

HP-UX Reference

Release 11i

User Commands

Section 1

Part 1 of 2 (A-M)

Volume 1 of 9

Edition 1

Customer Order Number: B2355-90688



Manufacturing Part Number: B2355-90689

E1200

Printed in: United States

© Copyright 1983-2000 Hewlett-Packard Company. All rights reserved.

Legal Notices

The information in this document is subject to change without notice.

Hewlett-Packard makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be held liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Warranty

A copy of the specific warranty terms applicable to your Hewlett-Packard product and replacement parts can be obtained from your local Sales and Service Office.

Restricted Rights Legend

Use, duplication or disclosure by the U.S. 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 for DOD agencies, and subparagraphs (c) (1) and (c) (2) of the Commercial Computer Software Restricted Rights clause at FAR 52.227-19 for other agencies.

HEWLETT-PACKARD COMPANY
3000 Hanover Street
Palo Alto, California 94304 U.S.A.

Use of this document and any supporting software media (CD-ROMs, flexible disks, and tape cartridges) supplied for this pack is restricted to this product only. Additional copies of the programs may be made for security and backup purposes only. Resale of the programs in their present form or with alterations is expressly prohibited.

Copyright Notices

Copyright © 1983-2000 Hewlett-Packard Company. All rights reserved.

Reproduction, adaptation, or translation of this document without prior written permission is prohibited, except as allowed under the copyright laws.

Copyright © 1979, 1980, 1983, 1985-93 Regents of the University of California. This software is based in part on the Fourth Berkeley Software Distribution under license from the Regents of the University of California.

Copyright © 1988 Carnegie Mellon University.

Copyright © 1990-1995 Cornell University.

Copyright © 1986 Digital Equipment Corporation.

Copyright © 1997 Isogon Corporation.

Copyright © 1985, 1986, 1988 Massachusetts Institute of Technology.

Copyright © 1991-1997 Mentat, Inc.

Copyright © 1996 Morning Star Technologies, Inc.

Copyright © 1990 Motorola, Inc.

Copyright © 1980, 1984, 1986 Novell, Inc.

Copyright © 1989-1993 The Open Software Foundation, Inc.

Copyright © 1996 Progressive Systems, Inc.

Copyright © 1989-1991 The University of Maryland

Copyright © 1986-1992 Sun Microsystems, Inc.

Trademark Notices

Apple® and Macintosh® are trademarks of Apple Computer, Inc., registered in the United States and other countries.

AppleShare® is a registered trademark of Apple Computer, Inc.

CHAMELEON™ is a trademark of NetManage, Inc.

DIGITAL™ and PATHWORKS™ are trademarks of Digital Equipment Corporation.

DiskAccess® is a registered trademark of Intergraph.

EXCURSION™ is a trademark of Digital Equipment Corporation.

Exeed® is a registered trademark of Hummingbird Communications Ltd.

eXodus™ is a trademark of White Pine Software, Inc.

MS-DOS® and Microsoft® are U.S. registered trademarks of Microsoft Corporation.

NTRIGUE™ is a trademark of Insignia Solutions, Inc.

OSF/Motif™ is a trademark of the Open Software Foundation, Inc. in the U.S. and other countries.

PC_Xware™ is a trademark, and WinCenter® is a registered trademark of Network Computing Devices, Inc.

REFLECTION® and WRQ® are registered trademarks of WRQ, Inc.

UNIX® is a registered trademark in the United States and other countries, licensed exclusively through The Open Group.

VERITAS® is a registered trademark of VERITAS Software Corporation.

VERITAS File System™ is a trademark of VERITAS Software Corporation.

WinDD™ is a trademark of Tektronix, Inc.

X Window System™ is a trademark of the Massachusetts Institute of Technology.

Publication History

The manual publication date and part number indicate its current edition. The printing date will change when a new edition is printed. Minor changes may be made at reprint without changing the printing date. The manual part number will change when extensive changes are made.

Manual updates may be issued between editions to correct errors or document product changes. To ensure that you receive the updated or new editions, you should subscribe to the appropriate product support service. See your HP sales representative for details.

First Edition: December 2000 (HP-UX Release 11i)

Volume One

Table of Contents

Introduction

Section 1

Volume One

Table of Contents

Introduction

Section 1

Table of Contents Volumes One and Two

Section 1: User Commands

| Entry Name(Section): name | Description |
|--|--|
| intro(1) | introduction to command utilities and application programs |
| adb(1): adb | absolute debugger |
| adjust(1): adjust | simple text formatter |
| admin(1): admin | create and administer SCCS files |
| alias: substitute command and/or file name | see sh-posix(1) |
| alias: substitute command and/or filename | see cs(1) |
| alias: substitute command and/or filename | see ksh(1) |
| alloc: show dynamic memory usage | see cs(1) |
| answer(1): answer | phone message transcription system |
| ar(1): ar | maintain portable archives and libraries |
| as(1): as | assembler (Precision Architecture) |
| asa(1): asa | interpret ASA carriage control characters |
| at(1): at, batch | execute batched commands immediately or at a later time |
| attributes(1): attributes | describe audio file |
| awk(1): awk | text pattern scanning and processing language |
| banner(1): banner | make posters in large letters |
| basename(1): basename, dirname | extract portions of path names |
| batch: execute batched commands immediately | see at(1) |
| bc(1): bc | arbitrary-precision arithmetic language |
| bdiff(1): bdiff | diff for large files |
| bfs(1): bfs | big file scanner |
| break: exit from enclosing for/next loop | see cs(1) |
| break: exit from enclosing for/next loop | see ksh(1) |
| break: exit from enclosing for/next loop | see sh-bourne(1) |
| break: exit from enclosing for/next loop | see sh-posix(1) |
| breaksw: break from switch and resume after endsw | see cs(1) |
| bs(1): bs | a compiler/interpreter for modest-sized programs |
| cal(1): cal | print calendar |
| calendar(1): calendar | reminder service |
| cancel: cancel requests on an LP printer or plotter | see lp(1) |
| case: label in a switch statement | see cs(1) |
| case: label in a switch statement | see ksh(1) |
| case: label in a switch statement | see sh-posix(1) |
| cat(1): cat | concatenate, copy, and print files |
| ccat: cat compacted files | see compact(1) |
| cc_bundled(1): cc_bundled | bundled C compiler |
| cd: change working directory | see cs(1) |
| cd: change working directory | see ksh(1) |
| cd: change working directory | see sh-bourne(1) |
| cd: change working directory | see sh-posix(1) |
| cd(1): cd | change working directory |
| cd(1): command | execute a simple command |
| cdc(1): cdc | change the delta commentary of an SCCS delta |
| chac(1): chac | add, modify, delete, copy, or summarize access control lists (ACLs) of files |
| chattr(1): chattr | change program's internal attributes |
| chdir: change current working directory | see cs(1) |
| checknr(1): checknr | check nroff/troff files |
| chfn(1): chfn | change user information in password file; used by finger |
| chgrp: change file group | see chown(1) |
| chkey(1): chkey | change user's secure RPC key |
| chmod(1): chmod | change file mode access permissions |
| chown(1): chown, chgrp | change file owner or group |
| chsh(1): chsh | change default login shell |
| ci(1): ci | check in RCS revisions |
| ckconfig(1): ckconfig | verify pathname of FTP configuration files |

Table of Contents

Volumes One and Two

| Entry Name(Section): name | Description |
|--|--|
| cksum(1): cksum | print file checksum and sizes |
| clear(1): clear | clear terminal screen |
| cmp(1): cmp | compare two files |
| co(1): co | check out RCS revisions |
| col(1): col | filter reverse line-feeds and backspaces |
| comb(1): comb | combine SCCS deltas |
| comm(1): comm | select or reject lines common to two sorted files |
| compact(1): compact, uncompact, ccat | compact and uncompact files, and cat them |
| compress(1): compress, compressdir, uncompress, uncompressdir, zcat | compress and expand data |
| compressdir: compress files in a directory | see compress(1) |
| continue: resume execution of nearest while or foreach | see cs(1) |
| continue: resume next iteration of enclosing for/next loop | see ksh(1) |
| continue: resume next iteration of enclosing for/next loop | see sh-bourne(1) |
| continue: resume next iteration of enclosing for/next loop | see sh-posix(1) |
| convert(1): convert | convert audio file |
| cp(1): cp | copy file, files, or directory subtree |
| cpio(1): cpio | copy file archives in and out |
| cpre(1): cpre | the C language preprocessor |
| crontab(1): crontab | user crontab file operations |
| crypt(1): crypt | encode/decode files |
| cs(1): cs | a shell (command interpreter) with C-like syntax |
| csplit(1): csplit | context split |
| ct(1): ct | spawn getty to a remote terminal (call terminal) |
| ctags(1): ctags | create a tags file |
| cu(1): cu | call another (UNIX) system; terminal emulator |
| cue(1): cue | HP Character-Terminal User Environment (CUE) |
| cut(1): cut | cut out (extract) selected fields of each line of a file |
| date(1): date | display or set the system-clock date and time |
| dc(1): dc | desk calculator |
| dd(1): dd | convert, reblock, translate, and copy a (tape) file |
| default: label default in switch statement | see cs(1) |
| delta(1): delta | make a delta (change) to an SCCS file |
| deroff(1): deroff | remove nroff, tbl, and neqn constructs |
| diff(1): diff, diffh | differential file comparator |
| diff3(1): diff3 | 3-way differential file comparison |
| diffh: differential file comparator | see diff(1) |
| diffmk(1): diffmk | mark differences between files |
| dircmp(1): dircmp | directory comparison |
| dirname: extract portions of path names | see basename(1) |
| dirs: print the directory stack | see cs(1) |
| disable: disable LP printers | see enable(1) |
| dmpxlt(1): dmpxlt | dump iconv translation tables to a readable format |
| domainname(1): domainname | set or display NIS domain name |
| dos2ux(1): dos2ux, ux2dos | convert ASCII file format |
| doschmod(1): doschmod | change attributes of a DOS file |
| doscp(1): doscp | copy to or from DOS files |
| dosdf(1): dosdf | report number of free disk clusters |
| dosll: list contents of DOS directories | see dosls(1) |
| dosls(1): dosls, dosll | list contents of DOS directories |
| dosmkdir(1): dosmkdir | make a DOS directory |
| dosrm(1): dosrm, dosrmdir | remove DOS files or directories |
| dosrmdir: remove DOS directories | see dosrm(1) |
| du(1): du | summarize disk usage |
| dumpmsg: extract messages from message catalog file | see findmsg(1) |
| echo: echo (print) arguments | see cs(1) |
| echo: echo (print) arguments | see ksh(1) |
| echo: echo (print) arguments | see sh-bourne(1) |
| echo: echo (print) arguments | see sh-posix(1) |
| echo(1): echo | echo (print) arguments |
| ed(1): ed, red | line-oriented text editor |
| edit: extended line-oriented text editor | see ex(1) |

Table of Contents Volumes One and Two

| Entry Name(Section): name | Description |
|---|--|
| egrep : search a file for a pattern | see grep(1) |
| elfdump(1) : elfdump | dump information contained in object files |
| elm(1) : elm | process electronic mail through a screen-oriented interface |
| elmalias(1) : elmalias | display/verify elm user and system aliases |
| enable(1) : enable, disable | enable/disable LP printers |
| end : terminate foreach or while loop | see cs(1) |
| endsw : terminate switch statement | see cs(1) |
| env(1) : env | set environment for command execution |
| eucset(1) : eucset | set and get EUC code widths for lterm |
| eval : read arguments as shell input and execute resulting | see cs(1) |
| eval : read arguments as shell input and execute resulting commands | see ksh(1) |
| eval : read arguments as shell input and execute resulting commands | see sh-bourne(1) |
| eval : read arguments as shell input and execute resulting commands | see sh-posix(1) |
| ex(1) : edit, ex | extended line-oriented text editor |
| exec : execute command without creating new process | see cs(1) |
| exec : execute command without creating new process | see ksh(1) |
| exec : execute command without creating new process | see sh-bourne(1) |
| exec : execute command without creating new process | see sh-posix(1) |
| exit : exit shell with exit status | see cs(1) |
| exit : exit shell with exit status | see ksh(1) |
| exit : exit shell with exit status | see sh-bourne(1) |
| exit : exit shell with exit status | see sh-posix(1) |
| expand(1) : expand, unexpand | expand tabs to spaces, and vice versa |
| expand_alias(1) : expand_alias | recursively expands the sendmail aliases |
| export : export variable names to environment of subsequent commands | see ksh(1) |
| export : export variable names to environment of subsequent commands | see sh-bourne(1) |
| export : export variable names to environment of subsequent commands | see sh-posix(1) |
| expr(1) : expr | evaluate arguments as an expression |
| factor(1) : factor, primes | factor a number, generate large primes |
| false : do nothing and return non-zero exit status | see true(1) |
| fastbind(1) : fastbind | prepare an incomplete executable for faster program start-up |
| fastmail(1) : fastmail | quick batch mail interface |
| fc : edit and execute previous command | see ksh(1) |
| fc : edit and execute previous command | see sh-posix(1) |
| fgrep : search a file for a string (fast) | see grep(1) |
| file(1) : file | determine file type |
| find(1) : find | find files |
| findmsg(1) : findmsg, dumpmsg | create message catalog file for modification |
| findstr(1) : findstr | find strings for inclusion in message catalogs |
| finger(1) : finger | user information lookup program |
| fmt(1) : fmt | format text |
| fold(1) : fold | fold long lines for finite width output device |
| for : execute a do list | see ksh(1) |
| for : execute a do list | see sh-posix(1) |
| forder(1) : forder | convert file data order |
| foreach : initiate repetitive loop | see cs(1) |
| from(1) : from | who is my mail from? |
| fruled(1) : fruled | turn on/off attention LEDs |
| ftio(1) : ftio | faster tape I/O |
| ftp(1) : ftp | file transfer program |
| ftpaccess(1) : ftpaccess | create shutdown message files to shutdown ftp servers |
| ftpcount(1) : ftpcount | show current number of users for each ftp class |
| ftprestart(1) : ftprestart | remove the shutdown message files created by ftpshut |
| ftpwho(1) : ftpwho | show process information on each ftp user |
| gencat(1) : gencat | generate a formatted message catalog file |
| genxlt(1) : genxlt | generate iconv translation |
| get(1) : get | get a version of an SCCS file |
| getaccess(1) : getaccess | list access rights to file(s) |
| getacl(1) : getacl | list access control lists for files, JFS only |
| getconf(1) : getconf | get POSIX configuration values |
| getopt(1) : getopt | parse command options |

Table of Contents

Volumes One and Two

| Entry Name(Section): name | Description |
|---|--|
| getopts(1): getopts | parse utility (command) options |
| getprivgrp(1): getprivgrp | get special attributes for group |
| glob: echo without '\\' escapes | see cs(1) |
| goto: continue execution on specified line | see cs(1) |
| gprof(1): gprof | display call graph profile data |
| grep(1): grep, egrep, fgrep | search a file for a pattern |
| grget: get password and group information | see pwget(1) |
| groups(1): groups | show group memberships |
| hash: remember command location in search path | see sh-bourne(1) |
| hashcheck: create hash codes from compressed spelling list | see spell(1) |
| hashmake: convert words to 9-digit hashcodes | see spell(1) |
| hashstat: print hash table effectiveness statistics | see cs(1) |
| head(1): head | print first few lines in a file |
| history: display event history list | see cs(1) |
| hostname(1): hostname | set or display name of current host system |
| hp(1): hp | handle special functions of HP 2640 and HP 2621-series terminals |
| hp-mc680x0: provide truth value about processor type | see machid(1) |
| hp-pa: provide truth value about processor type | see machid(1) |
| hp9000s200: provide truth value about processor type | see machid(1) |
| hp9000s300: provide truth value about processor type | see machid(1) |
| hp9000s500: provide truth value about processor type | see machid(1) |
| hp9000s800: provide truth value about processor type | see machid(1) |
| hyphen(1): hyphen | find hyphenated words |
| iconv(1): iconv | character code set conversion |
| id(1): id | print user and group IDs and names |
| ident(1): ident | identify files in RCS |
| idlookup(1): idlookup | identify the user of a particular TCP connection |
| ied(1): ied | input editor and command history for interactive programs |
| if: execute command if expression evaluates true | see cs(1) |
| if: execute command if previous command returns exit status 0 | see ksh(1) |
| if: execute command if previous command returns exit status 0 | see sh-posix(1) |
| insertmsg(1): insertmsg | use findstr(1) output to insert calls to catgets(3C) |
| inv: make unprintable and non-ASCII characters in a file invisible | see vis(1) |
| iostat(1): iostat | report I/O statistics |
| ipcrm(1): ipcrm | remove a message queue, semaphore set or shared memory id |
| ipcs(1): ipcs | report status of interprocess communication facilities |
| jobs: list active jobs | see cs(1) |
| jobs: list active jobs | see ksh(1) |
| jobs: list active jobs | see sh-posix(1) |
| join(1): join | relational database operator |
| kdestroy(1): kdestroy | destroy Kerberos tickets |
| kermi(1): kermi | communication software for serial and network connections |
| keylogin(1): keylogin | decrypt and store secret key |
| keylogout(1): keylogout | delete secret key stored with keyserv |
| keysh(1): keysh | context-sensitive softkey shell |
| kill: send termination or specified signal to a process | see cs(1) |
| kill: terminate job or process | see ksh(1) |
| kill: terminate job or process | see sh-posix(1) |
| kill(1): kill | send a signal to a process; terminate a process |
| kinit(1): kinit | obtain and cache Kerberos ticket-granting ticket |
| klist(1): klist | list cached Kerberos tickets |
| kpasswd(1): kpasswd | change a user's Kerberos password |
| ksh(1): ksh, rksh | shell, the standard/restricted command programming language |
| ktutil(1): ktutil | Kerberos keytab file maintenance utility |
| kvno(1): kvno | print key version numbers of Kerberos principals |
| l: list contents of directories | see ls(1) |
| last(1): last, lastb | indicate last logins of users and ttys |
| lastb: indicate last bad logins of users and ttys | see last(1) |
| lastcomm(1): lastcomm | show last commands executed in reverse order |
| lc: list contents of directories | see ls(1) |
| ld(1): ld | link editor |

Table of Contents Volumes One and Two

| Entry Name(Section): name | Description |
|--|---|
| ldd(1): ldd | list dynamic dependencies of executable files or shared libraries |
| leave(1): leave | remind you when you have to leave |
| let: evaluate arithmetic expression | see ksh(1) |
| let: evaluate arithmetic expression | see sh-posix(1) |
| liblu62.a: IBM APPC (Advanced Program-to-Program Communications) API | see sna(1) |
| lifcp(1): lifcp | copy to or from LIF files |
| lifinit(1): lifinit | write LIF volume header on file |
| lifs(1): lifs | list contents of a LIF directory |
| lifrename(1): lifrename | rename LIF files |
| lifrm(1): lifrm | remove a LIF file |
| limit: limit usage by current process | see csh(1) |
| line(1): line | read one line from user input |
| listusers(1M): listusers | display user login data |
| ll: list contents of directories | see ls(1) |
| ln(1): ln | link files and directories |
| locale(1): locale | get locale-specific (NLS) information |
| lock(1): lock | reserve a terminal |
| logger(1): logger | make entries in the system log |
| login: terminate login shell | see csh(1) |
| login(1): login | sign on; start terminal session |
| logname(1): logname | get login name |
| logout: terminate login shell | see csh(1) |
| lorder(1): lorder | find ordering relation for an object library |
| lp(1): cancel, lp, lpalt | print/alter/cancel requests on an LP printer or plotter |
| lpalt: alter requests on an LP printer or plotter | see lp(1) |
| lpfilter(1): lpfilter | filters used by the lp interface scripts |
| lpstat(1): lpstat | print LP status information |
| ls(1): ls, l, lc, ll, lsf, lsr, lsx | list contents of directories |
| lsacl(1): lsacl | list access control lists (ACLs) of files |
| lsf: list contents of directories | see ls(1) |
| lsr: list contents of directories | see ls(1) |
| lsx: list contents of directories | see ls(1) |
| m4(1): m4 | macro processor |
| machid(1): hp9000s200, hp9000s300, hp9000s500, hp9000s800, pdp11, u3b, u3b5, vax | provide truth value about processor type |
| mail(1): mail, rmail | send mail to users or read mail |
| mailfrom(1): mailfrom | summarize mail folders by subject and sender |
| mailq(1): mailq | prints the mail queue |
| mailstats(1): mailstats | print mail traffic statistics |
| mailx(1): mailx | interactive message processing system |
| make(1): make | maintain, update, and regenerate groups of programs |
| makekey(1): makekey | generate encryption key |
| man(1): man | find manual information by keywords; print out a manual entry |
| mediainit(1): mediainit | initialize disk or cartridge tape media, partition DDS tape |
| merge(1): merge | three-way file merge |
| mesg(1): mesg | permit or deny messages to terminal |
| mkdir(1): mkdir | make a directory |
| mkfifo(1): mkfifo | make FIFO (named pipe) special files |
| mkmf(1): mkmf | make a makefile |
| mkmsgs(1): mkmsgs | create message files for use by gettxt() |
| mkstr(1): mkstr | extract error messages from C source into a file |
| mktemp(1): mktemp | make a name for a temporary file |
| mkupath: manage the pathalias database | see uupath(1) |
| mm(1): mm, osdd | print/check documents formatted with the mm macros |
| model(1): model | print name of current HP-UX version |
| more(1): more, page | file perusal filter for crt viewing |
| mpsched(1): mpsched | control processor or locality domain on which a specific process executes |
| mt(1): mt | magnetic tape manipulating program |
| mv(1): mv | move or rename files and directories |
| neqn(1): neqn | format mathematical text for nroff |
| netstat(1): netstat | show network status |

Table of Contents

Volumes One and Two

| Entry Name(Section): name | Description |
|---|--|
| newalias(1): newalias | install new elm aliases for user or system |
| newform(1): newform | change or reformat a text file |
| newgrp: equivalent to exec newgrp | see sh-bourne(1) |
| newgrp: equivalent to exec newgrp | see cs(1) |
| newgrp: equivalent to exec newgrp | see ksh(1) |
| newgrp: equivalent to exec newgrp | see sh-posix(1) |
| newgrp(1): newgrp | switch to a new group |
| newmail(1): newmail | notify users of new mail in mailboxes |
| news(1): news | print news items |
| nice: alter command priority | see cs(1) |
| nice(1): nice | run a command at nondefault priority |
| nis+(1): nis+ | new version of the network information name service. |
| niscat(1): niscat | display NIS+ tables and objects |
| nischgrp(1): nischgrp | change the group owner of an NIS+ object |
| nischmod(1): nischmod | change access rights on an NIS+ object |
| nischown(1): nischown | change the owner of an NIS+ object |
| nischttl(1): nischttl | change the time to live value of an NIS+ object |
| nisdefaults(1): nisdefaults | display NIS+ default values |
| niserror(1): niserror | display NIS+ error messages |
| nisgrep(1): nisgrep | utility for searching NIS+ tables |
| nisgrpadm(1): nisgrpadm | administer NIS+ groups |
| nisln(1): nisln | symbolically link NIS+ objects |
| nisls(1): nisls | list the contents of an NIS+ directory |
| nismatch(1): nismatch | utility for searching NIS+ tables |
| nismkdir(1): nismkdir | create NIS+ directories |
| nispasswd(1): nispasswd | change NIS+ password information |
| nisrm(1): nisrm | remove NIS+ objects from the namespace |
| nisrmdir(1): nisrmdir | remove NIS+ directories |
| nistbladm(1): nistbladm | administer NIS+ tables |
| nistest(1): nistest | return the state of the NIS+ namespace using a conditional expression |
| nl(1): nl | line numbering filter |
| nljust(1): nljust | justify lines, left or right, for printing |
| nm(1): nm | print name list of common object file |
| nohup: ignore hangups during command execution | see cs(1) |
| nohup(1): nohup | run a command immune to hangups |
| notify: notify user of change in job status | see cs(1) |
| nroff(1): nroff | format text |
| nslookup(1): nslookup | query name servers interactively |
| nsquery(1): nsquery | query the Name Service Switch backend libraries |
| od(1): od, xd | octal and hexadecimal dump |
| odump(1): odump | dump information contained in SOM object files |
| on(1): on | execute a command on a remote host; environment similar to local environment |
| onintr: specify shell's treatment of interrupts | see cs(1) |
| osdd: print/check documents formatted with the mm macros | see mm(1) |
| pack(1): pack, pcat, unpack | compress and expand files |
| page: file perusal filter for crt viewing | see more(1) |
| parstatus(1): parstatus | display information about the Superdome complex |
| partition(1): partition | display information about Superdome commands |
| passwd(1): passwd | change login password and associated attributes |
| paste(1): paste | merge same lines of several files or subsequent lines of one file |
| patch(1): patch | applying a diff file to an original |
| pathalias(1): pathalias | electronic address router |
| pathchk(1): pathchk | check path names |
| pax(1): pax | portable archive exchange |
| pcat: compress and expand files | see pack(1) |
| pdclean(1): pdclean | remove jobs from specified object |
| pdcreate(1): pdcreate | creates print objects |
| pdddelete(1): pdddelete | deletes print objects |
| pddisable(1): pddisable | stops printers from accepting jobs and logs from logging |
| pdenable(1): pdenable | enables printers to accept jobs and logs to log |
| pdls(1): pdls | lists selected attribute values |

Table of Contents

Volumes One and Two

| Entry Name(Section): name | Description |
|---|--|
| pdmod(1): pdmod | modifies attributes of submitted print jobs |
| pdmsg(1): pdmsg | displays the text and description of a HPDPS message at the command line |
| pdpl1: provide truth value about processor type | see machid(1) |
| pdpause(1): pdpause | pauses jobs, physical printers, servers, or queues |
| pdpr(1): pdpr | submits print jobs |
| pdpromote(1): pdpromote | advances a job to the top of a queue |
| pdq(1): pdq | queries and lists the status of one or more print jobs |
| pdresubmit(1): pdresubmit | resubmits previously submitted print jobs |
| pdresume(1): pdresume | enables paused objects to resume operation |
| pdrm(1): pdrm | removes print jobs |
| pdset(1): pdset | modifies attributes of submitted print jobs |
| pdshutdown(1): pdshutdown | stop servers |
| pg(1): pg | file perusal filter for soft-copy terminals |
| popd: pop directory stack | see cs(1) |
| pppd(1): pppd | point to point protocol daemon |
| pr(1): pr | format and print files |
| praliases(1): praliases | print system-wide sendmail aliases |
| prealloc(1): prealloc | preallocate disk storage |
| primes: generate large prime numbers | see factor(1) |
| print: output from shell | see ksh(1) |
| print: output from shell | see sh-posix(1) |
| printenv(1): printenv | print out the environment |
| printf(1): echo | print formatted arguments |
| prmail(1): prmail | print out mail in the incoming mailbox file |
| prof(1): prof | display profile data |
| prs(1): prs | print and summarize an SCCS file |
| ps(1): ps | report process status |
| ptx(1): ptx | permuted index |
| pty: get the name of the pseudo-terminal | see tty(1) |
| pushd: push directory stack | see cs(1) |
| pwd: print current working directory | see ksh(1) |
| pwd: print current working directory | see sh-posix(1) |
| pwd: working directory name | see sh-bourne(1) |
| pwd(1): pwd | working directory name |
| pwget(1): pwget, grget | get password and group information |
| quota(1): quota | display disk usage and limits |
| ranlib(1): ranlib | regenerate archive symbol table |
| rcp(1): rcp | remote file copy |
| rcp(1): rcp | remote file copy |
| rcs(1): rcs | change RCS file attributes |
| rcsdiff(1): rcsdiff | compare RCS revisions |
| rcsmerge(1): rcsmerge | merge RCS revisions |
| rdist(1): rdist | remote file distribution |
| read: input and parse a line | see ksh(1) |
| read: input and parse a line | see sh-posix(1) |
| read: read line from standard input | see sh-bourne(1) |
| read(1): read | read a line from standard input |
| readmail(1): readmail | read mail from a mail folder or incoming mailbox |
| readonly: mark <i>name</i> as read-only | see sh-bourne(1) |
| readonly: mark names as unreddefinable | see ksh(1) |
| readonly: mark names as unreddefinable | see sh-posix(1) |
| red: restricted line-oriented text editor | see ed(1) |
| rehash: recompute internal hash table | see cs(1) |
| remsh(1): remsh | execute from a remote shell |
| repeat: execute command more than once | see cs(1) |
| return: exit function with return value | see sh-bourne(1) |
| return: shell function return to invoking script | see ksh(1) |
| return: shell function return to invoking script | see sh-posix(1) |
| rev(1): rev | reverse lines of a file |
| rksh: restricted Korn shell command programming language | see ksh(1) |
| rlog(1): rlog | print log messages and other information on RCS files |

Table of Contents

Volumes One and Two

| Entry Name(Section): name | Description |
|--|--|
| rlogin(1): rlogin | remote login |
| rm(1): rm | remove files or directories |
| rmail: send mail to users or read mail | see mail(1) |
| rmidel(1): rmidel | remove a delta from an SCCS file |
| rmdir(1): rmdir | remove directories |
| rmnl(1): rmnl | remove extra new-line characters from file |
| rpcgen(1): rpcgen | an RPC protocol compiler |
| rsh: restricted shell command programming language | see sh-bourne(1) |
| rtprio(1): rtprio | execute process with real-time priority |
| rtsched(1): rtsched | execute process with POSIX real-time priority |
| rup(1): rup | show host status of local machines (RPC version) |
| ruptime(1): ruptime | show status of local machines |
| rusers(1): rusers | determine who is logged onto machines on the local network |
| rwho(1): rwho | show who is logged in on local machines |
| sact(1): sact | print current SCCS file editing activity |
| samlog_viewer(1): samlog_viewer | tool for viewing and saving the SAM logfile |
| sccs(1): sccs | utility program for SCCS commands |
| sccsdiff(1): sccsdiff | compare two versions of an SCCS file |
| sccshelp(1): sccshelp | help for SCCS commands |
| script(1): script | make typescript of terminal session |
| sdiff(1): sdiff | side-by-side difference program |
| sed(1): sed | stream text editor |
| send_sound(1): send_sound | play audio file |
| serialize(1): serialize | force target process to run serially with other processes |
| set: set/define flags and arguments | see cs(1) |
| set: set/define options and arguments | see ksh(1) |
| set: set/define options and arguments | see sh-bourne(1) |
| set: set/define options and arguments | see sh-posix(1) |
| setacl(1): setacl | modify access control lists for files (JFS only) |
| setenv: define environment variable | see cs(1) |
| sh(1): sh | overview of various system shells |
| sh-bourne(1): sh, rsh | shell, the standard/restricted command programming language |
| sh-posix(1): sh | shell, the standard/restricted command programming language |
| shar(1): shar | make a shell archive package |
| shift: shift <i>argv</i> members one position to left | see cs(1) |
| shift: shift <i>argv</i> members one position to left | see ksh(1) |
| shift: shift <i>argv</i> members one position to left | see sh-posix(1) |
| shift: shift positional parameters to next lower position | see sh-bourne(1) |
| shl(1): shl | shell layer manager |
| size(1): size | print section sizes of object files |
| sleep(1): sleep | suspend execution for an interval |
| slp(1): slp | set printing options for a non-serial printer |
| sna(1): sna3179g, sna3270, sna3770 | IBM 3179G/3192G, 3270, 3777 terminal emulator |
| soelim(1): soelim | eliminate .so's from nroff input |
| sort(1): sort | sort or merge files |
| source: define source for command input | see cs(1) |
| spell(1): spell, hashmake, spellin, hashcheck | find spelling errors |
| spellin: create compressed spelling list from hash codes | see spell(1) |
| split(1): split | split a file into pieces |
| ssp(1): ssp | remove multiple line-feeds from output |
| strings(1): strings | find the printable strings in an object or other binary file |
| strip(1): strip | strip symbol and line number information from an object file |
| stty(1): stty | set the options for a terminal port |
| su(1): su | switch user |
| sum(1): sum | print checksum and block count of a file |
| switch: define switch statement | see cs(1) |
| tabs(1): tabs | set tabs on a terminal |
| tail(1): tail | deliver the last part of a file |
| talk(1): talk | talk to another user |
| tar(1): tar | tape file archiver |
| tbl(1): tbl | format tables for nroff |

Table of Contents Volumes One and Two

| Entry Name(Section): name | Description |
|---|--|
| tee(1): tee | pipe fitting |
| telnet(1): telnet | user interface to the TELNET protocol |
| test: evaluate conditional expression | see cs(1) |
| test: evaluate conditional expression | see ksh(1) |
| test: evaluate conditional expression | see sh-posix(1) |
| test(1): test | condition evaluation command |
| tftp(1): tftp | trivial file transfer program |
| time, times: print summary of time used by processes | see ksh(1) |
| time: print accumulated shell and children process times | see sh-bourne(1) |
| time: print summary of time used by shell and children | see cs(1) |
| time(1): time | time a command |
| time, times: print summary of time used by processes | see sh-posix(1) |
| times: print accumulated user and system process times | see sh-bourne(1) |
| timex(1): timex | time a command; report process data and system activity |
| top(1): top | display and update information about top processes on system |
| touch(1): touch | update access, modification, and/or change times of file |
| tput(1): tput | query terminfo database |
| tr(1): tr | translate characters |
| trap: execute command upon receipt of signal | see sh-bourne(1) |
| trap: trap specified signal | see ksh(1) |
| trap: trap specified signal | see sh-posix(1) |
| true(1): true, false | return zero or non-zero exit status |
| tset(1): tset | terminal-dependent initialization |
| tsm(1): tsm | Terminal Session Manager |
| tsm.command(1): tsm.command | send commands to Terminal Session Manager |
| tsm.info(1): tsm.info | get Terminal Session Manager state information |
| tsort(1): tsort | topological sort |
| tty(1): tty, pty | get the name of the terminal |
| ttytype(1): ttytype | terminal identification program |
| type: show interpretation of <i>name</i> as if a command | see sh-bourne(1) |
| typeset: control leading blanks and parameter handling | see ksh(1) |
| typeset: control leading blanks and parameter handling | see sh-posix(1) |
| u370: provide truth value about processor type | see machid(1) |
| u3b: provide truth value about processor type | see machid(1) |
| u3b10: provide truth value about processor type | see machid(1) |
| u3b2: provide truth value about processor type | see machid(1) |
| u3b5: provide truth value about processor type | see machid(1) |
| ul(1): ul | do underlining |
| ulimit: impose file size limit for child processes | see sh-bourne(1) |
| ulimit: set size or time limits | see ksh(1) |
| ulimit: set size or time limits | see sh-posix(1) |
| umask: set permissions mask for creating new files | see cs(1) |
| umask: set permissions mask for creating new files | see ksh(1) |
| umask: set permissions mask for creating new files | see sh-bourne(1) |
| umask: set permissions mask for creating new files | see sh-posix(1) |
| umask(1): umask | set file-creation mode mask |
| umodem(1): umodem | XMODEM-protocol file transfer program |
| unalias: discard specified alias | see cs(1) |
| unalias: discard specified alias | see ksh(1) |
| unalias: discard specified alias | see sh-posix(1) |
| uname(1): uname | display information about computer system; set node name (system name) |
| uncompact: uncompact files | see compact(1) |
| uncompress: expand compressed data | see compress(1) |
| uncompressdir: expand compressed files in a directory | see compress(1) |
| unexpand: convert spaces to tabs | see expand(1) |
| unget(1): unget | undo a previous get of an SCCS file |
| unhash: disable use of internal hash tables | see cs(1) |
| unifdef(1): unifdef | remove preprocessor lines |
| uniq(1): uniq | report repeated lines in a file |
| units(1): units | conversion program |
| unpack: compress and expand files | see pack(1) |

Table of Contents

Volumes One and Two

| Entry Name(Section): name | Description |
|---|--|
| unset: remove definition/setting of flags and arguments | see cs(1) |
| unset: remove definition/setting of options and arguments | see ksh(1) |
| unset: remove definition/setting of options and arguments | see sh-bourne(1) |
| unset: remove definition/setting of options and arguments | see sh-posix(1) |
| unsetenv: remove variable from environment | see cs(1) |
| until: execute commands until expression is non-zero | see ksh(1) |
| until: execute commands until expression is nonzero | see sh-posix(1) |
| uptime(1): uptime, w | show how long system has been up |
| users(1): users | compact list of users who are on the system |
| uucp(1): uucp, uulog, uuname, uutry | UNIX system to UNIX system copy |
| uudecode: decode a file encoded by uuencode | see uuencode(1) |
| uuencode(1): uuencode, uudecode | encode/decode a binary file for transmission by mailer |
| uulog: access UUCP summary logs | see uucp(1) |
| uuname: list known UUCP systems | see uucp(1) |
| uupath(1): uupath, mkuupath | access and manage the pathalias database |
| uupick: accept or reject incoming UUCP messages | see uuto(1) |
| uustat(1): uustat | uucp status inquiry and job control |
| uuto(1): uuto, uupick | public UNIX system to UNIX system file copy |
| uutry: test for successful login to remote system | see uucp(1) |
| uux(1): uux | UNIX system to UNIX system command execution |
| ux2dos: convert ASCII file format | see dos2ux(1) |
| vacation(1): vacation | return "I am not here" indication |
| val(1): val | validate SCCS file |
| vax: provide truth value about processor type | see machid(1) |
| vc(1): vc | version control |
| vedit: beginner's screen-oriented text editor | see vi(1) |
| vi(1): vedit, vi, view | extended screen-oriented text editor |
| view: read-only screen-oriented text editor | see vi(1) |
| vis(1): vis, inv | make unprintable and non-ASCII characters in a file visible or invisible |
| vmstat(1): vmstat | report virtual memory statistics |
| vt(1): vt | log in on another system over lan |
| wait: wait for background processes | see cs(1) |
| wait: wait for child process | see ksh(1) |
| wait: wait for child process | see sh-posix(1) |
| wait: wait for process and report termination status | see sh-bourne(1) |
| wait(1): wait | await completion of process |
| wc(1): wc | count words, lines, and bytes or characters in a file |
| what(1): what | get SCCS identification information |
| whence: define interpretation of name as a command | see ksh(1) |
| whence: define interpretation of name as a command | see sh-posix(1) |
| whereis(1): whereis | locate source, binary, and/or manual for program |
| which(1): which | locate a program file including aliases and paths |
| while: execute commands while expression is non-zero | see cs(1) |
| while: execute commands while expression is non-zero | see ksh(1) |
| while: execute commands while expression is nonzero | see sh-posix(1) |
| who(1): who | who is currently logged in on the system |
| whoami(1): whoami | print effective current user id |
| whois(1): whois | Internet user name directory service |
| write(1): write | interactively write (talk) to another user |
| xargs(1): xargs | construct argument list(s) and execute command |
| xd: hexadecimal dump | see od(1) |
| xstr(1): xstr | extract strings from C programs to implement shared strings |
| yes(1): yes | be repetitively affirmative |
| ypcat(1): ypcat | print all Network Information Service map values |
| ypmatch(1): ypmatch | print values of selected keys in Network Information Service map |
| yppasswd(1): yppasswd | change login password in Network Information System (NIS) |
| ypwhich(1): ypwhich | list Network Information System server or map master |
| zcat: expand and cat data | see compress(1) |

Introduction to HP-UX Reference

Introduction to HP-UX Reference

NAME

Introduction - an introduction to the HP-UX operating system and the HP-UX Reference

INTRODUCTION

HP-UX is the Hewlett-Packard Company's implementation of an operating system that is compatible with various industry standards. It is based on the UNIX® System V Release 4 operating system and includes important features from the Fourth Berkeley Software Distribution.

Improvements include enhanced capabilities and other features, developed by HP to make HP-UX a very powerful, useful, and reliable operating system, capable of supporting a wide range of applications ranging from simple text processing to sophisticated engineering graphics and design. It can readily be used to control instruments and other peripheral devices. Real-time capabilities further expand the flexibility of HP-UX as a powerful tool for solving tough problems in design, manufacturing, business, and other areas where responsiveness and performance are important.

Extensive international language support enables HP-UX to interact with users in any of dozens of human languages. HP-UX interfaces easily with local area networks and resource-sharing facilities. By using industry-standard protocols, HP-UX provides flexible interaction with other computers and operating systems. Optional software products extend HP-UX capabilities into a broad range of specialized needs.

The *HP-UX Reference* is not a learning tool for beginners. It is primarily a reference tool that is most useful for experienced users of UNIX or UNIX-like systems. If you are not already familiar with UNIX or HP-UX, refer to the series of Beginner's Guides, tutorial manuals, and other learning documents supplied with your system or available separately. System implementation and maintenance details are explained in the *Managing Systems and Workgroups* manual.

MANPAGE ORGANIZATION

The contents of the *HP-UX Reference* and its on-line counterpart are a number of independent entries called **manpages**. These are also called **manual entries** or **reference pages**.

For convenient reference, the manpages are divided into eight specialized sections. The printed manual also has a table of contents for each volume and a composite index.

Each manpage consists of one or more printed pages, with the manpage name and section number printed in the upper corners. Manpages are arranged alphabetically within each section of the reference, except for the *intro* page at the beginning of each section. Manpages are referred to by name and section number, in the form *pagename(section)*.

The manpages are available on-line through the **man** command if the manpages are present on the system. Refer to the *man(1)* manpage in Section 1 for more information.

Each page in the printed manual has two page numbers, printed at the bottom of the page. The center page number starts over with page 1 at the beginning of each new manpage; it is placed between two dashes in normal typeface. The number printed at the outside corner on each page sequences the printed pages within a section. Users usually locate manpages by the alphabetic headings at the top of the page as when reading a dictionary.

Some manpages describe two or more commands or routines. In such cases, the manpage is usually named for the first command or function that appears in the NAME section. Occasionally, a manpage name appears as a prefix to the NAME section. In such instances, the name describes the commands or functions in more general terms. For example, the *acct(1M)* manpage describes the **acctdisk**, **acctdusg**, **accton**, and **acctwtm** commands, while the *string(3C)* manpage describes many character string functions.

The various sections are described as follows:

Volume Table of Contents (Printed Manual)

A complete listing of all manpages in the order they appear in each section, as well as alphabetically intermixed lists of all command, function, and feature names that are the different from the manpage where they appear

Section 1: User Commands

Programs that are usually invoked directly by users or from command language procedures (scripts).

Section 1M: System Administration Commands

Commands used for system installation and maintenance, including boot processes, crash recovery, system integrity testing, and other needs. Most commands in this section require the superuser privilege.

Section 2: System Calls

Entries into the HP-UX kernel, including the C-language interface. These topics are primarily of interest to programmers.

Section 3: Library Functions

Available subroutines that reside (in binary form) in various system libraries. These topics are primarily of interest to programmers.

Section 4: File Formats

The structure of various types of files, primarily of interest to administrators and programmers. For example, the link editor output file format is described in *a.out*(4). Files that are used only by a single command (such as intermediate files used by assemblers) are not described. C-language declarations corresponding to the formats in Section 4 can be found in the directories `/usr/include` and `/usr/include/sys`.

Section 5: Miscellaneous

A variety of information, such as descriptions of header files, character sets, macro packages, and other topics.

Section 7: Device Special Files

The characteristics of special (device) files that provide the link between HP-UX and system I/O devices. The names for each topic usually refer to the type of I/O device rather than to the names of individual special files.

Section 9: Introduction and Glossary

A general introduction (this one) and definitions of terms used in the HP-UX environment.

Composite Index (Printed Manual)

An alphabetical listing of keywords and topics based on the NAME section near the beginning of each manpage as well as other information, cross-referenced to manpage names and sections. The index also contains references to built-in features in the various command interpreters ("shells").

MANPAGE FORMATS

All manpages follow an established section heading format, but not all section headings are included in each manpage. A few manpages have self-explanatory specialized headings.

NAME

Gives the names of the commands, functions, or features and briefly states the purpose.

SYNOPSIS

Summarizes the syntax of the command or program entity. A few conventions are used:

Constant-width characters indicate literal characters that should be entered exactly as they appear. These characters appear in bold in the online manpages.

Italic strings represent variable elements that should be replaced with appropriate values.

Roman square brackets ([]) indicate that the contents are optional.

Roman braces ({}) indicate a required element, usually in a choice.

Ellipses (...) indicate that the previous element can be repeated.

Note: An argument beginning with a dash (-), a plus sign (+), or an equal sign (=) is often defined as a command option, even if it appears in a position where a file name could appear. Therefore, it is unwise to have file names that begin with -, +, or =.

DESCRIPTION

Discusses the function and behavior of each entry.

EXTERNAL INFLUENCES

Information under this heading pertains to programming for various spoken languages. Typical entries indicate support for single- or multibyte characters, the effect of language-related environment variables on system behavior, and other related information.

NETWORKING FEATURES

Information under this heading is applicable only if you are using the network feature described there

(such as NFS).

RETURN VALUE

Describes the values returned by function calls or in the return code by commands.

DIAGNOSTICS

Describes diagnostic information that may be produced. Self-explanatory messages are not listed.

ERRORS

Lists function error conditions (set in **errno**) and their corresponding error messages.

EXAMPLES

Provides examples of typical usage.

WARNINGS

Describes potential problems and deficiencies.

DEPENDENCIES

Describes variations in HP-UX operation that are related to the use of specific hardware or combinations of hardware.

AUTHOR

Indicates the origin of the software documented by the manpage. Unless noted otherwise, the source of an entry is System V.

FILES

Lists file names that are used or affected by the program or command.

SEE ALSO

Provides pointers to related manpages and other documentation.

STANDARDS CONFORMANCE

For each command or subroutine entry point addressed by one or more of the following industry standards, this section lists the standard specifications to which that HP-UX component conforms.

The various standards are:

| | |
|------------|--|
| AES | OSF Application Environment Specification |
| ANSI C | ANSI X3.159-1989 |
| POSIX.1 | IEEE Standard 1003.1-1988 (IEEE Computer Society) (Portable Operating System Interface for Computer Environments) |
| POSIX.2 | IEEE Standard 1003.2-1990 (IEEE Computer Society) (Portable Operating System Interface for Computer Environments) |
| POSIX.4 | IEEE Standard 1003.1b-1993 (IEEE Computer Society) (Portable Operating System Interface for Computer Environments) |
| FIPS 151-1 | Federal Information Processing Standard 151-1 (National Institute of Standards and Technology) |
| FIPS 151-2 | Federal Information Processing Standard 151-2 (National Institute of Standards and Technology) |
| SVID2 | System V Interface Definition Issue 2 |
| SVID3 | System V Interface Definition Issue 3 |
| XPG2 | X/Open Portability Guide Issue 2 (X/Open, Ltd.) |
| XPG3 | X/Open Portability Guide Issue 3 (X/Open, Ltd.) |
| XPG4 | X/Open Portability Guide Issue 4 (X/Open, Ltd.) |
| XPG4.2 | X/Open Portability Guide Issue 4 (X/Open, Ltd.) Version 2 |

GETTING STARTED WITH HP-UX

This is a very brief overview of how to use the HP-UX system: how to log in and log out, how to communicate through your machine, and how to run a program.

HP-UX uses **control characters** to perform certain functions. Control characters are generally shown in the form **^x**, such as **^D** for Control-D. Hold down the **Control (Ctrl)** key while you press the character key.

Logging In

To log in you must have a valid user name and password, which can be obtained from your system administrator.

When a connection has been established, the system displays **login:** on your terminal. Type your user name and press the **Return** key. Enter your password (it is not echoed by the system) and press **Return**.

A list of copyright notices and a message-of-the-day may greet you before the first prompt.

It is important that you type your login name with lowercase letters, if possible. If you type uppercase letters, HP-UX assumes that your terminal cannot generate lowercase letters, and treats subsequent uppercase input as lowercase.

When you log in successfully, the system starts your login shell. The default is the POSIX shell, **/usr/bin/sh**. The POSIX shell (and its predecessors, the Korn and Bourne shells) use **\$** as the default prompt. The C shell uses **%**.

See *login(1)* for more on login, *passwd(1)* to change your password, *chsh(1)* to change your login shell.

Logging Out

You can log out of the shells by typing an **exit** command or the **eof** (end-of-file) character (see the *Special Interactive Characters* subsection below). The shell terminates and the **login:** prompt appears again. (If you are using the C, Korn, or POSIX shells, respectively, see *cs(1)*, *ksh(1)*, or *sh-posix(1)* for information about the **ignoreeof** special command.)

How to Communicate Through Your Terminal

HP-UX gathers keyboard input characters and saves them in a buffer. The accumulated characters are not passed to the shell or other program until you type **Return**.

HP-UX terminal input/output is full-duplex. It has full read-ahead, which means that you can type at any time, even while a program is printing on your display or terminal. Of course, if you type during output, the output display will have the input characters interspersed in it. However, whatever you type will be saved and interpreted in the correct sequence. There is a limit to the amount of read-ahead, but it is generous and not likely to be exceeded unless the system is severely overloaded or operating abnormally. When the read-ahead limit is exceeded, the system throws away *all* the saved characters.

stty(1) tells you how to describe the characteristics of your terminal to the system. *profile(4)* explains how to accomplish this task automatically every time you log in.

Special Interactive Characters

A number of special characters are used to control the input and output of your terminal. These characters have defaults and can be redefined with the **stty** command (see *stty(1)*).

| stty Name | Default At Login Character (ASCII Name; Key Names) | Common Redefinition |
|------------------|---|----------------------------|
| eof | ^D (EOT) | |
| erase | # | ^H (BS; Backspace) |
| kill | @ | ^U (NAK), ^X (CAN) |
| intr | ^? (DEL; Delete, Rub, Rubout) | ^C (ETX) |
| quit | ^\ (FS) | |
| start | ^Q (DC1; X-ON) | |
| stop | ^S (DC3; X-OFF) | |

The **eof** character terminates "file" input from the terminal, as read by programs and scripts. By extension, **eof** can also terminate the shell (see the *Logging Out* subsection above).

The **kill** character deletes all characters typed before it on a terminal input line. The **erase** character erases the last character typed. Successive uses of **erase** will erase characters back to, but not beyond, the beginning of the input line.

The **intr** character generates an interrupt signal that bypasses the input buffer. This signal generally causes whatever program you are running to terminate. It can be used to stop a long printout that you don't want. However, programs can arrange either to ignore this signal altogether, or to be notified when it happens (instead of being terminated). For example, the **vi** editor catches interrupts and stops what it is doing, instead of terminating, so that an interrupt can be used to halt an editing operation without losing

the file being edited.

The **quit** character generates a quit signal that bypasses the input buffer and most program traps and causes a running program to terminate. It can cause a core dump in the current directory.

The **stop** character can be used to pause output to the terminal. It is commonly used on video terminals to suspend output to the display while you read what is already being displayed. You can then resume output by typing the **start** character. When **stop** and **start** are used to suspend or resume output, they bypass the keyboard command-line buffer and are not passed to the program. However, any other characters typed on the keyboard are saved and used as input later in the program.

The **eof**, **erase**, and **kill** characters can be used as normal text characters if you escape them with a preceding ****, as in **\^D**. Therefore, to erase a ****, you need two **erases**.

The **intr**, **quit**, **start**, and **stop** characters cannot be escaped on the input line.

End-of-Line and Tab Characters

Besides adapting to the speed of the terminal, HP-UX tries to be intelligent as to whether you have a terminal with a **newline (line-feed)** key, or whether it must be simulated with a return/line-feed character pair. In the latter case, all incoming return characters are changed to line-feed characters (the standard line delimiter), and a return/line-feed pair is echoed to the terminal. If you get into the wrong mode, use the **stty** command to correct it (see *stty(1)*).

Tab characters are used freely in HP-UX source programs. If your terminal does not have the tab function, you can arrange to have tab characters changed into spaces during output, and echoed as spaces during input. The **stty** command sets or resets this mode. By default, the system assumes that tabs are set every eight character positions. The **tabs** command (see *tabs(1)*) can set tab stops on your terminal, if the terminal supports tabs.

How to Run a Program

When you have successfully logged into HP-UX, the shell monitors input from your terminal. The shell accepts typed lines from the terminal, splits them into command names and arguments, then executes the command. The command can be the name of a shell built-in, an executable script of commands, or an executable program. There is nothing special about system-provided commands, except that they are kept in directories where the shell can find them. You can also keep commands in your own directories and arrange for the shell to find them there.

The command name is the first word on an input line to the shell; the command and its arguments are separated from one another by blanks (one or more space and/or tab characters).

When a program terminates, the shell ordinarily regains control and prompts you to indicate that it is ready for another command. The shell has many other capabilities, which are described in detail in the appropriate manpages: *sh-posix(1)* for the POSIX shell, *ksh(1)* for the Korn shell, *sh-bourne(1)* for the Bourne shell, or *csh(1)* for the C shell.

The Current Directory

HP-UX has a file system arranged in a hierarchy of directories. When the system administrator gave you a user name, he or she also created a directory for you (ordinarily with the same name as your user name, and known as your *login* or *home* directory). When you log in, that directory becomes your *current* or *working* directory, and any file name you type is assumed to be in that directory by default. Because you are the owner of this directory, you have full permission to read, write, alter, or destroy its contents. The permissions you have for other directories and files will have been granted or denied to you by their respective owners, or by the system administrator. To change the current working directory use *cd(1)*.

Path Names

To refer to files not in the current directory, you must use a path name. Full (absolute) path names begin with **/**, which is the name of the *root* directory of the whole file system. After the slash comes the name of each directory containing the next subdirectory (followed by a **/**), until finally the file name is reached (for example, **/usr/ae/filex** refers to file **filex** in directory **ae**, while **ae** is itself a subdirectory of **usr**; **usr** is a subdirectory of the root directory). See *glossary(9)* for a formal definition of **path name**.

If your current directory contains subdirectories, the path names of files in them begin with the name of the corresponding subdirectory (without a prefixed **/**). Generally, a path name can be used anywhere a file name is required.

Important commands that modify the contents of directories are **cp**, **mv**, and **rm** which respectively copy, move (that is, rename, relocate, or both), and remove files. To determine the status of files or the contents

of directories, use the **ls** command. Use **mkdir** to make directories, **rmdir** to destroy them, and **mv** to rename them (see *cp(1)*, *ls(1)*, *mkdir(1)*, *mv(1)*, *rm(1)*, and *rmdir(1)*).

Writing a Program

To enter the text of a source program into an HP-UX file, use a text editing program such as **vi**, **ex**, or **ed** (see *vi(1)*, *ex(1)*, and *ed(1)*). The three principal languages available under HP-UX are C (see *cc_bundled(1)* and *cc(1)*), FORTRAN (see *f77(1)*), and Pascal (see *pc(1)*). After the program text has been entered with the editor and written into a file (whose name has the appropriate suffix), you can give the name of that file to the appropriate language processor as an argument. Normally, the output of the language processor will be left in a file named **a.out** in the current directory. Since the results of a subsequent compilation may also be placed in **a.out**, thus overwriting the current output, you may want to use **mv** to give the output a unique name. If the program is written in assembly language, you will probably need to link library sub-routines with it (see *ld(1)*). FORTRAN, C, and Pascal call the linker automatically.

When you have gone through this entire process without encountering any diagnostics, the resulting program can be run by giving its name to the shell in response to the prompt.

Your programs can receive arguments from the command line just as system programs do by using the *argc* and *argv* parameters. See the supplied C tutorial for details.

Text Processing

Almost all text is entered through a text editor. The editor preferred above all others provided with HP-UX is the **vi** editor. For batch-processing text files, the **sed** editor is very efficient. Other editors are used much less frequently. The **ex** editor is useful for handling certain situations while using **vi** but most other editors are rarely used except in various scripts.

The following editors are the same program masquerading under various names: **vi**, **view**, and **vedit** (see *vi(1)*) and **ex** and **edit** (see *ex(1)*). For information about the **sed** stream editor, see *sed(1)*. The **ed** line editor is described in *ed(1)*.

The commands most often used to display text on a terminal are **cat**, **more**, and **pr** (see *cat(1)*, *more(1)*, and *pr(1)*). The **cat** command simply copies ASCII text to the terminal, with no processing at all. The **more** command displays text on the terminal a screenful at a time, pausing for an acknowledgement from the user before continuing. The **pr** command paginates text, supplies headings, and has a facility for multicolumn output. **pr** is most commonly used in conjunction with the **lp** command (see *lp(1)*) to pipe formatted text to a line printer.

Interuser Communication

Certain commands provide interuser communication. Even if you do not plan to use them, it could be beneficial to learn about them, because someone else may direct them toward you. To communicate with another user that is currently logged in, you can use **write** to transfer text directly to that user's terminal display (if permission to do so has been granted by the other user). Otherwise, **elm**, **mailx**, or **mail** (in order of ease of use) can send a message to another user's mailbox. The user is then informed by HP-UX that mail has arrived (if currently logged in) or mail is present (when the user next logs in). Refer to *elm(1)*, *mail(1)*, *mailx(1)*, and *write(1)* for explanations of how these commands are used.

ACKNOWLEDGEMENTS

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited.

SEE ALSO

cat(1), *cc_bundled(1)*, *cd(1)*, *chsh(1)*, *cp(1)*, *cs(1)*, *ed(1)*, *ex(1)*, *ksh(1)*, *ld(1)*, *login(1)*, *lp(1)*, *ls(1)*, *mail(1)*, *mailx(1)*, *man(1)*, *mkdir(1)*, *more(1)*, *mv(1)*, *passwd(1)*, *pr(1)*, *rm(1)*, *rmdir(1)*, *sed(1)*, *sh(1)*, *sh-bourne(1)*, *sh-posix(1)*, *stty(1)*, *tabs(1)*, *vi(1)*, *write(1)*, *a.out(4)*, *profile(4)*, *glossary(9)*.

Web access to HP-UX documentation at <http://docs.hp.com>.

Section 1

Part 1

User Commands (A-M)

Section 1

Part 1

User Commands (A-M)

NAME

intro - introduction to command utilities and application programs

DESCRIPTION

This section describes commands accessible by users, as opposed to system calls in Section (2) or library routines in Section (3), which are accessible by user programs.

Command Syntax

Unless otherwise noted, commands described in this section accept options and other arguments according to the following syntax:

name [*option* (*s*)] [*cmd_arg* (*s*)]

where the elements are defined as follows:

name Name of an executable file.

option One or more *options* can appear on a command line. Each takes one of the following forms:

-no_arg_letter

A single letter representing an option without an argument.

-no_arg_letters

Two or more single-letter options combined into a single command-line argument.

-arg_letter<>opt_arg

A single-letter option followed by a required argument where:

arg_letter

is the single letter representing an option that requires an argument,

opt_arg

is an argument (character string) satisfying the preceding *arg_letter*,

<> represents optional white space.

cmd_arg Path name (or other command argument) *not* beginning with -, or - by itself indicating the standard input. If two or more *cmd_args* appear, they must be separated by white space.

Manual Entry Formats

All manual entries follow an established topic format, but not all topics are included in each entry.

NAME Gives the name(s) of the entry and briefly states its purpose.

SYNOPSIS Summarizes the use of the entry or program entity being described. A few conventions are used:

Computer font strings are literals, and are to be typed exactly as they appear in the manual (except for parameters in the SYNOPSIS section of entries in Sections 2 and 3).

Italic strings represent substitutable argument names and names of manual entries found elsewhere in the manual.

Square brackets [] around an argument name indicate that the argument is optional.

Ellipses (...) are used to show that the previous argument can be repeated.

A final convention is used by the commands themselves. An argument beginning with a dash (-), a plus sign (+), or an equal sign (=) is often taken to be some sort of option argument, even if it appears in a position where a file name could appear. Therefore it is unwise to have file names that begin with -, +, or =.

DESCRIPTION Discusses the function and behavior of each entry.

EXTERNAL INFLUENCES

Information under this heading pertains to programming for various spoken languages. Typical entries indicate support for single- and/or multi-byte characters, the effect of language-related environment variables on system behavior, and other related information.

NETWORKING FEATURES

Information under this heading is applicable only if you are using the networking feature described there (such as NFS).

RETURN VALUE Discusses various values returned upon completion of program calls.

DIAGNOSTICS Discusses diagnostics indications that may be produced. Self-explanatory messages are not listed.

ERRORS Lists error conditions and their corresponding error message or return value.

EXAMPLES Provides examples of typical usage, where appropriate.

WARNINGS Points out potential pitfalls.

DEPENDENCIES Points out variations in HP-UX operation that are related to the user or specific hardware or hardware combinations.

AUTHOR Indicate the origin of the software documented by the manual entry.

FILES Lists file names that are built into the program or command.

SEE ALSO Provides pointers to related topics.

BUGS Discusses known bugs and deficiencies, occasionally suggesting fixes.

STANDARDS CONFORMANCE

This section lists the standard specifications to which the HP-UX component conforms.

RETURN VALUE

Upon termination, each command returns two bytes of status, one supplied by the system giving the cause for termination, and (in the case of “normal” termination) one supplied by the program (for descriptions, see *wait(2)* and *exit(2)*). The system-supplied byte is 0 for normal termination. The byte provided by the program is customarily 0 for successful execution and non-zero to indicate errors or failure such as incorrect parameters in the command line, or bad or inaccessible data. Values returned are usually called variously “exit code”, “exit status”, “return code”, or “return value”, and are 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 when they encounter a null character (the string terminator) within a line.

SEE ALSO

getopt(1), *exit(2)*, *wait(2)*, *getopt(3C)*, *hier(5)*, *introduction(9)*.

Web access to HP-UX documentation at <http://docs.hp.com>.

NAME

adb, adb64 - absolute debugger

SYNOPSIS

adb [-w] [-I *dir*] [-k] [-m] [-P *pid*] *objfil* [*corfil*]

adb64 [-w] [-I *dir*] [-k] [-m] [-P *pid*] *objfil* [*corfil*]

DESCRIPTION

The **adb** command executes a general-purpose debugging program that is sensitive to the underlying architecture of the processor and operating system on which it runs. It can be used to examine files and provide a controlled environment for executing HP-UX programs.

adb calls **adb64** to process 64 bit files.

objfil is normally an executable program file, or an HP-UX kernel (vmunix), preferably containing a symbol table; if not, the symbolic features of **adb** cannot be used, although the file can still be examined. The default for *objfil* is **a.out**.

corfil is assumed to be a core image file produced after executing *objfil* or an HP-UX crash file produced from the *objfil*. The default for *corfil* is **core**.

Requests to **adb** are read from standard input and **adb** responds on standard output. If the **-w** flag is present, *objfil* is created (if necessary) and opened for reading and writing, to be modified using **adb**. The **-I** option specifies a directory where files read with **\$<** or **\$<<** (see below) are sought; the default is **/usr/lib/adb**. **adb** ignores QUIT; INTERRUPT causes return to the next **adb** command.

The following options are also supported:

-k Allows **adb** to read *objfil* as an HP-UX kernel file and *corfil* as an HP-UX crash dump. This also allows virtual-to-physical address translation, useful for kernel debugging. In this case, *corfil* should be an HP-UX crash dump or **/dev/mem**. Without **-k** or **-m**, **adb** treats *objfil* as an application program file and *corfil* as an application core file.

When **adb** is invoked with this option, it sets up the context of the currently running process using space registers four through seven. A user specified address is dereferenced by combining it with the appropriate space register, depending on the quadrant in which the 32-bit address lies.

When the current radix is not (decimal) ten, the **-k** option allows **adb** to support the notion of long pointers or addresses in the form *space.offset*. Once a space is specified, all subsequent addresses are dereferenced using that space until the user enters another long address. If a space equal to (hexadecimal) 0xffffffff is used, **adb** reverts to the previous context and uses space registers four through seven to dereference 32-bit addresses.

-m Must be specified instead of **-k** when a core dump is written to multiple files.

When **-m** is used, *corfil* must be specified as the path name of the directory that contains system core dump files. A command line using **-m** might look similar to the following:

```
adb -m /var/adm/crash/core.1/vmunix /var/adm/crash/core.1
```

Notice that when **-m** is specified on the command line, **-k** is not necessary.

-P *pid* Causes **adb** to adopt process *pid* as a "traced" process (see *ptrace(2)*). This option is helpful for debugging processes that were not originally run under the control of **adb**.

Requests to **adb** follow the form:

```
[address] [, count] [command] [;]
```

If *address* is present, *dot* is set to *address*. Initially *dot* is set to 0. For most commands, *count* specifies the number of times the command is to be executed. The default *count* is 1. *address* and *count* are expressions.

The interpretation of an address depends on the context in which it is used. If a subprocess is being debugged, addresses are interpreted in the address space of the subprocess. (For further details of address mapping see Addresses below.)

Expressions

Expressions are interpreted as follows:

| | |
|----------------------------------|---|
| . | The value of <i>dot</i> . |
| + | The value of <i>dot</i> increased by the current increment. |
| ^ | The value of <i>dot</i> decreased by the current decrement. |
| " | The last <i>address</i> typed. |
| <i>integer</i> | A number. The prefix 0 (zero) forces interpretation in octal radix; the prefixes 0d and 0D force interpretation in decimal radix; the prefixes 0x and 0X force interpretation in hexadecimal radix. Thus 020 = 0d16 = 0x10 = sixteen. If no prefix appears, the <i>default radix</i> is used; see the \$d command. The radix is initialized to the base used in the assembly language for the processor involved. Note that a hexadecimal number whose most significant digit would otherwise be an alphabetic character must have a 0x (or 0X) prefix. |
| <i>integer</i> . <i>fraction</i> | A 32-bit floating-point number. |
| ' <i>cccc</i> ' | The ASCII value of up to 4 characters. A backslash (\) can be used to escape a single quote ('). |
| < <i>name</i> | <i>name</i> can have the value of either a variable or a register. adb maintains a number of variables named by single letters or digits; see Variables below. If <i>name</i> is a register, the value of the register is obtained from the CORE_PROC segment in <i>corfil</i> (before the subprocess is initiated) or from the user area of the subprocess. Register names are implementation dependent; see the \$r command. |
| <i>symbol</i> | A <i>symbol</i> is a sequence of uppercase or lowercase letters, underscores, or digits, not starting with a digit. A backslash (\) can be used to escape other characters. The value of the <i>symbol</i> is taken from the symbol table in <i>objfil</i> . An initial underscore (_) is prefixed to <i>symbol</i> , if needed. |
| _ <i>symbol</i> | If the compiler prefixes _ to an external symbol, it may be necessary to cite this name to distinguish it from a symbol generated in assembly language. |
| (<i>exp</i>) | The value of the expression <i>exp</i> . |

The following are monadic operators:

| | |
|--------------|---|
| * <i>exp</i> | The contents of the location addressed by <i>exp</i> in <i>corfil</i> . |
| @ <i>exp</i> | The contents of the location addressed by <i>exp</i> in <i>objfil</i> . |
| - <i>exp</i> | Integer negation. |
| ~ <i>exp</i> | Bitwise complement. |

The following dyadic operators are left associative and are less binding than monadic operators:

| | |
|-----------------------|--|
| <i>e1</i> + <i>e2</i> | Integer addition. |
| <i>e1</i> - <i>e2</i> | Integer subtraction. |
| <i>e1</i> * <i>e2</i> | Integer multiplication. |
| <i>e1</i> % <i>e2</i> | Integer division. |
| <i>e1</i> & <i>e2</i> | Bitwise conjunction. |
| <i>e1</i> <i>e2</i> | Bitwise disjunction. |
| <i>e1</i> # <i>e2</i> | <i>e1</i> rounded up to the next multiple of <i>e2</i> . |

Commands

Most commands consist of an action character followed by a modifier or list of modifiers. The following action characters can take format specifiers. (The action characters ? and / can be followed by *; see Addresses for further details.)

| | |
|------------|--|
| ? <i>f</i> | Locations starting at <i>address</i> in <i>objfil</i> are printed according to the format <i>f</i> . <i>dot</i> is incremented by the sum of the increments for each format letter. If a subprocess has been initiated, <i>address</i> references a location in the address space of the subprocess instead of |
|------------|--|

objfil.

/f Locations starting at *address* in *corfil* are printed according to the format *f* and *dot* is increased like ?. If a subprocess has been initiated, *address* refers to a location in the address space of the subprocess instead of *corfil*.

=f The value of *address* is printed in the styles indicated by the format *f*. (For **i** format ? is printed for the parts of the instruction that refer to subsequent words.)

A *format* consists of one or more characters that specify a style of printing. Each format character can be preceded by an integer that indicates how many times the format is repeated. While stepping through a format, *dot* is increased by the amount given for each format character. If no format is given then the last format is used.

The following format characters are available:

- a** 0 Print the value of *dot* in symbolic form.
- b** 1 Print the addressed byte in hexadecimal.
- B** 1 Print the addressed byte in octal.
- c** 1 Print the addressed character (the sign bit is ignored).
- C** 1 Print the addressed character using the following escape convention. First, the sign bit is discarded, then character values 000 to 040 are printed as @ followed by the corresponding character in the range 0100 to 0140. The character @ is printed as @@.
- d** 2 Print 2 bytes in decimal.
- D** 4 Print 4 bytes in decimal.
- f** 4 Print the 32 bit value as a floating point number.
- F** 8 Print double floating point.
- i** *n* Print as machine instructions. The value of *n* is the number of bytes occupied by the instruction.
- n** 0 Print a new-line character.
- o** 2 Print 2 bytes in octal. All octal numbers output by **adb** are preceded by 0.
- O** 4 Print 4 bytes in octal.
- p** *n* Print the addressed value in symbolic form. The value of *n* is a machine-dependent constant.
- q** 2 Print 2 bytes in signed octal.
- Q** 4 Print 4 bytes in signed octal.
- r** 0 Print a space.
- s** *n* Print the addressed characters until a zero character is reached.
- S** *n* Print a string using the @ escape convention. The value *n* is the length of the string including its zero terminator.
- t** 0 When preceded by an integer, moves to the next appropriate tab stop. For example, **8t** moves to the next 8-space tab stop.
- u** 2 Print 2 bytes as an unsigned decimal number.
- U** 4 Print 4 bytes as an unsigned decimal number.
- x** 2 Print 2 bytes in hexadecimal.
- X** 4 Print 4 bytes in hexadecimal.
- Y** 4 Print 4 bytes in date format (see *ctime(3C)*).
- "..."** 0 Print the enclosed string.
- ^** *dot* is decreased by the current increment. Nothing is printed.
- +** *dot* is increased by 1. Nothing is printed.

- *dot* is decreased by 1. Nothing is printed.

new-line

Repeat the previous command with a *count* of 1. The value of *dot* continues from the end of the previous format, unlike a *[?/]* command with no *address*, which repeats the previous *address* value. New-line can also be used to repeat a *:s* or *:c* command; however, any arguments to the previous command are lost.

*[?/]*1 *value mask*

Words starting at *dot* are masked with *mask* and compared with *value* until a match is found. If *L* is used, *adb* looks to match 4 bytes at a time instead of 2. If no match is found, *dot* is left unchanged; otherwise *dot* is set to the matched location. If *mask* is omitted -1 is used.

*[?/]*w *value* ...

Write the 2-byte *value* into the addressed location. If the command is *W*, write 4 bytes. Odd addresses are not allowed when writing to the subprocess address space.

=m Toggle the address mapping of *corfil* between the initial map set up for a valid core file and the default mapping pair which the user can modify with */m*. If the *corfil* was invalid, only the default mapping is available.

*[?/]*m *b1 e1 f1* *[?/]*

Record new values for (*b1*, *e1*, *f1*). If fewer than three expressions are given, the remaining map parameters are left unchanged. If the *?* or */* is followed by ***, the second segment (*b2*, *e2*, *f2*) of the mapping is changed. If the list is terminated by *?* or */*, the file (*objfil* or *corfil*, respectively) is used for subsequent requests. (For example, */m?* causes */* to refer to *objfil*.) A */m* command switches the *corfil* mapping to the default mapping pair. For a valid core file, the =m command can be used to switch back to the initial mapping.

>name

Assign *dot* to the variable or register named.

! Call a shell to read the remainder of the line following !.

The following \$ commands take the form \$*modifier*:

\$<*f* Read commands from the file *f*. If this command is executed in a file, further commands in the file are not seen. If a *count* is given, and is zero, the command is ignored. The value of the count is placed in variable 9 before the first command in *f* is executed.

\$<<*f* Similar to \$< except it can be used in a file of commands without causing the file to be closed. Variable 9 is saved when the command executes and is restored when it completes. Only five \$<< files can be open at once.

\$>*f* Send output to the file *f*, which is created if it does not already exist.

\$new-line Print the process id and register values.

\$b Print all breakpoints and their associated counts and commands.

\$c C stack backtrace. If *address* is given, it is taken as the address of the current frame (instead of the normal stack frame pointer). If *count* is given, only the first *count* frames are printed.

\$d Set the default radix to *address* and report the new value. Note that *address* is interpreted in the (old) current radix. Thus 10\$d never changes the default radix. To make decimal the default radix, use 0d10\$d.

\$e The names and values of external variables are printed.

\$f Print the floating-point registers.

\$m Print the address map. This includes both the initial and default maps for a valid *corfil* with an indication of which is currently active.

\$N [*nodenumber*]

Print the number of nodes on V-class multinode machines and the current node number. To switch to another node, enter \$N *nodenumber*.

\$o The default for all integers input is octal.

\$q Exit from *adb*.

| | |
|------------|--|
| \$r | Print the general registers and the instruction addressed by the process counter. <i>dot</i> is set to the process counter contents. |
| \$s | Set the limit for symbol matches to <i>address</i> . The default is system dependent. |
| \$v | Print all non-zero variables in the current radix. |
| \$w | Set the page width for output to <i>address</i> (default 80). |
| \$x | The default for all integers input is hexadecimal. |
| \$z | Print a list of signals and how they are handled. See :z for information on changing signal handling. |

The available **:** commands manage subprocesses, and take the form **:modifier**:

| | |
|--------------------|---|
| :bc | Set breakpoint at <i>address</i> . The breakpoint is executed <i>count</i> -1 times before causing a stop. Each time the breakpoint is encountered, the command <i>c</i> is executed. If this command sets <i>dot</i> to zero, the breakpoint causes a stop. |
| :cs | Continue the subprocess with signal <i>s</i> (see <i>signal(5)</i>). If <i>address</i> is given, the subprocess continues at this address. If no signal is specified, the signal that caused the subprocess to stop is sent. Breakpoint skipping is the same as for :r . |
| :d | Delete breakpoint at <i>address</i> . :d* deletes all breakpoints. |
| :e | Set up a subprocess as in :r ; no instructions are executed. |
| :k | Terminate the current subprocess, if any. |
| :r | Run <i>objfil</i> as a subprocess. If <i>address</i> is given explicitly, the program is entered at this point; otherwise the program is entered at its standard entry point. The value <i>count</i> specifies how many breakpoints are ignored before stopping. Arguments to the subprocess may be supplied on the same line as the command. An argument starting with < or > causes the standard input or output to be established for the command. All signals are turned on when entering the subprocess. |
| :ss | As for :c except that the subprocess is single stepped <i>count</i> times. If there is no current subprocess, <i>objfil</i> is run as a subprocess as for :r . In this case no signal can be sent; the remainder of the line is treated as arguments to the subprocess. |
| :Ss | Same as :c except that a temporary breakpoint is set at the next instruction. Useful for stepping across subroutines. |
| :x a [b]... | Execute subroutine <i>a</i> with parameters <i>[b]...</i> |
| :zd | Change signal handling for a specified signal. Disposition <i>d</i> can be specified as: +s Stop process when signal is received. -s Do not stop process when signal is received. +r Report when signal is received. -r Do not report when signal is received. +d Deliver signal to the target process. -d Do not deliver signal to the target process. |

For example, **0x10:z+d** enables delivering of signal number **0x10** to the target process. Use **\$z** to display existing settings.

Variables

adb provides named and numbered variables. Named variables are set initially by **adb** but are not used subsequently. Numbered variables are reserved for communication as follows:

| | |
|----------|--|
| 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 \$< command. |

a

On entry, the following named variables are set from the coreheaders in the *corfil*. If *corfil* does not appear to be a **core** file, these values are set from *objfil*.

| | |
|----------|---------------------------------------|
| b | The base address of the data segment. |
| d | The data segment size. |
| s | The stack segment size. |
| t | The text segment size. |

The following variables are set from *objfil*.

| | |
|----------|---|
| e | The entry point. |
| m | The "magic" number as defined in <magic.h> . |

Addresses

The file address associated with a written address is determined by a mapping described below; see **\$m**. Both the *objfil* mapping and the default *corfil* mapping are represented by two triples (*b1*, *e1*, *f1*) and (*b2*, *e2*, *f2*). The initial mapping for a valid *corfil* contains a triple for each segment (coreheader).

The *file address* corresponding to a written *address* is calculated as follows:

If

$b1 \leq \text{address} < e1$, then $\text{file address} = \text{address} + f1 - b1$.

Otherwise, if

$b2 \leq \text{address} < e2$, then $\text{file address} = \text{address} + f2 - b2$.

Otherwise, the requested *address* is not valid. For a valid *corfil*, this pattern repeats as many times as there are segments (coreheaders) in the *corfil*, rather than twice. If **?** or **/** is followed by *****, only the second triple is used, or (when using the initial mapping of a valid *corfil*) only segments with a **CORE_STACK** coreheader.

The initial setting of both mappings is suitable for normal **a.out** and **core** files. If either file is not of the kind expected, **adb** sets *b1* to 0, *e1* to the maximum file size, and *f1* to 0; in this way the entire file can be examined with no address translation.

adb keeps all appropriate values as signed 32-bit integers so that it can be used on large files.

EXTERNAL INFLUENCES

International Code Set Support

Single- and multi-byte character code sets are supported.

RETURN VALUE

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

returned non-zero status.

DEPENDENCIES

- Setting breakpoints in shared libraries is not supported.

adb does not read the linker symbol table for shared libraries, and cannot access locations in shared libraries by name. In a stack backtrace (**\$c**), **adb** does not know the names of shared library procedures.

If the core file was created when the program was in a shared library function, the **\$c** command does not work. When a stack backtrace for the core file encounters a shared library procedure on the stack it aborts at that point.

- A leading zero by itself is not recognized as a radix indicator. Use the prefixes **0o** or **0O** (zero-oh) to force interpretation in octal radix. The prefixes **0t** and **0T** are also accepted to force interpretation in decimal radix. Thus **0o20** = **0t16** = sixteen. A hexadecimal number whose most significant digit would otherwise be an alphabetic character may begin with a leading zero instead of **0x** (or **0X**), if the default radix is hexadecimal.
- The **\$f** command prints floating point registers as 32-bit single precision and **\$F** prints these registers as 64-bit doubles.
- \$R** prints all registers available to **adb** users.
- The **:x** and **:S** commands are not currently supported.
- adb** can be used to inspect relocatable object files; it reads the symbol table and sets up the appropriate mappings for text and data. Note that relocatable object files do not necessarily contain an exact image of the initialized data; however, if this is the case, the data mapping is not set.

AUTHOR

adb was developed by AT&T and HP.

FILES

```
a.out
core
/dev/mem
/dev/kmem
/dev/swap
```

SEE ALSO

ptrace(2), crt0(3), ctime(3C), end(3C), a.out(4), core(4), signal(5).

ADB Tutorial

NAME

adjust - simple text formatter

SYNOPSIS

adjust [-b] [-c|-j|-r] [-m *column*] [-t *tabsize*] [*files* ...]

DESCRIPTION

The **adjust** command is a simple text formatter for filling, centering, left and right justifying, or only right justifying text paragraphs, and is designed for interactive use. It reads the concatenation of input files (or standard input if none are given) and produces on standard output a formatted version of its input, with each paragraph formatted separately. If - is given as an input filename, **adjust** reads standard input at that point (use -- as an argument to separate - from options.)

adjust reads text from input lines as a series of words separated by space characters, tabs, or newlines. Text lines are grouped into paragraphs separated by blank lines. By default, text is copied directly to the output, subject only to simple filling (see below) with a right margin of 72, and leading spaces are converted to tabs where possible.

Options

The **adjust** command recognizes the following command-line options:

- b Do not convert leading space characters to tabs on output; (output contains no tabs, even if there were tabs in input).
- c Center text on each line. Lines are pre- and post-processed, but no filling is performed.
- j Justify text. After filling, insert spaces in each line as needed to right justify it (except in the last line of each paragraph) while keeping the justified left margin.
- r After filling text, adjust the indentation of each line for a smooth right margin (ragged left margin).

-mcolumn

Set the right fill margin to the given column number, instead of 72. Text is filled, and optionally right justified, so that no output line extends beyond this column (if possible). If -m0 is given, the current right margin of the first line of each paragraph is used for that and all subsequent lines in the paragraph.

By default, text is centered on column 40. With -c, the -m option sets the middle column of the centering "window", but -m0 auto-sets the right side as before (which then determines the center of the "window").

- t*tabsize* Set the tab size to other than the default (eight columns).

Only one of the -c, -j, and -r options is allowed in a single command line.

Details

Before doing anything else to a line of input text, **adjust** first handles backspaces, rubbing out preceding characters in the usual way. Next, it ignores all non-printable characters except tab. It then expands all tabs to spaces.

For simple text filling, the first word of the first line of each paragraph is indented the same amount as in the input line. Each word is then carried to the output followed by one space. "Words" ending in *terminal_character*[*quote*][*closing_character*] are followed by two spaces, where *terminal_character* is any of ., :, ?, or !; *quote* is a single closing quote (') character or double-quote character ("), and *close* is any of),], or }. Here are some examples:

end. of? sentence.' sorts!" of.) words?"]

(**adjust** does not place two spaces after a pair of single closing quotes (' ') following a *terminal_character*).

adjust starts a new output line whenever adding a word (other than the first one) to the current line would exceed the right margin.

adjust understands indented first lines of paragraphs (such as this one) when filling. The second and subsequent lines of each paragraph are indented the same amount as the second line of the input paragraph if there is a second line, else the same as the first line.

- * **adjust** also has a rudimentary understanding of tagged paragraphs (such as this one) when filling. If the second line of a paragraph is indented more than the first, and the first line has a word beginning at the same indentation as the second line, the input column position of the tag word or words (prior to the one matching the second line indentation) is preserved.

Tag words are passed through without change of column position, even if they extend beyond the right margin. The rest of the line is filled or right justified from the position of the first non-tag word.

When **-j** is given, **adjust** uses an intelligent algorithm to insert spaces in output lines where they are most needed, until the lines extend to the right margin. First, all one space word separators are examined. One space is added to each separator, starting with the one having the most letters between it and the preceding and following separators, until the modified line reaches the right margin. If all one space separators are increased to two spaces and more spaces must be inserted, the algorithm is repeated with two space separators, and so on.

Output line indentation is held to one less than the right margin. If a single word is larger than the line size (right margin minus indentation), that word appears on a line by itself, properly indented, and extends beyond the right margin. However, if **-r** is used, such words are still right justified, if possible.

If the current locale defines class names **ekinsoku** and **bkinsoku** (see *iswctype(3C)*), **adjust** formats the text in accordance with the **ekinsoku/bkinsoku** character classification and margin settings (see **-r**, **-j**, and **-m** options).

EXTERNAL INFLUENCES

Environment Variables

LANG provides a default value for the internationalization variables that are unset or null. If **LANG** is unset or null, the default value of "C" (see *lang(5)*) is used. If any of the internationalization variables contains an invalid setting, **adjust** will behave as if all internationalization variables are set to "C". See *environ(5)*.

LC_ALL If set to a non-empty string value, overrides the values of all the other internationalization variables.

LC_CTYPE determines the interpretation of text as single and/or multi-byte characters, the classification of characters as printable, and the characters matched by character class expressions in regular expressions.

LC_MESSAGES determines the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.

NLSPATH determines the location of message catalogs for the processing of **LC_MESSAGES**.

International Code Set Support

Single- and multi-byte character code sets are supported.

DIAGNOSTICS

adjust complains to standard error and later returns a nonzero value if any input file cannot be opened (it skips the file). It does the same (but quits immediately) if the argument to **-m** or **-t** is out of range, or if the program is improperly invoked.

Input lines longer than **BUFSIZ** are silently split (before tab expansion) or truncated (afterwards). Lines that are too wide to center begin in column 1 (no leading spaces).

EXAMPLES

This command is useful for filtering text while in *vi*(1). For example,

```
!}adjust
```

reformats the rest of the current paragraph (from the current line down), evening the lines.

The *vi* command:

```
:map ^X {!}adjust -j^V^M
```

(where **^** denotes control characters) sets up a useful "finger macro". Typing **^X** (Ctrl-X) reformats the entire current paragraph.

adjust -m1 is a simple way to break text into separate words without white space, except for tagged-paragraph tags.

WARNINGS

This program is designed to be simple and fast. It does not recognize backslash to escape white space or other characters. It does not recognize tagged paragraphs where the tag is on a line by itself. It knows that lines end in newline or null, and how to deal with tabs and backspaces, but it does not do anything special with other characters such as form feed (they are simply ignored). For complex operations, standard text processors are likely to be more appropriate.

This program could be implemented instead as a set of independent programs, fill, center, and justify (with the **-r** option). However, this would be much less efficient in actual use, especially given the program's special knowledge of tagged paragraphs and last lines of paragraphs.

AUTHOR

adjust was developed by HP.

SEE ALSO

nroff(1).


a

NAME

admin - create and administer SCCS files

SYNOPSIS

```
admin -i[name] [-n] [-b] [-a login] ... [-d flag[flag-val]] ... [-f flag[flag-val]] ...
      [-m mrlist] ... [-r rel] [-t[name]] [-y[comment]] file ...

admin -n [-a login] ... [-d flag[flag-val]] ... [-f flag[flag-val]] ... [-m mrlist] ...
      [-t[name]] [-y[comment]] file ...

admin [-a login] ... [-e login] ... [-d flag[flag-val]] ... [-m mrlist] ...
      [-r rel] [-t[name]] file ...

admin -h file ...

admin -z file ...
```

DESCRIPTION

The **admin** command is used to create new SCCS files and change the parameters of existing ones. Arguments to **admin**, which may appear in any order, (unless **--** is specified as an argument, in which case all arguments after **--** are treated as files) consist of option arguments, beginning with **-**, and named *files* (note that SCCS file names must begin with the characters **s.**). If a named *file* does not exist, it is created and its parameters are initialized according to the specified option arguments. Parameters not initialized by an option argument are assigned a default value. If a named *file* does exist, parameters corresponding to specified option arguments are changed, and other parameters are left unaltered.

If *directory* is named instead of *file*, **admin** acts on each *file* in *directory*, except that non-SCCS files (the last component of the path name does not begin with **s.**) and unreadable files are silently ignored. If a name of **-** is given, the standard input is read, and each line of the standard input is assumed to be the name of an SCCS file to be processed. Again, non-SCCS files and unreadable files are silently ignored.

The **admin** option arguments apply independently to all named *files*, whether one file or many. In the following discussion, each option is explained as if only one file is specified, although they affect single or multiple files identically.

Options

The **admin** command supports the following options and command-line arguments:

- n** This option indicates that a new SCCS file is to be created.
- i[name]** The *name* of a file from which the contents for a new SCCS file is to be taken. (if name is a binary file, then you must specify the **-b** option) The contents constitutes the first delta of the file (see the **-r** option for the delta numbering scheme). If the **-i** option is used but the file name is omitted, the text is obtained by reading the standard input until an end-of-file is encountered. If this option is omitted, the SCCS file is created with an empty initial delta. Only one SCCS file can be created by an **admin** command on which the **-i** option is supplied. Using a single **admin** to create two or more SCCS files requires that they be created empty (no **-i** option). Note that the **-i** option implies the **-n** option.
- b** Encode the contents of name, specified to the **-i** option. This keyletter must be used if name is a binary file; otherwise, a binary file will not be handled properly by SCCS commands.
- r rel** The release (*rel*) into which the initial delta is inserted. This option can be used only if the **-i** option is also used. If the **-r** option is not used, the initial delta is inserted into release 1. The level of the initial delta is always 1 (by default initial deltas are named 1.1).
- t[name]** The *name* of a file from which descriptive text for the SCCS file is to be taken. If the **-t** option is used and **admin** is creating a new SCCS file (the **-n** and/or **-i** options are also used), the descriptive text file name must also be supplied. In the case of existing SCCS files:
 - A **-t** option without a file name causes removal of descriptive text (if any) currently in the SCCS file.
 - A **-t** option with a file name causes text (if any) in the named file to replace the descriptive text (if any) currently in the SCCS file.

-f *flag*

This option specifies a *flag*, and possibly a value for the *flag*, to be placed in the SCCS file. Several **-f** options can be supplied on a single **admin** command line. The allowable *flags* and their values are:

- b** Allows use of the **-b** option on a **get** command (see *get(1)*) to create branch deltas.
- cceil** The highest release (i.e., "ceiling"), a number less than or equal to 9999, which can be retrieved by a **get** command for editing. The default value for an unspecified **c** flag is 9999.
- ffloor** The lowest release (i.e., "floor"), a number greater than 0 but less than 9999, which may be retrieved by a **get** command for editing. The default value for an unspecified **f** flag is 1.
- dSID** The default delta number *SID* to be used by a **get** command (see *get(1)*).
- istr** Causes the message:

No id keywords (cm7)

issued by **get** or **delta** to be treated as a fatal error (see *delta(1)*). In the absence of this flag, the message is only a warning. The message is issued if no SCCS identification keywords (see *get(1)*) are found in the text retrieved or stored in the SCCS file. If a value is supplied, the keywords must exactly match the given string. However, the string must contain a keyword, but must not contain embedded newlines.

- j** Allows concurrent **get** commands for editing on the same *SID* of an SCCS file. This allows multiple concurrent updates to the same version of the SCCS file.
- Only one user can perform concurrent edits. Access by multiple users is usually accomplished by using a common login or a set user ID program (see *chmod(1)* and *exec(2)*).
- l list** A *list* of releases to which deltas can no longer be made. (A **get -e** against one of these locked releases fails). The *list* has the following syntax:

list ::= *range* / *list* , *range*
range ::= *RELEASE NUMBER* | **a**

The character **a** in the *list* is equivalent to specifying *all releases* for the named SCCS file. Omitting any list is equivalent to **a**.

- n** Causes **delta** to create a null delta in each of those releases being skipped (if any) when a delta is made in a *new* release (such as when making delta 5.1 after delta 2.7, release 3 and release 4 are skipped). These null deltas serve as **anchor points** so that branch deltas can be created from them later. The absence of this flag causes skipped releases to be nonexistent in the SCCS file, preventing branch deltas from being created from them in the future.
- qtext** User-definable text substituted for all occurrences of the **%Q%** keyword in SCCS file text retrieved by **get**.
- mmod** The *module* name of the SCCS file substituted for all occurrences of the **%M%** keyword in SCCS file text retrieved by **get**. If the **m** flag is not specified, the value assigned is the name of the SCCS file with the leading **s.** removed.
- ttype** The *type* of module in the SCCS file substituted for all occurrences of **%Y%** keyword in SCCS file text retrieved by **get**.

- v**[*pgm*] Causes **delta** to prompt for Modification Request (*MR*) numbers as the reason for creating a delta. The optional value specifies the name of a (*MR*) number validity checking program (see *delta*(1)). (If this flag is set when creating an SCCS file, the **m** option must also be used even if its value is null).
- x** Causes **get** to create files with execute permissions.
- d** *flag* Causes removal (deletion) of the specified *flag* from an SCCS file. The **-d** option can be specified only when processing existing SCCS files. Several **-d** options can be supplied on a single **admin** command line. See the **-f** option for allowable *flag* names.
- l***list* A *list* of releases to be unlocked. See the **-f** option for a description of the **l** flag and the syntax of a *list*.
- a** *login* A *login* name, or numerical HP-UX group ID, to be added to the list of users allowed to make deltas (changes) to the SCCS file. A group ID is equivalent to specifying all *login* names common to that group ID. Several **a** options can be used on a single **admin** command line. As many *logins* or numerical group IDs as desired can be on the list simultaneously. If the list of users is empty, anyone can add deltas. A *login* or group ID preceded by a **!** denies permission to make deltas.
- e** *login* A *login* name or numerical group ID to be erased from the list of users allowed to make deltas (changes) to the SCCS file. Specifying a group ID is equivalent to specifying all *login* names common to that group ID. Several **e** options can be used on a single **admin** command line.
- y**[*comment*] The *comment* text is inserted into the SCCS file as a comment for the initial delta in a manner identical to that of *delta*(1). Omission of the **-y** option results in a default comment line being inserted in the form:
- date and time created YY / MM / DD / HH / MM / SS by login**
- The **-y** option is valid only if the **-i** and/or **-n** options are specified (i.e., a new SCCS file is being created).
- m** *mrlist* The list of Modification Request (*MR*) numbers is inserted into the SCCS file as the reason for creating the initial delta, in a manner identical to *delta*(1). The **v** flag must be set and the (*MR*) numbers are validated if the **v** flag has a value (the name of an (*MR*) number validation program). Diagnostic messages occur if the **v** flag is not set or (*MR*) validation fails.
- h** Causes **admin** to check the structure of the SCCS file (see *sccsfile*(4)), and to compare a newly computed checksum (the sum of all of the characters in the SCCS file except those in the first line) with the checksum that is stored in the first line of the SCCS file. Appropriate error diagnostics are produced.
- This option inhibits writing on the file, thus canceling the effect of any other options supplied, and therefore is only meaningful when processing existing files.
- z** The SCCS file checksum is recomputed and stored in the first line of the SCCS file (see **-h**, above).
- Note that use of this option on a truly corrupted file can prevent future detection of the corruption.

Access Control Lists (ACLs)

Do not add optional ACL entries to SCCS files. SCCS removes them, possibly causing unexpected and undesirable access modes.

EXTERNAL INFLUENCES

Environment Variables

LC_CTYPE determines the interpretation of text as single- and/or multi-byte characters.

LC_MESSAGES determines the language in which messages are displayed.

If **LC_CTYPE** or **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of **C** (see *lang*(5)) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **admin** behaves as if all internationalization variables are set to **C**. See

environ(5).

International Code Set Support

Single-byte and multi-byte character code sets are supported.

a

DIAGNOSTICS

Use *sccshelp*(1) for explanations.

WARNINGS

SCCS files can be any length, but the number of lines in the text file itself cannot exceed 99 999 lines.

FILES

The last component of all SCCS file names must be of the form *s.filename*. New SCCS files are given mode 444 (see *chmod*(1)). Write permission in the pertinent directory is required to create a file. All writing done by **admin** is to a temporary x-file, called *x.filename*, (see *get*(1)), created with mode 444 if the **admin** command is creating a new SCCS file, or with the same mode as the SCCS file if it exists. After successful execution of **admin**, the SCCS file is removed (if it exists), and the x-file is renamed to the name of the SCCS file. This ensures that changes are made to the SCCS file only if no errors occurred.

It is recommended that directories containing SCCS files be mode 755 and that SCCS files themselves be mode 444. The mode of any given directory allows only the owner to modify SCCS files contained in that directory. The mode of the SCCS files prevents any modification at all except by SCCS commands.

If it should be necessary to patch an SCCS file for any reason, the mode can be changed to 644 by the owner, thus allowing the use of **vi** or any other suitable editor. *Care must be taken!* The edited file should *always* be processed by an **admin -h** to check for corruption followed by an **admin -z** to generate a proper checksum. Another **admin -h** is recommended to ensure the SCCS file is valid.

admin also makes use of a transient lock file called *z.filename*, which is used to prevent simultaneous updates to the SCCS file by different users. See *get*(1) for further information.

SEE ALSO

delta(1), *ed*(1), *get*(1), *sccshelp*(1), *prs*(1), *what*(1), *sccsfile*(4), *acl*(5).

STANDARDS CONFORMANCE

admin: SVID2, SVID3, XPG2, XPG3, XPG4

NAME

answer - phone message transcription system

SYNOPSIS

answer [-pu]

DESCRIPTION

The **answer** interactive program helps you to transcribe telephone (and other) messages into electronic mail.

The program uses your personal **elm** alias database and the system **elm** alias database, allowing you to use aliases to address the messages.

Options

answer supports the following options:

- p Prompt for phone-slip-type message fields.
- u Allow addresses that are not aliases.

Operation

answer begins with the **Message to:** prompt. Enter a one-word alias or a two-word user name ("Words" are separated by spaces.) The user name is converted to an alias in the form *f_lastword*, where *f* is the first character of the first word, *lastword* is the second word, and all letters are shifted to lowercase. For example, **Dave Smith** is converted to the alias **d_smith**.

Without the **-u** option, the specified or converted alias must exist in the alias databases. With the **-u** option, if the processed "alias" is not in the alias databases, it is used for the address as is.

The fully expanded address is displayed.

With the **-p** option, you are asked for typical message slip data:

```

Caller:
of:
Phone:

TELEPHONED      -
CALLED TO SEE YOU -
WANTS TO SEE YOU -
RETURNED YOUR CALL -
PLEASE CALL     -
WILL CALL AGAIN -
*****URGENT***** -

```

Enter the appropriate data. You can put just an **X** or nothing after the pertinent dash prompts, or you can type longer comments. Whatever you enter becomes part of the message. Lines with no added text are omitted from the message.

Finally, you are prompted for a message. Enter a message, if any, ending with a blank line. The message is sent and the **Message to:** prompt is repeated.

To end the program, enter any one of **bye**, **done**, **exit**, or **quit**, at the **Message to:** prompt.

EXAMPLES

User input is in normal type.

With No Options

This example shows a valid alias, an invalid user name, and a valid user name. In the invalid case, the converted alias is displayed in square brackets.

```

-----

Message to: oswald
address 'oswaldr@mycompany.com (Oswald Rarebit)'
```

Enter message for oswald ending with a blank line.

a

```
> And here is the message.
>
```

```
-----

Message to: albert einstein
Sorry, could not find 'albert einstein' [a_einstein] in list!
```

```
Message to: john jones
address 'mrjones@companybee.com (John P. Jones)'
```

```
Enter message for john jones ending with a blank line.
```

```
> A nice message
>
```

```
-----

Message to: quit
```

With the -u Option

If you enter **answer -u**, the previous error is treated differently.

```
-----

Message to: albert einstein
address 'a_einstein'
```

```
Enter message for albert einstein ending with a blank line.
```

```
> About your theory ...
>
```

With the -p Option

If you enter **answer -p**, the phone-slip prompts are added. The three lines with no added text are deleted from the message.

```
-----

Message to: George Dancer
address 'g_dancer@cup.hp.com (George B. Dancer)'
```

```
Caller: Harold Maudlin
of:      Terpsichore Inc.
Phone:   123 456 7890
```

```
TELEPHONED      - at 4:30pm
CALLED TO SEE YOU -
WANTS TO SEE YOU - X
RETURNED YOUR CALL -
PLEASE CALL     - X
WILL CALL AGAIN -
*****URGENT***** - very very!
```

```
Enter message for George Dancer ending with a blank line.
```

```
> He said that you would ...
>
```

FILES

| | |
|--------------------------------------|-------------------------------------|
| <code>\$HOME/.elm/aliases</code> | User alias database data table |
| <code>\$HOME/.elm/aliases.dir</code> | User alias database directory table |

| | |
|--|---------------------------------------|
| <code>\$HOME/.elm/aliases.pag</code> | User alias database hash table |
| <code>\$HOME/.elm/aliases.text</code> | User alias source text |
| <code>/var/mail/.elm/aliases</code> | System alias database data table |
| <code>/var/mail/.elm/aliases.dir</code> | System alias database directory table |
| <code>/var/mail/.elm/aliases.pag</code> | System alias database hash table |
| <code>/var/mail/.elm/aliases.text</code> | System alias source text |
| <code>/tmp/snd.pid</code> | Outbound mail message edit buffer |

AUTHOR

`answer` was developed by HP.

SEE ALSO

`elm(1)`, `newaliases(1)`.



a

NAME

ar - create and maintain portable archives and libraries

SYNOPSIS

ar [-*key* [-][*modifier* ...] [*posname*] *afile* [*name* ...]

DESCRIPTION

The **ar** command maintains groups of files combined into a single archive file. Its main use is to create and update library files as used by the link editor (see *ld(1)*). It can be used, however, for any similar purpose. The magic string and file headers used by **ar** consist of printable ASCII characters. If an archive is composed of printable files, the entire archive is printable.

Individual files are inserted without conversion into the archive file. When **ar** creates an archive, it creates headers in a format that is portable across all machines. See *ar(4)* for a detailed description of the portable archive format and structure. The archive symbol table (described in *ar(4)*) is used by the link editor to search repeatedly and efficiently through libraries of object files. An archive symbol table is created and maintained by **ar** only when the archive contains at least one object file. 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 **ar** is used to create or update the contents of an archive, the symbol table is rebuilt (unless the **z** modifier is used). The **s** modifier described below forces the symbol table to be rebuilt.

One *key* operation character from the set, **drqtpmx**, is required and can be optionally preceded by a hyphen (-). The required *key* operation character can be specified with one or more *modifier* characters from the set **abcfFilsuvzACT**. *posname* is used with the **r** and **m** key operations and the **a**, **b**, and **i** modifiers to specify a position in the archive. *afile* is the archive file. Constituent files in the archive file are specified by *name* arguments.

The following list describes the *key* operation characters:

- d** Delete the named files from the archive file.
- r** Replace the named files, or add a new file to the archive:
 - If the **u** modifier is used with the operation character **r**, only those files with modification dates later than those of the corresponding member files are replaced.
 - If an optional positioning character from the set **abi** is used, the *posname* argument must be present and specifies that new files are to be placed after (**a**) or before (**b** or **i**) *posname*. In the absence of a positioning character, new files are placed at the end.
 - **ar** creates *afile* if it does not already exist.
 - If no *name* is specified and:
 - the specified archive file does not exist, **ar** creates an empty archive file containing only the archive header (see *ar(4)*).
 - the archive contains one or more files whose names match names in the current directory, each matching archive file is replaced by the corresponding local file without considering which file may be newer unless the **u** modifier is also specified.
- q** Quickly append the named files to the end of the archive file. Positioning characters are invalid. The operation does not check to determine whether the added members are already in the archive. **ar** creates *afile* if it does not already exist.
- t** Print a table of contents of the archive file to the standard output. If no names are given, all files in the archive are described. If names are given, information about only those files appears.
- p** Print the named files in the archive to the standard output. If no names are specified, the contents of all files are printed in the order that they appear in the archive.
- m** Move the named files. By default, the files are moved to the end of the archive. If a positioning character is present, the *posname* argument must be present and, as in the **r** operation, *posname* specifies where the files are to be moved. Note that, when used with a positioning character, the files are moved in the same order that they currently appear in the archive, *not* in the order specified on the command line. See EXAMPLES.
- x** Extract the named files. If no names are given, all files in the archive are extracted. In neither case does **x** alter entries from the archive file.

The following list describes the optional *modifier* characters:

- a** Position the files after the existing positioning file specified by *posname* .
- b** Place the new files before the existing positioning file specified by *posname* .
- c** Suppress the message normally produced when *afile* is created. For **r** and **q** operations, **ar** normally creates *afile* if it does not already exist.
- f** Truncate the named file names to 14 bytes before performing operations on an archive. This modifier has been provided for compatibility with previous releases where file names up to a maximum of 14 bytes were supported. Longer file names were truncated. When used with the **r** operation, the first existing file that matches the truncated file name is replaced. The **f** modifier can also be used with other operations to allow the full file names to be specified, rather than the truncated file names. Also see the description of the **F** modifier.
- i** Place the new files before the existing positioning file specified by *posname* . Identical to the **b** modifier.
- l** Place temporary files in the local current working directory rather than in the directory specified by the environment variable **TMPDIR** or in the default directory **/var/tmp** . Only the **d**, **m**, **q**, and **r** operations and the **s** and **F** modifiers use temporary files.
- s** Regenerate the archive symbol table even if **ar** is not invoked with an operation that modifies the archive contents. This modifier is useful for restoring the archive symbol table after the **strip** command has been used on the archive (see *strip(1)*) or after the archive has been modified using the **z** modifier.
- u** Update the archive. (**r** operations only) Do not copy the local file to the archive unless the local file is newer than the corresponding existing file in the archive.
- v** Give a verbose file-by-file description of the creation or modification of an archive file to the standard output. When used with **t**, **v** gives a long listing of all information about the files. When used with the **d**, **m**, **p**, **q**, or **x** operations, the verbose modifier causes **ar** to print each *key* operation character and the file name associated with that operation. For the **r** operation, **ar** shows an **a** if it adds a new file or an **r** if it replaces an existing one. For the **p** operation, **ar** prints the name of the file to the standard output before the contents of the file are printed.
- z** Suppress the rebuilding of the symbol table when the archive is modified. This modifier is useful only to avoid long build times when creating a large archive piece-by-piece. If an existing archive contains a symbol table, the **z** modifier will cause it to be invalidated. If a file name longer than 15 bytes is given the entire archive is rewritten. To rebuild the symbol table, either use the **ranlib** command (see *ranlib(1)*), or invoke **ar** again with the **s** modifier.
- A** Suppress warning messages regarding optional access control list entries. **ar** does not archive optional access control list entries in a file's access control list (see *acl(5)*). Normally, a warning message is printed for each file having optional access control list entries.
- C** Prevent extracted files from replacing files with the same name. The **C** modifier can only be used with the **x** operation.
- F** Truncate the entire archive. The **F** modifier causes the entire archive to be rewritten such that all file names within the archive are truncated to 14 bytes, even when **ar** does not modify the archive contents. The long name table will be removed (see *ar(4)*). This modifier has been provided for compatibility with previous releases where file names up to a maximum of 14 bytes were supported. Also see the description of the **f** modifier.
- T** Truncate file names whose archive names are longer than those supported by the file system. By default, files with names longer than those supported by the file system will not be extracted and will cause an error. The **T** modifier can only be used with the **x** operation.

Only the following combinations are meaningful; no other combination of modifiers with operations have any effect on the operation:

- d:** **v**, **f**, **F**, **l**
- m:** **v**, **f**, **F**, **l**, and **a** | **b** | **i**
- r:** **v**, **f**, **F**, **l**, **c**, **A**, **u**, and **a** | **b** | **i**
- q:** **v**, **f**, **F**, **l**, **c**, **A**, **z**, **s**
- t:** **v**, **f**, **F**, **s**

p: v, f, F, s
x: v, f, F, s, C, T

EXTERNAL INFLUENCES

Environment Variables

The following internationalization variables affect the execution of **ar**:

LANG

Determines the locale category for native language, local customs and coded character set in the absence of **LC_ALL** and other **LC_*** environment variables. If **LANG** is not specified or is set to the empty string, a default of **C** (see *lang(5)*) is used instead of **LANG**.

LC_ALL

Determines the values for all locale categories and has precedence over **LANG** and other **LC_*** environment variables.

LC_CTYPE

Determines the locale category for character handling functions.

LC_MESSAGES

Determines the locale that should be used to affect the format and contents of diagnostic messages written to standard error.

LC_NUMERIC

Determines the locale category for numeric formatting.

LC_TIME

Determines the format and contents of date and time formatting.

NLSPATH

Determines the location of message catalogues for the processing of **LC_MESSAGES**.

If any internationalization variable contains an invalid setting, **ar** behaves as if all internationalization variables are set to **C**. See *environ(5)*.

In addition, the following environment variable affects **ar**:

TMPDIR

Specifies a directory for temporary files (see *tmpnam(3S)*). The **l** modifier overrides the **TMPDIR** variable, and **TMPDIR** overrides **/var/tmp**, the default directory.

DIAGNOSTICS

phase error on file name

The named file was modified by another process while **ar** was copying it into the archive. When this happens, **ar** exits and the original archive is left untouched.

ar write error: file system error message

ar could not write to a temporary file or the final output file. If **ar** was trying to write the final output file, the original archive is lost.

ar reports **cannot create file.a**, where *file.a* is an **ar**-format archive file, even if *file.a* already exists. This message is triggered when *file.a* is write-protected or inaccessible.

EXAMPLES

Create a new file (if one does not already exist) in archive format with its constituents entered in the order indicated:

```
ar r newlib.a f3 f2 f1 f4
```

Replace files **f2** and **f3** such that the new copies follow file **f1**, and **f3** follows **f2**:

```
ar ma f1 newlib.a f2 f3
ar ma f2 newlib.a f3
ar r newlib.a f2 f3
```

The archive is then ordered:

```
newlib.a: f1 f2' f3' f4
```

where the single quote marks indicate updated files. The first command says "move **f2** and **f3** after **f1** in *newlib.a*", thus creating the order:

```
f1 f3 f2 f4
```

Note that the relative order of **f2** and **f3** has not changed. The second command says "move **f3** after **f2** in *newlib.a*", creating the order:

```
f1 f2 f3 f4
```

The third command then replaces files **f2** and **f3**. Since files **f2** and **f3** both already existed in the archive, this sequence of commands could not be simply replaced by:

```
ar ra f1 newlib.a f2 f3
```

because the previous position and relative order of **f2** and **f3** in the archive are preserved (no matter how the files are specified on the command line), producing the following archive:

```
newlib.a: f3' f2' f1 f4
```

WARNINGS

If you are a user who has appropriate privileges, **ar** can alter any archive file, even if it is write-protected.

If the same file is mentioned twice in an argument list, it might be put in the archive twice.

If multiple copies of a file exist in an archive, **ar** matches the first occurrence of the file in the archive.

ar automatically creates an archive symbol table, a task performed in early HP-UX versions by **ranlib**. Use of the **z** modifier either suppresses generation of the symbol table, or invalidates it if it exists. The **ranlib** command can be used to rebuild the symbol table if an archive was built with the **z** modifier.

FILES

/var/tmp/ar* Temporary files

SEE ALSO

System Tools:

ld(1) Invoke the link editor

Miscellaneous:

| | |
|-------------------|--|
| acl(5) | Access control lists |
| a.out(4) | Assembler, compiler, and linker output |
| ar(4) | Archive format |
| lorder(1) | Find the ordering relation for object files or archive libraries |
| ranlib(1) | Regenerate an archive symbol table |
| strip(1) | Strip symbol and line number information from an object file |
| tmpnam(3S) | Create a name for a temporary file |

Texts and Tutorials:

HP-UX Linker and Libraries Online User Guide
(See the **+help** option)
HP-UX Linker and Libraries User's Guide
(See *manuals(5)* for ordering information)

STANDARDS CONFORMANCE

ar: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

as - assembler (Precision Architecture)

SYNOPSIS

as *[[option]... [file] ...]...*

DESCRIPTION

The **as** command assembles source text from files or standard input and produces a relocatable object file suitable for the link editor, **ld** (see *ld(1)*).

Source text is read from standard input only if no *file* argument is given. Standard input cannot be a device file, such as a terminal. The *option* and *file* arguments can be intermingled on the command line. Every specified *option* applies to every specified *file*, or standard input. The source files are concatenated to form a single input stream.

If the **-o** *objfile* option is not specified, the **.s** suffix (if any) is stripped from the end of the last source file name and **.o** is appended to the name to form the name of the default object code output file.

as output is not optimized. **as** creates a relocatable object file that must be processed by **ld** before it can be successfully executed (see *ld(1)*).

The **cc** compiler normally runs the C preprocessor **cpre** (see *cpre(1)*), then invokes **as** to assemble the **.s** files together with **/usr/lib/pcc_prefix.s**, and subsequently invokes **ld**.

Arguments

as recognizes the following arguments.

| | |
|---------------|--|
| <i>file</i> | A text file contain assembler source code. |
| <i>option</i> | An option described below in Options. |

Options

as recognizes the following values for the *option* argument.

| | |
|--------------------------------|--|
| -e | Permit an unlimited number of errors to be tolerated before the assembly process is abandoned. By default, one hundred errors are allowed before the assembler aborts. |
| -f | Set the default value for the .CALLINFO directive to CALLS . The normal default value for a .CALLINFO that omits the CALLER , CALLS or NO_CALLS parameter is NO_CALLS . |
| -l | Write a listing of the program to standard output after assembly. This listing shows the offsets of instructions and actual values for fields. |
| -o <i>objfile</i> | Name the output object file <i>objfile</i> instead of using the default .o suffix on the file name of the last <i>file</i> specified. |
| -p <i>number</i> | Set the default privilege level for an .EXPORT directive to <i>number</i> . By default, all user-level procedures are exported at privilege level 3. |
| -s | Set the output file suffix to .ss instead of .o . The file will have a format suitable for conversion to the ROM burning programs. |
| -u | Do not create unwind descriptors. To avoid the need for the .CALLINFO directive, the .ENTER and .LEAVE directives must not have been used. |
| -v <i>xrfile</i> | Write cross-reference data to the file named <i>xrfile</i> . |
| -V | Print the version number of the assembler program to standard error before assembling the source text. |
| -w [<i>number</i>] | Either suppress all warning messages if no <i>number</i> is supplied or suppress just the warning <i>number</i> provided. Multiple -wnumber options can be used to suppress additional warning messages. |
| +DA <i>architecture</i> | Assemble code for the <i>architecture</i> specified. The use of this option is discouraged. The preferred method for selecting the <i>architecture</i> is to have a .LEVEL directive contained within the assembly source file. |

The assembler uses the following precedence to determine the target architecture.

1. Use the **.LEVEL** directive within the assembly source file.
2. Use the **+DA** command-line specification.
3. Use the default architecture of **PA_RISC_1.0**.

+z, +Z

Both of these options are used in the building of shared libraries. For a more complete discussion regarding these options, see the manual *HP-UX Linker and Libraries User's Guide*.

EXTERNAL INFLUENCES

International Code Set Support

Single- and multibyte character code sets are supported.

DIAGNOSTICS

When syntactic or semantic errors occur, a single-line diagnostic is written to standard error, that includes the file name and the line number where it occurred. The format is as follows:

```
as: "filename",line line: error error: message
source = source-line
```

WARNINGS

as does not invoke *cpp*(1) or *m4*(1) to perform macro processing.

FILES

| | |
|--|--|
| <code>/usr/include/hard_reg.hr</code> | Hardware register definitions |
| <code>/usr/include/soft_reg.h</code> | Software calling convention register definitions |
| <code>/usr/include/std_space.h</code> | Standard space and subspace definition |
| <code>/usr/lib/nls/msg/C/as.cat</code> | Assembler error message catalog |
| <code>/usr/lib/pcc_prefix.s</code> | Space, subspace and register definitions |
| <code>file.o</code> | Object file |

SEE ALSO

`adb`(1), `cc_bundled`(1), `ld`(1), `nm`(1), `crt0`(3).

Assembly Language Reference Manual,
Precision Architecture and Instruction Reference Manual,
Procedure Calling Conventions Reference Manual.

NAME

asa - interpret ASA carriage control characters

SYNOPSIS

asa [*files*]

DESCRIPTION

asa interprets the output of FORTRAN programs that utilize ASA carriage control characters. It processes either the *files* whose names are given as arguments, or the standard input if - is specified or if no file names are given. The first character of each line is assumed to be a control character. The following control characters are interpreted as indicated:

- (blank) Output a single new-line character before printing.
- (space) (XPG4 only.) The rest of the line will be output without change.
- 0 Output two new-line characters before printing.
- 0 (XPG4 only.) A <newline> shall be output, then the rest of the input line.
- 1 Output a new-page character before printing.
- + Overprint previous line.
- + (XPG4 only.) The <newline> of the previous line shall be replaced with one or more implementation-defined characters that causes printing to return to column position 1, followed by the rest of the input line. If the + is the first character in the input, it shall have the same effect as <space>.

Lines beginning with other than the above characters are treated the same as lines beginning with a blank. The first character of a line is *not* printed. If any such lines appear, an appropriate diagnostic is sent to standard error. This program forces the first line of each input file to start on a new page.

(XPG4 only.) The action of the **asa** utility is unspecified upon encountering any character other than those listed above as the first character in a line.

To view the output of FORTRAN programs which use ASA carriage control characters and have them appear in normal form, **asa** can be used as a filter:

```
a.out | asa | lp
```

The output, properly formatted and paginated, is then directed to the line printer. FORTRAN output previously sent to a file can be viewed on a user terminal screen by using:

```
asa file
```

EXTERNAL INFLUENCES**Environment Variables**

LC_CTYPE determines the interpretation of text within file as single- and/or multi-byte characters.

LC_MESSAGES determines the language in which messages are displayed.

If **LC_CTYPE** or **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang*(5)) is used instead of **LANG**.

If any internationalization variable contains an invalid setting, **asa** behaves as if all internationalization variables are set to "C". See *environ*(5).

International Code Set Support

Single- and multi-byte character code sets are supported.

SEE ALSO

efl(1), *f77*(1), *fsplit*(1), *ratfor*(1).

STANDARDS CONFORMANCE

asa: XPG4, POSIX.2

NAME

at, batch - execute batched commands immediately or at a later time

SYNOPSIS

Enter commands from standard input to run at a specified time:

```
at [-m] [-q queue] -t spectime
commands
eof
```

```
at [-m] [-q queue] time [date] [next timeunit | +count timeunit]
commands
eof
```

Enter commands from a file to run at a specified time:

```
at -f job-file [-m] [-q queue] -t spectime
```

```
at -f job-file [-m] [-q queue] time [date] [next timeunit | +count timeunit]
```

List scheduled jobs:

```
at -d job-id ...
```

```
at -l [job-id ...]
```

```
at -l -q queue
```

Cancel (remove) a scheduled job:

```
at -r job-id ...
```

Enter commands from standard input to run as a batch process:

```
batch
commands
eof
```

Enter commands from a file to run as a batch process:

```
batch < job-file
```

DESCRIPTION

The **at** and **batch** commands schedule jobs for execution by the **cron** daemon (see *cron(1M)*).

at schedules a job for execution at a specified time. **at** can also list (-l) or remove (-r) existing scheduled **at** and **batch** jobs.

batch schedules a job for execution immediately, or as soon as system load levels permit.

You can enter commands into a job in one of the following ways:

- From the keyboard on separate lines immediately after the **at** or **batch** command line, followed by the currently defined *eof* (end-of-file) character to end the input. The default *eof* is Ctrl-D. It can be redefined in your environment (see *stty(1)*).
- With the -f option of the **at** command to read input from a script file.
- From output piped from a preceding command.

Options and Arguments

at recognizes the following options and arguments.

| | |
|----------------------|--|
| <i>commands</i> | One or more HP-UX commands that can be executed as a shell script by at or batch . |
| <i>eof</i> | End-of-file character. The default is Ctrl-D unless defined otherwise in your environment. |
| <i>job-file</i> | The path name of an existing file. |
| <i>job-id</i> | The job identifier reported by at or batch when the job was originally scheduled. |
| -d <i>job-id</i> ... | Displays the contents of the specified job. An unprivileged user is restricted to display information only on jobs that the user owns. A user with the appropriate |

privileges is able to display information about all jobs.

- f** *job-file* Read in the commands contained in *job-file* instead of using standard input.
- l** [*job-id* ...] List the jobs specified. If no *job-ids* are given, all jobs are listed.
- m** Send mail to the invoking user after the job has run, announcing its completion. Unless redirected elsewhere within the job, standard output and standard error produced by the job are automatically mailed to the user as well.
- q** *queue* Submit the specified job to the *queue* indicated (see *queuedefs*(4)). Queues **a**, **b**, and **d** through **y** can be used. **at** uses queue **a** by default. **batch** always uses queue **b**. All queues except **b** require a *time* or a **-t** specification. **at -qb** is equivalent to **batch**.
- r** *job-id* ... Remove the jobs specified by each *job-id*.
- t** *spectime* Define the absolute time to start the job.

spectime A date and time in the format:

[[*CC*]*YY*]*MMDDhhmm* [. *ss*]

where the decimal digit pairs are as follows:

CC The first two digits of the year (19, 20).

YY The second two digits of the year (69–99, 00–68). See WARNINGS.

MM The month of the year (01–12).

DD The day of the month (01–31).

hh The hour of the day (00–23).

mm The minute of the hour (00–59).

ss The second of the minute (00–61).

If both *CC* and *YY* are omitted, the default is the current year.

If *CC* is omitted and *YY* is in the range 69–99, *CC* defaults to 19. Otherwise, *CC* defaults to 20.

The range for *ss* provides for two leap seconds. If *ss* is 60 or 61, and the resulting time, as affected by the **TZ** environment variable, does not refer to a leap second, the time is set to the whole minute following *mm*.

If *ss* is omitted, it defaults to 00.

time [*date*] Define the base time for starting the job.

time A time specified as one, two, or four digits. One- and two-digit numbers represent hours; four digits represent hours and minutes.

Alternately, *time* can be specified as two numbers separated by a colon (:), a single quote ('), the letter *h* (**h**), a period (.), or a comma (,). Spaces may be present between the separator and digits representing minutes. If defined in *langinfo*(5), special time unit characters can be used.

am or **pm** can be appended to indicate morning or afternoon. Otherwise, a 24-hour clock is understood. For example, 0815, 8:15, 8'15, 8h15, 8.15, and 8,15 are read as 15 minutes after eight in the morning. The suffixes **zulu** and **utc** can be used to specify Coordinated Universal Time (UTC), equivalent to Greenwich Mean Time (GMT).

The special names **midnight**, **noon**, and **now** are also recognized.

date A day of the week (fully spelled out or abbreviated) or a date consisting of a day, a month, and optionally a year. The day and year fields must be numeric, and the month can be either fully spelled out, abbreviated, or numeric. The fields in the date string are separated by punctuation marks such as slash (/), hyphen (-), period (.), and comma (,). If defined in *langinfo*(5), special date unit characters can

be present. A field having a value greater than 31 is treated as the year field and the remaining two fields in the date string are treated as month and day fields. Otherwise, if a given date is ambiguous (such as 2/5 or 2/5/10), the `D_T_FMT` string (if defined in *langinfo(5)*) is used to resolve the ambiguity.

Two special days, **today** and **tomorrow**, are also recognized. If no *date* is given, **today** is assumed if the given time is greater than the current time; **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. Two-digit years in the range 69 to 99 are expanded to 1969 to 1999; in the range 00 to 68, to 2000 to 2068.

`next timeunit | + count timeunit`

Delay the execution date and time by a specific number of time units after the base time specified by *time* [*date*].

count A decimal number. **next** is equivalent to +1.

timeunit A time unit, one of the following: **minutes**, **hours**, **days**, **weeks**, **months**, or **years**, or their singular forms.

How Jobs Are Processed

When a job is accepted, **at** and **batch** print a message to standard error in the form:

```
job job-id at execution-date
```

where *job-id* is the job identifier in the form *jobnumber.queue*, such as **756284400.a**, and *execution-date* is the date and time when the job will be released for execution.

If your login shell is not the POSIX shell (`/usr/bin/sh`), the commands also print a warning message:

```
warning: commands will be executed using /usr/bin/sh
```

at jobs default to queue **a**. **batch** jobs always go in queue **b**. See the **-q** option.

An **at** or **batch** job consists of a two-part script stored in `/var/spool/cron/atjobs` that can be executed by the POSIX shell.

The first part sets up the environment to match the environment when the **at** or **batch** command was issued. This includes the current shell environment variables, current directory, **umask**, and **ulimit** (see *ulimit(2)*, *umask(1)*, and *proto(4)*). Open file descriptors, traps, and priority are lost.

The second part consists of the commands that you entered.

When **cron** dispatches the job, it starts a POSIX shell to execute the script.

The number of jobs executing from a queue at any time is controlled by parameters in the file `/var/adm/cron/queuedefs` (see *queuedefs(4)*).

Standard output and standard error from the job are mailed to the user unless they are redirected elsewhere within the job.

Scheduled jobs are immune to the **SIGHUP** hangup signal, and remain scheduled if the user logs off.

Users are permitted to use the **at** and **batch** commands if their user names appear in the file `/usr/lib/cron/at.allow`. If that file does not exist, users can use **at** and **batch** if their names *do not* appear in the file `/usr/lib/cron/at.deny`. If neither file exists, only superuser is allowed to submit jobs. If only `at.deny` exists but is empty, all users can use **at** and **batch**. The *allow/deny* files consist of one user name per line.

All users can list and remove their own jobs. Users with appropriate privileges can list and remove jobs other than their own.

EXTERNAL INFLUENCES

Environment Variables

LC_TIME determines the format and contents of date and time strings.

LC_MESSAGES determines the language in which messages are displayed.

`LC_MESSAGES` also determines the language in which the words `days`, `hours`, `midnight`, `minutes`, `months`, `next`, `noon`, `now`, `today`, `tomorrow`, `weeks`, `years`, and their singular forms can also be specified.

If `LC_TIME` or `LC_MESSAGES` is not specified in the environment or is set to the empty string, the value of `LANG` is used as a default for each unspecified or empty variable. If `LANG` is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of `LANG`.

If any internationalization variable contains an invalid setting, all internationalization variables default to "C" (see *environ(5)*).

International Code Set Support

Single- and multi-byte character code sets are supported.

RETURN VALUE

The exit code is set to one of the following:

- 0 Successful completion
- 1 Failure

DIAGNOSTICS

`at` produces self-explanatory messages for syntax errors and out-of-range times.

warning: commands will be executed using /usr/bin/sh

If your login shell is not the POSIX shell (`/usr/bin/sh`), `at` and `batch` produce a warning message as a reminder that `at` and `batch` jobs are executed using `/usr/bin/sh`.

EXAMPLES

The following commands show three different ways to run a POSIX shell script file named `delayed-job` five minutes from now:

```
at -f delayed-job now + 5 minutes
cat delayed-job | at now + 5 minutes
at now + 5 minutes <delayed-job
```

Run a typical HP-UX command (`nroff` in this case) when system load levels permit, and redirect standard output and standard error to files:

```
batch
nroff source-file >output-file 2>error-file
eof      (the default is Ctrl-D)
```

Run a job contained in `future` in the home directory at 12:20 a.m. on December 27, 2013:

```
at -f $HOME/future -t201312271220.00
```

Redirect standard error to a pipe (useful in a shell procedure). Note that the sequence of the output redirection specifications is significant. Standard error is redirected to where standard output is going; standard output is redirected to a file; the original "standard output" (which now consists of the former standard error) is piped to the `mail` program.

```
batch <<!!      (sets eof temporarily to !!)
nroff input-file 2>&1 1> output-file | mail loginid
!!
```

Run a job contained in `jobfile` in the home directory at 5:00 a.m. next Tuesday:

```
at -f $HOME/jobfile 5am tuesday next week
```

Run the same job at 5:00 a.m. one week from next Tuesday (i.e., 2 Tuesdays in advance):

```
at -f $HOME/jobfile 5am tuesday + 2 weeks
```

Add a command to the file named `weekly-run` in directory `jobs` in the home directory so that it automatically reschedules itself every time it runs. This example reschedules itself every Thursday at 1900 (7:00 p.m.):

```
echo "sh $HOME/jobs/weekly-run" | at 1900 thursday next week
```

The following commands show several forms recognized by `at` and include native language usage:

```

at 0815 Jan 24
at 8:15 Jan 24
at 9:30am tomorrow
at now + 1 day
at -f job 5 pm Friday
at 17:40 Tor.           # in Danish
at 17h46 demain        # in French
at 5:30 26. Feb. 1988  # in German
at 12:00 26-02         # in Finnish

```

WARNINGS

If the *date* argument begins with a number and the *time* argument is also numeric without a suffix, the *time* argument should be a four-digit number that can be correctly interpreted as hours and minutes.

If you use both **next** and **+count** within a single **at** command, the first operator is accepted and the trailing operator is silently ignored.

If you use both **-t** and *time* ... in the same command, the first specified is accepted and the second is silently ignored.

If the FIFO used to communicate with **cron** fills up, **at** is suspended until **cron** has read sufficient messages from the FIFO to make room for the message **at** is trying to write. This condition can occur if **at** is writing messages faster than **cron** can process them or if **cron** is not executing.

Scheduled processes are run in the background. Any script file that calls itself will cause the user or the system to run out of available processes.

If the execution-time request for a job duplicates the execution time of a currently scheduled job, the new job time is set to the next available second.

at will not schedule jobs whose start time precedes the current Epoch (00:00:00 January 1, 1970 UTC).

at will not schedule jobs beyond the year 2037.

DEPENDENCIES

HP Process Resource Manager

If the optional HP Process Resource Management (PRM) software is installed and configured, jobs are launched in the initial process resource group of the user that scheduled the job. The user's initial group is determined at the time the job is started, not when the job is scheduled. If the user's initial group is not defined, the job runs in the user default group (**PRMID=1**). See *prmconfig(1)* for a description of how to configure HP PRM, and *prmconf(4)* for a description of how the user's initial process resource group is determined.

AUTHOR

at was developed by AT&T and HP.

FILES

| | |
|--------------------------------|---|
| /usr/bin/sh | POSIX shell |
| /var/adm/cron | Main cron directory |
| /var/adm/cron/.proto | This file contains a set of shell commands which are added to the at job file to make the environment for the at job same as the current environment. See <i>proto(4)</i> . |
| /usr/lib/cron/at.allow | List of allowed users |
| /usr/lib/cron/at.deny | List of denied users |
| /var/adm/cron/queuedefs | Scheduling information |
| /var/spool/cron/atjobs | Spool area |

SEE ALSO

crontab(1), kill(1), mail(1), nice(1), ps(1), sh(1), stty(1), cron(1M), proto(4), queuedefs(4).

HP Process Resource Manager:

prmconfig(1), prmconf(4) in *HP Process Resource Manager User's Guide*.

at(1)

at(1)

STANDARDS CONFORMANCE

at: SVID2, SVID3, XPG2, XPG3, XPG4

batch: SVID2, SVID3, XPG2, XPG3, XPG4

a

NAME

`attributes` - describe an audio file

SYNOPSIS

`/opt/audio/bin/attributes filename`

DESCRIPTION

This command provides information about an audio file, including file format, data format, sampling rate, number of channels, data length and header length.

EXAMPLE

The following is an example of using `attributes` on an audio file supplied with HP-UX.

```
$ /opt/audio/bin/attributes /opt/audio/sounds/welcome.au
File Name: /opt/audio/sounds/welcome.au
File Type: NeXT/Sun
Data Format: Mu-law
Sampling Rate: 22050
Channels: Mono
Duration: 1.972 seconds
Bits per Sample: 8
Header Length: 40 bytes
Data Length: 43492 bytes
```

AUTHOR

`attributes` was developed by HP.

Sun is a trademark of Sun Microsystems, Inc.

NeXT is a trademark of NeXT Computers, Inc.

SEE ALSO

`audio(5)`, `asecure(1M)`, `aserver(1M)`, `convert(1)`, `send_sound(1)`.

Using the Audio Developer's Kit

NAME

awk - pattern-directed scanning and processing language

SYNOPSIS

awk [-F*fs*] [-v *var=value*] [*program* | -f *progfile* ...] [*file* ...]

DESCRIPTION

awk scans each input *file* for lines that match any of a set of patterns specified literally in *program* or in one or more files specified as -f *progfile*. With each pattern there can be an associated action that is to be performed when a line in a *file* matches the pattern. Each line is matched against the pattern portion of every pattern-action statement, and the associated action is performed for each matched pattern. The file name - means the standard input. Any *file* of the form *var=value* is treated as an assignment, not a filename. An assignment is evaluated at the time it would have been opened if it were a filename, unless the -v option is used.

An input line is made up of fields separated by white space, or by regular expression **FS**. The fields are denoted \$1, \$2, ...; \$0 refers to the entire line.

Options

awk recognizes the following options and arguments:

- F *fs* Specify regular expression used to separate fields. The default is to recognize space and tab characters, and to discard leading spaces and tabs. If the -F option is used, leading input field separators are no longer discarded.
- f *progfile* Specify an awk program file. Up to 100 program files can be specified. The pattern-action statements in these files are executed in the same order as the files were specified.
- v *var=value* Cause *var=value* assignment to occur before the **BEGIN** action (if it exists) is executed.

Statements

A pattern-action statement has the form:

```
pattern { action }
```

A missing { *action* } means print the line; a missing pattern always matches. Pattern-action statements are separated by new-lines or semicolons.

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

```
if( expression ) statement [ else statement ]
while( expression ) statement
for( expression ; expression ; expression ) statement
for( var in array ) statement
do statement while( expression )
break
continue
{ [statement ...] }
expression # commonly var = expression
print [expression-list] [ > expression ]
printf format [, expression-list] [ > expression ]
return [expression]
next # skip remaining patterns on this input line.
delete array [expression] # delete an array element.
exit [expression] # exit immediately; status is expression.
```

Statements are terminated by semicolons, newlines or right braces. An empty *expression-list* stands for \$0. String constants are quoted (" "), with the usual C escapes recognized within. Expressions take on string or numeric values as appropriate, and are built using the operators +, -, *, /, %, ^ (exponentiation), and concatenation (indicated by a blank). The operators ++, --, +=, -=, *=, /=, %=, ^=, **=, >, >=, <, <=, ==, !=, and ? : are also available in expressions. Variables can be scalars, array elements (denoted x[i]) or fields. Variables are initialized to the null string. Array subscripts can 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**.

The **print** statement prints its arguments on the standard output (or on a file if *>file* or *>>file* is present or on a pipe if *|cmd* is present), separated by the current output field separator, and terminated by the output record separator. *file* and *cmd* can be literal names or parenthesized expressions. Identical string values in different statements denote the same open file. The **printf** statement formats its expression list according to the format (see *printf(3)*).

Built-In Functions

The built-in function **close**(*expr*) closes the file or pipe *expr* opened by a **print** or **printf** statement or a call to **getline** with the same string-valued *expr*. This function returns zero if successful, otherwise, it returns non-zero.

The customary functions **exp**, **log**, **sqrt**, **sin**, **cos**, **atan2** are built in. Other built-in functions are:

| | |
|---|---|
| length [(<i>s</i>)] | Length of its associated argument (in bytes) taken as a string, or of \$0 if no argument. |
| length [(<i>s</i>)] | Length of its associated argument (in characters) taken as a string, or of \$0 if no argument. |
| rand () | Returns a random number between zero and one. |
| srand ([<i>expr</i>]) | Sets the seed value for <i>rand</i> , and returns the previous seed value. If no argument is given, the time of day is used as the seed value; otherwise, <i>expr</i> is used. |
| int (<i>x</i>) | Truncates to an integer value |
| substr (<i>s</i> , <i>m</i> [, <i>n</i>]) | Return the at most <i>n</i> -character substring of <i>s</i> that begins at position <i>m</i> , numbering from 1. If <i>n</i> is omitted, the substring is limited by the length of string <i>s</i> . |
| index (<i>s</i> , <i>t</i>) | Return the position, in characters, numbering from 1, in string <i>s</i> where string <i>t</i> first occurs, or zero if it does not occur at all. |
| match (<i>s</i> , <i>ere</i>) | Return the position, in characters, numbering from 1, in string <i>s</i> where the extended regular expression <i>ere</i> occurs, or 0 if it does not. The variables RSTART and RLENGTH are set to the position and length of the matched string. |
| split (<i>s</i> , <i>a</i> [, <i>fs</i>]) | Splits the string <i>s</i> into array elements <i>a</i> [1], <i>a</i> [2], ..., <i>a</i> [<i>n</i>], and returns <i>n</i> . The separation is done with the regular expression <i>fs</i> , or with the field separator FS if <i>fs</i> is not given. |
| sub (<i>ere</i> , <i>repl</i> [, <i>in</i>]) | Substitutes <i>repl</i> for the first occurrence of the extended regular expression <i>ere</i> in the string <i>in</i> . If <i>in</i> is not given, \$0 is used. |
| gsub | Same as sub except that all occurrences of the regular expression are replaced; sub and gsub return the number of replacements. |
| sprintf (<i>fnt</i> , <i>expr</i> , ...) | String resulting from formatting <i>expr</i> ... according to the <i>printf</i> (3S) format <i>fnt</i> |
| system (<i>cmd</i>) | Executes <i>cmd</i> and returns its exit status |
| toupper (<i>s</i>) | Converts the argument string <i>s</i> to uppercase and returns the result. |
| tolower (<i>s</i>) | Converts the argument string <i>s</i> to lowercase and returns the result. |

The built-in function **getline** sets **\$0** to the next input record from the current input file; **getline <file** sets **\$0** to the next record from *file*. **getline x** sets variable *x* instead. Finally, *cmd* | **getline** pipes the output of *cmd* into **getline**; each call of **getline** returns the next line of output from *cmd*. In all cases, **getline** returns 1 for a successful input, 0 for end of file, and -1 for an error.

Patterns

Patterns are arbitrary Boolean combinations (with **!** | **&&**) of regular expressions and relational expressions. **awk** supports Extended Regular Expressions as described in *regex(5)*. Isolated regular expressions in a pattern apply to the entire line. Regular expressions can also occur in relational expressions, using the operators **~** and **!~**. */re/* is a constant regular expression; any string (constant or variable) can be used as a regular expression, except in the position of an isolated regular expression in a pattern.

a

A pattern can consist of two patterns separated by a comma; in this case, the action is performed for all lines from an occurrence of the first pattern though an occurrence of the second.

A relational expression is one of the following:

expression matchop regular-expression
expression relop expression
expression in array-name
(expr,expr,...) in array-name

where a relop is any of the six relational operators in C, and a matchop is either `~` (matches) or `!~` (does not match). A conditional is an arithmetic expression, a relational expression, or a Boolean combination of the two.

The special patterns **BEGIN** and **END** can be used to capture control before the first input line is read and after the last. **BEGIN** and **END** do not combine with other patterns.

Special Characters

The following special escape sequences are recognized by **awk** in both regular expressions and strings:

| Escape | Meaning |
|--------------------|--------------------------------------|
| <code>\a</code> | alert character |
| <code>\b</code> | backspace character |
| <code>\f</code> | form-feed character |
| <code>\n</code> | new-line character |
| <code>\r</code> | carriage-return character |
| <code>\t</code> | tab character |
| <code>\v</code> | vertical-tab character |
| <code>\nnn</code> | 1- to 3-digit octal value <i>nnn</i> |
| <code>\xhhh</code> | 1- to n-digit hexadecimal number |

Variable Names

Variable names with special meanings are:

| | |
|-----------------|--|
| FS | Input field separator regular expression; a space character by default; also settable by option <code>-Ffs</code> . |
| NF | The number of fields in the current record. |
| NR | The ordinal number of the current record from the start of input. Inside a BEGIN action the value is zero. Inside an END action the value is the number of the last record processed. |
| FNR | The ordinal number of the current record in the current file. Inside a BEGIN action the value is zero. Inside an END action the value is the number of the last record processed in the last file processed. |
| FILENAME | A pathname of the current input file. |
| RS | The input record separator; a newline character by default. |
| OFS | The print statement output field separator; a space character by default. |
| ORS | The print statement output record separator; a newline character by default. |
| OFMT | Output format for numbers (default <code>%.6g</code>). If the value of OFMT is not a floating-point format specification, the results are unspecified. |
| CONVFMT | Internal conversion format for numbers (default <code>%.6g</code>). If the value of CONVFMT is not a floating-point format specification, the results are unspecified. |
| SUBSEP | The subscript separator string for multi-dimensional arrays; the default value is <code>"\034"</code> |
| ARGC | The number of elements in the ARGV array. |
| ARGV | An array of command line arguments, excluding options and the <i>program</i> argument numbered from zero to ARGC -1. |

The arguments in **ARGV** can be modified or added to; **ARGC** can be altered. As each input file ends, **awk** will treat the next non-null element of **ARGV**, up to the current value of **ARGC**-1, inclusive, as the name of the next input file. Thus, setting

an element of **ARGV** to null means that it will not be treated as an input file. The name **-** indicates the standard input. If an argument matches the format of an *assignment* operand, this argument will be treated as an assignment rather than a *file* argument.

| | |
|----------------|---|
| ENVIRON | Array of environment variables; subscripts are names. For example, if environment variable V=thing , ENVIRON["V"] produces thing . |
| RSTART | The starting position of the string matched by the match function, numbering from 1. This is always equivalent to the return value of the match function. |
| RLENGTH | The length of the string matched by the match function. |

Functions can be defined (at the position of a pattern-action statement) as follows:

```
function foo(a, b, c) { ...; return x }
```

Parameters are passed by value if scalar, and by reference if array name. Functions can be called recursively. Parameters are local to the function; all other variables are global.

Note that if pattern-action statements are used in an HP-UX command line as an argument to the **awk** command, the pattern-action statement must be enclosed in single quotes to protect it from the shell. For example, to print lines longer than 72 characters, the pattern-action statement as used in a script (**-f progfile** command form) is:

```
length > 72
```

The same pattern action statement used as an argument to the **awk** command is quoted in this manner:

```
awk 'length > 72'
```

EXTERNAL INFLUENCES

Environment Variables

| | |
|--------------------|---|
| LANG | Provides a default value for the internationalization variables that are unset or null. If LANG is unset or null, the default value of "C" (see <i>lang(5)</i>) is used. If any of the internationalization variables contains an invalid setting, awk will behave as if all internationalization variables are set to "C". See <i>environ(5)</i> . |
| LC_ALL | If set to a non-empty string value, overrides the values of all the other internationalization variables. |
| LC_CTYPE | Determines the interpretation of text as single and/or multi-byte characters, the classification of characters as printable, and the characters matched by character class expressions in regular expressions. |
| LC_NUMERIC | Determines the radix character used when interpreting numeric input, performing conversion between numeric and string values and formatting numeric output. Regardless of locale, the period character (the decimal-point character of the POSIX locale) is the decimal-point character recognized in processing awk programs (including assignments in command-line arguments). |
| LC_COLLATE | Determines the locale for the behavior of ranges, equivalence classes and multi-character collating elements within regular expressions. |
| LC_MESSAGES | Determines the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output. |
| NLSPATH | Determines the location of message catalogues for the processing of LC_MESSAGES . |
| PATH | Determines the search path when looking for commands executed by system(cmd) , or input and output pipes. |

In addition, all environment variables will be visible via the **awk** variable **ENVIRON**.

International Code Set Support

Single- and multi-byte character code sets are supported except that variable names must contain only ASCII characters and regular expressions must contain only valid characters.

DIAGNOSTICS

awk supports up to 199 fields (\$1, \$2, ..., \$199) per record.

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:

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

Print all lines between start/stop pairs:

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

Simulate **echo** command (see *echo(1)*):

```
BEGIN {
    # Simulate echo(1)
    for (i = 1; i < ARGV; i++) printf "%s ", ARGV[i]
    printf "\n"
    exit }
```

AUTHOR

awk was developed by AT&T, IBM, OSF, and HP.

SEE ALSO

lex(1), sed(1).

A. V. Aho, B. W. Kernighan, P. J. Weinberger: *The AWK Programming Language*, Addison-Wesley, 1988.

STANDARDS CONFORMANCE

awk: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

banner - make posters in large letters

SYNOPSIS

banner *strings*

DESCRIPTION

banner prints its arguments (each up to 10 characters long) in large letters on the standard output.

Each argument is printed on a separate line. Note that multiple-word arguments must be enclosed in quotes in order to be printed on the same line.

EXAMPLES

Print the message "Good luck Susan" in large letters on the screen:

```
banner "Good luck" Susan
```

The words **Good luck** are displayed on one line, and **Susan** is displayed on a second line.

WARNINGS

This command is likely to be withdrawn from X/Open standards. Applications using this command might not be portable to other vendors' platforms.

SEE ALSO

echo(1).

STANDARDS CONFORMANCE

banner: SVID2, SVID3, XPG2, XPG3

b

NAME

basename, dirname - extract portions of path names

SYNOPSIS

basename *string* [*suffix*]

dirname [*string*]

DESCRIPTION

basename deletes any prefix ending in `/` and the *suffix* (if present in *string*) from *string*, and prints the result on the standard output. If *string* consists entirely of slash characters, *string* is set to a single slash character. If there are any trailing slash characters in *string*, they are removed. If the suffix operand is present but not identical to the characters remaining in *string*, but it is identical to a suffix of the characters remaining in *string*, the suffix is removed from *string*. **basename** is normally used inside command substitution marks (``...``) within shell procedures.

dirname delivers all but the last level of the path name in *string*. If *string* does not contain a directory component, **dirname** returns `.`, indicating the current working directory.

EXTERNAL INFLUENCES

Environment Variables

LC_CTYPE determines the interpretation of string and, in the case of **basename**, suffix as single and/or multi-byte characters.

If **LC_CTYPE** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **basename** and **dirname** behave as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single- and multi-byte character code sets are supported.

EXAMPLES

The following shell script, invoked with the argument `/usr/src/cmd/cat.c`, compiles the named file and moves the output to a file named `cat` in the current directory:

```
cc $1
mv a.out `basename $1 .c`
```

The following example sets the shell variable **NAME** to `/usr/src/cmd`:

```
NAME=`dirname /usr/src/cmd/cat.c`
```

RETURNS

basename and **dirname** return one of the following values:

- 0** Successful completion.
- 1** Incorrect number of command-line arguments.

SEE ALSO

expr(1), *sh(1)*.

STANDARDS CONFORMANCE

basename: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

dirname: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

bc - arbitrary-precision arithmetic language

SYNOPSIS

bc [-**c**] [-**l**] [*file ...*]

DESCRIPTION

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.

Options:

bc recognizes the following command-line options:

- c** Compile only. **bc** is actually a preprocessor for **dc** which **bc** invokes automatically (see *dc(1)*). Specifying **-c** prevents invoking *dc*, and sends the *dc* input to standard output.
- l** causes an arbitrary-precision math library to be predefined. As a side effect, the scale factor is set.

Program Syntax:

L a single letter in the range **a** through **z**;
E expression;
S statement;
R relational expression.

Comments:

Comments are enclosed in **/*** and ***/**.

Names:

Names include:

simple variables: **L**
 array elements: **L [E]**
 The words **ibase**, **obase**, and **scale**
 stacks: **L**

Other Operands

Other operands include:

Arbitrarily long numbers with optional sign and decimal point.

(**E**)

sqrt (**E**)

length (**E**) number of significant decimal digits

scale (**E**) number of digits right of decimal point

L (**E** , ... , **E**)

Strings of ASCII characters enclosed in quotes (").

Arithmetic Operators:

Arithmetic operators yield an **E** as a result and include:

+ **-** ***** **/** **%** **^** (**%** is remainder (not mod, see below); **^** is power).

++ **--** (prefix and append; apply to names)

= **+=** **-=** ***=** **/=** **%=** **^=**

Relational Operators

Relational operators yield an **R** when used as **E op E**:

== **<=** **>=** **!=** **<** **>**

Statements

```

E
{ S ; ... ; S }
if ( R ) S
while ( R ) S
for ( E ; R ; E ) S
null statement
break
quit

```

Function Definitions:

```

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

```

Functions in -l Math Library:

Functions in the -l math library include:

| | |
|--------|-----------------|
| 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. Trigonometric angles are in radians where 2 pi radians = 360 degrees.

The value of a statement that is an expression is printed unless the main operator is an assignment. No operators are defined for strings, but the string is printed if it appears in a context where an expression result would be printed. Either semicolons or new-lines can separate statements. Assignment to *scale* influences the number of digits to be retained on arithmetic operations in the manner of *dc(1)*. Assignments to *ibase* or *obase* set the input and output number radix respectively, again as defined by *dc(1)*.

The same letter can be used simultaneously as an array, a function, and a simple variable. 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.

The % operator yields the remainder at the current scale, not the integer modulus. Thus, at scale 1, 7 % 3 is .1 (one tenth), not 1. This is because (at scale 1) 7 / 3 is 2.3 with .1 as the remainder.

EXAMPLES

Define 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
    }
}

```

Print approximate values of the exponential function of the first ten integers.

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

WARNINGS

There are currently no **&&** (AND) or **||** (OR) comparisons.

The **for** statement must have all three expressions.

quit is interpreted when read, not when executed.

bc's parser is not robust in the face of input errors. Some simple expression such as 2+2 helps get it back into phase.

The assignment operators: **=+** **=-** **=*** **=/** **=%** and **=^** are obsolete. Any occurrences of these operators cause a syntax error with the exception of **=-** which is interpreted as **=** followed by a unary minus.

Neither entire arrays nor functions can be passed as function parameters.

FILES

| | |
|-----------------------|------------------------------------|
| /usr/bin/dc | desk calculator executable program |
| /usr/lib/lib.b | mathematical library |

SEE ALSO

bs(1), **dc(1)**.

bc tutorial in *Number Processing Users Guide*

STANDARDS CONFORMANCE

bc: XPG4, POSIX.2

NAME

bdiff - diff for large files

SYNOPSIS

bdiff *file1 file2* [*n*] [-s]

DESCRIPTION

bdiff compares two files and produces output identical to what would be produced by **diff** (see *diff(1)*), specifying changes that must be made to make the files identical. **bdiff** is designed for handling files that are too large for **diff**, but it can be used on files of any length.

bdiff processes files as follows:

- Ignore lines common to the beginning of both files.
- Split the remainder of each file into *n*-line segments, then execute **diff** on corresponding segments. The default value of *n* is 3500.

Command-Line Arguments

bdiff recognizes the following command-line arguments:

| | |
|------------------------------|---|
| <i>file1</i> <i>file2</i> | Names of two files to be compared by bdiff . If <i>file1</i> or <i>file2</i> (but not both) is -, standard input is used instead. |
| <i>n</i> | If a numeric value is present as the third argument, the files are divided into <i>n</i> -line segments before processing by diff . Default value for <i>n</i> is 3500. This option is useful when 3500-line segments are too large for processing by diff . |
| -s | Silent option suppresses diagnostic printing by bdiff , but does not suppress possible error messages from diff . If the <i>n</i> and -s arguments are both used, the <i>n</i> argument must precede the -s option on the command line or it will not be properly recognized. |

EXTERNAL INFLUENCES**Environment Variables**

LC_MESSAGES determines the language in which messages are displayed.

If **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**.

If any internationalization variable contains an invalid setting, **bdiff** behaves as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single- and multi-byte character code sets are supported.

DIAGNOSTICS

both files standard input (bd2)

Standard input was specified for both files. Only one file can be specified as standard input.

non-numeric limit (bd4)

A non-numeric value was specified for the *n* (third) argument.

EXAMPLES

Find differences between two large files: **file1** and **file2**, and place the result in a new file named **diffs_1.2**.

```
bdiff file1 file2 >diffs_1.2
```

Do the same, but limit file length to 1400 lines; suppress error messages:

```
bdiff file1 file2 1400 -s >diffs_1.2
```

WARNINGS

bdiff produces output identical to output from **diff**, and makes the necessary line-number corrections so that the output looks like it was processed by **diff**. However, depending on where the files are split, **bdiff** may or may not find a fully minimized set of file differences.

FILES

/tmp/bd??????

SEE ALSO

diff(1).



b

NAME

bfs - big file scanner

SYNOPSIS

bfs [-] *name*

DESCRIPTION

bfs is similar to **ed** except that it is read-only (see *ed(1)*) **bfs** can handle files with up to 32K - 1 lines; each line can contain up to 512 characters, including the new-line character. **bfs** is usually more efficient than **ed** for scanning a file, since the file is not copied to a buffer. Historically, this command was most useful for identifying sections of a large file where **csplit** could be used to divide it into more manageable pieces for editing (see *csplit(1)*). However, most editors now support files larger than the above-mentioned limits.

Normally, the size of the file being scanned is printed, as is the size of any file written with the **w** command. The optional **-** suppresses printing of sizes. Input is prompted with ***** if **P** and a carriage-return are typed, as in **ed**. Prompting can be turned off again by inputting another **P** and pressing Return. Note that messages are given in response to errors if prompting is turned on.

bfs supports the Basic Regular Expression (RE) syntax (see *regex(5)*) with the addition that a null RE (e.g., **/**) is equivalent to the last RE encountered. All address expressions described under **ed** are supported. In addition, regular expressions can be surrounded with two symbols besides **/** and **?**: **>** indicates downward search without wrap-around, and **<** indicates upward search without wrap-around. There is a slight difference in mark names: only the letters **a** through **z** can be used, and all 26 marks are remembered.

The **e**, **g**, **v**, **k**, **n**, **p**, **q**, **w**, **=**, **!** and null commands operate as described under **ed**. Commands such as **--**, **+++**, **+++=**, **-12**, and **+4p** are accepted. Note that **1,10p** and **1,10** both print the first ten lines. The **f** command only prints the name of the file being scanned; there is no *remembered* file name. The **w** command is independent of output diversion, truncation, or crunching (see the **xo**, **xt**, and **xc** commands, below). The following additional commands are available:

xf file Further commands are taken from the named *file*. When an end-of-file is reached, an interrupt signal is received or an error occurs, reading resumes with the file containing the **xf**. **xf** commands may be nested to a depth of 10.

xo [file] Further output from the **p** and null commands is diverted to the named *file*, which, if necessary, is created mode 666. If *file* is missing, output is diverted to the standard output. Note that each diversion causes truncation or creation of the file.

: label This positions a *label* in a command file. *label* is terminated by a new-line, and blanks between the **:** and the start of *label* are ignored. This command can also be used to insert comments into a command file, since labels need not be referenced.

(. , .)xb / regular expression / label

A jump (either upward or downward) is made to *label* if the command succeeds. It fails under any of the following conditions:

1. Either address is not between **1** and **\$**.
2. The second address is less than the first.
3. The regular expression does not match at least one line in the specified range, including the first and last lines.

On success, **.** is set to the line matched and a jump is made to *label*. This command is the only one that does not issue an error message on bad addresses. Thus it can be used to test whether addresses are bad before other commands are executed. Note that the command

xb / label

is an unconditional jump.

The **xb** command is allowed only if it is read from someplace other than a terminal. If it is read from a pipe only a downward jump is possible.

xn List the marks currently in use (marks are set by the **k** command).

xt number Output from the **p** and null commands is truncated to at most *number* characters. The initial number is 255.

xv[*digit*][*spaces*][*value*]

The variable name is the specified *digit* following the **xv**. **xv5100** or **xv5 100** both assign the value 100 to the variable 5. **xv61,100p** assigns the value 1,100p to the variable 6. To reference a variable, put a % in front of the variable name. For example, using the above assignments for variables 5 and 6:

```
1,%5p
1,%5
%6
```

all print the first 100 lines.

```
g/%5/p
```

globally searches for the characters 100 and prints each line containing a match. To escape the special meaning of %, a \ must precede it. For example, to match and list lines in a program file that contain **printf()** format strings specifying characters, decimal integers, or strings, the following could be used:

```
g/".*\%[cds]/p
```

Another feature of the **xv** command is that the first line of output from an HP-UX command can be stored into a variable. The only requirement is that the first character of *value* be an !. For example:

```
.w junk
xv5!cat junk
!rm junk
!echo "%5"
xv6!expr %6 + 1
```

each put the current line into variable 5, print it, and increment the variable 6 by one. To escape the special meaning of ! as the first character of *value*, precede it with a \.

```
xv7\!date
```

stores the value **!date** into variable 7.

xbz *label*
xbn *label*

These two commands test the last saved *return code* from the execution of an HP-UX system command (*!command*) for a zero or non-zero value, respectively, and cause a branch to the specified label. The two examples below both search for the next five lines containing the string **size**.

First example:

```
xv55
: 1
/size/
xv5!expr %5 - 1
!if [ %5 != 0 ] ; then exit 2 ; fi
xbn 1
```

Second Example:

```
xv45
: 1
/size/
xv4!expr %4 - 1
!if [ %4 = 0 ] ; then exit 2 ; fi
xbz 1
```

xc [*switch*] If *switch* is 1, output from the **p** and null commands is crunched; if *switch* is 0 it isn't. Without an argument, **xc** reverses *switch*. Initially *switch* is set for no crunching. Crunched output has strings of tabs and blanks reduced to one blank, and blank lines suppressed.

EXTERNAL INFLUENCES

Environment Variables

LC_COLLATE determines the collating sequence used in evaluating regular expressions.

LC_CTYPE determines the classification of characters as letters, and the characters matched by character class expressions in regular expressions.

If **LC_COLLATE** or **LC_CTYPE** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **bfs** behaves as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single-byte character code sets are supported.

DIAGNOSTICS

? for errors in commands, if prompting is turned off. Self-explanatory error messages when prompting is on.

SEE ALSO

csplit(1), ed(1), lang(5), regexp(5).

NAME

bs - a compiler/interpreter for modest-sized programs

SYNOPSIS

bs [*file* [*args*]]

DESCRIPTION

bs is a remote descendant of BASIC and SNOBOL4 with some C language added. **bs** is designed for programming tasks where program development time is as important as the resulting speed of execution. Formalities of data declaration and file/process manipulation are minimized. Line-at-a-time debugging, the **trace** and **dump** statements, and useful run-time error messages all simplify program testing. Furthermore, incomplete programs can be debugged; *inner* functions can be tested before *outer* functions have been written, and vice versa.

If *file* is specified on the command-line, it is used for input before any input is taken from the keyboard. By default, statements read from *file* are compiled for later execution. Likewise, statements entered from the keyboard are normally executed immediately (see **compile** and **execute** below). Unless the final operation is assignment, the result of an immediate expression statement is printed.

bs programs are made up of input lines. If the last character on a line is a \, the line is continued. **bs** accepts lines of the following form:

```
statement
label statement
```

A label is a *name* (see below) followed by a colon. A label and a variable can have the same name.

A **bs** statement is either an expression or a keyword followed by zero or more expressions. Some keywords (**clear**, **compile**, **!**, **execute**, **include**, **ibase**, **obase**, and **run**) are always executed as they are compiled.

Statement Syntax:

- | | |
|--------------------------------------|--|
| expression | The expression is executed for its side effects (value, assignment, or function call). The details of expressions follow the description of statement types below. |
| break | break exits from the innermost for/while loop. |
| clear | Clears the symbol table and compiled statements. clear is executed immediately. |
| compile [<i>expression</i>] | Succeeding statements are compiled (overrides the immediate execution default). The optional expression is evaluated and used as a file name for further input. A clear is associated with this latter case. compile is executed immediately. |
| continue | continue transfers to the loop-continuation of the current for/while loop. |
| dump [<i>name</i>] | The name and current value of every non-local variable is printed. Optionally, only the named variable is reported. After an error or interrupt, the number of the last statement is displayed. The user-function trace is displayed after an error or stop that occurred in a function. |
| edit | A call is made to the editor selected by the EDITOR environment variable if it is present, or ed(1) if EDITOR is undefined or null. If the <i>file</i> argument is present on the command line, <i>file</i> is passed to the editor as the file to edit (otherwise no file name is used). Upon exiting the editor, a compile statement (and associated clear) is executed giving that file name as its argument. |
| exit [<i>expression</i>] | Return to system level. The expression is returned as process status. |
| execute | Change to immediate execution mode (an interrupt has a similar effect). This statement |

does not cause stored statements to execute (see **run** below).

```
for name = expression expression statement
for name = expression expression
```

```
...
next
```

```
for expression , expression , expression statement
for expression , expression , expression
```

next The **for** statement repetitively executes a statement (first form) or a group of statements (second form) under control of a named variable. The variable takes on the value of the first expression, then is incremented by one on each loop, not to exceed the value of the second expression. The third and fourth forms require three expressions separated by commas. The first of these is the initialization, the second is the test (true to continue), and the third is the loop-continuation action (normally an increment).

```
fun f( [a, ...] ) [v, ...]
```

nuf **fun** defines the function name, arguments, and local variables for a user-written function. Up to ten arguments and local variables are allowed. Such names cannot be arrays, nor can they be I/O associated. Function definitions cannot be nested. Calling an undefined function is permissible; see function calls below.

freturn A way to signal the failure of a user-written function. See the interrogation operator (?) below. If interrogation is not present, **freturn** merely returns zero. When interrogation is active, **freturn** transfers to that expression (possibly by-passing intermediate function returns).

goto name Control is passed to the internally stored statement with the matching label.

ibase n **ibase** sets the input base (radix) to *n*. The only supported values for *n* are the constants 8, 10 (the default), and 16. Hexadecimal values 10-15 are entered as **a-f**. A leading digit is required (i.e., **f0a** must be entered as **0f0a**). **ibase** (and **obase** discussed below) are executed immediately.

```
if expression statement
if expression
```

```
...
[else...]
```

fi The statement (first form) or group of statements (second form) is executed if the expression evaluates to non-zero. The strings 0 and "" (null) evaluate as zero. In the second form, an optional **else** provides for a second group of statements to be executed when the first group is not. The only statement permitted on the same line with an **else** is an **if**; only other **fis** can be on the same line with a **fi**. The concatenation of **else** and **if** into an **elif** is supported. Only a single **fi** is required to close an **if ... elif ... [else ...]** sequence.

include expression *expression* must evaluate to a file name. The file must contain **bs** source statements. Such statements become part of the program being compiled. **include** statements cannot be nested.

obase n **obase** sets the output base to *n* (see **ibase** above).

```
onintr label
```

onintr **onintr** provides program control of interrupts. In the first form, control passes to the label given, just as if a **goto** had been executed at the time **onintr** was executed. The effect of the statement is cleared after each interrupt. In the second form, an interrupt causes **bs** to terminate.

return [expression] The expression is evaluated and the result is passed back as the value of a function call. If no expression is given, zero is returned.

run The random number generator is reset. Control is passed to the first internal statement. If the **run** statement is contained in a file, it should be the last statement.

stop Execution of internal statements is stopped. **bs** reverts to immediate mode.

trace [*expression*]
 The **trace** statement controls function tracing. If the expression is null (or evaluates to zero), tracing is turned off. Otherwise, a record of user-function calls/returns is printed. Each **return** decrements the **trace** expression value.

while *expression statement*
while *expression*
 ...

next **while** is similar to **for** except that only the conditional expression for loop-continuation is given.

! *shell command*
 An immediate escape to the shell.

... This statement is ignored (treated as a comment).

Expression Syntax:

name A name is used to specify a variable. Names are composed of a letter (uppercase or lowercase) optionally followed by letters and digits. Only the first six characters of a name are significant. Except for names declared in *fun* statements, all names are global to the program. Names can take on numeric (double float) values, string values, or can be associated with input/output (see the built-in function **open**() below).

name ([*expression* [, *expression*] ...])
 Functions can be called by a name followed by the arguments in parentheses separated by commas. Except for built-in functions (listed below), the name must be defined with a *fun* statement. Arguments to functions are passed by value. If the function is undefined, the call history to the call of that function is printed, and a request for a return value (as an expression) is made. The result of that expression is taken to be the result of the undefined function. This permits debugging programs where not all the functions are yet defined. The value is read from the current input file.

name [*expression* [, *expression*] ...]
 This syntax is used to reference either arrays or tables (see built-in *table* functions below). For arrays, each expression is truncated to an integer and used as a specifier for the name. The resulting array reference is syntactically identical to a name; **a**[1,2] is the same as **a**[1][2]. The truncated expressions are restricted to values between 0 and 32767.

number A number is used to represent a constant value. A number is written in Fortran style, and contains digits, an optional decimal point, and possibly a scale factor consisting of an **e** followed by a possibly signed exponent.

string Character strings are delimited by " characters. The \ escape character allows the double quote (\"), new-line (\n), carriage return (\r), backspace (\b), and tab (\t) characters to appear in a string. Otherwise, \ stands for itself.

(*expression*) Parentheses are used to alter the normal order of evaluation.

(*expression* , *expression* [, *expression* ...]) [*expression*]
 The bracketed expression is used as a subscript to select a comma-separated expression from the parenthesized list. List elements are numbered from the left, starting at zero.

The expression:
 (**False** , **True**) [**a == b**]
 has the value **True** if the comparison is true.

? expression The interrogation operator tests for the success of the expression rather than its value. At the moment, it is useful for testing end-of-file (see examples in the *Programming Tips* section below), the result of the **eval** built-in function, and for checking the return from user-written functions (see **freturn**). An interrogation "trap" (end-of-file, etc.) causes an immediate transfer to the most recent interrogation, possibly skipping assignment statements or intervening function levels.

- expression The result is the negation of the expression.

| | |
|---------------------|---|
| ++ name | Increments the value of the variable (or array reference). The result is the new value. |
| -- name | Decrements the value of the variable. The result is the new value. |
| ! expression | The logical negation of the expression. Watch out for the shell escape command. |
| expression | <i>operator</i> expression Common functions of two arguments are abbreviated by the two arguments separated by an operator denoting the function. Except for the assignment, concatenation, and relational operators, both operands are converted to numeric form before the function is applied. |

b

Binary Operators (in increasing precedence):

| | |
|------------------------------------|--|
| = | = is the assignment operator. The left operand must be a name or an array element. The result is the right operand. Assignment binds right to left, all other operators bind left to right. |
| _ | _ (underscore) is the concatenation operator. |
| & | & (logical AND) has result zero if either of its arguments are zero. It has result one if both of its arguments are non-zero; (logical OR) has result zero if both of its arguments are zero. It has result one if either of its arguments is non-zero. Both operators treat a null string as a zero. |
| < <= > >= == != | The relational operators (< : less than, <= : less than or equal, > : greater than, >= : greater than or equal, == : equal to, != : not equal to) return one if their arguments are in the specified relation, or return zero otherwise. Relational operators at the same level extend as follows: a>b>c is equivalent to a>b & b>c . A string comparison is made if both operands are strings. |
| + - | Add and subtract. |
| * / % | Multiply, divide, and remainder. |
| ^ | Exponentiation. |

Built-in Functions:**Dealing with arguments**

| | |
|-----------------|---|
| arg(i) | is the value of the <i>i</i> -th actual parameter on the current level of function call. At level zero, <i>arg</i> returns the <i>i</i> -th command-line argument (<i>arg</i> (0) returns bs). |
| narg() | returns the number of arguments passed. At level zero, the command argument count is returned. |

Mathematical

| | |
|-------------------|---|
| abs(x) | is the absolute value of <i>x</i> . |
| atan(x) | is the arctangent of <i>x</i> . Its value is between $-\pi/2$ and $\pi/2$. |
| ceil(x) | returns the smallest integer not less than <i>x</i> . |
| cos(x) | is the cosine of <i>x</i> (radians). |
| exp(x) | is the exponential function of <i>x</i> . |
| floor(x) | returns the largest integer not greater than <i>x</i> . |
| log(x) | is the natural logarithm of <i>x</i> . |
| rand() | is a uniformly distributed random number between zero and one. |
| sin(x) | is the sine of <i>x</i> (radians). |
| sqrt(x) | is the square root of <i>x</i> . |

String operations

| | |
|-----------------------|---|
| size(s) | the size (length in bytes) of <i>s</i> is returned. |
| format(f, a) | returns the formatted value of <i>a</i> . <i>f</i> is assumed to be a format specification in the style of <i>printf</i> (3S). Only the % ... f , % ... e , and % ... s types are safe. Since it is not always |

possible to know whether **a** is a number or a string when the **format** call is coded, coercing **a** to the type required by **f** by either adding zero (for **e** or **f** format) or concatenating (.) the null string (for **s** format) should be considered.

index(*x*, *y*)

returns the number of the first position in *x* that any of the characters from *y* matches. No match yields zero.

trans(*s*, *f*, *t*)

Translates characters of the source *s* from matching characters in *f* to a character in the same position in *t*. Source characters that do not appear in *f* are copied to the result. If the string *f* is longer than *t*, source characters that match in the excess portion of *f* do not appear in the result.

substr(*s*, *start*, *width*)

returns the sub-string of *s* defined by the *starting* position and *width*.

match(*string*, *pattern*)

mstring(*n*) The *pattern* is a regular expression according to the Basic Regular Expression definition (see *regex*(5)). **mstring** returns the *n*-th (1 ≤ *n* ≤ 10) substring of the subject that occurred between pairs of the pattern symbols \ (and \) for the most recent call to *match*. To succeed, patterns must match the beginning of the string (as if all patterns began with ^). The function returns the number of characters matched. For example:

```
match("a123ab123", ".*\[a-z]\)") == 6
mstring(1) == "b"
```

File handling

open(*name*, *file*, *function*)

close(*name*)

name argument must be a **bs** variable name (passed as a string). For the **open**, the *file* argument can be:

1. a 0 (zero), 1, or 2 representing standard input, output, or error output, respectively;
2. a string representing a file name; or
3. a string beginning with an ! representing a command to be executed (via **sh -c**). The *function* argument must be either **r** (read), **w** (write), **W** (write without new-line), or **a** (append). After a **close**, *name* reverts to being an ordinary variable. If *name* was a pipe, a **wait**() is executed before the close completes (see *wait*(2)). The **bs exit** command does not do such a wait. The initial associations are:

```
open("get", 0, "r")
open("put", 1, "w")
open("puterr", 2, "w")
```

Examples are given in the following section.

access(*s*, *m*)

executes **access**() (see *access*(2)).

fctype(*s*)

returns a single character file type indication: **f** for regular file, **p** for FIFO (i.e., named pipe), **d** for directory, **b** for block special, or **c** for character special.

Tables

table(*name*, *size*)

A table in **bs** is an associatively accessed, single-dimension array. "Subscripts" (called keys) are strings (numbers are converted). The *name* argument must be a **bs** variable name (passed as a string). The *size* argument sets the minimum number of elements to be allocated. **bs** prints an error message and stops on table overflow. The result of *table* is *name*.

item(*name*, *i*)

key()

The **item** function accesses table elements sequentially (in normal use, there is no orderly progression of key values). Where the **item** function accesses values, the **key** function accesses the "subscript" of the previous **item** call. It fails (or in the absence of an

interrogate operator, returns null) if there was no valid subscript for the previous **item** call. The *name* argument should not be quoted. Since exact table sizes are not defined, the interrogation operator should be used to detect end-of-table; for example:

```
table("t", 100)
```

```
...
# If word contains "party", the following expression adds one
# to the count of that word:
++t[word]
```

```
...
# To print out the the key/value pairs:
for i = 0, ?(s = item(t, i)), ++i if key() put = key()_"_"s
```

If the interrogation operator is not used, the result of **item** is null if there are no further elements in the table. Null is, however, a legal "subscript".

iskey(*name*, *word*)

iskey tests whether the key *word* exists in the table *name* and returns one for true, zero for false.

Odds and ends

eval(*s*) The string argument is evaluated as a **bs** expression. The function is handy for converting numeric strings to numeric internal form. **eval** can also be used as a crude form of indirection, as in:

```
name = "xyz" eval("++" _ name)
```

which increments the variable *xyz*. In addition, **eval** preceded by the interrogation operator permits the user to control **bs** error conditions. For example:

```
?eval("open(\"x\", \"XXX\", \"r\")")
```

returns the value zero if there is no file named **XXX** (instead of halting the user's program). The following executes a **goto** to the label **L** (if it exists):

```
label="L"
if !(?eval("goto " _ label)) puterr = "no label"
```

plot(*request*, *args*)

If the **tplot** command is available, the **plot** function produces output on devices recognized by **tplot**. The *requests* are as follows:

| Call | Function |
|---|---|
| plot (0, <i>term</i>) | causes further <i>plot</i> output to be piped into <i>tplot</i> with an argument of -Tterm . <i>term</i> can be up to 40 characters in length. |
| plot (1) | "erases" the plotter. |
| plot (2, <i>string</i>) | labels the current point with <i>string</i> . |
| plot (3, <i>x1</i> , <i>y1</i> , <i>x2</i> , <i>y2</i>) | draws the line between (<i>x1</i> , <i>y1</i>) and (<i>x2</i> , <i>y2</i>). |
| plot (4, <i>x</i> , <i>y</i> , <i>r</i>) | draws a circle with center (<i>x</i> , <i>y</i>) and radius <i>r</i> . |
| plot (5, <i>x1</i> , <i>y1</i> , <i>x2</i> , <i>y2</i> , <i>x3</i> , <i>y3</i>) | draws an arc (counterclockwise) with center (<i>x1</i> , <i>y1</i>) and endpoints (<i>x2</i> , <i>y2</i>) and (<i>x3</i> , <i>y3</i>). |
| plot (6) | is not implemented. |
| plot (7, <i>x</i> , <i>y</i>) | makes the current point (<i>x</i> , <i>y</i>). |
| plot (8, <i>x</i> , <i>y</i>) | draws a line from the current point to (<i>x</i> , <i>y</i>). |
| plot (9, <i>x</i> , <i>y</i>) | draws a point at (<i>x</i> , <i>y</i>). |
| plot (10, <i>string</i>) | sets the line mode to <i>string</i> . |
| plot (11, <i>x1</i> , <i>y1</i> , <i>x2</i> , <i>y2</i>) | makes (<i>x1</i> , <i>y1</i>) the lower left corner of the plotting area and (<i>x2</i> , <i>y2</i>) the upper right corner of the plotting area. |

```
plot(12, x1, y1, x2, y2)
```

causes subsequent x (y) coordinates to be multiplied by *x1* (*y1*) and then added to *x2* (*y2*) before they are plotted. The initial scaling is `plot(12, 1.0, 1.0, 0.0, 0.0)`.

Some requests do not apply to all plotters. All requests except zero and twelve are implemented by piping characters to *tplot*.

Each statement executed from the keyboard re-invokes *tplot*, making the results unpredictable if a complete picture is not done in a single operation. Plotting should thus be done either in a function or a complete program, so all the output can be directed to *tplot* in a single stream.

`last()` in immediate mode, `last` returns the most recently computed value.

EXTERNAL INFLUENCES

Environment Variables

`LC_COLLATE` determines the collating sequence used in evaluating regular expressions.

`LC_CTYPE` determines the characters matched by character class expressions in regular expressions.

If `LC_COLLATE` or `LC_CTYPE` is not specified in the environment or is set to the empty string, the value of `LANG` is used as a default for each unspecified or empty variable. If `LANG` is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of `LANG`. If any internationalization variable contains an invalid setting, `bs` behaves as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single-byte character code sets are supported.

EXAMPLES

Using `bs` as a calculator (\$ is the shell prompt):

```
$ bs
# Distance (inches) light travels in a nanosecond.
186000 * 5280 * 12 / 1e9
11.78496
...
# Compound interest (6% for 5 years on $1,000).
int = .06 / 4
bal = 1000
for i = 1 5*4 bal = bal + bal*int
bal - 1000
346.855007
...
exit
```

The outline of a typical `bs` program:

```
# initialize things:
var1 = 1
open("read", "infile", "r")
...
# compute:
while ?(str = read)
...
next
# clean up:
close("read")
...
# last statement executed (exit or stop):
exit
# last input line:
run
```

Input/Output examples:

```
# Copy file oldfile to file newfile.
open("read", "oldfile", "r")
open("write", "newfile", "w")
...
while ?(write = read)
...
# close "read" and "write":
close("read")
close("write")

# Pipe between commands.
open("ls", "!ls *", "r")
open("pr", "!pr -2 -h 'List'", "w")
while ?(pr = ls) ...
...
# be sure to close (wait for) these:
close("ls")
close("pr")
```

WARNINGS

The graphics mode (`plot ...`) is not particularly useful unless the `tplot` command is available on your system.

bs is not tolerant of some errors. For example, mistyping a **fun** declaration is difficult to correct because a new definition cannot be made without doing a **clear**. The best solution in such a case is to start by using the **edit** command.

SEE ALSO

`ed(1)`, `sh(1)`, `access(2)`, `printf(3S)`, `stdio(3S)`, `lang(5)`, `regex(5)`.

See Section (3M) for a further description of the mathematical functions.

`pow()` is used for exponentiation — see `exp(3M)`;

bs uses the Standard I/O package.

NAME

cal - print calendar

SYNOPSIS

cal *[[month] year]*

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. *year* can be between 1 and 9999. *month* is a decimal number between 1 and 12. The calendar produced is a Gregorian calendar.

EXTERNAL INFLUENCES**Environment Variables**

LANG determines the locale to use for the locale categories when both **LC_ALL** and the corresponding environment variable (beginning with **LC_**) do not specify a locale. If **LANG** is not set or is set to the empty string, a default of "C" (see *lang(5)*) is used.

LC_CTYPE determines the locale for interpretation of sequences of bytes of text data as characters (e.g., single- versus multibyte characters in arguments and input files).

LC_TIME determines the format and contents of the calendar.

TZ determines the timezone used to calculate the value of the current month.

If any internationalization variable contains an invalid setting, **cal** behaves as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single- and multi-byte character code sets are supported.

EXAMPLES

The command:

```
cal 9 1850
```

prints the calendar for September, 1850 on the screen as follows:

```

      September 1850
  S   M   Tu   W   Th   F   S
  1   2   3   4   5   6   7
  8   9  10  11  12  13  14
 15  16  17  18  19  20  21
 22  23  24  25  26  27  28
 29  30
```

However, for XPG4 the output looks like below:

```

      Sep 1850
  Sun  Mon  Tue  Wed  Thu  Fri  Sat
  1    2    3    4    5    6    7
  8    9   10   11   12   13   14
 15   16   17   18   19   20   21
 22   23   24   25   26   27   28
 29   30
```

WARNINGS

The year is always considered to start in January even though this is historically naive.

Beware that **cal 83** refers to the early Christian era, not the 20th century.

STANDARDS CONFORMANCE

cal: SVID2, SVID3, XPG2, XPG3, XPG4

NAME

calendar - reminder service

SYNOPSIS

calendar [-]

DESCRIPTION

calendar consults the file **calendar** in the current directory and prints out lines containing today's or tomorrow's date anywhere in the line. On weekends, "tomorrow" extends through Monday.

When a **-** command-line argument is present, **calendar** searches for the file **calendar** in each user's home directory, and sends any positive results to the user by **mail** (see *mail(1)*). Normally this is done daily in the early morning hours under the control of **cron** (see *cron(1M)*). When invoked by **cron**, **calendar** reads the first line in the **calendar** file to determine the user's environment.

Language-dependent information such as spelling and date format (described below) are determined by the user-specified **LANG** statement in the **calendar** file. This statement should be of the form **LANG=language** where *language* is a valid language name (see *lang(5)*). If this line is not in the **calendar** file, the action described in the EXTERNAL INFLUENCES Environment Variable section is taken.

calendar is concerned with two fields: month and day. A month field can be expressed in three different formats: a string representing the name of the month (either fully spelled out or abbreviated), a numeric month, or an asterisk (representing any month). If the month is expressed as a string representing the name of the month, the first character can be either upper-case or lower-case; other characters must be lower-case. The spelling of a month name should match the string returned by calling `nl_langinfo()` (see *nl_langinfo(3C)*). The day field is a numeric value for the day of the month.

Month-Day Formats

If the month field is a string, it can be followed by zero or more blanks. If the month field is numeric, it must be followed by either a slash (/) or a hyphen (-). If the month field is an asterisk (*), it must be followed by a slash (/). The day field can be followed immediately by a blank or non-digit character.

Day-Month Formats

The day field is expressed as a numeral. What follows the day field is determined by the format of the month. If the month field is a string, the day field must be followed by zero or one dot (.) followed by zero or more blanks. If the month field is a numeral, the day field must be followed by either a slash (/) or a hyphen (-). If the month field is an asterisk, the day field must be followed by a slash (/).

EXTERNAL INFLUENCES**Environment Variables**

LC_TIME determines the format and contents of date and time strings when no **LANG** statement is specified in the **calendar** file.

LANG determines the language in which messages are displayed.

If **LC_TIME** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **calendar** behaves as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single- and multi-byte character code sets are supported.

EXAMPLES

The following **calendar** file illustrates several formats recognized by *calendar*:

```
LANG=en_US.roman8
Friday, May 29th: group coffee meeting
meeting with Boss on June 3.
3/30/87 - quarter end review
4-26 Management council meeting at 1:00 pm
It is first of the month ( */1 ); status report due.
```

In the following **calendar** file, dates are expressed according to European English usage:

```
LANG=en_GB.roman8
On 20 Jan. code review
```

```
Jim's birthday is on the 3. February
30/3/87 - quarter end review
26-4 Management council meeting at 1:00 pm
It is first of the month ( 1/* ); status report due.
```

WARNINGS

To get reminder service, either your calendar must be public information or you must run **calendar** from your personal **crontab** file, independent of any **calendar** - run systemwide. Note that if you run **calendar** yourself, the calendar file need not reside in your home directory.

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

This command is likely to be withdrawn from X/Open standards. Applications using this command might not be portable to other vendors' platforms.

AUTHOR

calendar was developed by AT&T and HP.

FILES

```
calendar
/tmp/cal*
/usr/sbin/calprog to figure out today's and tomorrow's dates
/usr/bin/crontab
/etc/passwd
```

SEE ALSO

cron(1M), nl_langinfo(3C), mail(1), environ(5).

STANDARDS CONFORMANCE

calendar: SVID2, SVID3, XPG2, XPG3

C

NAME

cat - concatenate, copy, and print files

SYNOPSIS

cat [-benrstuv] file ...

DESCRIPTION

cat reads each *file* in sequence and writes it on the standard output. Thus:

```
cat file
```

prints *file* on the default standard output device;

```
cat file1 file2 > file3
```

concatenates *file1* and *file2*, and places the result in *file3*.

If **-** appears as a *file* argument, **cat** uses standard input. To combine standard input and other files, use a combination of **-** and *file* arguments.

Options

cat recognizes the following options:

- b Omit line numbers from blank lines when **-n** option is specified. If this option is specified, the **-n** option is automatically selected.
- e Print a **\$** character at the end of each line (prior to the new-line). If this option is specified, the **-v** option is automatically selected.
- n Display output lines preceded by line numbers, numbered sequentially from 1.
- r Replace multiple consecutive empty lines with one empty line, so that there is never more than one empty line between lines containing characters.
- s Silent option. **cat** suppresses error messages about non-existent files, identical input and output, and write errors. Normally, input and output files cannot have identical names unless the file is a special file.
- t Print each tab character as **^I** and form feed character as **^L**. If this option is specified, the **-v** option is automatically selected.
- u Do not buffer output (handle character-by-character). Normally, output is buffered.
- v Cause non-printing characters (with the exception of tabs, new-lines and form-feeds) to be printed visibly. Control characters are printed using the form **^X** (Ctrl-*X*), and the DEL character (octal 0177) is printed as **^?** (see *ascii*(5)). Single-byte control characters whose most significant bit is set, are printed using the form **M-^x**, where *x* is the character specified by the seven low order bits. All other non-printing characters are printed as **M-x**, where *x* is the character specified by the seven low order bits. This option is influenced by the **LC_CTYPE** environment variable and its corresponding code set.

EXTERNAL INFLUENCES**Environment Variables**

LANG provides a default value for the internationalization variables that are unset or null. If **LANG** is unset or null, the default value of "C" (see *lang*(5)) is used. If any of the internationalization variables contains an invalid setting, **cat** will behave as if all internationalization variables are set to "C". See *environ*(5).

LC_ALL If set to a non-empty string value, overrides the values of all the other internationalization variables.

LC_CTYPE determines the interpretation of text as single and/or multi-byte characters, the classification of characters as printable, and the characters matched by character class expressions in regular expressions.

LC_MESSAGES determines the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.

NLSPATH determines the location of message catalogues for the processing of **LC_MESSAGES**.

International Code Set Support

Single- and multi-byte character code sets are supported.

RETURN VALUE

Exit values are:

| | |
|--------------|---------------------------|
| 0 | Successful completion. |
| >0 | Error condition occurred. |

EXAMPLES

To create a zero-length file, use any of the following:

```
cat /dev/null > file
cp /dev/null file
touch file
```

The following prints `^I` for all the occurrences of tab character in *file1*

```
cat -t file1
```

To suppress error messages about files that do not exist, use:

```
cat -s file1 file2 file3 > file
```

If *file2* does not exist, the above command concatenates *file1* and *file3* without reporting the error on *file2*. The result is the same if `-s` option is not used, except that **cat** displays the error message.

To view non-printable characters in *file2*, use:

```
cat -v file2
```

WARNINGS

Command formats such as

```
cat file1 file2 > file1
```

overwrites the data in *file1* before the concatenation begins, thus destroying the file. Therefore, be careful when using shell special characters.

SEE ALSO

cp(1), more(1), pg(1), pr(1), rmnl(1), ssp(1).

STANDARDS CONFORMANCE

cat: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

(Bundled C Compiler - Limited Functionality)

NAME

cc - bundled C compiler

SYNOPSIS

cc [*options*] *files*

DESCRIPTION

This manual page describes the Bundled C compiler. See *cc(1)*, online only, for a description of the ANSI-compliant HP-UX manual page.

This *cc* accepts several types of arguments as *files*:

.c Suffix

Arguments whose names end with **.c** are understood to be C source files. Each is compiled and the resulting object file is left in a file having the corresponding base name, **.o** instead of **.c**. However, if a single C file is compiled and linked, all in one step, the **.o** file is deleted.

.s Suffix

Arguments whose names end with **.s** are understood to be assembly source files and are assembled, producing a **.o** file for each **.s** file.

.i Suffix

Arguments whose names end with **.i** are assumed to be the output of **cpp** (see the **-P** option below). They are compiled without invoking **cpp** (see *cpp(1)*). Each object file is left in a file having the corresponding base name, but suffixed with **.o** instead of **.i**.

-l_x Form

Arguments of the form **-l_x** cause the linker to search the library **lib_x.sl** or **lib_x.a** in an attempt to resolve currently unresolved external references. Because a library is searched when its name is encountered, placement of a **-l** is significant. If a file contains an unresolved external reference, the library containing the definition must be placed *after* the file on the command line. See *ld(1)* for further details.

-l:lib_x.suffix Form

Arguments of the form **-l:lib_x.suffix** cause the linker to search the library **lib_x.sl** or **lib_x.a** (depending on *suffix*) in an attempt to resolve currently unresolved external references. It is similar to the **-l** option except the current state of the **-Wl,-a** option is not important.

Other Suffixes

All other arguments, such as those whose names end with **.o** or **.a**, are taken to be relocatable object files that are to be included in the link operation.

Arguments and options can be passed to the compiler through the **CCOPTS** environment variable as well as on the command line. The compiler reads the value of **CCOPTS** and divides these options into two sets; options that appear before a vertical bar (**|**), and options that appear after the vertical bar. The first set of options are placed before any of the command-line parameters to *cc*; the second set of options are placed after the command-line parameters to *cc*. If the vertical bar is not present, all options are placed before the command-line parameters. For example (in *sh(1)* notation),

```
CCOPTS="-v | -lmalloc"
export CCOPTS
cc -w prog.c
```

is equivalent to

```
cc -v -w prog.c -lmalloc
```

When set, the **TMPDIR** environment variable specifies a directory to be used by the compiler for temporary files, overriding the default directory **/var/tmp**.

Options

The following options are the only options which are recognized by the bundled C compiler.

- c** Suppress the link edit phase of the compilation, and force an object (**.o**) file to be produced for each **.c** file, even if only one program is compiled. Object files produced from C programs must be linked before being executed.
- C** Prevent the preprocessor from stripping C-style comments. See *cpp(1)* for details.

(Bundled C Compiler - Limited Functionality)

- Dname=def** Define *name* to the preprocessor, as if by '#define'. See *cpp*(1) for details.
- Dname**
- E** Run only **cpp** on the named C or assembly files, and send the result to the standard output.
- I dir** Change the algorithm used by the preprocessor for finding include files to also search in directory *dir*. See *cpp*(1) for details.
- l x** Refer to the **DESCRIPTION** section.
- L dir** Change the algorithm used by the linker to search for **libx.sl** or **libx.a**. The **-L** option causes **cc** to search in *dir* before searching in the default locations. See *ld*(1) for details.
- o outfile** Name the output file from the linker *outfile*. The default name is **a.out**.
- P** Run only **cpp** on the named C files and leave the result on corresponding files suffixed **.i**. The **-P** option is also passed along to **cpp**.
- +R num** Allow only the first *num* **register** variables to actually have the **register** class. Use this option when the register allocator issues the message:
- out of general registers**
- s** Cause the output of the linker to be stripped of symbol table information. See *strip*(1) for more details. The use of this option prevents the use of a symbolic debugger on the resulting program. See *ld*(1) for more details.
- S** Compile the named C files, and leave the assembly language output on corresponding files suffixed **.s**.
- tx, name** Substitute subprocess *x* with *name* where *x* is one or more of a set of identifiers indicating the subprocess(es). This option works in two modes: 1) if *x* is a single identifier, *name* represents the full path name of the new subprocess; 2) if *x* is a set of identifiers, *name* represents a prefix to which the standard suffixes are concatenated to construct the full path names of the new subprocesses.
- The *x* can take one or more of the values:
- p** Preprocessor (standard suffix is **cpp**)
 - c** Compiler (standard suffix is **ccom**)
 - a** Assembler (standard suffix is **as**)
 - l** Linker (standard suffix is **ld**)
- Uname** Remove any initial definition of *name* in the preprocessor. See *cpp*(1) for details.
- v** Enable verbose mode, which produces a step-by-step description of the compilation process on the standard error.
- V** Cause each invoked subprocess to print its version information to stdout.
- w** Suppress warning messages.
- Wx, arglist** Pass the comma-separated argument(s) in *arglist* to subprocess *x*. The **-W** option specification allows additional, implementation-specific options to be recognized by the compiler driver. For example,
- Wl, -a, archive**
- causes the linker to link with archive libraries instead of with shared libraries. See *ld*(1) for details.

The *x* can assume one of the following values:

- p** Preprocessor
- a** Assembler
- l** Linker

Any other options encountered generate a warning to standard error.

Other arguments are assumed to be C-compatible object programs, typically produced by an earlier **cc** run, or perhaps libraries of C-compatible routines. These programs, together with the results of any compilations specified, are linked (in the order given) to produce an executable program with the

(Bundled C Compiler - Limited Functionality)

name **a.out**.

EXTERNAL INFLUENCES**Environment Variables**

LANG determines the language in which messages are displayed.

If **LC_CTYPE** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of **C** (see *lang(5)*) is used. If any internationalization variable contains an invalid setting, **cc** behaves as if all internationalization variables are set to **C**. See *environ(5)*.

International Code Set Support

Single byte and multibyte character code sets are supported.

DIAGNOSTICS

The diagnostics produced by **C** itself are intended to be self-explanatory. Occasionally, messages may be produced by the preprocessor, assembler, or the link editor.

If any errors occur before **cc** is completed, a nonzero value is returned. Otherwise, zero is returned.

EXAMPLES

The example below compiles the C file **prog.c** to create a **prog.o** file, then invokes the **ld** link editor to link **prog.o** and **procedure.o** with all of the C startup routines in **/usr/ccs/lib/crt0.o** and library routines from the C library **libc.sl** or **libc.a**. The resulting executable program is placed in file **prog**:

```
cc prog.c procedure.o -o prog
```

WARNINGS

Options not recognized by **cc** are not passed on to the link editor. The option **-Wl, arg** can be used to pass any such option to the link editor.

By default, the return value from a C program is completely random. The only two guaranteed ways to return a specific value are to explicitly call **exit()** (see *exit(2)*) or leave the function **main()** with a **return expression;** construct.

FILES

| | |
|----------------------------|--|
| file.c | Input file |
| file.o | Object file |
| a.out | Linked executable output file |
| /var/tmp/ctm* | Temporary files used by the compiler |
| /usr/ccs/bin/as | Assembler (see <i>as(1)</i>) |
| /usr/ccs/bin/ld | Link editor (see <i>ld(1)</i>) |
| /usr/ccs/lib/crt0.o | Runtime startup |
| /usr/lib/libc.a | Standard C library (archive version), see <i>HP-UX Reference Section (3)</i> |
| /usr/lib/libc.sl | Standard C library (shared version), see <i>HP-UX Reference Section (3)</i> |
| /usr/include | Standard directory for #include files |

Bundled C Compiler Files

| | |
|---------------------------------------|----------------------------|
| /usr/ccs/bin/cc | C driver |
| /usr/ccs/lib/cocom | C compiler |
| /usr/lib/nls/msg/\$LANG/cc.cat | C compiler message catalog |
| /usr/ccs/lib/cpp | C preprocessor |

SEE ALSO**System Tools**

| | |
|---------------|--|
| as(1) | Translate assembly code to machine code. |
| cpp(1) | Invoke the C language preprocessor. |
| ld(1) | Invoke the link editor. |
| cc(1) | The ANSI-compliant C compiler on HP-UX. |

(Bundled C Compiler - Limited Functionality)**Miscellaneous**

| | |
|----------------|---|
| matherr(3M) | Trap math errors. |
| fpgetround(3M) | Floating-point mode control functions. |
| strip(1) | Strip symbol and line number information from an object file. |
| crt0(3) | Execution startup routine. |
| end(3C) | Symbol of the last locations in program. |
| exit(2) | Termination of a process. |

Tutorials and Standards Documents

B. W. Kernighan and D. M. Ritchie, *The C Programming Language*, Prentice-Hall, 1978.


C

NAME

cd - change working directory

SYNOPSIS

cd [*directory*]

DESCRIPTION

If *directory* is not specified, the value of shell parameter **HOME** is used as the new working directory. If *directory* specifies a complete path starting with /, ., or . ., *directory* 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 *directory*.

cd exists only as a shell built-in command because a new process is created whenever a command is executed, making **cd** useless if written and processed as a normal system command. Moreover, different shells provide different implementations of **cd** as a built-in utility. Features of **cd** as described here may not be supported by all the shells. Refer to individual shell manual entries for differences.

If **cd** is called in a subshell or a separate utility execution environment such as:

```
find . -type d -exec cd {}; -exec foo {};
```

(which invokes **foo** on accessible directories) **cd** does not affect the current directory of the caller's environment. Another usage of **cd** as a stand-alone command is to obtain the exit status of the command.

EXTERNAL INFLUENCES**International Code Set Support**

Single- and multi-byte character code sets are supported.

EXAMPLES

Change the current working directory to the **HOME** directory from any location in the file system:

```
cd
```

Change to new current working directory **foo** residing in the current directory:

```
cd foo
```

or

```
cd ../foo
```

Change to directory **foobar** residing in the current directory's parent directory:

```
cd ../foobar
```

Change to the directory whose absolute pathname is **/usr/local/lib/work.files**:

```
cd /usr/local/lib/work.files
```

Change to the directory **proj1/schedule/staffing/proposals** relative to home directory:

```
cd $HOME/proj1/schedule/staffing/proposals
```

VARIABLES

The following environment variables affect the execution of **cd**:

| | |
|---------------|--|
| HOME | The name of the home directory, used when no directory operand is specified. |
| CDPATH | A colon-separated list of pathnames that refer to directories. If the directory operand does not begin with a slash (/) character, and the first component is not dot or dot-dot, cd searches for <i>directory</i> relative to each directory named in the CDPATH variable, in the order listed. The new working directory is set to the first matching directory found. An empty string in place of a directory pathname represents the current directory. If CDPATH is not set, it is treated as if it was an empty string. |

RETURN VALUE

Upon completion, **cd** exits with one of the following values:

0 The directory was successfully changed.
>0 An error occurred. The working directory remains unchanged.

SEE ALSO

cd(1), pwd(1), ksh(1), sh-posix(1), sh(1), chdir(2).

STANDARDS CONFORMANCE

cd: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

cdc - change the delta commentary of an SCCS delta

SYNOPSIS

```
cdc -r SID [-m[mrlist]] [-y[comment]] files
```

DESCRIPTION

The **cdc** command changes the **delta commentary**, for the *SID* specified by the **-r** option, of each named SCCS file.

Delta commentary is defined to be the Modification Request (MR) and comment information normally specified via the *delta(1)* command (**-m** and **-y** options).

If a directory is named, **cdc** behaves as if each file in the directory were specified as a named file, except that non-SCCS files (last component of the path name does not begin with **s.**) and unreadable files are silently ignored. If a name of **-** is given, the standard input is read (see **WARNINGS**); each line of the standard input is taken to be the name of an SCCS file to be processed.

Options

Arguments to **cdc**, which can appear in any order, consist of *option* arguments and file names.

All of the described *option* arguments apply independently to each named file:

- rSID** Used to specify the *SCCS ID*entification (SID) string of a delta for which the delta commentary is to be changed.
- m[mrlist]** If the SCCS file has the **v** option set (see *admin(1)*), a list of MR numbers to be added and/or deleted in the delta commentary of the *SID* specified by the **-r** option *may* be supplied. A null MR list has no effect.

MR entries are added to the list of MRs in the same manner as that of *delta(1)*. To delete an MR, precede the MR number with the character **!** (see **EXAMPLES**). If the MR to be deleted is currently in the list of MRs, it is removed and changed into a "comment" line. A list of all deleted MRs is placed in the comment section of the delta commentary and preceded by a comment line stating that they were deleted.

If **-m** is not used and the standard input is a terminal, the prompt **MRS?** is issued on the standard output before the standard input is read; if the standard input is not a terminal, no prompt is issued. The **MRS?** prompt always precedes the **comments?** prompt (see **-y** option).

MRs in a list are separated by blanks and/or tab characters. An unescaped new-line character terminates the MRs list.

Note that if the **v** option has a value (see *admin(1)*), it is treated as the name of a program (or shell procedure) that validates the correctness of the MR numbers. If a non-zero exit status is returned from the MR number validation program, **cdc** terminates and the delta commentary remains unchanged.
- y[comment]** Arbitrary text used to replace the *comment* or *comments* already existing for the delta specified by the **-r** option. Previous comments are kept and preceded by a comment line stating that they were changed. A null *comment* has no effect.

If **-y** is not specified and the standard input is a terminal, the prompt **comments?** is issued on the standard output before standard input is read; if standard input is not a terminal, no prompt is issued. An unescaped new-line character terminates the *comment* text.

The exact permissions necessary to modify the SCCS file are documented in *get(1)*. Simply stated, they are either:

- If you made the delta, you can change its delta commentary, or
- If you own the file and directory, you can modify the delta commentary.

EXTERNAL INFLUENCES**Environment Variables**

LANG determines the language in which messages are displayed.

International Code Set Support

Single- and multi-byte character code sets are supported.

DIAGNOSTICS

Use *sccshelp(1)* for explanations.

EXAMPLES

Add **b178-12345** and **b179-00001** to the MR list, remove **b177-54321** from the MR list, and add the comment **trouble** to delta 1.6 of **s.file**:

```
cdc -r1.6 -m"b178-12345 !b177-54321 b179-00001" -ytrouble s.file
```

The following does the same thing:

```
cdc -r1.6 s.file
MRs? !b177-54321 b178-12345 b179-00001
comments? trouble
```

WARNINGS

If SCCS file names are supplied to the **cdc** command via the standard input (**-** on the command line), the **-m** and **-y** options must also be used.

FILES

| | |
|---------------|-----------------------|
| <i>x-file</i> | See <i>delta(1)</i> . |
| <i>z-file</i> | See <i>delta(1)</i> . |

SEE ALSO

admin(1), *delta(1)*, *get(1)*, *sccshelp(1)*, *prs(1)*, *sccsfile(4)*, *rcsfile(4)*, *acl(5)*, *rcsintro(5)*.

NAME

chac1 - add, modify, delete, copy, or summarize access control lists (ACLs) of files

SYNOPSIS

```
/usr/bin/chac1 acl file ...
chac1 -r acl file ...
chac1 -d aclpatt file ...
chac1 -f fromfile tofile ...
chac1 -[z | Z | F] file...
```

DESCRIPTION

chac1 extends the capabilities of *chmod*(1), by enabling the user to grant or restrict file access to additional specific users and/or groups. Traditional file access permissions, set when a file is created, grant or restrict access to the file's owner, group, and other users. These file access permissions (eg., *rwrxrw-r--*) are mapped into three base access control list entries: one entry for the file's owner (*u.%*, *mode*), one for the file's group (*%g*, *mode*), and one for other users (*%.%*, *mode*).

chac1 enables a user to designate up to thirteen additional sets of permissions (called optional access control list (ACL) entries) which are stored in the access control list of the file.

To use *chac1*, the owner (or superuser) constructs an *acl*, a set of (*user.group*, *mode*) mappings to associate with one or more files. A specific user and group can be referred to by either name or number; any user (*u*), group (*g*), or both can be referred to with a *%* symbol, representing *any* user or group. The *@* symbol specifies the file's owner or group.

Read, write, and execute/search (*rwrx*) *modes* are identical to those used by *chmod*; symbolic operators (*op*) add (+), remove (-), or set (=) access rights. The entire *acl* should be quoted if it contains whitespace or special characters. Although two variants for constructing the *acl* are available (and fully explained in *acl*(5)), the following syntax is suggested:

```
entry[, entry] ...
```

where the syntax for an *entry* is

```
u.g op mode[ op mode] ...
```

By default, **chac1** modifies existing ACLs. It adds ACL entries or modifies access rights in existing ACL entries. If *acl* contains an ACL entry already associated with a file, the entry's mode bits are changed to the new value given, or are modified by the specified operators. If the file's ACL does not already contain the specified entry, that ACL entry is added. **chac1** can also remove all access to files. Giving it a null *acl* argument means either "no access" (when using the *-r* option) or "no changes."

For a summary of the syntax, run **chac1** without arguments.

If *file* is specified as *-*, **chac1** reads from standard input.

Options

chac1 recognizes the following options:

-r Replace old ACLs with the given ACL. All optional ACL entries are first deleted from the specified file's ACLs, their base permissions are set to zero, and the new ACL is applied. If *acl* does not contain an entry for the owner (*u.%*), the group (*%g*), or other (*%.%*) users of a file, that base ACL entry's mode is set to zero (no access). The command affects all of the file's ACL entries, but does not change the file's owner or group ID.

In *chmod*(1), the "modify" and "replace" operations are distinguished by the syntax (string or octal value). There is no corollary for ACLs because they have a variable number of entries. Hence **chac1** modifies specific entries by default, and optionally replaces all entries.

-d Delete the specified entries from the ACLs on all specified files. The *aclpatt* argument can be an exact ACL or an ACL pattern (see *acl*(5)). **chac1 -d** updates each file's ACL only if entries are deleted from it.

If you attempt to delete a base ACL entry from any file, the entry remains but its access mode is set to zero (no access). If you attempt to delete a non-existent ACL entry from a file (that is, if an ACL entry pattern matches no ACL entry), **chac1** informs you of the error,

continues, and eventually returns non-zero.

-f *fromfile tofile*

Copy the ACL from *fromfile* to the specified *tofile*, transferring ownership, if necessary (see *acl*(5), *chown*(2), or *chownacl*(3C)). *fromfile* can be **-** to represent standard input.

This option implies the **-r** option. If the owner and group of *fromfile* are identical to those of *tofile*, **chacl -f** is identical to:

```
chacl -r 'lsacl fromfile' tofile ...
```

To copy an ACL without transferring ownership, the above command is suggested instead of **chacl -f**.

-z Delete ("zap") all optional entries in the specified file's ACLs, leaving only base entries.

-Z Delete ("zap") all optional entries in the specified file's ACLs, and set the access modes in all base entries to zero (no access). This is identical to replacing the old ACL with a null ACL:

```
chacl -r '' file ...
```

or using *chmod*(1), which deletes optional entries as a side effect:

```
chmod 0 file ...
```

-F Incorporate ("fold") optional ACL entries into base ACL entries. The base ACL entry's permission bits are altered, if necessary, to reflect the caller's effective access rights to the file; all optional entries, if any, are deleted.

For ordinary users, only the access mode of the owner base ACL entry can be altered. Unlike *getaccess*, the write bit is not turned off for a file on a read-only file system or a shared-text program being executed (see *getaccess*(1)).

For super-users, only the execute mode bit in the owner base ACL entry might be changed, only if the file is not an regular file or if an execute bit is not already set in a base ACL entry mode, but is set in an optional ACL entry mode.

acl also can be obtained from a string in a file:

```
chacl 'cat file' files ...
```

Using **@** in *acl* to represent "file owner or group" can cause **chacl** to run more slowly because it must reparse the ACL for each file (except with the **-d** option).

EXTERNAL INFLUENCES

Environment Variables

LANG determines the language in which messages are displayed.

If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang*(5)) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **chacl** behaves as if all internationalization variables are set to "C". See *environ*(5).

RETURN VALUE

If **chacl** succeeds, it returns a value of zero.

If **chacl** encounters an error before it changes any file's ACL, it prints an error message to standard error and returns 1. Such errors include invalid invocation, invalid syntax of *acl* (*aclpatt*), a given user name or group name is unknown, or inability to get an ACL from *fromfile* with the **-f** option.

If **chacl** cannot execute the requested operation, it prints an error message to standard error, continues, and later returns 2. This includes cases when a file does not exist, a file's ACL cannot be altered, more ACL entries would result than are allowed, or an attempt is made to delete a non-existing ACL entry.

EXAMPLES

The following command adds read access for user **jpc** in any group, and removes write access for any user in the file's groups, for files **x** and **y**.

```
chacl "jpc.%r, %.-w" x y
```

This command replaces the ACL on the file open as standard input and on file **test** with one which only allows the file owner read and write access.

```
chacl -r '(@.%,rw-)' - test
```

Delete from file **myfile** the specific access rights, if any, for user 165 in group 13. Note that this is different from adding an ACL entry that restricts access for that user and group. The user's resulting access rights depend on the entries remaining in the ACL. The command also deletes all entries for user **jpc** that have a read bit turned on (the asterisk can be used as a wildcard in the ACL pattern for user, group, or access mode):

```
chacl -d '165.13, jpc.*+r' myfile
```

Copy the ACL from **oldfile** to **slow/hare** and **fast/tortoise**.

```
chacl -f oldfile slow/hare fast/tortoise
```

Delete the optional ACL entries, if any, on the file open as standard input.

```
chacl -z -
```

Deny all access to all files in the current directory whose names start with **a**, **b**, or **c**:

```
chacl -Z [a-c]*
```

Incorporate the optional ACL entries of a file (**fun.stuff**) into the base ACL entries:

```
chacl -F fun.stuff
```

WARNINGS

An ACL string cannot contain more than 16 unique entries, even though converting @ symbols to user or group names and combining redundant entries might result in fewer than 16 entries for some files.

DEPENDENCIES

chacl will fail when the target file resides on a file system which does not support ACLs.

NFS

Only the **-F** option is supported on remote files.

AUTHOR

chacl was developed by HP.

SEE ALSO

chmod(1), **getaccess**(1), **lsacl**(1), **getacl**(2), **setacl**(2), **acl**(5), **glossary**(9).

NAME

chatr - change program's internal attributes

SYNOPSIS

PA32 SOM chatr

```
chatr [-nqsmN] [-l library] [-B mode] [+b flag] [+es flag] [+gst flag] [+gstbuckets size]
      [+gstsize size] [+k flag] [+l library] [+pd size] [+pi size] [+plabel_cache flag]
      [+q3p flag] [+q4p flag] [+r flag] [+s flag] [+z flag] file ...
```

PA64 ELF chatr

There are two possible syntactic forms that can be used to invoke PA64 **chatr**.

FORMAT 1: The first syntactic form, which is compatible with the SOM **chatr**, is used for backward compatibility, and for easy manipulation of ordinary files that only have a single text and a single data segment:

```
chatr [-nqs] [-l library] [-B mode] [+b flag] [+cd flag] [+ci flag] [+es flag] [+gst flag]
      [+gstsize size] [+k flag] [+l library] [+md flag] [+mi flag] [+pd size] [+pi size] [+s
      flag] [+z flag] file ...
```

FORMAT 2: The second syntactic form provides the ability to explicitly specify segments to be modified:

```
chatr [-s] [-B mode] [+c flag] [+dz flag] [+k flag] [+m flag] [+p size] [+r flag] [+s flag]
      [+si index] [+sa address] [+sall] [+z flag] file ...
```

DESCRIPTION

chatr allows you to change a program's internal attributes for 32-bit mode SOM and 64-bit mode ELF files.

Upon completion, **chatr** prints the file's old and new values to standard output unless **-s** is specified.

The **+pd** and **+pi** options only provide a hint for the virtual memory page size. The actual page sizes may vary. Under certain conditions, page size hints of **L** may result in better performance, depending on the specific memory requirements of the application.

The performance of some applications may benefit from static branch prediction, others may not. The **+r** option provides a hint for using or avoiding this feature.

The **+gst** and related options provide performance enhancements through use of global symbol table which improves searching for exported symbols. See *dld.sl(5)* and the *HP-UX Linker and Libraries Online User Guide* for more information.

COMMON OPTIONS FOR PA32 SOM AND PA64 ELF (FORMAT 1) chatr

chatr, by default, prints each *file*'s magic number and file attributes to the standard output.

- l library** Indicate that the specified shared library is subject to run-time path lookup if directory path lists are provided (see **+s** and **+b**).
- n** Change *file* from demand-loaded (**DEMAND_MAGIC**) to shared (**SHARE_MAGIC**) (Ignored in PA64 FORMAT 1.)
- q** Change *file* from shared (**SHARE_MAGIC**) to demand-loaded (**DEMAND_MAGIC**). (Ignored in PA64 FORMAT 1.)
- s** Perform its operation silently. (Available with the PA64 FORMAT 2 command.)
- B mode** Select run-time binding behavior mode of a program using shared libraries. You must specify one of the major binding modes **immediate** or **deferred**. One or more of the binding modifiers **nonfatal**, **verbose**, or **restricted** can also be specified, each with a separate option. See the *HP-UX Linker and Libraries User's Guide* manual for a description of binding modes. (Available with the PA64 FORMAT 2 command.)
- +b flag** Control whether the embedded path list stored when the program (if any) was built can be used to locate shared libraries needed by the program. The two flag values, **enable** and **disable**, respectively enable and disable use of the embedded path list. See the **+s** option. You can use the **+b** option to enable the embedded path for filter libraries.
- +es flag** Control the ability of user code to execute from stack with the *flag* values, **enable** and **disable**. See the *Restricting Execute Permission on Stacks* section below for additional

information related to security issues.

- +gst flag** Control whether the global symbol table hash mechanism is used to look up values of symbol import/export entries. The two flag values, **enable** and **disable**, respectively enable and disable use of the global symbol table hash mechanism. The default is **disable**.
- +gstsize size** Request a particular hash array *size* using the global symbol table hash mechanism. The value can vary between 1 and **MAXINT**. The default value is 1103. Use this option with **+gst enable**.
- +k flag** Request kernel assisted branch prediction. The flags **enable** and **disable** turn this request on and off, respectively. (Available with the PA64 FORMAT 2 command.)
- +l library** Indicate that the specified shared library is not subject to run-time path lookup if directory path lists are provided (see **+s** and **+b**).
- +pd size** Request a particular virtual memory page size that should be used for data. Sizes of **4K**, **16K**, **64K**, **256K**, **1M**, **4M**, **16M**, **64M**, **256M**, and **L** are supported. A size of **L** will result in using the largest page size available. The actual page size may vary if the requested size cannot be fulfilled.
- +pi size** Request a particular virtual memory page size that should be used for instructions. See the **+pd** option for additional information.
- +r flag** Request static branch prediction when executing this program. The flags **enable** and **disable** turn this request on and off, respectively. (Available with the PA64 FORMAT 2 command.)
- +s flag** Control whether the directory path list specified with the **SHLIB_PATH** environment variable can be used to locate shared libraries needed by the program. The two flag values, **enable** and **disable**, respectively enable and disable use of the environment variable. If both **+s** and **+b** are used, their relative order on the command line indicates which path list will be searched first. See the **+b** option. (Available with the PA64 FORMAT 2 command.)
- +z** Enable lazy swap on all data segments (using PA32 chatr or PA64 chatr FORMAT 1) or on a specific segment (using PA64 ELF chatr FORMAT 2). May not be used with non-data segments.

OPTIONS FOR PA32 SOM chatr ONLY

- M** Change *file* from **EXEC_MAGIC** to **SHMEM_MAGIC**. (This option is an interim solution until 64-bit addressability is available with a true 64-bit kernel. See the "chatr and MAGIC Numbers" and "Using SHMEM_MAGIC" sections below.)
- N** Change *file* from **SHMEM_MAGIC** to **EXEC_MAGIC**. (This option is an interim solution until 64-bit addressability is available with a true 64-bit kernel. See the "chatr and MAGIC Numbers" and "Using SHMEM_MAGIC" sections below.)
- +gstbuckets size** Request a particular number of buckets per entry using the global symbol table hash mechanism. The value can vary between 1 and **MAXINT**. The default value is 3. Use this option with **+gst enable**.
- +plabel_cache flag** Control the use of the plabel caching mechanism. The flags **enable** and **disable** turn this request on and off, respectively. The default is **disable**. Use this option with **+gst enable**.

This option is effective with C++. In C++ applications, the dynamic loader needs to repetitively access PLABEL information (import stub). In order to make this access faster, the dynamic loader uses the global symbol table structure to also contain PLABEL entries. This behavior is enabled when the PLABEL_CACHE flag is set in the dl_header structure (enabled **ld +plabel_cache enable a.out** or **chatr +plabel_cache enable a.out**).
- +q3p flag** Control the flag bit setting to indicate how 32-bit processes use the third quadrant as data space.

The **enable** flag sets the flag bit to indicate that 32-bit processes use the third quadrant as a private data space. By setting the bit, the private data space increases from 1.9GB to 2.85GB for 32-bit processes.

The **disable** flag unsets the bit, which returns the third quadrant to the default state, in which it is used for shared memory.

This flag mechanism differs from how to set usage for the first and second quadrants. Set these values by using the magic number of the executable. (See the **-M** and **-N** options.)

+q4p flag

Control the flag bit setting to indicate how 32-bit processes use the third and fourth quadrant as data space.

The **enable** flag sets the flag bit to indicate that 32-bit processes use the fourth quadrant as a private data space. By setting the **+q4p** flag bit, the private data space increases from 1.9GB to 3.8GB for 32-bit processes. When you set the fourth quadrant for private data space, the third quadrant is automatically set for use as private data space, ignoring the current **+q3p** value.

The **disable** flag unsets the flag bit, which returns the fourth quadrant to the default state, in which it is used for shared memory. With **+q4p disable**, the value of the **+q3p** flag controls whether the third quadrant is used as a private data space or for shared memory.

This flag mechanism differs from how to set usage for the first and second quadrants. Set these values by using the magic number of the executable. (See the **-M** and **-N** options.)

OPTIONS FOR PA64 ELF chatr

PA64 ELF **chatr** is similar to SOM **chatr** but supports new options (and obsoletes others).

New options:

OPTIONS FOR PA64 ELF chatr (FORMAT 1)

- +cd** Set the code bit for the file's data segment(s).
- +ci** Set the code bit for the file's text segments(s).
- +md** Set the modification bit for the file's data segment(s).
- +mi** Set the modification bit for the file's text segment(s).

OPTIONS FOR PA64 ELF chatr (FORMAT 2)

With common options: **-s**, **-B mode**, **+k flag**, **+r flag**, **+s flag**, **+z flag**.

- +c** Set the code bit for a specified segment.
- +dz** Enable or disable lazy swap allocation for dynamically allocated segments (such as the stack or heap).
- +m** Set the modification bit for a specified segment.
- +p** Set the page size for a specified segment.
- +sa** Specify a segment using an address for a set of attribute modifications.
- +sall** Use all segments in the file for a set of attribute modifications.
- +si** Specify a segment using a segment index number for a set of attribute modifications.

chatr and MAGIC Numbers

The term **shared** applies to the magic number **SHARE_MAGIC** while the term **demand-loaded** applies to the magic number **DEMAND_MAGIC**. See *magic(4)* and the *HP-UX Linker and Libraries Online User Guide* for more information.

chatr labels the following type of executables in output.

- SHARE_MAGIC:** shared executable
- DEMAND_MAGIC:** demand load executable
- EXEC_MAGIC:** normal executable

SHMEM_MAGIC: normal **SHMEM_MAGIC** executable

The linker produces **SHARE_MAGIC** executables by default.

Using SHMEM_MAGIC

SHMEM_MAGIC is an interim solution until 64-bit addressability is available with a true 64-bit kernel.

SHMEM_MAGIC will not be supported on future HP implementations of 64-bit architectures (beyond PA2.0). Programs that need larger than 1.75 GB of shared memory on those architectures will have to be recompiled (as 64-bit executables) for those architectures.

Programs that are compiled as 64-bit executables on any 64-bit HP implementation (including PA 2.0) cannot be marked as **SHMEM_MAGIC** nor do they need to be as they will already have access to more than 1.75 GB of shared memory.

The additional 1 GB of shared memory that is available over other types of executables can be availed of only for system V shared memory and not other forms of shared memory (like memory mapped files).

Restricting Execute Permission on Stacks

A frequent or common method of breaking into systems is by maliciously overflowing buffers on a program's stack, such as passing unusually long, carefully chosen command line arguments to a privileged program that does not expect them. Malicious unprivileged users can use this technique to trick a privileged program into starting a superuser shell for them, or to perform similar unauthorized actions.

One simple yet highly effective way to reduce the risk from this type of attack is to remove the execute permission from a program's stack pages. This improves system security without sacrificing performance and has no negative effects on the vast majority of legitimate applications. The changes described in this section only affect the very small number of programs that try to execute (or are tricked into executing) instructions located on the program's stack(s).

If the stack protection feature described in this section is enabled for a program and that program attempts to execute code from its stack(s), the HP-UX kernel will terminate the program with a **SIGKILL** signal, display a message referring to this manual page section, and log an error message to the system message log (use **dmesg** to view the error message). The message logged by the kernel is:

```
WARNING: UID # may have attempted a buffer overflow attack. PID #
(program_name) has been terminated. See the '+es enable' option of
chatr(1).
```

If you see one of these messages, check with the program's owner to determine whether this program is legitimately executing code from its stack. If it is, you can use one or both of the methods described below to make the program functional again. If the program is not legitimately executing code from its stack, you should suspect malicious activity and take appropriate action.

HP-UX provides two options to permit legitimate execution from a program's stack(s). Combinations of these two options help make site-specific tradeoffs between security and compatibility.

The first method is the use of the **+es** option of **chatr** and affects individual programs. It is typically used to specify that a particular binary must be able to execute from its stack, regardless of the system default setting. This allows a restrictive system default while not preventing legitimate programs from executing code on their stack(s). Ideally this option should be set (if needed) by the program's provider, to minimize the need for manual intervention by whomever installs the program.

An alternate method is setting the kernel tunable parameter, **executable_stack**, to set a system-wide default for whether stacks are executable. Setting the **executable_stack** parameter to 1 (one) with **sam** (see **sam(1M)**) tells the HP-UX kernel to allow programs to execute on the program stack(s). Use this setting if compatibility with older releases is more important than security. Setting the **executable_stack** parameter to 0 (zero), the recommended setting, is appropriate if security is more important than compatibility. This setting significantly improves system security with minimal, if any, negative effects on legitimate applications.

Combinations of these settings may be appropriate for many applications. For example, after setting **executable_stack** to 0, you may find that one or two critical applications no longer work because they have a legitimate need to execute from their stack(s). Programs such as simulators or interpreters that use self-modifying code are examples you might encounter. To obtain the security benefits of a restrictive system default while still letting these specific applications run correctly, set **executable_stack** to 0, and run **chatr +es enable** on the specific binaries that need to execute code from their stack(s). These binaries can be easily identified when they are executed, because they will print error messages referring to

this manual page.

The possible settings for **executable_stack** are as follows:

executable_stack = 0

A setting of 0 causes stacks to be non-executable and is strongly preferred from a security perspective.

executable_stack = 1 (default)

A setting of 1 (the default value) causes all program stacks to be executable, and is safest from a compatibility perspective but is the least secure setting for this parameter.

executable_stack = 2

A setting of 2 is equivalent to a setting of 0, except that it gives non-fatal warnings instead of terminating a process that is trying to execute from its stack. Using this setting is helpful for users to gain confidence that using a value of 0 will not hurt their legitimate applications. Again, there is less security protection.

The table below summarizes the results from using the possible combinations of **chatr +es** and **executable_stack** when executing from the program's stack. Running **chatr +es disable** relies solely on the setting of the **executable_stack** kernel tunable parameter when deciding whether or not to grant execute permission for stacks and is equivalent to not having run **chatr +es** on the binary.

| chatr +es | executable_stack | ACTION |
|-----------------------------|------------------|--|
| enable | 1 | program runs normally |
| disable or chatr is not run | 1 | program runs normally |
| enable | 0 | program runs normally |
| disable or chatr is not run | 0 | program is killed |
| enable | 2 | program runs normally |
| disable or chatr is not run | 2 | program runs normally with warning displayed |

RETURN VALUE

chatr returns zero on success. If the command line contents is syntactically incorrect, or one or more of the specified files cannot be acted upon, **chatr** returns information about the files whose attributes could not be modified. If no files are specified, **chatr** returns decimal 255.

Illegal options

For PA32 **chatr**, if you use an illegal option, **chatr** returns the number of words in the command line. For example,

chatr +b enable +xyz enable returns 5 (because of illegal option **+xyz**).

chatr +b enable +xyz enable +mno file1 file2 returns 8.

For PA64 **chatr**, if you use an illegal option, **chatr** returns the number of *non-option words* present after the first illegal option.

chatr +b enable +xyz enable +mno enable +pqr enable file returns 4.

Invalid arguments

If you use an invalid argument with a valid option and you do not specify a filename, both PA32 and PA64 **chatr** return 0.

chatr +b <no argument> returns 0.

For PA32 **chatr**, if you specify a file name (regardless of whether or not the file exists), **chatr** returns number of words in the command line.

chatr +b <no argument> file returns 4.

For PA64 **chatr**, if you specify a file name (regardless of whether or not the file exists), **chatr** returns the number of files specified.

chatr +b <no argument> file1 file2 file3 returns 3.

Invalid files

For both PA32 and PA64 **chatr**, if the command cannot act on any of the files given, it returns the total number of files specified (if some option is specified). Otherwise it returns the number of files upon which it could not act.

chatr +b enable a1 a2 a3 a4 (where **a2** does not have read/write permission) returns 4.

chatr a1 a2 a3 a4 returns 1.

EXTERNAL INFLUENCES

Environment Variables

The following internationalization variables affect the execution of **chatr**:

| | |
|--------------------|--|
| LANG | Determines the locale category for native language, local customs and coded character set in the absence of LC_ALL and other LC_* environment variables. If LANG is not specified or is set to the empty string, a default of C (see <i>lang(5)</i>) is used instead of LANG . |
| LC_ALL | Determines the values for all locale categories and has precedence over LANG and other LC_* environment variables. |
| LC_CTYPE | Determines the locale category for character handling functions. |
| LC_MESSAGES | Determines the locale that should be used to affect the format and contents of diagnostic messages written to standard error. |
| LC_NUMERIC | Determines the locale category for numeric formatting. |
| NLSPATH | Determines the location of message catalogues for the processing of LC_MESSAGES . |

If any internationalization variable contains an invalid setting, **chatr** behaves as if all internationalization variables are set to **C**. See *environ(5)*.

In addition, the following environment variable affects **chatr**:

| | |
|---------------|---|
| TMPDIR | Specifies a directory for temporary files (see <i>tmpnam(3S)</i>). |
|---------------|---|

EXAMPLES

Change **a.out** to demand-loaded

```
chatr -q a.out
```

Change binding mode of program file that uses shared libraries to immediate and nonfatal. Also enable usage of **SHLIB_PATH** environment variable:

```
chatr -B immediate -B nonfatal +s enable a.out
```

Disallow run-time path lookup for the shared library **/usr/lib/libc.sl** that the shared library **libfoo.sl** depends on:

```
chatr +l /usr/lib/libc.sl libfoo.sl
```

Given segment index number 5 from a previous run of **chatr**, change the page size to 4 kilobytes:

```
chatr +si 5 +p 4K average64
```

AUTHOR

chatr was developed by HP.

SEE ALSO

System Tools:

ld(1) invoke the link editor

Miscellaneous:

a.out(4) assembler, compiler, and linker output
magic(4) magic number for HP-UX implementations
sam(1M) system administration manager

Texts and Tutorials:

HP-UX Linker and Libraries Online User Guide

(See the **+help** option)

HP-UX Linker and Libraries User's Guide

(See *manuals(5)* for ordering information)

NAME

checknr - check nroff/troff files

SYNOPSIS

checknr [-s] [-f] [-a.x1.y1.x2.y2xn.yn] [-c.x1.x2.x3...c .xn] [file ...]

DESCRIPTION

checknr searches a list of **nroff** or **troff** input files for certain kinds of errors involving mismatched opening and closing delimiters and unknown commands. If no files are specified, **checknr** searches the standard input. **checknr** looks for the following:

- Font changes using **\fx** ... **\fP**.
- Size changes using **\sx** ... **\s0**.
- Macros that come in *open ... close* forms, such as the **.TS** and **.TE** macros, which must appear in matched pairs.

checknr knows about the **ms** and **me** macro packages.

Options

checknr recognizes the following options:

- a Define additional macro pairs in the list. -a is followed by groups of six characters, each group defining a pair of macros. Each six characters consist of a period, the first macro name, another period, and the second macro name. For example, to define the pairs **.BS** and **.ES**, and **.XS** and **.XE**, use:

```
-a.BS.ES.XS.XE
```

No spaces are allowed between the option and its arguments.

- c Define commands that **checknr** would otherwise interpret as undefined.
- f Ignore **\fx** font changes.
- s Ignore **\sx** size changes.

EXTERNAL INFLUENCES**International Code Set Support**

Single-byte character code sets are supported.

DIAGNOSTICS

checknr complains about unmatched delimiters, unrecognized commands, and bad command syntax.

EXAMPLES

Check file **sorting** for errors that involve mismatched opening and closing delimiters and unknown commands, but disregard errors caused by font changes:

```
checknr -f sorting
```

WARNINGS

checknr is designed for use on documents prepared with the intent of using **checknr**, much the same as **lint** is used. It expects a certain document writing style for **\f...** and **\s...** commands, in which each **\fx** is terminated with **\fP** and each **\sx** is terminated with **\s0**. Although text files format properly when the next font or point size is coded directly instead of using **\fP** or **\s0**, such techniques produce complaints from **checknr**. If files are to be examined by **checknr**, the **\fP** and **\s0** delimiting conventions should be used.

-a cannot be used to define single-character macro names.

checknr does not recognize certain reasonable constructs such as conditionals.

AUTHOR

checknr was developed by the University of California, Berkeley.

SEE ALSO

checkeq(1), lint(1), nroff(1).

NAME

chfn - change user information; used by **finger**

SYNOPSIS

```
chfn [login-name]
chfn -r files [login-name]
chfn -r nis [login-name]
chfn -r nisplus [login-name]
chfn -r dce [login-name]
```

DESCRIPTION

The **chfn** command changes the user information that is stored in the **repository** for the current logged-in user or for the user specified by *login-name* (see *passwd(1)*).

The information is organized as four comma-separated subfields within the reserved (5th) field of the password file entry. It consists of the user's full name, location code, office phone number, and home phone number, in that order. This information is used by the **finger** command and other programs (see *finger(1)*).

chfn prompts you for each subfield. The prompt includes a default value, which is enclosed in brackets. Accept the default value by pressing the Return key. To enter a blank subfield, type the word **none**.

The DCE repository (**-r dce**) is only available if Integrated Login has been configured, see *auth.adm(1M)*. If Integrated Login has been configured, other considerations apply. A user with appropriate DCE privileges is capable of modifying a user's finger (gecos) information; this is not dependent upon superuser privileges.

If the repository is not specified; i.e., **chfn [login-name]**, the finger information is changed in the *passwd* file only.

Run **finger** after running **chfn** to make sure the information was processed correctly.

Options

The following option is recognized:

-r Specify the repository to which the operation is to be applied. Supported repositories include **files**, **nis**, **nisplus**, and **dce**.

Subfield Values

| | |
|---------------------|--|
| Name | Up to 1022 printing characters. The finger command and other utilities expand an & found anywhere in this subfield by substituting the login name for it and shifting the first letter of the login name to uppercase. (chfn does not alter the input & .) |
| Location | Up to 1022 printing characters. |
| Office Phone | Up to 25 printing characters. finger inserts appropriate hyphens if the value is all digits. |
| Home Phone | Up to 25 printing characters. finger inserts appropriate hyphens if the value is all digits. |

Security Restrictions

You must have appropriate privileges to use the optional *login-name* argument to change another user's information.

EXAMPLES

The following is a sample run. The user's input is shown in regular type.

```
Name [Tracy Simmons]:
Location (Ex: 47U-P5) []: 42L-P1
Office Phone (Ex: 1632) [77777]: 71863
Home Phone (Ex: 9875432) [4085551546]: none
```

WARNINGS

The encoding of office and extension information is installation-dependent.

For historical reasons, the user's name, etc., are stored in the `/etc/passwd` file. This is an inappropriate place to store the information.

Because two users may try to write the `passwd` file at once, a synchronization method was developed. On rare occasions, **chfn** prints a message that the password file is busy. When this occurs, **chfn** sleeps for a short time, then tries to write to the `passwd` file again.

C**AUTHOR**

chfn was developed by the University of California, Berkeley.

FILES

`/etc/passwd`
`/etc/ptmp`

NOTES

The **chfn** command is a hard link to `passwd` command. When **chfn** is executed, actually the `passwd` command gets executed with appropriate arguments to change the user *gecos* information in the *repository* specified in command line. If no *repository* is specified the *gecos* information is changed in `/etc/passwd` file.

SEE ALSO

`chsh(1)`, `finger(1)`, `passwd(1)`, `passwd(4)`.

NAME

chkey - change user's secure RPC key pair

SYNOPSIS

chkey [-p] [-s nisplus | nis | files]

DESCRIPTION

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.

chkey ensures that the login password and the secure-rpc password are kept the same.

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.

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.

Non root users are not allowed to change their key pair in the */etc/publickey* file.

Options

- p** Re-encrypt the existing secret key with the user's login password.
- s nisplus** Update the NIS+ database.
- s nis** Update the NIS database.
- s files** Update the **files** database.

AUTHOR

chkey was developed by Sun Microsystems, Inc.

FILES

/etc/nsswitch.conf
/etc/publickey

SEE ALSO

keylogin(1), keylogout(1), keyserv(1M), newkey(1M), nisaddcred(1M), nsswitch.conf(4), publickey(4).

NAME

chmod - change file mode access permissions

SYNOPSIS

/usr/bin/chmod [-A] [-R] *symbolic_mode_list* file ...

Obsolescent form:

/usr/bin/chmod [-A] [-R] *numeric_mode* file ...

DESCRIPTION

The **chmod** command changes the permissions of one or more *files* according to the value of *symbolic_mode_list* or *numeric_mode*. You can display the current permissions for a file with the **ls -l** command (see **ls(1)**).

Symbolic Mode List

A *symbolic_mode_list* is a comma-separated list of operations in the following form. Whitespace is not permitted.

[*who*]*op*[*permission*][, ...]

The variable fields can have the following values:

who One or more of the following letters:

- u** Modify permissions for user (owner).
- g** Modify permissions for group.
- o** Modify permissions for others.
- a** Modify permissions for all users (**a** is equivalent to **ugo**).

op Required; one of the following symbols:

- +** Add *permission* to the existing file mode bits of *who*.
- Delete *permission* from the existing file mode bits of *who*.
- =** Replace the existing mode bits of *who* with *permission*.

permission One or more of the following letters:

- r** Add or delete the read permission for *who*.
- w** Add or delete the write permission for *who*.
- x** Add or delete the execute file (search directory) permission for *who*.
- s** Add or delete the set-owner-ID-on-file-execution or set-group-ID-on-file-execution permission for *who*. Useful only if **u** or **g** is expressed or implied in *who*.
- t** Add or delete the save-text-image-on-file-execution (sticky bit) permission. Useful only if **u** is expressed or implied in *who*. See **chmod(2)**.
- X** Conditionally add or delete the execute/search permission as follows:
 - If *file* is a directory, add or delete the search permission to the existing file mode for *who*. (Same as **x**.)
 - If *file* is not a directory, and the current file permissions include the execute permission (**ls -l** displays an **x** or an **s**) for at least one of user, group, or other, then add or delete the execute file permission for *who*.
 - If *file* is not a directory, and no execute permissions are set in the current file mode, then do not change any execute permission.

Or one only of the following letters:

- u** Copy the current user permissions to *who*.
- g** Copy the current group permissions to *who*.
- o** Copy the current other permissions to *who*.

The operations are performed in the order specified, and can override preceding operations specified in the same command line.

If *who* is omitted, the **r**, **w**, **x**, and **X** permissions are changed for all users if the changes are permitted by the current file mode creation mask (see **umask(1)**). The **s** and **t** permissions are changed as if **a** was specified in *who*.

Omitting *permission* is useful only when used with **=** to delete all permissions.

Numeric Mode (Obsolescent)

Absolute permissions can be set by specifying a *numeric_mode*, an octal number constructed from the logical OR (sum) of the following mode bits:

Miscellaneous mode bits:

```
4000  (= u=s)  Set user ID on file execution (file only)
2000  (= g=s)  Set group ID on file execution (file only)
1000  (= u=t)  Set sticky bit; see below and chmod(2)
```

Permission mode bits:

```
0400  (= u=r)  Read by owner
0200  (= u=w)  Write by owner
0100  (= u=x)  Execute (search in directory) by owner
0040  (= g=r)  Read by group
0020  (= g=w)  Write by group
0010  (= g=x)  Execute/search by group
0004  (= o=r)  Read by others
0002  (= o=w)  Write by others
0001  (= o=x)  Execute/search by others
```

Options

- A Preserve any optional access control list (ACL) entries associated with the file (HFS file systems only). By default, in conformance with the IEEE Standard POSIX 1003.1-1988, optional HFS ACL entries are deleted. For JFS ACLs, this option has no effect, because optional JFS ACL entries are always preserved. For information about access control lists, see *acl(5)* and *aclv(5)*.
- R Recursively change the file mode bits. For each *file* operand that names a directory, **chmod** alters the file mode bits of the named directory and all files and subdirectories in the file hierarchy below it.

Only the owner of a file, or a user with appropriate privileges, can change its mode.

Only a user having appropriate privileges can set (or retain, if previously set) the sticky bit of a regular file.

If the sticky bit is set on a directory, files inside the directory may be renamed or removed only by the owner of the file, the owner of the directory, or the superuser (even if the modes of the directory would otherwise allow such an operation).

In order to set the set-group-ID bit, the group of the file must correspond to your current group ID.

If **chmod** is used on a symbolic link, the mode of the file referred to by the link is changed.

EXTERNAL INFLUENCES**Environment Variables**

LC_MESSAGES determines the language in which messages are displayed.

If **LC_MESSAGES** is not specified or is null, it defaults to the value of **LANG**. If **LANG** is not specified or is null, it defaults to **C** (see *lang(5)*).

If any internationalization variable contains an invalid setting, all internationalization variables default to **C**. See *environ(5)*.

International Code Set Support

Single- and multibyte character code sets are supported.

RETURN VALUE

Upon completion, **chmod** returns one of the following values:

```
0    Successful completion.
>0  An error condition occurred.
```

EXAMPLES

Deny write permission to others:

```
chmod o-w file
```

Make a file executable by everybody:

```
chmod a+x file
```

Assign read and execute permission to everybody, and set the set-user-ID bit:

```
chmod a=rx,u+s file
```

Assign read and write permission to the file owner, and read permission to everybody else:

```
chmod u=rw,go=r file
```

or the obsolescent form:

```
chmod 644 file
```

Traverse a directory subtree making all regular files readable by user and group only, and all executables and directories executable (searchable) by everyone:

```
chmod -R ug+r,o-r,a+X pathname
```

If the current value of **umask** is 020 (**umask -S** displays **u=rwx,g=rx,o=rwx**; do not change write permission for group) and the current permissions for file **mytest** are 444 (**a=r**), displayed by **ls -l** as **-r--r--r--**, then the command

```
chmod +w mytest
```

sets the permissions to 646 (**uo=rw,g=r**), displayed by **ls -l** as **-rw-r--rw-**.

If the current value of **umask** is 020 (**umask -S** displays **u=rwx,g=rx,o=rwx**; do not change write permission for group) and the current permissions for file **mytest** are 666 (**a=rw**), displayed by **ls -l** as **-rw-rw-rw-**, then the command

```
chmod -w mytest
```

sets the permissions to 464 (**uo=r,g=rw**), displayed by **ls -l** as **-r--rw-r--**.

DEPENDENCIES

The **-A** option causes **chmod** to fail on file systems that do not support ACLs.

AUTHOR

chmod was developed by AT&T and HP.

SEE ALSO

chacl(1), **ls(1)**, **umask(1)**, **chmod(2)**, **acl(5)**, **aclv(5)**.

STANDARDS CONFORMANCE

chmod: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

chown, chgrp - change file owner or group

SYNOPSIS

chown [-h] [-R] *owner[:group]* *file* ...

chgrp [-h] [-R] *group* *file* ...

DESCRIPTION

The **chown** command changes the owner ID of each specified *file* to *owner* and optionally the group ID of each specified *file* to *group*.

The **chgrp** command changes the group ID of each specified *file* to *group*.

owner can be either a decimal user ID or a login name found in the */etc/passwd* file.

group can be either a decimal group ID or a group name found in the */etc/group* file.

In order to change the owner or group, you must own the file and have the CHOWN privilege (see *setprivgrp*(1M)). If either command is invoked on a regular file by other than the superuser, the set-user-ID and set-group-ID bits of the file mode (04000 and 02000 respectively) are cleared. Note that a given user's or group's ability to use this command can be restricted by **setprivgrp** (see *setprivgrp*(1M)).

Access Control Lists – HFS File Systems Only

Users can permit or deny specific individuals and groups to access a file by setting optional ACL entries in the file's access control list (see *aclv*(5)). When using **chown** in conjunction with HFS ACLs, if the new owner and/or group of a file does not have an optional ACL entry corresponding to *user.%* and/or *%.group* in the file's access control list, the file's access permission bits remain unchanged. However, if the new owner and/or group is already designated by an optional ACL entry of *user.%* and/or *%.group* in the file's ACL, **chown** sets the corresponding file access permission bits (and the corresponding base ACL entries) to the permissions contained in that entry.

Access Control Lists – JFS File Systems Only

Users can permit or deny specific individuals and groups to access a file by setting optional ACL entries in the file's access control list (see *aclv*(5)). When using **chown** in conjunction with JFS ACLs, if the new owner and/or group of a file have optional ACL entries corresponding to **user:uid:perm** and/or **group:gid:perm** in the file's access control list, those entries remain in the ACL but no longer have any effect, being superseded by the file's **user::perm** and/or **group::perm** entries.

Options

chown and **chgrp** recognize the following options:

- h Change the owner or group of a symbolic link.

By default, the owner or group of the target file that a symbolic link points to is changed. With **-h**, the target file that the symbolic link points to is not affected. If the target file is a directory, and you specify **-h** and **-R**, recursion does not take place.

- R Recursively change the owner or group. For each *file* operand that names a directory, the owner or group of the directory and all files and subdirectories in the file hierarchy below it are changed.

EXTERNAL INFLUENCES**Environment Variables**

LC_MESSAGES determines the language in which messages are displayed.

If **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang*(5)) is used instead of **LANG**.

If any internationalization variable contains an invalid setting, **chown** behaves as if all internationalization variables are set to "C". See *environ*(5).

International Code Set Support

Single- and multi-byte character code sets are supported.

RETURN VALUE

chown and **chgrp** return the following values:

- 0 Successful completion.
- >0 An error condition occurred.

EXAMPLES

The following command changes the owner of the file **jokes** to **sandi**:

```
chown sandi jokes
```

The following command searches the directory **design_notes** and changes each file in that directory to owner **mark** and group **users**:

```
chown -R mark:users design_notes
```

WARNINGS

The default operation of **chown** and **chgrp** for symbolic links has changed as of HP-UX release 10.0. Use the **-h** option to get the former default operation.

FILES

```
/etc/group  
/etc/passwd
```

SEE ALSO

chmod(1), **setprivgrp(1M)**, **chown(2)**, **group(4)**, **passwd(4)**, **acl(5)**, **aclv(5)**.

STANDARDS CONFORMANCE

chown: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

chgrp: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

chsh - change default login shell

SYNOPSIS

```
chsh login-name [shell]
chsh -r files login-name [shell]
chsh -r nisplus login-name [shell]
chsh -r nis login-name [shell]
chsh -r dce login-name [shell]
```

DESCRIPTION

The **chsh** command changes the login-shell for a user's login name in the repository (see *passwd(1)*).

The DCE repository (**-r dce**) is only available if Integrated Login has been configured, see *auth.adm(1M)*. If Integrated Login has been configured, other considerations apply. A user with appropriate DCE privileges is capable of modifying a user's shell; this is not dependent upon superuser privileges.

If the repository is not specified; i.e., **chsh [login-name]**, the login shell is changed in *passwd* file only.

Run **finger** after running **chsh** to make sure the information was processed correctly.

Arguments

login-name A login name of a user.

shell The absolute path name of a shell. If the file **/etc/shells** exists, the new login shell must be listed in that file. Otherwise, you can specify one of the standard shells listed in the *getusershell(3C)* manual entry. If *shell* is omitted, it defaults to the POSIX shell, **/usr/bin/sh**.

Options

The following option is recognized:

-r Specify the repository to which the operation is to be applied. Supported repositories include **files**, **nis**, **nisplus**, and **dce**.

Security Restrictions

You must have the **syslog** sensitivity label to execute **chsh**. You must have the **owner** kernel authorization to change another user's login shell.

NETWORKING FEATURES**NFS**

File **/etc/passwd** can be implemented as a Network Information Service (NIS) database.

EXAMPLES

To change the login shell for user **voltaire** to the default:

```
chsh voltaire
```

To change the login shell for user **descartes** to the C shell:

```
chsh descartes /usr/bin/csh
```

To change the login shell for user **aristotle** to the Korn shell in the DCE registry:

```
chsh -r dce aristotle /usr/bin/ksh
```

WARNINGS

As many users may try to write the **/etc/passwd** file simultaneously, a *passwd* locking mechanism was devised. If this locking fails after subsequent retrying, **chsh** terminates.

AUTHOR

chsh was developed by HP and the University of California, Berkeley.

NOTES

The **chsh** command is a hard link to **passwd** command. When **chsh** is executed actually the **passwd** command gets executed with appropriate arguments to change the user login shell in the *repository* specified in command line. If no *repository* is specified the login shell is changed in **/etc/passwd** file.

FILES

/etc/shells
/etc/ptmp

SEE ALSO

chfn(1), csh(1), ksh(1), passwd(1), sh(1), sh-bourne(1), sh-posix(1), getusershell(3C), pam(3), passwd(4), shells(4).

C

NAME

ci - check in RCS revisions

SYNOPSIS

ci [*options*] *file...*

DESCRIPTION

ci stores new revisions into RCS files. Each file name ending in ,v is treated as an RCS file; all others are assumed to be working files. ci deposits the contents of each working file into the corresponding RCS file (see *rcsintro*(5)).

If the RCS file does not exist, ci creates it and deposits the contents of the working file as the initial revision. The default number is "1.1". The access list is initialized to empty. Instead of the log message, ci requests descriptive text (see the -t option below).

An RCS file created by ci inherits the read and execute permissions from the working file. If the RCS file exists, ci preserves its read and execute permissions. ci always turns off all write permissions of RCS files.

The caller of the command must have read/write permission for the directories containing the RCS file and the working file, and read permission for the RCS file itself. A number of temporary files are created. A semaphore file is created in the directory containing the RCS file. ci always creates a new RCS file and unlinks the old one; therefore links to RCS files are useless.

For ci to work, the user's login must be in the access list unless the access list is empty, the user is the owner of the file, or the user is super-user.

Normally, ci checks whether the revision to be deposited is different from the preceding one. If it is not different, ci either aborts the deposit (if -q is given) or asks whether to abort (if -q is omitted). A deposit can be forced with the -f option.

If sufficient memory is not available for checking the difference between the revision to be deposited and the preceding one, then either swap or maxdsiz values can be increased.

For each revision deposited, ci prompts for a log message. The log message should summarize the change and must be terminated with a line containing a single "." or a control-D. If several files are being checked in, ci asks whether or not to reuse the log message from the previous file. If the standard input is not a terminal, ci suppresses the prompt and uses the same log message for all files (see -m option below).

The number of the deposited revision can be given with any of the options -r, -f, -k, -l, -u, or -q (see -r option below).

To add a new revision to an existing branch, the head revision on that branch must be locked by the caller. Otherwise, only a new branch can be created. This restriction is not enforced for the owner of the file, unless locking is set to strict (see *rcs*(1)). A lock held by someone else can be broken with the *rcs* command (see *rcs*(1)).

Options

- f[*rev*] Forces a deposit. The new revision is deposited even if it is not different from the preceding one.
- k[*rev*] Searches the working file for keyword values to determine its revision number, creation date, author, and state (see *co*(1)), and assigns these values to the deposited revision, rather than computing them locally. A revision number given with a command option overrides the number in the working file. This option is useful for software distribution. A revision that is sent to several sites should be checked in with the -k option at these sites to preserve its original number, date, author, and state.
- l[*rev*] Works like -r, except it performs an additional co -l for the deposited revision. Thus, the deposited revision is immediately checked out again and locked. This is useful for saving a revision although one wants to continue editing it after the check-in.
- m" *msg*" Uses the string *msg* as the log message for all revisions checked in.
- n" *name*" Assigns the symbolic name *name* to the checked-in revision. ci prints an error message if *name* is already assigned to another number.
- N" *name*" Same as -n, except that it overrides a previous assignment of *name*.

- q[*rev*]** Quiet mode; diagnostic output is not printed. A revision that is not different from the preceding one is not deposited unless **-f** is given.
- r[*rev*]** Assigns the revision number *rev* to the checked-in revision, releases the corresponding lock, and deletes the working file. This is the default.
- If *rev* is omitted, **ci** derives the new revision number from the caller's last lock. If the caller has locked the head revision of a branch, the new revision is added to the head of that branch and a new revision number is assigned to the new revision. The new revision number is obtained by incrementing the head revision number. If the caller locked a non-head revision, a new branch is started at the locked revision, and the number of the locked revision is incremented. The default initial branch and level numbers are 1. If the caller holds no lock, but is the owner of the file and locking is not set to *strict*, the revision is added to the head of the trunk.
- If *rev* indicates a revision number, it must be higher than the latest one on the branch to which *rev* belongs, or must start a new branch.
- If *rev* indicates a branch instead of a revision, the new revision is added to the head of that branch. The level number is obtained by incrementing the head revision number of that branch. If *rev* indicates a non-existing branch, that branch is created with the initial revision numbered *rev*.1.
- NOTE: On the trunk, revisions can be added to the head, but not inserted.
- s "*state*"** Sets the state of the checked-in revision to the identifier *state*. The default is **Exp**.
- t[*txtfile*]** Writes descriptive text into the RCS file (deletes the existing text). If *txtfile* is omitted, **ci** prompts the user for text from standard input that is terminated with a line containing a single **.** or Ctrl-D. Otherwise, the descriptive text is copied from the file *txtfile*. During initialization, descriptive text is requested even if **-t** is not given. The prompt is suppressed if standard input is not a terminal.
- u[*rev*]** Similar to **-l**, except that the deposited revision is not locked. This is useful if one wants to process (e.g., compile) the revision immediately after check in.

Access Control Lists (ACLs)

Optional ACL entries should not be added to RCS files, because they might be deleted.

DIAGNOSTICS

For each revision, **ci** prints the RCS file, the working file, and the number of both the deposited and the preceding revision. The exit status always refers to the last file checked in, and is 0 if the operation was successful, 1 if unsuccessful.

EXAMPLES

If the current directory contains a subdirectory **RCS** with an RCS file **io.c,v**, all of the following commands deposit the latest revision from **io.c** into **RCS/io.c,v**:

```
ci io.c
ci RCS/io.c,v
ci io.c,v
ci io.c RCS/io.c,v
ci io.c io.c,v
ci RCS/io.c,v io.c
ci io.c,v io.c
```

Check in version 1.2 of RCS file **foo.c,v**, with the message **Bug fix**:

```
ci -r1.2 -m"Bug Fix" foo.c,v
```

WARNINGS

The names of RCS files are generated by appending **,v** to the end of the working file name. If the resulting RCS file name is too long for the file system on which the RCS file should reside, **ci** terminates with an error message.

The log message cannot exceed 2046 bytes.

A file with approximately 240 revisions may cause a hash table overflow. **ci** cannot add another revision to the file until some of the old revisions have been removed. Use the **rcs -o** (obsolete) command option

to remove old revisions.

RCS is designed to be used with TEXT files only. Attempting to use RCS with non-text (binary) files results in data corruption.

AUTHOR

ci was developed by Walter F. Tichy.

SEE ALSO

co(1), **ident(1)**, **rcs(1)**, **rcsdiff(1)**, **rcsmerge(1)**, **rlog(1)**, **rcsfile(4)**, **acl(5)**, **rcsintro(5)**.


C

NAME

ckconfig - verify the path names of all the FTP configuration files.

SYNOPSIS

`/usr/bin/ckconfig`

DESCRIPTION

The `ckconfig` utility is used to verify the path names of the FTP configuration files, `/etc/ftpd/ftpusers`, `/etc/ftpd/ftpaccess`, `/etc/ftpd/ftpconversions`, `/etc/ftpd/ftpgroups`, `/etc/ftpd/ftphosts`, `/var/adm/syslog/xferlog`, and `/etc/ftpd/pids/*`.

This utility checks to see that all the FTP configuration files are in the path specified. If it is not able to find the configuration files in the path, it will give out an error message to the system administrator about the same.

FILES

`/usr/bin/ckconfig`

AUTHOR

`ckconfig` was developed by the Washington University, St. Louis, Missouri.

SEE ALSO

`ftpusers(4)`, `ftpconversions(4)`, `ftpaccess(4)`, `ftphosts(4)`, `ftpgroups(4)`, `xferlog(5)`.

NAME

cksum - print file checksum and sizes

SYNOPSIS

cksum [*file* ...]

DESCRIPTION

The **cksum** command calculates and prints to standard output a checksum for each named file, the number of octets in the file and the filename.

cksum uses a portable algorithm based on a 32-bit Cyclic Redundancy Check. This algorithm finds a broader spectrum of errors than the 16-bit algorithms used by **sum** (see *sum(1)*). The CRC is the sum of the following expressions, where *x* is each byte of the file.

$$x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^7 + x^5 + x^4 + x^2 + x^1 + x^0$$

The results of the calculation are truncated to a 32-bit value. The number of bytes in the file is also printed.

Standard input is used if no file names are given.

cksum is typically used to verify data integrity when copying files between systems.

EXTERNAL INFLUENCES**Environment Variables**

LANG determines the locale to use for the locale categories when both **LC_ALL** and the corresponding environment variable (beginning with **LC_**) do not specify a locale. If **LANG** is not set or is set to the empty string, a default of "C" (see *lang(5)*) is used.

LC_CTYPE determines the locale for interpretation of sequences of bytes of text data as characters (e.g., single- versus multibyte characters in arguments and input files).

LC_MESSAGES determines the language in which messages are displayed.

If any internationalization variable contains an invalid setting, **cksum** behaves as if all internationalization variables are set to "C". See *environ(5)*.

RETURN VALUE

Upon completion, **cksum** returns one of the following values:

- 0 All files were processed successfully.
- >0 One or more files could not be read or another error occurred.

If an inaccessible file is encountered, **cksum** continues processing any remaining files, but the final exit status is affected.

SEE ALSO

sum(1), *wc(1)*, *pdf(4)*.

STANDARDS CONFORMANCE

cksum: XPG4, POSIX.2

NAME

clear - clear terminal screen

SYNOPSIS

clear

DESCRIPTION

clear clears the terminal screen if it is possible to do so. It reads the **TERM** environment variable for the terminal type, then reads the appropriate **terminfo** database to determine how to clear the screen.

FILES

/usr/share/lib/terminfo/?/* terminal database files

AUTHOR

clear was developed by the University of California, Berkeley.

SEE ALSO

terminfo(4).

C

NAME

cmp - compare two files

SYNOPSIS

cmp [-1] [-s] *file1 file2* [*skip1* [*skip2*]]

DESCRIPTION

cmp compares two files (if *file1* or *file2* 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 number at which the difference occurred. If one file is an initial subsequence of the other, that fact is noted. *skip1* and *skip2* are initial byte offsets into *file1* and *file2*, respectively; and maybe octal or decimal; the form of the number is determined by the environment variable **LC_NUMERIC** (in the C locale, a leading 0 denotes an octal number. See **LANG** on *environ*(5) and *strtol*(3C)).

cmp recognizes the following options:

- 1 Print the byte number (decimal) and the differing bytes (octal) for each difference (byte numbering begins at 1 rather than 0).
- s Print nothing for differing files; return codes only.

EXTERNAL INFLUENCES**Environment Variables**

LANG determines the language in which messages are displayed. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang*(5)) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **cmp** behaves as if all internationalization variables are set to "C". See *environ*(5).

International Code Set Support

Single- and multi-byte character code sets are supported.

DIAGNOSTICS

cmp returns the following exit values:

- 0 Files are identical.
- 1 Files are not identical.
- 2 Inaccessible or missing argument.

cmp prints the following warning if the comparison succeeds till the end of file of *file1*(*file2*) is reached.

cmp: EOF on file1(file2)

SEE ALSO

comm(1), *diff*(1).

STANDARDS CONFORMANCE

cmp: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

co - check out RCS revisions

SYNOPSIS

co [*options*] *file* ...

DESCRIPTION

co retrieves revisions from RCS files. Each file name ending in *,v* is taken to be an RCS file. All other files are assumed to be working files. co retrieves a revision from each RCS file and stores it in the corresponding working file (see also *rcsintro*(5)).

Revisions of an RCS file can be checked out locked or unlocked. Locking a revision prevents overlapping updates. A revision checked out for reading or processing (e.g., compiling) need not be locked. A revision checked out for editing and later checked in must normally be locked. Locking a revision currently locked by another user fails (a lock can be broken with the *rcs* command, but poses inherent risks when independent changes are being made simultaneously (see *rcs*(1)). co with locking requires the caller to be on the access list of the RCS file unless: he is the owner of the file, a user with appropriate privileges, or the access list is empty. co without locking is not subject to access list restrictions.

A revision is selected by number, check-in date/time, author, or state. If none of these options are specified, the latest revision on the trunk is retrieved. When the options are applied in combination, the latest revision that satisfies all of them is retrieved. The options for date/time, author, and state retrieve a revision on the selected branch. The selected branch is either derived from the revision number (if given), or is the highest branch on the trunk. A revision number can be attached to the options *-l*, *-p*, *-q*, or *-r*.

The caller of the command must have write permission in the working directory, read permission for the RCS file, and either read permission (for reading) or read/write permission (for locking) in the directory that contains the RCS file.

The working file inherits the read and execute permissions from the RCS file. In addition, the owner write permission is turned on, unless the file is checked out unlocked and locking is set to *strict* (see *rcs*(1)).

If a file with the name of the working file exists already and has write permission, co aborts the check out if *-q* is given, or asks whether to abort if *-q* is not given. If the existing working file is not writable, it is deleted before the check out.

A number of temporary files are created. A semaphore file is created in the directory of the RCS file to prevent simultaneous update.

A co command applied to an RCS file with no revisions creates a zero-length file. co always performs keyword substitution (see below).

Options

- l[rev]* Locks the checked out revision for the caller. If omitted, the checked out revision is not locked. See option *-r* for handling of the revision number *rev*.
- p[rev]* Prints the retrieved revision on the standard output rather than storing it in the working file. This option is useful when co is part of a pipe.
- q[rev]* Quiet mode; diagnostics are not printed.
- ddate* Retrieves the latest revision on the selected branch whose check in date/time is less than or equal to *date*. The date and time may be given in free format and are converted to local time. Examples of formats for *date*:

| | |
|--------------------------------|-------------------------------------|
| Tue-PDT, 1981, 4pm Jul 21 | (free format) |
| Fri April 16 15:52:25 EST 1982 | (output of <i>ctime</i> (3C)) |
| 4/21/86 10:30am | (format: <i>mm/dd/yy hh:mm:ss</i>) |

Most fields in the date and time can be defaulted. co determines the defaults in the order year, month, day, hour, minute, and second (from most- to least-significant). At least one of these fields must be provided. For omitted fields that are of higher significance than the highest provided field, the current values are assumed. For all other omitted fields, the lowest possible values are assumed. For example, the date *20, 10:30* defaults to 10:30:00 of the 20th of the current month and current year. Date/time fields can be delimited by spaces or commas. If spaces are used, the string must be surrounded by double quotes.

For 2-digit year input (*yy*) without the presence of the century field, the following interpretation is taken: [70-99, 00-69 (1970-1999, 2000-2069)].

- r[*rev*]** Retrieves the latest revision whose number is less than or equal to *rev*. If *rev* indicates a branch rather than a revision, the latest revision on that branch is retrieved. *rev* is composed of one or more numeric or symbolic fields separated by **.**. The numeric equivalent of a symbolic field is specified with the **ci -n** and **rcs -n** commands (see **ci(1)** and **rcs(1)**).
- s *state*** Retrieves the latest revision on the selected branch whose state is set to *state*.
- w[*login*]** Retrieves the latest revision on the selected branch that was checked in by the user with login name *login*. If the argument *login* is omitted, the caller's login is assumed.
- j *joinlist*** Generates a new revision that is the result of the joining of the revisions on *joinlist*. *joinlist* is a comma-separated list of pairs of the form *rev2:rev3*, where *rev2* and *rev3* are (symbolic or numeric) revision numbers. For the initial pair, *rev1* denotes the revision selected by the options **-l**, **...**, **-w**. For all other pairs, *rev1* denotes the revision generated by the previous pair. (Thus, the output of one join becomes the input to the next.)

For each pair, **co** joins revisions *rev1* and *rev3* with respect to *rev2*. This means that all changes that transform *rev2* into *rev1* are applied to a copy of *rev3*. This is particularly useful if *rev1* and *rev3* are the ends of two branches that have *rev2* as a common ancestor. If *rev1* < *rev2* < *rev3* on the same branch, joining generates a new revision that is similar to *rev3*, but with all changes that lead from *rev1* to *rev2* undone. If changes from *rev2* to *rev1* overlap with changes from *rev2* to *rev3*, **co** prints a warning and includes the overlapping sections, delimited as follows:

```
<<<<<<<
rev1
=====
rev3
>>>>>>>
```

For the initial pair, *rev2* can be omitted. The default is the common ancestor. If any of the arguments indicate branches, the latest revisions on those branches are assumed. If the **-l** option is present, the initial *rev1* is locked.

Keyword Substitution

Strings of the form **\$keyword\$** and **\$keyword:...\$** embedded in the text are replaced with strings of the form **\$keyword: value \$**, where *keyword* and *value* are pairs listed below. Keywords may be embedded in literal strings or comments to identify a revision.

Initially, the user enters strings of the form **\$keyword\$**. On check out, **co** replaces these strings with strings of the form **\$keyword: value \$**. If a revision containing strings of the latter form is checked back in, the value fields are replaced during the next checkout. Thus, the keyword values are automatically updated on checkout.

Keywords and their corresponding values:

- \$Author\$** The login name of the user who checked in the revision.
- \$Date\$** The date and time the revision was checked in.
- \$Header\$** A standard header containing the RCS file name, the revision number, the date, the author, and the state.
- \$Locker\$** The login name of the user who locked the revision (empty if not locked).
- \$Log\$** The log message supplied during checkin, preceded by a header containing the RCS file name, the revision number, the author, and the date. Existing log messages are *not* replaced. Instead, the new log message is inserted after **\$Log:...\$**. This is useful for accumulating a complete change log in a source file.
- \$Revision\$** The revision number assigned to the revision.
- \$Source\$** The full pathname of the RCS file.
- \$State\$** The state assigned to the revision with **rcs -s** or **ci -s**.

Access Control Lists (ACLs)

Optional ACL entries should not be added to RCS files because they might be deleted.

DIAGNOSTICS

The RCS file name, the working file name, and the revision number retrieved are written to the diagnostic output. The exit status always refers to the last file checked out, and is 0 if the operation was successful, 1 if unsuccessful.

EXAMPLES

Assume the current directory contains a subdirectory named **RCS** with an RCS file named **io.c,v**. Each of the following commands retrieves the latest revision from **RCS/io.c,v** and stores it into **io.c**:

```
co io.c
co RCS/io.c,v
co io.c,v
co io.c RCS/io.c,v
co io.c io.c,v
co RCS/io.c,v io.c
co io.c,v io.c
```

Check out version 1.1 of RCS file **foo.c,v**:

```
co -r1.1 foo.c,v
```

Check out version 1.1 of RCS file **foo.c,v** to the standard output:

```
co -p1.1 foo.c,v
```

Check out the version of file **foo.c,v** that existed on September 18, 1992:

```
co -d"09/18/92" foo.c,v
```

WARNINGS

The **co** command generates the working file name by removing the **,v** from the end of the RCS file name. If the given RCS file name is too long for the file system on which the RCS file should reside, **co** terminates with an error message.

There is no way to suppress the expansion of keywords, except by writing them differently. In **nroff** and **troff**, this is done by embedding the null-character **\&** into the keyword.

The **-d** option gets confused in some circumstances, and accepts no date before 1970.

The **-j** option does not work for files containing lines consisting of a single **.**.

RCS is designed to be used with *text* files only. Attempting to use RCS with non-text (binary) files results in data corruption.

AUTHOR

co was developed by Walter F. Tichy.

SEE ALSO

ci(1), **ident(1)**, **rcs(1)**, **rcsdiff(1)**, **rcsmerge(1)**, **rlog(1)**, **rcsfile(4)**, **acl(5)**, **rcsintro(5)**.

NAME

col - filter reverse line-feeds and backspaces

SYNOPSIS

col [-blfxp]

DESCRIPTION

col 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**). **col** is particularly useful for filtering multi-column output made with the **nroff .rt** command, and output resulting from use of the **tbl** preprocessor (see **nroff(1)** and **tbl(1)**).

If the **-b** option is given, **col** 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 is output.

If the **-l** option is given, **col** assumes the output device is a line printer (rather than a character printer) and removes backspaces in favor of multiply overstruck full lines. It generates the minimum number of print operations necessary to generate the required number of overstrikes. (All but the last print operation on a line are separated by carriage returns (**\r**); the last print operation is terminated by a newline (**\n**).)

Although **col** 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 **-f** (fine) option; in this case, the output from **col** may contain forward half-line feeds (**ESC-9**), but will still never contain either kind of reverse line motion.

Unless the **-x** option is given, **col** converts white space to tabs on output wherever possible to shorten printing time.

The ASCII control characters **SO (\016)** and **SI (\017)** are assumed by **col** 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.

On input, the only control characters accepted are space, backspace, tab, return, new-line, **SI**, **SO**, and **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.

Normally, **col** ignores any unrecognized escape sequences found in its input; the **-p** option can be used to cause **col** 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.

EXTERNAL INFLUENCES**Environment Variables**

LANG provides a default value for the internationalization variables that are unset or null. If **LANG** is unset or null, the default value of "C" (see *lang(5)*) is used. If any of the internationalization variables contains an invalid setting, **col** will behave as if all internationalization variables are set to "C". See *environ(5)*.

LC_ALL If set to a non-empty string value, overrides the values of all the other internationalization variables.

LC_CTYPE determines the interpretation of text as single and/or multi-byte characters, the classification of characters as printable, and the characters matched by character class expressions in regular expressions.

LC_MESSAGES determines the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.

NLSPATH determines the location of message catalogues for the processing of **LC_MESSAGES**.

International Code Set Support

Single- and multi-byte character code sets are supported.

EXAMPLES

col is used most often with **nroff** and **tbl**. A common usage is:

```
tbl filename | nroff -man | col | more -s
```

(very similar to the usual *man*(1) command). This command allows vertical bars and outer boxes to be printed for tables. The file is run through the **tbl** preprocessor, and the output is then piped through **nroff**, formatting the output using the **-man** macros. The formatted output is then piped through **col**, which sets up the vertical bars and aligns the columns in the file. The file is finally piped through the **more** command, which prints the output to the screen with underlining and highlighting substituted for italic and bold typefaces. The **-s** option deletes excess space from the output so that multiple blank lines are not printed to the screen.

C

SEE ALSO

nroff(1), tbl(1), ul(1), man(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** is a device that can interpret half-line motions, and **-Tlp** otherwise.

BUGS

Cannot back up more than 128 lines. Cannot back up across page boundaries.

There is a maximum limit for the number of characters, including backspaces and overstrikes, on a line. The maximum limit is at least 800 characters.

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.

WARNINGS

This command is likely to be withdrawn from X/Open standards. Applications using this command might not be portable to other vendors' systems.

STANDARDS CONFORMANCE

col: SVID2, SVID3, XPG2, XPG3

NAME

comb - combine SCCS deltas

SYNOPSIS

comb [-p *SID*] [-c *list*] [-o] [-s] *file* ...

DESCRIPTION

The **comb** command generates a shell procedure (see *sh(1)*) which, when run, reconstructs the given SCCS files. The reconstructed files are usually smaller than the original files. Arguments can be specified in any order, but all options apply to all named SCCS files. If a directory is named, **comb** behaves as though each file in the directory were specified as a named file, except that non-SCCS files (last component of the path name does not begin with **s.**) and unreadable files are silently ignored. If a name of - is given, the standard input is read; each line of the standard input is taken to be the name of an SCCS file to be processed; non-SCCS files and unreadable files are silently ignored. The generated shell procedure is written on the standard output.

Options

comb recognizes the following options. Each is explained as if only one named file is to be processed, but the effects of any option apply independently to each named file.

- p *SID* The SCCS *I*Dentification string (SID) of the oldest delta to be preserved. All older deltas are discarded in the reconstructed file.
- c *list* A *list* of deltas to be preserved (see *get(1)* for the syntax of a *list*). All other deltas are discarded.
- o For each **get -e** generated, this option causes the reconstructed file to be accessed at the release of the delta to be created, otherwise the reconstructed file would be accessed at the most recent ancestor. Use of the -o option can decrease the size of the reconstructed SCCS file. It can also alter the shape of the delta tree of the original file.
- s This option causes **comb** to generate a shell procedure which, when run, produces a report giving, for each file: the file name, size (in blocks) after combining, original size (also in blocks), and percentage change computed by:

$$100 \times (\text{original} - \text{combined}) / \text{original}$$

It is recommended that this option be used before any SCCS files are actually combined to determine exactly how much space is saved by the combining process.

If no options are specified, **comb** preserves only leaf deltas and the minimal number of ancestors needed to preserve the tree.

EXTERNAL INFLUENCES**International Code Set Support**

Single- and multi-byte character code sets are supported.

DIAGNOSTICS

Use *scscshelp(1)* for explanations.

EXAMPLES

The command:

```
comb -c1.1,1.3,1.6 s.document > save_file
```

creates a shell script named **save_file**, which if executed, creates a new **s.document** using only the deltas **1.1**, **1.3**, and **1.6** from the old **s.document**. The script overwrites the old **s.document**; thus, it might be wise to copy the original elsewhere. Here is an example of typical technique:

```
cp s.document s.save
comb -c1.1,1.3,1.6 s.document > save_file
sh save_file
```

WARNINGS

comb may rearrange the shape of the tree of deltas. Combining files may or may not save space; in fact, it is possible for the reconstructed file to actually be larger than the original.

FILES

| | |
|--------------------|----------------|
| s.COMB????? | Temporary file |
| comb????? | Temporary file |

SEE ALSO

admin(1), delta(1), get(1), sccshelp(1), prs(1), sh(1), sccsfile(4).

C

NAME

comm - select or reject lines common to two sorted files

SYNOPSIS

comm [-[123]] *file1 file2*

DESCRIPTION

comm reads *file1* and *file2*, which should be ordered in increasing collating sequence (see *sort(1)* and Environment Variables below), and produces a three-column output:

Column 1: Lines that appear only in *file1*,
 Column 2: Lines that appear only in *file2*,
 Column 3: Lines that appear in both files.

If - is used for *file1* or *file2*, the standard input is used.

Options 1, 2, or 3 suppress printing of the corresponding column. Thus **comm -12** prints only the lines common to the two files; **comm -23** prints only lines in the first file but not in the second; **comm -123** does nothing useful.

EXTERNAL INFLUENCES**Environment Variables**

LC_COLLATE determines the collating sequence **comm** expects from the input files.

LC_MESSAGES determines the language in which messages are displayed.

If **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** determines the language in which messages are displayed. If **LC_COLLATE** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **comm** behaves as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single- and multi-byte character code sets are supported.

EXAMPLES

The following examples assume that *file1* and *file2* have been ordered in the collating sequence defined by the **LC_COLLATE** or **LANG** environment variable.

Print all lines common to *file1* and *file2* (in other words, print column 3):

```
comm -12 file1 file2
```

Print all lines that appear in *file1* but not in *file2* (in other words, print column 1):

```
comm -23 file1 file2
```

Print all lines that appear in *file2* but not in *file1* (in other words, print column 2):

```
comm -13 file1 file2
```

SEE ALSO

cmp(1), *diff(1)*, *sdiff(1)*, *sort(1)*, *uniq(1)*.

STANDARDS CONFORMANCE

comm: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

`command` - execute a simple command

SYNOPSIS

command *command_name* [*argument* ...]

DESCRIPTION

command enables the shell to treat the arguments as a simple command, suppressing the shell function lookup.

If *command_name* is not the name of the function, the effect of **command** is the same as omitting *command*.

OPERANDS

command recognizes the following operands:

| | |
|---------------------|---|
| <i>command_name</i> | The name of a HP-UX command or a shell built-in command. |
| <i>argument</i> | One or more strings to be interpreted as arguments to <i>command_name</i> . |

The **command** command is necessary to allow functions that have the same name as a command to call the command (instead of a recursive call to the function).

Nothing in the description of **command** is intended to imply that the command line is parsed any differently than any other simple command. For example,

```
command a | b ; c
```

is not parsed in any special way that causes `|` or `;` to be treated other than a pipe operator or semicolon or that prevents function lookup on *b* or *c*.

EXTERNAL INFLUENCE**Environment Variables**

PATH determines the search path used during the command search.

RETURN VALUE

command exits with one of the following values:

- If **command** fails:
 - 126** The utility specified by the *command_name* is found but not executable.
 - 127** An error occurred in the **command** utility or the utility specified by *command_name* is not found.
- If **command** does not fail:

The exit status of **command** is the same as that of the simple command specified by the arguments:

```
command_name[argument ...]
```

EXAMPLES

Create a version of the `cd` command that always prints the name of the new working directory whenever it is used:

```
cd() {
    command "$@" >/dev/null
    pwd
}
```

Circumvent the redefined `cd` command above, and change directories without printing the name of the new working directory:

```
command cd
```

SEE ALSO

`getconf(1)`, `sh-posix(1)`, `confstr(3C)`.

command(1)

command(1)

STANDARDS CONFORMANCE

command: XPG4, POSIX.2



C

NAME

compact, uncompact, ccat - compact and uncompact files, and cat them

SYNOPSIS

compact [*name* ...]

uncompact [*name* ...]

ccat [*file* ...]

DESCRIPTION

compact compresses the named files using an adaptive Huffman code. If no file names are given, standard input is compacted and sent to the standard output. **compact** operates as an on-line algorithm. Each time a byte is read, it is encoded immediately according to the current prefix code. This code is an optimal Huffman code for the set of frequencies seen so far. It is unnecessary to attach a decoding tree in front of the compressed file because the encoder and the decoder start in the same state and stay synchronized. Furthermore, **compact** and **uncompact** can operate as filters. In particular,

... | **compact** | **uncompact** | ...

operates as a (very slow) no-op.

When an argument *file* is given, it is compacted, the resulting file is placed in *file.C*, and *file* is unlinked. The first two bytes of the compacted file code the fact that the file is compacted. These bytes are used to prohibit recompaction.

The amount of compression to be expected depends on the type of file being compressed. Typical file size reduction (in percent) through compression are: Text, 38%; Pascal Source, 43%; C Source, 36%; and Binary, 19%.

uncompact restores the original file from a file compressed by **compact**. If no file names are specified, standard input is uncompactd and sent to the standard output.

ccat cats the original file from a file compressed by **compact**, without uncompressing the file.

Access Control Lists (ACLs)

On systems that implement access control lists, when a new file is created with the effective user and group ID of the caller, the original file's ACL is copied to the new file after being altered to reflect any change in ownership (see *acl(5)* and *aclv(5)*). In JFS file systems, files created by **compact**, **uncompact** or **ccat** do not inherit their parent directory's default ACL entries (if any), but instead retain their original ACLs. When a file being compacted or uncompactd resides on a JFS file system, and the compacted or uncompactd file resides on an HFS file system (or vice versa), as the result of **ccat** or the use of **compact** or **uncompact** as a filter, optional ACL entries are lost.

WARNINGS

On short-filename systems, the last segment of the file name must contain 12 or fewer characters to allow space for the appended *.C*.

DEPENDENCIES**NFS**

Access control list entries of networked files are summarized (as returned in *st_mode* by *stat()*), but not copied to the new file (see *stat(2)*).

FILES

.C* compacted file created by **compact, removed by **uncompact**

SEE ALSO

compress(1), *pack(1)*, *acl(5)*, *aclv(5)*.

Gallager, Robert G., "Variations on a Theme of Huffman," *I.E.E.E. Transactions on Information Theory*, vol. IT-24, no. 6, November 1978, pp. 668 - 674.

AUTHOR

compact was developed by Colin L. Mc Master.

NAME

compress, uncompress, zcat, compressdir, uncompressdir - compress and expand data

SYNOPSIS

Compress Files

compress [-d] [-f] [-z] [-z] [-v] [-c] [-V] [-b *maxbits*] [*file* ...]

uncompress [-f] [-v] [-c] [-V] [*file* ...]

zcat [-V] [*file* ...]

Compress Entire Directory Subtrees

compressdir [*options*] [*directory* ...]

uncompressdir [*options*] [*directory* ...]

DESCRIPTION

The following commands compress and uncompress files and directory subtrees as indicated:

| | |
|----------------------|---|
| compress | Reduce the size of the named <i>files</i> using adaptive Lempel-Ziv coding. If reduction is possible, each <i>file</i> is replaced by a new file of the same name with the suffix .Z added to indicate that it is a compressed file. Original ownership, modes, access, and modification times are preserved. If no <i>file</i> is specified, or if - is specified, standard input is compressed to the standard output. |
| uncompress | Restore the compressed <i>files</i> to original form. Resulting files have the original filename, ownership, and permissions, and the .Z filename suffix is removed. If no <i>file</i> is specified, or if - is specified, standard input is uncompressed to the standard output. |
| zcat | Restore the compressed <i>files</i> to original form and send the result to standard output. If no <i>file</i> is specified, or if - is specified, standard input is uncompressed to the standard output. |
| compressdir | Front-end processor. Recursively descend each specified <i>directory</i> subtree and use compress to compress each file in <i>directory</i> . Existing files are replaced by a compressed file having the same name plus the suffix .Z , provided the resulting file is smaller than the original. If no directories are specified, compression is applied to all files starting with the current directory. <i>options</i> may include any valid compress command options (they are passed through to compress). To force compression of all files, even when the result is larger than the original file, use the -f option. |
| uncompressdir | Opposite of compressdir . Restore compressed files to their original form. <i>options</i> may include any valid uncompress command options (they are passed through to uncompress). |

The amount of compression obtained depends on the size of the input, the maximum number of bits (*maxbits*) per code, and the distribution of common substrings. Typically, text such as source code or English is reduced by 50-60 percent. Compression is generally much better than that achieved by Huffman coding (as used in **pack**), or adaptive Huffman coding (**compact**), and takes less time to compute.

Options

These commands recognize the following options in the combinations shown above in SYNOPSIS:

| | |
|-----------|---|
| -d | Decompress <i>file</i> . compress -d is equivalent to uncompress . |
| -f | Force compression of <i>file</i> . This is useful for compressing an entire directory, even if some of the files do not actually shrink. If -f is not given and compress is run in the foreground, the user is prompted as to whether an existing file should be overwritten. |
| -z | This is the same as the -f option except that it does not force compression when there is null compression. |
| -v | Print a message describing the percentage of reduction for each file compressed. |
| -c | Force compress and uncompress to write to the standard output; no files are changed. The nondestructive behavior of zcat is identical to that of |

- uncompress -c.**
- v** Print the current version and compile options onto the standard error.
- b maxbits** Specify the maximum number of bits the **compress** algorithm will use. The default is 16 and the range can be any integer between 9 and 16.

compress uses the modified Lempel-Ziv algorithm popularized in *A Technique for High Performance Data Compression*, Terry A. Welch, *IEEE Computer*, vol. 17, no. 6 (June 1984), pages 8-19. Common substrings in the file are first replaced by 9-bit codes 257 and up. When code 512 is reached, the algorithm switches to 10-bit codes and continues to use more bits until the limit specified by the **-b** flag is reached (default 16).

After the *maxbits* limit is attained, **compress** periodically checks the compression ratio. If it is increasing, **compress** continues to use the existing code dictionary. However, if the compression ratio is decreasing, **compress** discards the table of substrings and rebuilds it from scratch. This allows the algorithm to adapt to the next "block" of the file.

Note that the **-b** flag is omitted for **uncompress** since the *maxbits* parameter specified during compression is encoded within the output, along with a magic number to ensure that neither decompression of random data nor recompression of compressed data is attempted.

Access Control Lists

compress retains a file's access control list when compressing and expanding data.

EXTERNAL INFLUENCES

Environment Variables

LC_MESSAGES determines the language in which messages are displayed.

If **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**.

If any internationalization variable contains an invalid setting, **compress**, **uncompress**, and **zcat** behave as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single- and multi-byte character code sets are supported.

RETURN VALUE

These commands return the following values upon completion:

- 0 Completed successfully.
- 2 Last file is larger after (attempted) compression.
- 1 An error occurred.

DIAGNOSTICS

Usage: **compress [-f|-z] [-dvcV] [-b maxbits] [file ...]**
Invalid options were specified on the command line.

Missing maxbits
maxbits must follow **-b**.

file: not in compressed format
The file specified to **uncompress** has not been compressed.

file: compressed with xxbits, can only handle yybits
file was compressed by a program that could deal with a higher value of *maxbits* than the **compress** code on this machine. Recompress the file with a lower value of *maxbits*.

file: already has .Z suffix -- no change
The file is assumed to be already compressed. Rename the file and try again.

file: filename too long to tack on .Z
The output file name, which is the source file name with a **.Z** extension, is too long for the file system on which the source file resides. Make the source file name shorter and try again.

file already exists; do you wish to overwrite (y or n)?
Respond **y** if you want the output to replace the existing file; otherwise, respond **n**.

uncompress: corrupt input

A SIGSEGV violation was detected which usually means that the input file has been 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 (a directory for example), it is left unaltered.

-- has *xx*other links: unchanged

The input file has links which are not symbolic links and has been left unchanged. See *ln(1)* for more information.

-- has symbolic links: unchanged

The input file has symbolic links and has been left unchanged. See *ln(1)* for more information.

-- file unchanged

No savings is achieved by compression. The input remains unaltered.

EXAMPLES

Compress the file named **zenith** and print compression information to the terminal:

```
compress -v zenith
```

The terminal display shows either a line resembling

```
zenith: Compression: 23.55% -- replaced with zenith.Z
```

indicating that the compressed file is 23.55% smaller than the original, or a line resembling

```
zenith: Compression: -12.04% -- file unchanged
```

indicating that an additional 12.04% space must be used to compress the file.

Undo the compression by typing either of the following commands:

```
uncompress zenith.Z  
compress -d zenith.Z
```

This restores file **zenith.Z** to its original uncompressed form and name.

uncompress will perform on standard input if no files are specified. For example, to list a compressed tar file:

```
uncompress -c arch.tar.Z | tar -tvf -
```

WARNINGS

Although compressed files are compatible between machines with large memory, **-b12** should be used for file transfer to architectures with a small process data space (64K bytes or less).

NFS

Access control lists of networked files are summarized (as returned in **st_mode** by **stat()**), but not copied to the new file (see *stat(2)*).

AUTHOR

compress was developed by Joseph M. Orst, Kenneth E. Turkowski, Spencer W. Thomas, and James A. Woods.

FILES

***.Z** Compressed file created by **compress** and removed by **uncompress**.

SEE ALSO

compact(1), *pack(1)*, *acl(5)*.

STANDARDS CONFORMANCE

compress: XPG4

NAME

convert - convert an audio file

SYNOPSIS

```
/opt/audio/bin/convert [source_file] [target_file] [-sfmt format] [-dfmt format]
[-ddata data_type] [-srate rate] [-drate rate]
[-schannels number] [-dchannels number]
```

DESCRIPTION

This command converts audio files from one supported file format, data format, sampling rate, and number of channels to another. The unconverted file is retained as a source file.

-sfmt format -dfmt format

are the file formats for the source and destination files. Each *format* can be one of these:

| | |
|------------|---------------------------------------|
| au | Sun file format |
| snd | NeXT file format |
| wav | Microsoft RIFF Waveform file format |
| u | MuLaw format |
| a1 | ALaw |
| 116 | linear 16-bit format |
| 108 | offset (unsigned) linear 8-bit format |
| 18 | linear 8-bit format |

If you omit **-sfmt**, **convert** uses the header or filename extension in the source file. You can omit **-dfmt** if you supply a filename extension for the destination file.

-ddata data_type

is the data type for the destination files. *data_type* can be one of these:

| | |
|------------|-------------------------------------|
| u | MuLaw |
| a1 | ALaw |
| 116 | linear 16-bit |
| 108 | offset (unsigned) linear 8-bit data |
| 18 | linear 8-bit data |

If you omit **-ddata**, **convert** uses an appropriate data type, normally the data type of the source file.

-srate rate -drate rate

are the number of samples per second for the source and destination file. Typical sampling rates range from 8 to 11k (for voice quality) to 44,100 (for CD quality). You can use **k** to indicate thousands. For example, **8k** means 8,000 samples per second.

If you omit **-srate**, **convert** uses a rate defined by the source file header or its filename extension. For a raw file with no extension, 8,000 is used. By playing the file, you can determine if 8,000 samples is too fast or too slow.

If you omit **-drate**, **convert** uses a sampling rate appropriate for the destination file format; if possible, it matches the sampling rate of the source file.

-schannels number -dchannels number

are the number of channels in the source and destination files. Use **1** for mono; **2** for stereo. If **-schannels** is omitted, **convert** uses the information in the header; for raw data files, it uses mono.

If **-dchannels** is omitted, **convert** matches what was used for the source file (through the header or **-schannels** option); for raw data files, it uses mono.

EXAMPLES

Convert a raw data file to a headered file.

```
cd /opt/audio/bin
convert beep.116 beep.au
```

Convert a raw data file to a headered file when the source has no extension, was sampled at 11,025 per second, and has stereo data.

```
cd /opt/audio/bin
convert beep beep.au -sfmt 116 -srate 11025 -schannels 2
```

To save disk space, convert an audio file with CD quality sound to voice quality sound.

```
cd /opt/audio/bin
convert idea.au idea2.au -ddata u -drate 8k -dchannels 1
```

AUTHOR

convert was developed by HP.

Sun is a trademark of Sun Microsystems, Inc.

NeXT is a trademark of NeXT Computers, Inc.

Microsoft is a trademark of Microsoft Corporation.

SEE ALSO

audio(5), asecure(1M), aserver(1M), attributes(1), send_sound(1).

Using the Audio Developer's Kit

NAME

cp - copy files and directory subtrees

SYNOPSIS

cp [-f|-i|-p] [-e *extarg*] *file1 new_file*

cp [-f|-i|-p] [-e *extarg*] *file1 [file2 ...] dest_directory*

cp [-f|-i|-p] [-R|-r] [-e *extarg*] *directory1 [directory2 ...] dest_directory*

DESCRIPTION

cp copies:

- *file1* to new or existing *new_file*,
- *file1* to existing *dest_directory*,
- *file1*, *file2*, ... to existing *dest_directory*,
- directory subtree *directory1*, to new or existing *dest_directory*. or
- multiple directory subtrees *directory1*, *directory2*, ... to new or existing *dest_directory*.

cp fails if *file1* and *new_file* are the same (be cautious when using shell metacharacters). When destination is a directory, one or more files are copied into that directory. If two or more files are copied, the destination must be a directory. When copying a single file to a new file, if *new_file* exists, its contents are destroyed.

If the access permissions of the destination *dest_directory* or existing destination file *new_file* forbid writing, cp aborts and produces an error message "cannot create file".

To copy one or more directory subtrees to another directory, the -r option is required. The -r option is ignored if used when copying a file to another file or files to a directory.

If *new_file* is a link to an existing file with other links, cp overwrites the existing file and retains all links. If copying a file to an existing file, cp does not change existing file access permission bits, owner, or group.

When copying files to a directory or to a new file that does not already exist, cp creates a new file with the same file permission bits as *file1*, modified by the file creation mask of the user if the -p option was not specified, and then bitwise inclusively ORed with S_IRWXU. The owner and group of the new file or files are those of the user. The last modification time of *new_file* (and last access time, if *new_file* did not exist) and the last access time of the source *file1* are set to the time the copy was made.

Options

- i (interactive copy) Cause cp to write a prompt to standard error and wait for a response before copying a file that would overwrite an existing file. If the response from the standard input is affirmative, the file is copied if permissions allow the copy. If the -i (interactive) and -f (forced-copy) options are both specified, the -i option is ignored.
- f Force existing destination pathnames to be removed before copying, without prompting for confirmation. This option has the effect of destroying and replacing any existing file whose name and directory location conflicts with the name and location of the new file created by the copy operation.
- p (preserve permissions) Causes cp to preserve in the copy as many of the modification time, access time, file mode, user ID, and group ID as allowed by permissions.
- r (recursive subtree copy) Cause cp to copy the subtree rooted at each source directory to *dest_directory*. If *dest_directory* exists, it must be a directory, in which case cp creates a directory within *dest_directory* with the same name as *file1* and copies the subtree rooted at *file1* to *dest_directory/file1*. An error occurs if *dest_directory/file1* already exists. If *dest_directory* does not exist, cp creates it and copies the subtree rooted at *file1* to *dest_directory*. Note that cp -r cannot merge subtrees.

Usually normal files and directories are copied. Character special devices, block special devices, network special files, named pipes, symbolic links, and sockets are copied, if the user has access to the file; otherwise, a warning is printed stating that the file cannot be created, and the file is skipped.

dest_directory should not reside within *directory1*, nor should *directory1* have a cyclic directory structure, since in both cases cp attempts to copy an infinite amount of data.

- R** (recursive subtree copy) The **-R** option is identical to the **-r** option with the exception that directories copied by the **-R** option are created with read, write, and search permission for the owner. User and group permissions remain unchanged.

With the **-R** and **-r** options, in addition to regular files and directories, **cp** also copies FIFOs, character and block device files and symbolic links. Only superusers can copy device files. All other users get an error. Symbolic links are copied so the target points to the same location that the source did.

Warning: While copying a directory tree that has device special files, use the **-r** option; otherwise, an infinite amount of data is read from the device special file and is duplicated as a special file in the destination directory occupying large file system space.

-e *extarg*

Specifies the handling of any extent attributes of the file[s] to be copied. *extarg* takes one of the following values.

- warn** Issues a warning message if extent attributes cannot be copied, but copies the file anyway.
- ignore** Does not copy the extent attributes.
- force** Fails to copy the file if the extent attribute can not be copied.

Extent attributes can not be copied if the files are being copied to a file system which does not support extent attributes or if that file system has a different block size than the original. If **-e** is not specified, the default value for *extarg* is **warn**.

Access Control Lists (ACLs)

If *new_file* is a new file, or if a new file is created in *dest_directory*, it inherits the access control list of the original *file1*, *file2*, etc., altered to reflect any difference in ownership between the two files (see *acl*(5) and *aclv*(5)). In JFS file systems, new files created by **cp** do not inherit their parent directory's default ACL entries (if any), but instead retain the ACLs of the files being copied. When copying files from a JFS file system to an HFS file system or vice versa, optional ACL entries are lost.

EXTERNAL INFLUENCES

Environment Variables

LC_CTYPE determines the interpretation of text as single and/or multi-byte characters.

LANG and **LC_CTYPE** determine the local language equivalent of y (for yes/no queries).

LANG determines the language in which messages are displayed.

If **LC_CTYPE** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang*(5)) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **cp** behaves as if all internationalization variables are set to "C". See *environ*(5).

International Code Set Support

Single- and multi-byte character code sets are supported.

EXAMPLES

The following command moves the directory *sourcedir* and its contents to a new location (*targetdir*) in the file system. Since **cp** creates the new directory, the destination directory *targetdir* should not already exist.

```
cp -r sourcedir targetdir && rm -rf sourcedir
```

The **-r** option copies the subtree (files and subdirectories) in directory *sourcedir* to directory *targetdir*. The double ampersand (**&&**) causes a conditional action. If the operation on the left side of the **&&** is successful, the right side is executed (and removes the old directory). If the operation on the left of the **&&** is not successful, the old directory is not removed.

This example is equivalent to:

```
mv sourcedir targetdir
```

To copy all files and directory subtrees in the current directory to an existing *targetdir*, use:

```
cp -r * targetdir
```

To copy all files and directory subtrees in *sourcedir* to *targetdir*, use:

```
cp -r sourcedir/* targetdir
```

Note that directory pathnames can precede both *sourcedir* and *targetdir*.

To create a zero-length file, use any of the following:

```
cat /dev/null >file  
cp /dev/null file  
touch file
```

C**DEPENDENCIES****NFS**

Access control lists of networked files are summarized (as returned in `st_mode` by `stat()`), but not copied to the new file. When using `mv` or `ln` on such files, a `+` is not printed after the mode value when asking for permission to overwrite a file.

AUTHOR

`cp` was developed by AT&T, the University of California, Berkeley, and HP.

SEE ALSO

`cpio(1)`, `ln(1)`, `mv(1)`, `rm(1)`, `link(1M)`, `lstat(2)`, `readlink(2)`, `stat(2)`, `symlink(2)`, `symlink(4)`, `acl(5)`, `aclv(5)`.

STANDARDS CONFORMANCE

`cp`: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

cpio - copy file archives in and out; duplicate directory trees

SYNOPSIS

```
cpio -o [-e extarg] [achvxABC]
cpio -i|bcdfmrstuvxBPRSU6] [pattern...]
cpio -p [-e extarg] [adlmruvxU] directory
```

DESCRIPTION

The **cpio** command saves and restores archives of files on magnetic tape, other devices, or a regular file, and copies files from one directory to another while replicating the directory tree structure. When **cpio** completes processing the files, it reports the number of blocks written.

cpio -o (copy out, export) Read standard input to obtain a list of path names, and copy those files to standard output together with path name and status information. The output is padded to a 512-byte boundary.

cpio -i (copy in, import) Extract files from standard input, which is assumed to be the result of a previous **cpio -o**.

If *pattern...*, is specified, only the files with names that match a *pattern* according to the rules of Pattern Matching Notation (see *regex*(5)) are selected. A leading **!** on a pattern indicates that only those names that do *not* match the remainder of the pattern should be selected. Multiple patterns can be specified. The patterns are additive. If no *pattern* is specified, the default is ***** (select all files). See the **x** option, as well.

Extracted files are conditionally created and copied into the current directory tree, as determined by the options described below. The permissions of the files match the permissions of the original files when the archive was created by **cpio -o** unless the **U** option is used. File owner and group are that of the current user unless the user has appropriate privileges, in which case **cpio** retains the owner and group of the files of the previous **cpio -o**.

cpio -p (passthrough) Read standard input to obtain a list of path names of files which are then conditionally created and copied into the destination *directory* tree as determined by the options described below. *directory* must exist. Destination path names are interpreted relative to *directory*.

With the **-p** option, when handling a link, only the link is passed and no data blocks are actually read or written. This is especially noteworthy with **cpio -pl**, where it is very possible that all the files are created as links, such that no blocks are written and "0 blocks" is reported by **cpio**. (See below for a description of the **-l** option.)

Options

cpio recognizes the following options, which can be appended as appropriate to **-i**, **-o**, and **-p**. Whitespace and hyphens are not permitted between these options and **-i**, **-o**, or **-p**.

- a** Reset access times of input files after they are copied.
- b** Swap both bytes and half-words. Use only with **-i**. See the **P** option for details; see also the **s** and **S** options.
- c** Write or read header information in ASCII character form for portability.
- d** Create directories as needed.

-e *extarg*

Specifies the handling of any extent attributes of the file(s) to be archived or copied. *extarg* takes one of the following values.

- warn** Archive or copy the file and issue a warning message if extent attributes cannot be preserved.
- ignore** Do not issue a warning message even if extent attributes cannot be preserved.
- force** Any file(s) with extent attributes will not be archived and a warning message will be issued.

When using the **-o** option, extent attributes are not preserved in the archive. Furthermore, the **-p** option will not preserve extent attributes if the files are being copied to a file system that does not support extent attributes. If **-e** is not specified, the default value for *extarg* is

warn.

- f** Copy in all files except those selected by *pattern*....
- h** Follow symbolic links as though they were normal files or directories. Normally, **cpio** archives the link.
- l** Whenever possible, link files rather than copying them. This option does not destroy existing files. Use only with **-p**.
- m** Retain previous file modification time. This option does not affect directories that are being copied.
- r** Rename files interactively. If the user types a null line, the file is skipped.
- s** Swap all bytes of the file. Use only with **-i**. See the **P** option for details; see also the **s** and **S** options.
- t** Print only a table of contents of the input. No files are created, read, or copied.
- u** Copy unconditionally (normally, an older file does not replace a newer file with the same name).
- v** Print a list of file names as they are processed. When used with the **t** option, the table of contents has the format:

numeric-mode owner-name blocks date-time filename

where *numeric-mode* is the file privileges in numeric format, *owner-name* is the name of the file owner, *blocks* is the size of the file in 512-byte blocks, *date-time* is the date and time the file was last modified, and *filename* is the path name of the file as recorded in the archive.

- x** Save or restore device special files. Since **mknod()** is used to recreate these files on a restore, **-ix** and **-px** can be used only by users with appropriate privileges (see *mknod(2)*). This option is intended for intrasystem (backup) use only. Restoring device files from previous versions of the OS, or from different systems can be very dangerous. **cpio** may prevent the restoration of certain device files from the archive.
- A** Suppress warning messages regarding optional access control list entries. **cpio** does not back up optional access control list entries in a file's access control list (see *acl(5)*). Normally, a warning message is printed for each file that has optional access control list entries.
- B** Block input/output at 5120 bytes to the record (does not apply to **cpio -p**). This option is meaningful only with data directed to or from devices that support variable-length records such as magnetic tape.
- C** Have **cpio** checkpoint itself at the start of each volume. If **cpio** is writing to a streaming tape drive with immediate-report mode enabled and a write error occurs, it normally aborts and exits with return code 2. With this option specified, **cpio** instead automatically restarts itself from the checkpoint and rewrites the current volume. Alternatively, if **cpio** is not writing to such a device and a write error occurs, **cpio** normally continues with the next volume. With this option specified, however, the user can choose to either ignore the error or rewrite the current volume.
- P** Read a file written on a PDP-11 or VAX system (with byte-swapping) that did not use the **c** option. Use only with **-i**. Files copied in this mode are not changed. Non-ASCII files are likely to need further processing to be readable. This processing often requires knowledge of file contents, and thus cannot always be done by this program. The **b**, **s**, and **S** options can be used when swapping all the bytes on the tape (rather than just in the headers) is appropriate. In general, text is best processed with **P** and binary data with one of the other options.
(PDP-11 and VAX are registered trademarks of Digital Equipment Corporation.)
- R** Resynchronize automatically when **cpio** goes "out of phase", (see *DIAGNOSTICS*).
- S** Swap all half-words in the file. Use only with **-i**. See the **P** option for details; see also the **b** and **s** options.
- U** Use the process's file-mode creation mask (see *umask(2)*) to modify the mode of files created, in the same manner as *creat(2)*.
- 6** Process a UNIX Sixth-Edition-format file. Use only with **-i**.

Note that **cpio** archives created using a raw device file must be read using a raw device file.

When the end of the tape is reached, **cpio** prompts the user for a new special file and continues.

If you want to pass one or more metacharacters to **cpio** without the shell expanding them, be sure to precede each of them with a backslash (\).

Device files written with the **-ox** option (such as **/dev/tty03**) do not transport to other implementations of HP-UX.

EXTERNAL INFLUENCES

Environment Variables

LC_COLLATE determines the collating sequence used in evaluating pattern matching notation for file name generation.

LC_CTYPE determines the interpretation of text as single and/or multi-byte characters, and the characters matched by character class expressions in pattern matching notation.

LC_TIME determines the format and content of date and time strings output when listing the contents of an archive with the **v** option.

LANG determines the language in which messages are displayed.

If **LC_COLLATE**, **LC_CTYPE**, or **LC_TIME** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **cpio** behaves as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single- and multi-byte character code sets are supported.

RETURN VALUE

cpio returns the following exit codes:

- 0 Successful completion. Review standard error for files that could not be transferred.
- 1 Error during resynchronization. Some files may not have been recovered.
- 2 Out-of-phase error. A file header is corrupt or in the wrong format.

DIAGNOSTICS

Out of phase--get help

Perhaps the "c" option should[n't] be used

cpio -i could not read the header of an archived file. The header is corrupt or it was written in a different format. Without the **R** option, **cpio** returns an exit code of 2.

If no file name has been displayed yet, the problem may be the format. Try specifying a different header format option: **n** for standard format; **c** for ASCII; **b**, **s**, **P**, or **S**, for one of the byte-swapping formats; or **6** for UNIX Sixth Edition.

Otherwise, a header may be corrupt. Use the **R** option to have **cpio** attempt to resynchronize the file automatically. Resynchronizing means that **cpio** tries to find the next good header in the archive file and continues processing from there. If **cpio** tries to resynchronize from being out of phase, it returns an exit code of 1.

Other diagnostic messages are self-explanatory.

EXAMPLES

Copy the contents of a directory into a tape archive:

```
ls | cpio -o > /dev/rmt/c0t0d0BEST
```

Duplicate a directory hierarchy:

```
cd olddir
find . -depth -print | cpio -pd newdir
```

The trivial case

```
find . -depth -print | cpio -oB >/dev/rmt/c0t0d0BEST
```

can be handled more efficiently by:

```
find . -cpio /dev/rmt/c0t0d0BEST
```

WARNINGS

Because of industry standards and interoperability goals, **cpio** does not support the archival of files larger than 2GB or files that have user/group IDs greater than 60K. Files with user/group IDs greater than 60K are archived and restored under the user/group ID of the current process.

Do not redirect the output of **cpio** to a named **cpio** archive file residing in the same directory as the original files belonging to that **cpio** archive. This can cause loss of data.

cpio strips any leading ./ characters in the list of filenames piped to it.

Path names are restricted to **PATH_MAX** characters (see **<limits.h>** and *limits(5)*). If there are too many unique linked files, the program runs out of memory to keep track of them. Thereafter, linking information is lost. Only users with appropriate privileges can copy special files.

cpio tapes written on HP machines with the **-ox[c]** options can sometimes mislead (non-HP) versions of **cpio** that do not support the **x** option. If a non-HP (or non-AT&T) version of **cpio** happens to be modified so that the (HP) **cpio** recognizes it as a device special file, a spurious device file might be created.

If **/dev/tty** is not accessible, **cpio** issues a complaint and exits.

The **-pd** option does not create the directory typed on the command line.

The **-idr** option does not make empty directories.

The **-plu** option does not link files to existing files.

POSIX defines a file named **TRAILER!!!** as an end-of-archive marker. Consequently, if a file of that name is contained in a group of files being written by **cpio -o**, the file is interpreted as end-of-archive, and no remaining files are copied. The recommended practice is to avoid naming files anything that resembles an end-of-archive file name.

To create a POSIX-conforming **cpio** archive, the **c** option must be used. To read a POSIX-conforming **cpio** archive, the **c** option must be used and the **b**, **s**, **S**, and **6** options should not be used. If the user does not have appropriate privileges, the **U** option must also be used to get POSIX-conforming behavior when reading an archive. Users with appropriate privileges should not use this option to get POSIX-conforming behavior.

DEPENDENCIES

If the path given to **cpio** contains a symbolic link as the last element, this link is traversed and pathname resolution continues. **cpio** uses the symbolic link's target, rather than that of the link.

SEE ALSO

ar(1), find(1), tar(1), cpio(4), acl(5), environ(5), lang(5), regexp(5).

STANDARDS CONFORMANCE

cpio: SVID2, SVID3, XPG2, XPG3

NAME

cpp - the C language preprocessor

SYNOPSIS

/usr/ccs/bin/cpp [*option ...*] [*ifile* [*ofile*]]

DESCRIPTION

cpp is the C language preprocessor which is invoked as the first pass of any C compilation using the **cc** command (see **cc(1)**). Its purpose is to process **#include** and conditional compilation instructions and macros. Thus the output of **cpp** is designed to be in a form acceptable as input to the next pass of the C compiler. As the C language evolves, **cpp** and the rest of the C compilation package will be modified to follow these changes. Therefore, the use of **cpp** in other than this framework is not suggested. The preferred way to invoke **cpp** is through the **cc** command, since the functionality of **cpp** may someday be moved elsewhere. See **m4(1)** for a general macro processor.

cpp optionally accepts two file names as arguments. *ifile* and *ofile* are respectively the input and output for the preprocessor. They default to standard input and standard output if not specified.

Options

The following options are recognized by **cpp**:

- A** Remove all predefined symbols that begin with a letter and **_HPUX_SOURCE**. The user is expected to define **_POSIX_SOURCE** or **_XOPEN_SOURCE** when using this option.
- C** By default, **cpp** strips C-style comments. If the **-C** option is specified, all comments (except those found on **cpp** directive lines) are passed along.
- Dname** Define *name* as if by a **#define** directive. If no **=def** is given, *name* is defined as 1. The **-D** option has lower precedence than the **-U** option. Thus, if the same name is used in both a **-U** option and a **-D** option, the name is undefined regardless of the order of the options.
- Dname=def** Define *name* as if by a **#define** directive. If no **=def** is given, *name* is defined as 1. The **-D** option has lower precedence than the **-U** option. Thus, if the same name is used in both a **-U** option and a **-D** option, the name is undefined regardless of the order of the options.
- Hnnn** Change the internal macro definition table to be *nnn* bytes in size. The default buffer size is at least 8188 bytes. This option serves to eliminate "Macro param too large", "Macro invocation too large", "Macro param too large after substitution", "Quoted macro param too large", "Macro buffer too small", "Input line too long", and "Catenated input line too long" errors.
- h[*inclfile*]** Generates included files and sends the results to the file *inclfile*. If the argument *inclfile* is omitted, the result is sent to the standard error.
- I *dir*** Change the algorithm for searching for **#include** files whose names do not begin with **/** to look in *dir* before looking in the directories on the standard list. Thus, **#include** files whose names are enclosed in double quotes ("**"**") are searched for first in the directory of the file containing the **#include** line, then in directories named in **-I** options in left-to-right order, and last in directories on a standard list. For **#include** files whose names are enclosed in angle brackets (<**>#include** line is not searched. However, directory *dir* is still searched.
- M[*makefile*]** Generates makefile dependencies and sends the results to the file *makefile*. If the argument *makefile* is omitted, the result is sent to the standard error.
- P** Preprocess the input without producing the line-control information used by the next pass of the C compiler.
- T** HP-UX no longer restricts preprocessor symbols to eight characters. The **-T** option forces **cpp** to use only the first eight characters for distinguishing different preprocessor names. This behavior is the same as preprocessors on some other systems with respect to the length of names, and is included for backward compatibility.
- Uname** Remove any initial definition of *name*, where *name* is a reserved symbol that is predefined by the particular preprocessor. The current list of these symbols includes:

| | | | |
|-------------------|-------------------|-------------------|---------------------|
| Operating system: | unix | __unix | |
| Hardware: | hp9000s200 | hp9000s300 | __hp9000s300 |
| | hp9000s500 | hp9000s800 | __hp9000s800 |
| | hp9000ipc | hppa | __hppa |

| | | | | |
|-----------------------|--------------------------|--------------------------|---------------------------|--------------------|
| | <code>_PA_RISC1_0</code> | <code>_PA_RISC1_1</code> | <code>_SIO</code> | <code>_WSIO</code> |
| UNIX systems variant: | <code>hpux</code> | <code>_hpux</code> | <code>_HPUX_SOURCE</code> | |
| | <code>PWB</code> | <code>_PWB</code> | | |
| <code>lint(1):</code> | <code>lint</code> | <code>__lint</code> | | |

In addition, all symbols that begin with an underscore and either an upper-case letter or another underscore are reserved. Other symbols may be defined by the `CCOPTS` variable or other command-line options to the C compiler at compile time (see `cc(1)`). All HP-UX systems have the symbols `PWB`, `hpux`, `unix`, `_PWB`, `__hpux`, and `__unix` defined. Each system defines at least one hardware variant, as appropriate. The lint symbols are defined when `lint(1)` is running. See DEPENDENCIES.

Two special names are understood by `cpp`. `__LINE__` is defined as the current line number (as a decimal integer) as counted by `cpp`. `__FILE__` is defined as the current file name (as a C string) as known by `cpp`. They can be used anywhere (including in macros) just as any other defined names.

Directives

All `cpp` directives start with lines begun by `#`. Any number of blanks and tabs are allowed between the `#` and the directive. The directives are:

#define *name token-string*

Replace subsequent instances of *name* with *token-string*. *token-string* can be null.

#define *name(arg, ... , arg) token-string*

Replace subsequent instances of *name* followed by a (, a list of comma-separated set of arguments, and a) by *token-string*, where each occurrence of an *arg* in the *token-string* is replaced by the corresponding set of tokens 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` restarts its scan for names to expand at the beginning of the newly created *token-string*.

Notice that there can be no space between *name* and the (.

#endif [*text*]

Ends a section of lines begun by a test directive (`#if`, `#ifdef`, or `#ifndef`). Each test directive must have a matching `#endif`. Any *text* occurring on the same line as the `#endif` is ignored and thus may be used to mark matching `#if`–`#endif` pairs. This makes it easier, when reading the source, to match `#if`, `#ifdef`, and `#ifndef` directives with their associated `#endif` directive.

#elif *constant-expression*

Equivalent to:

```
#else
#if constant-expression
```

#else

Reverses the notion of the test directive that matches this directive. Thus, if lines previous to this directive are ignored, the following lines appear in the output, and vice versa.

#if *constant-expression*

The lines following appear in the output if and only if the *constant-expression* evaluates to nonzero. All binary nonassignment C operators, the `?:` operator, the unary `-`, `!`, and `~` operators are all legal in *constant-expression*. The precedence of the operators is the same as defined by the C language.

There is also a unary operator, `defined`, which can be used in *constant-expression* in these two forms: `defined(name)` or `defined name`. This allows the use of `#ifdef` and `#ifndef` in an `#if` directive.

Only these operators, integer constants, and names that are known by `cpp` should be used in *constant-expression*. In particular, the `sizeof` operator is not available.

#ifndef *name*

The lines following appear in the output if and only if *name* has been the subject of a previous **#define** without being the subject of an intervening **#undef**.

```
#ifndef name
```

The lines following do not appear in the output if and only if *name* has been the subject of a previous **#define** without being the subject of an intervening **#undef**.

```
#include "filename"
```

```
#include <filename>
```

Include at this point the contents of *filename* (which are then run through **cpp**). See the **-I** option above for more detail.

```
#line integer-constant "filename"
```

Causes **cpp** to generate line-control information for the next pass of the C compiler. *integer-constant* is the line number of the next line and *filename* is the file where it comes from. If *filename* and the quotation marks are omitted, the current file name is unchanged.

```
#undef name
```

Cause the definition of *name* (if any) to be forgotten from now on.

The test directives and the possible **#else** directives can be nested. **cpp** supports names up to 255 characters in length.

EXTERNAL INFLUENCES

Environment Variables

LC_CTYPE determines the interpretation of comments and string literals as single- or multibyte characters.

LANG determines the language in which messages are displayed.

If **LC_CTYPE** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, it defaults to "C" (see *lang(5)*). If any internationalization variable contains an invalid setting, **cpp** behaves as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single- and multibyte character code sets are supported.

DIAGNOSTICS

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.

WARNINGS

When newline characters were found in argument lists for macros to be expanded, previous versions of **cpp** put out the newlines as they were found and expanded. The current version of **cpp** replaces these newlines with blanks to alleviate problems that the previous versions had when this occurred.

DEPENDENCIES

Series 700

The symbols **hp9000s700** and **__hp9000s700** are not reserved symbols recognized by the **-U** option. They are supplied to **cpp** either automatically by the compiler, or by the use of a compiler option. For example, on a Series 700 system, the command:

```
cc -v t.c
```

produces:

```
/usr/ccs/libin/cpp      t.c      /var/tmp/ctmAAAa29220      -D__hp9000s700
-D__hp9000s800 ...
```

(Also see the **-D** option of the **cc** command.)

FILES

/usr/include Standard directory for **#include** files

SEE ALSO

m4(1).

NOTES

The macro substitution scheme has been changed. Previous versions of **cpp** saved macros in a macro definition table whose table size is 128 000 bytes by default. The current version of **cpp** replaces this macro definition table with several small buffers. The default size of the small buffers is 8 188 bytes.

STANDARDS CONFORMANCE

cpp: SVID2, SVID3, XPG2

C

NAME

crontab - user job file scheduler

SYNOPSIS

```
crontab [file]
crontab -e [username]
crontab -l [username]
crontab -r [username]
```

DESCRIPTION

The **crontab** command manages a crontab file for the user. You can use a crontab file to schedule jobs that are executed automatically by **cron** (see *cron(1M)*) on a regular basis. The command has four forms:

```
crontab [file]      Create or replace your crontab file by copying the specified file, or standard
                    input if file is omitted or - is specified as file, into the crontab directory,
                    /var/spool/cron/crontabs. The name of your crontab file in the
                    crontab directory is the same as your effective user name.

crontab -e [username] Edit a copy of the user's crontab file, or create an empty file to edit if the
                    crontab file does not exist. When editing is complete, the file will be copied
                    into the crontab directory as the user's crontab file.

crontab -l [username] Lists the user's crontab file.

crontab -r [username] Remove the user's crontab file from the crontab directory.
```

Only a privileged user can use username following the -e, -l, or -r options, to edit, list, or remove the crontab file of the specified user.

The entries in a crontab file are lines of six fields each. The fields are separated by spaces or tabs. The lines have the following format:

```
minute hour monthday month weekday command
```

The first five are integer patterns that specify when the sixth field, *command*, should be executed. They can have the following ranges of values:

```
minute      The minute of the hour, 0-59
hour        The hour of the day, 0-23
monthday    The day of the month, 1-31
month       The month of the year, 1-12
weekday     The day of the week, 0-6, 0=Sunday
```

Each pattern can be either an asterisk (*), meaning all legal values, or a list of elements separated by commas. An element is either a number in the ranges shown above, or two numbers in the range separated by a hyphen (meaning an inclusive range). Note that the specification of days can be made in two fields: *monthday* and *weekday*. If both are specified in an entry, they are cumulative. For example,

```
0 0 1,15 * 1 command
```

runs *command* at midnight on the first and fifteenth of each month, as well as every Monday. To specify days in only one field, set the other field to asterisk (*). For example,

```
0 0 * * 1 command
```

runs *command* only on Mondays.

The sixth field, *command* (the balance of a line including blanks 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 a backslash (\)) is translated to a newline character, dividing the field into "lines". Only the first "line" (up to a % or end-of-line) of the command field is executed by the shell. Any other "lines" are made available to the command as standard input.

Blank lines and those whose first non-blank character is # will be ignored.

cron invokes the command from the user's **HOME** directory with the POSIX shell, (**/usr/bin/sh**). It runs in the **c** queue (see *queuedefs(4)*).

cron supplies a default environment for every shell, defining:

```
HOME= user's-home-directory
LOGNAME= user's-login-id
PATH=/usr/bin:/usr/sbin:.
SHELL=/usr/bin/sh
```

Users who desire to have their **.profile** executed must explicitly do so in the crontab entry or in a script called by the entry.

You can execute **crontab** if your name appears in the file **/usr/lib/cron/cron.allow**. If that file does not exist, you can use **crontab** if your name does not appear in the file **/usr/lib/cron/cron.deny**. If only **cron.deny** exists and is empty, all users can use **crontab**. If neither file exists, only the **root** user can use **crontab**. The **allow/deny** files consist of one user name per line.

EXTERNAL INFLUENCES

Environment Variables

LC_CTYPE determines the interpretation of text within file as single- and/or multi-byte characters.

LC_MESSAGES determines the language in which messages are displayed.

If **LC_CTYPE** or **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**.

If any internationalization variable contains an invalid setting, **crontab** behaves as if all internationalization variables are set to "C". See *environ(5)*. **EDITOR** determines the editor to be invoked when **-e** option is specified. The default editor is **vi**.

International Code Set Support

Single-byte and multi-byte character code sets are supported.

WARNINGS

Be sure to redirect the standard output and standard error from commands. If this is not done, any generated standard output or standard error is mailed to the user.

FILES

| | |
|---------------------------------|--|
| /var/adm/cron | Main cron directory |
| /var/adm/cron/cron.allow | List of allowed users |
| /var/adm/cron/cron.deny | List of denied users |
| /var/adm/cron/log | Accounting information |
| /var/spool/cron/crontabs | Directory containing the crontab files |

SEE ALSO

sh(1), **cron(1M)**, **queuedefs(4)**.

STANDARDS CONFORMANCE

crontab: SVID2, SVID3, XPG2, XPG3, XPG4

NAME

crypt - encode/decode files

SYNOPSIS

crypt [*password*]

DESCRIPTION

crypt reads from the standard input and writes on the standard output. *password* is a key that selects a particular transformation. If no *password* 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:

```
crypt key <clear >cypher
crypt key <cypher|pr
```

The latter command decrypts the file and prints the clear version.

Files encrypted by **crypt** are compatible with those treated by the **ed** editor in encryption mode (see *ed(1)*).

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 clear text can become visible must be minimized.

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.

The transformation of a key into the internal settings of the machine is deliberately designed to be expensive; i.e., to take a substantial fraction of a second to compute. However, if keys are restricted to, for example, three lowercase letters, then encrypted files can be read by expending only a substantial fraction of five minutes of machine time.

Since the key is an argument to the **crypt** command, it is potentially visible to users executing the **ps** or a derivative (see *ps(1)*). The choice of keys and key security are the most vulnerable aspect of **crypt**.

EXAMPLES

The following example demonstrates the use of **crypt** to edit a file that the user wants to keep strictly confidential:

```
$ crypt <plans >plans.x
key: violet
$ rm plans
...
$ vi -x plans.x
key: violet
...
:wq
$
...
$ crypt <plans.x | pr
key: violet
```

Note that the **-x** option is the encryption mode of **vi**, and prompts the user for the same key with which the file was encrypted.

WARNINGS

If output is piped to **nroff** and the encryption key is *not* given on the command line, **crypt** can leave terminal modes in a strange state (see *nroff(1)* and *stty(1)*).

If two or more files encrypted with the same key are concatenated and an attempt is made to decrypt the result, only the the first of the original files is decrypted correctly.

FILES

/dev/tty for typed key

SEE ALSO

ed(1), makekey(1), stty(1).


C

NAME

cs - a shell (command interpreter) with C-like syntax

SYNOPSIS

cs [-**cefinstvXTVX**] [*command_file*] [*argument_list* ...]

DESCRIPTION

cs is a command language interpreter that incorporates a command history buffer, C-like syntax, and job control facilities.

Command Options

Command options are interpreted as follows:

- c Read commands from the (single) following argument which must be present. Any remaining arguments are placed in **argv**.
- e C shell exits if any invoked command terminates abnormally or yields a non-zero exit status.
- f Suppress execution of the **.cshrc** file in your home directory, thus speeding up shell start-up time.
- i Force **cs** to respond interactively when called from a device other than a computer terminal (such as another computer). **cs** normally responds non-interactively. If **cs** is called from a computer terminal, it always responds interactively, regardless of which options are selected.
- n Parse but do *not* execute commands. This is useful for checking syntax in shell scripts. All substitutions are performed (history, command, alias, etc.).
- s Take command input from the standard input.
- t Read and execute a single line of input.
- v Set the **verbose** shell variable, causing command input to be echoed to the standard output device after history substitutions are made.
- x Set the **echo** shell variable, causing all commands to be echoed to the standard error immediately before execution.
- T Disable the tenex features which use the ESC key for command/file name completion and CTRL-D for listing available files (see the *CSH UTILITIES* section below)
- V Set the **verbose** variable before **.cshrc** is executed so that all **.cshrc** commands are also echoed to the standard output.
- X Set the **echo** variable before **.cshrc** is executed so that all **.cshrc** commands are also echoed to the standard output.

After processing the command options, if arguments remain in the argument list, and the **-c**, **-i**, **-s**, or **-t** options were not specified, the first remaining argument is taken as the name of a file of commands to be executed.

COMMANDS

A simple command is a sequence of words, the first of which specifies the command to be executed. A sequence of simple commands separated by vertical bar (|) characters forms a pipeline. The output of each command in a pipeline becomes the input for the next command in the pipeline. Sequences of pipelines can be separated by semicolons (;) which causes them to be executed sequentially. A sequence of pipelines can be executed in background mode by adding an ampersand character (&) after the last entry.

Any pipeline can be placed in parentheses to form a simple command which, in turn, can be a component of another pipeline. Pipelines can also be separated by || or && indicating, as in the C language, that the second pipeline is to be executed only if the first fails or succeeds, respectively.

Jobs

cs associates a **job** with each pipeline and keeps a table of current jobs (printed by the **jobs** command) and assigns them small integer numbers. When a job is started asynchronously using &, the shell prints a line resembling:

[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 want to do something else, you can type the currently defined *suspend* character (see *termio(7)*) which sends a stop signal to the current job. **cs**h then normally indicates that the job has been 'Stopped', and prints another prompt. You can then manipulate the state of this job, putting it in the background with the **bg** command, run some other commands, and then eventually bring the job back into the foreground with the foreground command **fg**. A *suspend* takes effect immediately and is like an interrupt in that pending output and unread input are discarded when it is typed. There is a delayed *suspend* character which does not generate a stop signal until a program attempts to *read(2)* it. This can usefully be typed ahead when you have prepared some commands for a job which you want to stop after it has read them.

A job being run in the background stops 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** (see *stty(1)*). If you set this tty option, background jobs stop when they try to produce output, just as they do when they try to read input. Keyboard signals and line-hangup signals from the terminal interface are not sent to background jobs on such systems. This means that background jobs are immune to the effects of logging out or typing the interrupt, quit, suspend, and delayed suspend characters (see *termio(7)*).

There are several ways to refer to jobs in the shell. The character % introduces a job name. If you wish to refer to job number 1, you can name it as %1. Just naming a job brings it to the foreground; thus %1 is a synonym for **fg %1**, bringing job 1 back into the foreground. Similarly, typing %1 & resumes job 1 in the background. Jobs can also be named by prefixes of the string typed in to start them if these prefixes are unambiguous; thus %ex normally restarts a suspended *ex(1)* job, if there is only one suspended job whose name begins with the string *ex*. It is also possible to say %?string which specifies a job whose text contains *string*, if there is only one such job.

csh maintains a notion of the current and previous jobs. In output pertaining to jobs, the current job is marked with a + and the previous job with a -. The abbreviation %+ refers to the current job and %- refers to the previous job. For close analogy with the syntax of the **history** mechanism (described below), %% is also a synonym for the current job.

csh 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 printing a prompt. This is done so that it does not otherwise disturb your work. If, however, you set the shell variable **notify**, **cs**h notifies you immediately of changes in status of background jobs. There is also a **cs**h built-in command called **notify** which marks a single process so that any status change is immediately reported. By default, **notify** marks the current process. Simply type **notify** after starting a background job to mark it.

If you try to leave the shell while jobs are stopped, **cs**h sends the warning message: **You have stopped jobs**. Use the **jobs** command to see what they are. If you do this or immediately try to exit again, **cs**h does not warn you a second time, and the suspended jobs are terminated (see *exit(2)*).

Built-In Commands

Built-in commands are executed within the shell without spawning a new process. If a built-in command occurs as any component of a pipeline except the last, it is executed in a subshell. The built-in commands are:

```
alias
alias name
alias name wordlist
```

The first form prints all aliases. The second form prints the alias for *name*. The third form assigns the specified *wordlist* as the alias of *name*. Command and file name substitution are performed on *wordlist*. *name* cannot be **alias** or **unalias**.

```
bg [%job ...]
```

Put the current (*job* not specified) or specified jobs into the background, continuing them if they were stopped.

```
break
```

Causes execution to resume after the **end** of the nearest enclosing **foreach** or **while**. The remaining commands on the current line are executed. Multi-level breaks are thus possible by writing them all on one line.

breaksw

Causes a break from a **switch**, resuming after the **endsw**.

case *label*:

A label in a **switch** statement as discussed below.

cd

cd *directory_name*

chdir

chdir *directory_name*

Change the shell's current working directory to *directory_name*. If not specified, *directory_name* defaults to your home directory.

If *directory_name* is not found as a subdirectory of the current working directory (and does not begin with */*, *./*, or *../*), each component of the variable *cdpath* is checked to see if it has a subdirectory *directory_name*. Finally, if all else fails, **cs** treats *directory_name* as a shell variable. If its value begins with */*, this is tried to see if it is a directory. See also *cd*(1).

continue

Continue execution of the nearest enclosing **while** or **foreach**. The rest of the commands on the current line are executed.

default:

Labels the default case in a **switch** statement. The default should come after all other **case** labels.

dirs

Prints the directory stack; the top of the stack is at the left; the first directory in the stack is the current directory.

echo *wordlist*

echo **-n** *wordlist*

The specified words are written to the shell's standard output, separated by spaces, and terminated with a new-line unless the **-n** option is specified. See also *echo*(1).

else

end

endif

endsw See the descriptions of the **foreach**, **if**, **switch**, and **while** statements below.

eval *arguments ...*

(Same behavior as *sh*(1).) *arguments* are read as input to the shell and the resulting command(s) executed. This is usually used to execute commands generated as the result of command or variable substitution, since parsing occurs before these substitutions.

exec *command*

The specified command is executed in place of the current shell.

exit

exit (*expression*)

cs exits either with the value of the **status** variable (first form) or with the value of the specified *expression* (second form).

fg [%*job* ...]

Brings the current (*job* not specified) or specified jobs into the foreground, continuing them if they were stopped.

foreach *name* (*wordlist*)

end . . .

The variable *name* is successively set to each member of *wordlist* and the sequence of commands between this command and the matching **end** are executed. (Both **foreach** and **end** must appear alone on separate lines.)

The built-in command **continue** can be used to continue the loop prematurely; the built-in command **break** to terminate it prematurely. When this command is read from the terminal, the loop is read once, prompting with *?* before any statements in the loop are executed. If you make a mistake while typing in a loop at the terminal, use the erase or line-kill character as appropriate to recover.

glob *wordlist*

Like **echo** but no ** escapes are recognized and words are delimited by null characters in

the output. Useful in programs that use the shell to perform file name expansion on a list of words.

goto *word*

The specified *word* is file name and command expanded to yield a string of the form **label**. The shell rewinds its input as much as possible and searches for a line of the form **label:** possibly preceded by blanks or tabs. Execution continues after the specified line.

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 [-h] [-r] [*n*]

Displays the history event list. If *n* is given, only the *n* most recent events are printed. The **-r** option reverses the order of printout to be most recent first rather than oldest first. The **-h** option prints the history list without leading numbers for producing files suitable for the **source** command.

if (*expression*) *command*

If *expression* evaluates 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, a parenthesized command list, or an aliased command. Input/output redirection occurs even if *expression* is false, meaning that *command* is *not* executed (this is a bug).

if (*expression1*) **then**

...

else if (*expression2*) **then**

...

else

...

endif If *expression1* is true, all commands down to the first **else** are executed; otherwise if *expression2* is true, all commands from the first **else** down to the second **else** are executed, etc. Any number of **else-if** pairs are possible, but only one **endif** is needed. The **else** part is likewise optional. (The words **else** and **endif** must appear at the beginning of input lines. The **if** must appear alone on its input line or after an **else**.)

jobs [-l]

Lists active jobs. The **-l** option lists process IDs in addition to the usual information.

kill % *job*

kill - *sig* % *job* ...

kill *pid*

kill - *sig* *pid*...

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 **/usr/include/signal.h**, stripped of the SIG prefix (see *signal(2)*). The signal names are listed by **kill -l**. There is no default, so **kill** used alone does not send a signal to the current job. If the signal being sent is TERM (terminate) or HUP (hangup), the job or process is sent a CONT (continue) signal as well. See also *kill(1)*.

limit [-h] [*resource*] [*maximum_use*]

Limits the usage by the current process and each process it creates not to (individually) exceed *maximum_use* on the specified *resource*. If *maximum_use* is not specified, then the current limit is displayed; if *resource* is not specified, then all limitations are given.

If the **-h** flag is specified, the hard limits are used instead of the current limits. The hard limits impose a ceiling on the values of the current limits. Only the superuser can raise the hard limits, but a user can lower or raise the current limits within the legal range.

Controllable resources currently include:

addressspace Maximum address space in bytes for a process

| | |
|---------------------|--|
| coredumpsize | Size of the largest core dump that is created |
| cputime | Maximum number of CPU seconds to be used by each process |
| datasize | Maximum growth of the data region allowed beyond the end of the program text |
| descriptors | Maximum number of open files for each process |
| filesize | Largest single file that can be created |
| memoryuse | Maximum size to which a process's resident set size can grow |
| stacksize | Maximum size of the automatically extended stack region |

The *maximum_use* argument can be specified as a floating-point or integer number followed by a scale factor: *k* or *kilobytes* (1024 bytes), *m* or *megabytes*, or *b* or *blocks* (the units used by the *ulimit* system call). For both *resource* names and scale factors, unambiguous prefixes of the names can be used. *filesize* can be lowered by an instance of *cs*, but can only be raised by an instance whose effective user ID is *root*. For more information, refer to the documentation for the *ulimit* system call.

login Terminates a login shell, replacing it with an instance of */usr/bin/login*. This is one way to log off, included for compatibility with *sh*(1).

logout Terminates a login shell. Especially useful if *ignoreeof* is set. A similar function, **bye**, which works for sessions that are not login shells, is provided for historical reasons. Its use is not recommended because it is not part of the standard BSD *cs* and may not be supported in future releases.

newgrp Changes the group identification of the caller; for details see *newgrp*(1). A new shell is executed by **newgrp** so that the current shell environment is lost.

nice

nice *+number*

nice *command*

nice *+number command*

The first form sets the *nice* (run command priority) for this shell to 4 (the default). The second form sets the priority to the given *number*. The final two forms run *command* at priority 4 and *number* respectively. The user with appropriate privileges can raise the priority by specifying negative niceness using **nice** *-number ... command* is always executed in a sub-shell, and restrictions placed on commands in simple *if* statements apply. See also *nice*(1).

nohup [*command*]

Without an argument, **nohup** can be used in shell scripts to cause hangups to be ignored for the remainder of the script. With an argument, causes the specified *command* to be run with hangups ignored. All processes executed in the background with **&** are effectively **nohup**ed as described under Jobs in the *COMMANDS* section.

notify [*job ...*]

Causes the shell to notify the user asynchronously when the status of the current (*job* not specified) or specified jobs changes; normally notification is presented before a prompt. This is automatic if the shell variable *notify* is set.

onintr [-] [*label*]

Controls the action of the shell on interrupts. With no arguments, *onintr* restores the default action of the shell on interrupts, which action is to terminate shell scripts or return to the terminal command input level. If **-** is specified, all interrupts are ignored. If a *label* is given, the shell executes a **goto** *label* when an interrupt is received or a child process terminates because it was interrupted.

If the shell is running in the background and interrupts are being ignored, *onintr* has no effect; interrupts continue to be ignored by the shell and all invoked commands.

popd [*+n*]

Pops the directory stack, returning to the new top directory. With an argument, discards the *n*th entry in the stack. The elements of the directory stack are numbered from 0 starting at

the top. A synonym for **popd**, called **rd**, is provided for historical reasons. Its use is not recommended because it is not part of the standard BSD **cs** and may not be supported in future releases.

pushd [*name*] [+*n*]

With no arguments, *pushd* exchanges the top two elements of the directory stack. Given a *name* argument, *pushd* changes to the new directory (using *cd*) and pushes the old current working directory (as in *csw*) onto the directory stack. With a numeric argument, *pushd* rotates the *n*th argument of the directory stack around to be the top element and changes to that directory. The members of the directory stack are numbered from the top starting at 0. A synonym for *pushd*, called *gd*, is provided for historical reasons. Its use is not recommended since it is not part of the standard BSD **cs** and may not be supported in future releases.

rehash

Causes the internal hash table of the contents of the directories in the *path* variable to be recomputed. This is needed if new commands are added to directories in the *path* while you are logged in. This should only be necessary if you add commands to one of your own directories or if a systems programmer changes the contents of one of the system directories.

repeat *count command*

The specified *command* (which is subject to the same restrictions as the *command* in the one-line **if** statement above) is executed *count* times. I/O redirections occur exactly once, even if *count* is 0.

set

set *name*

set *name=word*

set *name[index]=word*

set *name=(wordlist)*

The first form of **set** shows the value of all shell variables. Variables whose value is other than a single word print as a parenthesized word list. The second form sets *name* to the null string. The third form sets *name* to the single *word*. The fourth form sets the *index*th component of *name* to *word*; this component must already exist. The final form sets *name* to the list of words in *wordlist*. In all cases the value is command and file-name expanded.

These arguments can be repeated to set multiple values in a single **set** command. Note, however, that variable expansion happens for all arguments before any setting occurs.

setenv *name value*

Sets the value of environment variable *name* to be *value*, a single string. The most commonly used environment variables, **USER**, **TERM**, and **PATH**, are automatically imported to and exported from the **cs** variables *user*, *term*, and *path*; there is no need to use **setenv** for these.

shift [*variable*]

If no argument is given, the members of **argv** are shifted to the left, discarding **argv[1]**. An error occurs if **argv** is not set or has less than two strings assigned to it. When *variable* is specified, *shift* performs the same function on the specified *variable*.

source [-h] *name*

cs reads commands from *name*. **source** commands can be nested, but if nested too deeply the shell may run out of file descriptors. An error in a **source** at any level terminates all nested **source** commands. Normally, input during **source** commands is not placed on the history list. The **-h** option can be used to place commands in the history list without being executing them.

stop [%*job* ...]

Stops the current (no argument) or specified jobs executing in the background.

suspend

Causes **cs** to stop as if it had been sent a *suspend* signal. Since **cs** normally ignores *suspend* signals, this is the only way to suspend the shell. This command gives an error message if attempted from a login shell.

switch (*string*)

case *str1*:

```

    . . .
breaksw
    . . .
default:
    . . .
breaksw
endsw

```

Each **case** label (*str1*) is successively matched against the specified *string* which is first command and file name expanded. The form of the **case** labels is the Pattern Matching Notation with the exception that non-matching lists in bracket expressions are not supported (see *regexp(5)*). If none of the labels match before a **default** label is found, the execution begins after the **default** label. Each **case** label and the **default** label must appear at the beginning of a line. The **breaksw** command causes execution to continue after the *endsw*. Otherwise, control may fall through **case** labels and **default** labels as in C. If no label matches and there is no default, execution continues after the **endsw**.

time [*command*]

When *command* is not specified, a summary of time used by this shell and its children is printed. If specified, the simple *command* is timed and a time summary as described under the **time** variable is printed. If necessary, an extra shell is created to print the time statistic when the command completes.

umask [*value*]

The current file creation mask is displayed (*value* not specified) or set to the specified *value*. The mask is given in octal. Common values for the mask are 002, which gives all permissions to the owner and group and read and execute permissions to all others, or 022, which gives all permissions to the owner, and only read and execute permission to the group and all others. See also *umask(1)*.

unalias *pattern*

All aliases whose names match the specified *pattern* are discarded. Thus, all aliases are removed by **unalias ***. No error occurs if *pattern* does not match an existing alias.

unhash

Use of the internal hash table to speed location of executed programs is disabled.

unset *pattern*

All variables whose names match the specified *pattern* are removed. Thus, all variables are removed by **unset ***; this has noticeably undesirable side-effects. No error occurs if *pattern* matches nothing.

unsetenv *pattern*

Removes all variables whose names match the specified *pattern* from the environment. See also the **setenv** command above and *printenv(1)*.

wait

Waits for all background jobs to terminate. If the shell is interactive, an interrupt can disrupt the wait, at which time the shell prints names and job numbers of all jobs known to be outstanding.

while (*expression*)

```

    . . .
end

```

While the specified *expression* evaluates non-zero, the commands between the **while** and the matching **end** are evaluated. **break** and **continue** can be used to terminate or continue the loop prematurely. (The **while** and **end** must appear alone on their input lines.) If the input is a terminal (i.e., not a script), prompting occurs the first time through the loop as for the **foreach** statement.

%job Brings the specified job into the foreground.

%job & Continues the specified job in the background.

@

@ name=expression

@ name[index]=expression

The first form prints the values of all the shell variables. The second form sets the specified *name* to the value of *expression*. If the expression contains **<**, **>**, **&**, or **|**, at least this part of the expression must be placed within parentheses. The third form assigns the value of *expression* to the *indexth* argument of *name*. Both *name* and its *indexth* component must already exist.

The operators `*`, `+`, etc., are available as in C. White space can optionally separate the *name* from the assignment operator. However, spaces are mandatory in separating components of *expression* which would otherwise be single words.

Special postfix `++` and `--` operators increment and decrement *name*, respectively (e.g., `@ i++`).

Non-Built-In Command Execution

When a command to be executed is not a built-in command, **cs**h attempts to execute the command via *exec*(2). Each word in the variable *path* names a directory in which the shell attempts to find the command (if the command does not begin with `/`). If neither `-c` nor `-t` is given, the shell hashes the names in these directories into an internal table so that an *exec* is attempted only in those directories where the command might possibly reside. This greatly speeds command location when a large number of directories are present in the search path. If this mechanism has been turned off (via *unhash*), or if `-c` or `-t` was given, or if any directory component of *path* does not begin with a `/`, the shell concatenates the directory name and the given command name to form a path name of a file which it then attempts to execute.

Commands placed inside parentheses are always executed in a subshell. Thus

```
(cd ; pwd)
```

prints the *home* directory then returns to the current directory upon completion, whereas:

```
cd ; pwd
```

remains in the *home* directory upon completion.

When commands are placed inside parentheses, it is usually to prevent *chdir* from affecting the current shell.

If the file has execute permissions but is not an executable binary file, it is assumed to be a script file, which is a file of data for an interpreter that is executed as a separate process.

csh first attempts to load and execute the script file (see *exec*(2)). If the first two characters of the script file are `#!`, *exec*(2) expects an interpreter path name to follow and attempts to execute the specified interpreter as a separate process to read the entire script file.

If no `#!` **interpreter** is named, and there is an *alias* for the shell, the words of the *alias* are inserted at the beginning of the argument list to form the shell command. The first word of the *alias* should be the full path name of the command to be used. Note that this is a special, late-occurring case of *alias* substitution, which inserts words into the argument list without modification.

If no `#!` **interpreter** is named and there is no shell *alias*, but the first character of the file is `#`, the interpreter named by the `$shell` variable is executed (note that this normally would be `/usr/bin/csh`, unless the user has reset `$shell`). If `$shell` is not set, `/usr/bin/csh` is executed.

If no `!#` **interpreter** is named, and there is no shell *alias*, and the first character of the file is not `#`, `/usr/bin/sh` is executed to interpret the script file.

History Substitutions

History substitutions enable you to repeat commands, use words from previous commands as portions of new commands, repeat arguments of a previous command in the current command, and fix spelling or typing mistakes in an earlier command.

History substitutions begin with an exclamation point (`!`). Substitutions can begin anywhere in the input stream, but *cannot* be nested. The exclamation point can be preceded by a backslash to cancel its special meaning. For convenience, an exclamation point is passed to the parser unchanged when it is followed by a blank, tab, newline, equal sign, or left parenthesis. Any input line that contains history substitution is echoed on the terminal before it is executed for verification.

Commands input from the terminal that consist of one or more words are saved on the history list. The history substitutions reintroduce sequences of words from these saved commands into the input stream. The number of previous commands saved is controlled by the **history** variable. The previous command is always saved, regardless of its value. Commands are numbered sequentially from 1.

You can refer to previous events by event number (such as `!10` for event 10), relative event location (such as `!-2` for the second previous event), full or partial command name (such as `!d` for the last event using a command with initial character `d`), and string expression (such as `!?mic?` referring to an event containing the characters `mic`).

These forms, without further modification, simply reintroduce the words of the specified events, each separated by a single blank. As a special case, `!!` is a re-do; it refers to the previous command.

To select words from a command, use a colon (`:`) and a designator for the desired words after the event specification. The words of an input line are numbered from zero. The basic word designators are:

- `0` First word (i.e., the command name itself).
- `n` *n*th word.
- `^` First argument. (This is equivalent to `1`.)
- `$` Last word.
- `a-b` Range of words from *a* through *b*. Special cases are `-y`, an abbreviation for “word 0 through word *y*”; and `x-`, which means “word *x* up to, but not including, word `$`”.
- `*` Range from the second word through the last word.
- `%` Used with a search sequence to substitute the immediately preceding matching word.

The colon separating the command specification from the word designator can be omitted if the argument selector begins with a `^`, `$`, `*`, `-`, or `%`.

After word designator can be followed by a sequence of modifiers, each preceded by a colon. The following modifiers are defined:

- `h` Use only the first component of a path name by removing all following components.
- `r` Use the root file name by removing any trailing suffix (`.xxx`).
- `e` Use the file name's trailing suffix (`.xxx`) by removing the root name.
- `s /l/r`
substitute the value of *r* for the value *l* in the indicated command.
- `t` Use only the final file name of a path name by removing all leading path name components.
- `&` Repeat the previous substitution.
- `p` Print the new command but do not execute it.
- `q` Quote the substituted words, preventing further substitutions.
- `x` Like `q`, but break into words at blanks, tabs and newlines.
- `g` Use a `g`lobal command as a prefix to another modifier to cause the specified change to be made globally. All words in the command are changed, one change per word, and each string enclosed in single quotes (`'`) or double quotes (`"`) is treated as a single word.

Unless preceded by a `g`, the modification is applied only to the first modifiable word. An error results if a substitution is attempted and cannot be completed (i.e., if you ask for a substitution of `!11` on a history buffer containing only 10 commands).

The left hand side of substitutions are strings; not regular expressions in the sense of HP-UX editors. Any character can be used as the delimiter in place of a slash (`/`). Use a backslash to quote a delimiter character if it is used in the *l* or *r* string. The character `&` in the right-hand side is replaced by the text from the left. A `\` also quotes `&`. A null *l* string uses the previous string either from an *l* or from a contextual scan string *s* in `! ? s ?`. The trailing delimiter in the substitution can be omitted if a new-line character follows immediately, as may the trailing `?` in a contextual scan.

A history reference can be given without an event specification (as in `!$`). In this case, the reference is to the previous command unless a previous history reference occurred on the same line, in which case this form repeats the previous reference. Thus

```
! ?foo? ^ ! $
```

gives the first and last arguments from the command matching `?foo?`.

A special abbreviation of a history reference occurs when the first non-blank character of an input line is a circumflex (`^`). This is equivalent to `! : s ^`, providing a convenient shorthand for substitutions on the text of the previous line. Thus `^lb^lib` fixes the spelling of `lib` in the previous command.

Finally, a history substitution can be enclosed within curly braces `{ }` if necessary to insulate it from the characters which follow. Thus, after

```
ls -ld ~paul
one could execute !{1}a to do
ls -ld ~paula
while !1a would look for a command starting with 1a.
```

Quoting with Single and Double Quotes

The quotation of strings by single quotes (') and double quotes (") can be used to prevent all or some of the remaining substitutions. Strings enclosed in single quotes are protected from any further interpretation. Strings enclosed in double quotes are still variable- and command-expanded as described below.

In both cases the resulting text becomes (all or part of) a single word. Only in one special case (see *Command Substitution* below) does a double-quoted string yield parts of more than one word; single-quoted strings never do.

Alias Substitution

cs maintains a list of aliases that can be established, displayed, and modified by the **alias** and **unalias** commands. After a command line is scanned, it is parsed into distinct commands and the first word of each command, left-to-right, is checked to see if it has an alias. If it does, the text which is the alias for that command is reread with the history mechanism available as if that command was the previous input line. The resulting words replace the command and argument list. If no reference is made to the history list, the argument list is left unchanged.

Thus, if the alias for **ls** is **ls -l**, the command **ls /usr** maps to **ls -l /usr**, leaving the argument list undisturbed. Similarly, if the alias for **lookup** was **grep !^ /etc/passwd**, **lookup bill** maps to **grep bill /etc/passwd**.

If an alias is found, the word transformation of the input text is performed and the aliasing process begins again on the re-formed input line. Looping is prevented if the first word of the new text is the same as the old by flagging it to prevent further aliasing. Other loops are detected and cause an error.

Note that the mechanism allows aliases to introduce parser metasyntax. Thus:

```
alias print 'pr \!* | lp'
```

makes a command that uses *pr*(1) to print its arguments on the line printer.

Expressions

Some of the built-in commands take expressions in which the operators are similar to those of C, with the same precedence. These expressions appear in the **@**, **exit**, **if**, and **while** commands. The following operators are available (shown in order of increasing precedence):

```
|| && | ^ & == != =~ !~ <= >= < > << >> + - * / % ! ~ ( )
```

The following list shows the grouping of these operators. The precedence decreases from top to bottom in the list:

```
* / %
+ -
<< >>
<= >= < >
== != =~ !~
```

The operators **==**, **!=**, **=~**, and **!~** compare their arguments as strings; all others operate on numbers. The operators **=~** and **!~** are similar to **!=** and **==**, except that the right-hand side is a *pattern* (containing *****, **?**, **s**, and instances of **[...]**) against which the left hand operand is matched. This reduces the need for use of the **switch** statement in shell scripts when all that is really needed is pattern matching.

Strings beginning with **0** are considered octal numbers. Null or missing arguments are considered **0**. The result of all expressions are strings that represent decimal numbers. It is important to note that no two components of an expression can appear in the same word. These components should be surrounded by spaces except when adjacent to components of expressions that are syntactically significant to the parser: **-**, **&**, **|**, **<**, **>**, **(**, and **)**.

Also available in expressions as primitive operands are command executions enclosed in curly braces (**{ }**) and file enquiries of the form **-l filename**, where *l* is one of:

| | |
|---|----------------|
| r | read access |
| w | write access |
| x | execute access |
| e | existence |
| o | ownership |
| z | zero size |
| f | plain file |
| d | directory |

The specified *filename* is command- and file-name expanded then tested to see if it has the specified relationship to the real user. If the file does not exist or is inaccessible, all inquiries return false (0). Command executions succeed, returning true, if the command exits with status 0; otherwise they fail, returning false. If more detailed status information is required, the command should be executed outside of an expression and the **status** variable examined.

Control of the Flow

csh contains a number of commands that can be used to regulate the flow of control in command files (shell scripts) and (in limited but useful ways) from terminal input. These commands all operate by forcing the shell to reread or skip parts of its input and, due to the implementation, restrict the placement of some of the commands.

The **foreach**, **switch**, and **while** statements, as well as the **if-then-else** form of the **if** statement require that the major keywords appear in a single simple command on an input line as shown below.

If the shell's input is not seekable, the shell buffers input whenever a loop is being read and performs seeks in this internal buffer to accomplish the rereading implied by the loop. (To the extent that this allows, backward **gotos** succeed on non-seekable inputs.)

Signal Handling

csh normally ignores *quit* signals. Jobs running in background mode are immune to signals generated from the keyboard, including hangups. Other signals have the values which the shell inherited from its parent. **cs**h's handling of interrupts and terminate signals in shell scripts can be controlled by *onintr*. Login shells catch the *terminate* signal; otherwise this signal is passed on to children from the state in the shell's parent. In no case are interrupts allowed when a login shell is reading the file *.logout*.

Command Line Parsing

csh splits input lines into words at blanks and tabs. The following exceptions (parser metacharacters) are considered separate words:

| | |
|----|---------------------------|
| & | ampersand; |
| | vertical bar; |
| ; | semicolon; |
| < | less-than sign; |
| > | greater-than sign; |
| (| left parenthesis; |
|) | right parenthesis; |
| && | double ampersand; |
| | double vertical bar; |
| << | double less-than sign; |
| >> | double greater-than sign; |
| # | comment delimiter |

The backslash (\) removes the special meaning of these parser metacharacters. A parser metacharacter preceded by a backslash is interpreted as its ASCII value. A newline character (ASCII 10) preceded by a backslash is equivalent to a blank.

Strings enclosed in single or double quotes form parts of a word. Metacharacters in these strings, including blanks and tabs, do not form separate words. Within pairs of backslashes or quotes, a newline preceded by a backslash gives a true newline character.

When **cs**h's input is not a terminal, the **#** character introduces a comment terminated by a newline.

CSH VARIABLES

csh maintains a set of variables. Each variable has a value equal to zero or more strings (words). Variables have names consisting of up to 80 letters and digits starting with a letter. The underscore character is considered a letter. The value of a variable may be displayed and changed by using the **set** and

unset commands. Some of the variables are Boolean, that is, the shell does not care what their value is, only whether they are set or not.

Some operations treat variables numerically. The **at** sign (@) command permits numeric calculations to be performed and the result assigned to a variable. The null string is considered to be zero, and any subsequent words of multi-word values are ignored.

After the input line is aliased and parsed, and before each command is executed, variable expansion is performed keyed by the dollar sign (\$) **character**. Variable expansion can be prevented by preceding the dollar sign with a backslash character (\) except within double quotes (") where substitution **always** occurs. Variables are never expanded if enclosed in single quotes. Strings quoted by single quotes are interpreted later (see *Command Substitution*) so variable substitution does not occur there until later, if at all. A dollar sign is passed unchanged if followed by a blank, tab, or end-of-line.

Input/output redirections are recognized before variable expansion, and are variable expanded separately. Otherwise, the command name and entire argument list are expanded together.

Unless enclosed in double quotes or given the **:q** modifier, the results of variable substitution may eventually be command and file name substituted. Within double quotes, a variable whose value consists of multiple words expands to a portion of a single word, with the words of the variable's value separated by blanks. When the **:q** modifier is applied to a substitution, the variable expands to multiple words with each word separated by a blank and quoted to prevent later command or file name substitution.

The following metasequences are provided for introducing variable values into the shell input. Except as noted, it is an error to reference a variable that is not set.

\$variable_name
\${variable_name}

When interpreted, this sequence is replaced by the words of the value of the variable *variable_name*, each separated by a blank. Braces insulate *variable_name* from subsequent characters that would otherwise be interpreted to be part of the variable name itself.

If *variable_name* is not a **cs** variable, but is set in the environment, that value is used. Non-**cs** variables cannot be modified as shown below.

\$variable_name[selector]
\${variable_name[selector]}

This modification selects only some of the words from the value of *variable_name*. The selector is subjected to variable substitution, and can consist of a single number or two numbers separated by a dash. The first word of a variable's value is numbered 1. If the first number of a range is omitted it defaults to 1. If the last member of a range is omitted it defaults to the total number of words in the variable (**\$#variable_name**). An asterisk meta-character used as a selector selects all words.

\$#variable_name
\${#variable_name}

This form gives the number of words in the variable, and is useful for forms using a [*selector*] option.

\$0 This form substitutes the name of the file from which command input is being read. An error occurs if the file name is not known.

\$number
\${number}

This form is equivalent to an indexed selection from the variable **argv** (**\$argv[number]**).

\$* This is equivalent to selecting all of *argv* (**\$argv[*]**).

The modifiers **:h**, **:t**, **:r**, **:q**, and **:x** can be applied to the substitutions above, as can CR **:gh**, CR **:gt**, and CR **:gr**. If curly braces ({}) appear in the command form, the modifiers must appear within the braces. *The current implementation allows only one : modifier on each \$d expansion.*

The following substitutions cannot be modified with **:** modifiers:

\$?variable_name
\${?variable_name}

Substitutes the string 1 if *variable_name* is set, 0 if it is not.

\$?0 Substitutes 1 if the current input file name is known, 0 if it is not.

\$\$ Substitutes the (decimal) 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 shell script.

Pre-Defined and Environment Variables

The following variables have special meaning to the shell. Of these **autologout**, **argv**, **cwd**, **home**, **path**, **prompt**, **shell**, and **status** are always set by the shell. Except for **cwd** and **status**, this setting occurs only at initialization (initial execution of **cs**). These variables are not modified unless modified explicitly by the user.

cs copies the HP-UX environment variable **USER** into the shell variable **user**, the environment variable **TERM** into **term**, the environment variable **HOME** into **home**, and **PATH** into **path**. **cs** copies these values back into the environment whenever the **cs** variables are reset.

In a windowed environment, if **cs** detects that the window has changed size, **cs** sets the environment variables **LINES** and **COLUMNS** to match the new window size.

| | |
|------------------|--|
| argv | This variable is set to the arguments of the cs command statement. It is from this variable that positional parameters are substituted; i.e., \$1 is replaced by \$argv[1] , etc. |
| cdpath | This variable gives a list of alternate directories searched to find subdirectories in <i>chdir</i> commands. |
| cwd | This variable contains the absolute path name of the current working directory. Whenever changing directories (using <i>cd</i>), this variable is updated. |
| echo | This variable is set by the -x command line option. If set, all built-in commands and their arguments are echoed to the standard output device just before being executed. Built-in commands are echoed before command and file name substitution, since these substitutions are then done selectively. For non-built-in commands, all expansions occur before echoing. |
| history | This variable is used to create the command history buffer and to set its size. If this variable is not set, no command history is maintained and history substitutions cannot be made. Very large values of history can cause shell memory overflow. Values of 10 or 20 are normal. All commands, executable or not, are saved in the command history buffer. |
| home | This variable contains the absolute path name to your home directory. The variable home is initialized from the HP-UX environment. File name expansion of tilde (~) refers to this variable. |
| ignoreeof | If set, cs ignores end-of-file characters from input devices that are terminals. cs exits normally when it encounters the end-of-file condition (CTRL-D typed as the first character on a command line). Setting <i>ignoreeof</i> prevents the current shell from being killed by an accidental (CTRL-D). However, to prevent an infinite loop of EOF input, cs terminates if it receives 26 consecutive EOFs. |
| mail | This variable contains a list of the files where cs checks for your mail. cs periodically (default is 10 minutes) checks this variable before producing a prompt upon command completion. If the variable contains a file name that has been modified since the last check (resulting from mail being put in the file), cs prints You have new mail . If the first word of the value of mail is numeric, that number specifies a different mail checking interval in seconds. If multiple mail files are specified, the shell says New mail in file_name , where <i>file_name</i> is the file containing the mail. |
| noclobber | This variable places restrictions on output redirection to ensure that files are not accidentally destroyed, and that commands using append redirection (>>) refer to existing files. |
| noglob | If set, file name expansion is inhibited. This is most useful in shell scripts that are not dealing with file names, or after a list of file names has been obtained and further expansions are not desirable. |
| nonomatch | If set, it is no longer an error for a file name expansion to not match any existing files. If there is no match, the primitive pattern is returned. It is still an error for the |

C

| | |
|-----------------|---|
| | primitive pattern to be malformed. For example, <code>'echo ['</code> still gives an error. |
| notify | If set, cs h notifies you immediately (through your standard output device) of background job completions. The default is unset (indicate job completions just before printing a prompt). |
| path | Each word of the path variable specifies a directory in which commands are to be sought for execution. A null word specifies your current working directory. If there is no path variable, only full path names can be executed. When path is not set and when users do not specify full path names, cs h searches for the command through the directories . (current directory) and /usr/bin . A cs h which is given neither the -c nor the -t option normally hashes the contents of the directories in the path variable after reading .cshrc , and each time the path variable is reset. If new commands are added to these directories while the shell is active, it is necessary to execute rehash for cs h to access these new commands. |
| prompt | This variable lets you select your own prompt character string. The prompt is printed before each command is read from an interactive terminal input. If a ! appears in the string, it is replaced by the current command history buffer event number unless a preceding \ is given. The default prompt is the percent sign (%) for users and the # character for the super-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 cs h during startup. |
| shell | This variable contains the name of the file in which the cs h program resides. This variable is used in forking shells to interpret files that have their execute bits set but which are not executable by the system. (See the description of <i>Non-Built-In Command Execution</i>). |
| status | This variable contains the status value returned by the last command. If the command terminated abnormally, 0200 is added to the status variable's value. Built-in commands which terminated abnormally return exit status 1, and all other built-in commands set status to 0. |
| time | This variable contains a numeric value that controls the automatic timing of commands. If set, cs h prints, for any command taking more than the specified number of cpu seconds, a line of information to the standard output device giving user, system, and real execution times plus a utilization percentage. The utilization percentage is the ratio of user plus system times to real time. This message is printed after the command finishes execution. |
| verbose | This variable is set by the -v command line option. If set, the words of each command are printed on the standard output device after history substitutions have been made. |

Command and File name Substitution

The remaining substitutions, command and file name substitution, are applied selectively to the arguments of built-in commands. This means that portions of expressions that are not evaluated are not subjected to these expansions. For commands which are not internal to the shell, the command name is substituted separately from the argument list. This occurs very late, after input-output redirection is performed, and in a child of the main shell.

Command Substitution

Command substitution is indicated by a command enclosed in grave accents (``...``). The output from such a command is normally broken into separate words at blanks, tabs and newlines, with null words being discarded; this text then replacing the original string. Within double quotes, only newlines force new words; blanks and tabs are preserved.

In any case, the single final newline does not force a new word. Note that it is thus possible for a command substitution to yield only part of a word, even if the command outputs a complete line.

File name Substitution

Each command *word* is processed as a pattern for file name substitution, also known as **globbing**, and replaced with a sorted list of file names which match the pattern. The form of the patterns is the Pattern Matching Notation defined by *regex*(5) with the following exceptions:

- Non-matching lists in bracket expressions are not supported.
- In a list of words specifying file name substitution it is an error for no pattern to match an existing file name, but it is not required for each pattern to match.
- The metanotation **a**(b,c,d)**e** is a shorthand for "abe ace ade". Left to right order is preserved, with results of matches being sorted separately at a low level to preserve this order. This construct may be nested. Thus:

```
~source/s1/{oldls,ls}.c
```

expands to

```
/home/source/s1/oldls.c /home/source/s1/ls.c
```

whether or not these files exist, without any chance of error if the home directory for **source** is **/home/source**. Similarly,

```
../{memo,*box}
```

might expand to

```
../memo ../box ../mbox
```

(Note that **memo** was not sorted with the results of matching ***box**.) As a special case, {, }, and { } are passed undisturbed.

Input/Output

The standard input and standard output of a command can be redirected with the following syntax:

| | |
|------------------|--|
| < <i>name</i> | Open file <i>name</i> (which is first variable, command and file name expanded) as the standard input. |
| << <i>word</i> | Read the shell input up to a line which is identical to <i>word</i> . <i>word</i> is not subjected to variable, file name or command substitution, and each input line is compared to <i>word</i> before any substitutions are done on this input line. Unless a quoting \, ', or ` appears in <i>word</i> , variable and command substitution is performed on the intervening lines, allowing \ to quote \$, \ and `. Commands which are substituted have all blanks, tabs, and newlines preserved, except for the final newline which is dropped. The resultant text is placed in an anonymous temporary file which is given to the command as standard input. |
| > <i>name</i> | |
| >! <i>name</i> | |
| >& <i>name</i> | |
| >&! <i>name</i> | The file <i>name</i> is used as standard output. If the file does not exist, it is created; if the file exists, it is truncated, and its previous contents are lost. If the variable noclobber is set, the file must not exist or be a character special file (e.g., a terminal or /dev/null) or an error results. This helps prevent accidental destruction of files. In this case the exclamation point (!) forms can be used to suppress this check. The forms involving the ampersand character (&) route the standard error into the specified file as well as the standard output. <i>name</i> is expanded in the same way as < input file names are. |
| >> <i>name</i> | |
| >>& <i>name</i> | |
| >>! <i>name</i> | |
| >>&! <i>name</i> | Uses file <i>name</i> as standard output the same as >, but appends output to the end of the file. If the variable noclobber is set, it is an error for the file not to exist unless one of the ! forms is given. Otherwise, it is similar to >. |

A command receives the environment in which the shell was invoked as modified by the input-output parameters and the presence of the command in a pipeline. Thus, unlike some previous shells, commands executed from a shell script have no access to the text of the commands by default; rather they receive the original standard input of the shell. The << mechanism should be used to present inline data. This permits shell scripts to function as components of pipelines and allows the shell to block-read its input.

Diagnostic output can be directed through a pipe with the standard output. Simply use the form `|&` rather than `|` by itself.

CSH UTILITIES

File Name Completion

In typing file names as arguments to commands, it is no longer necessary to type a complete name, only a unique abbreviation is necessary. When you want the system to try to match your abbreviation, press the ESC key. The system then completes the file name for you, echoing the full name on your terminal. If the abbreviation does not match an available file name, the terminal's bell is sounded. The file name may be partially completed if the prefix matches several longer file names. In this case, the name is extended up to the ambiguous deviation, and the bell is sounded.

File name completion works equally well when other directories are addressed. In addition, the tilde (~) convention for home directories is understood in this context.

Viewing a File or Directory List

At any point in typing a command, you can request "what files are available" or "what files match my current specification". Thus, when you have typed:

```
% cd ~speech/data/bench/fritz/
```

you may wish to know what files or subdirectories exist (in `~speech/data/bench/fritz`), without aborting the command you are typing. Typing CTRL-D at this point lists the files available. Files are listed in multicolumn format, sorted by column. Directories and executable files are identified by a trailing / and *, respectively. Once printed, the command is re-echoed for you to complete. Additionally, you may want to know which files match a prefix, the current file specification so far. If you had typed:

```
% cd ~speech/data/bench/fr
```

followed by a CTRL-D, all files and subdirectories whose prefix was `fr` in the directory `~speech/data/bench` would be printed. Notice that the example before was simply a degenerate case of this with a null trailing file name. (The null string is a prefix of all strings.) Notice also that a trailing slash is required to pass to a new sub-directory for both file name completion and listing. Note that the degenerate case

```
% ~^D
```

prints a full list of login names on the current system.

Command Name Recognition

Command name recognition and completion works in the same manner as file name recognition and completion above. The current value of the environment variable `PATH` is used in searching for the command. For example

```
% newa [Escape]
```

might expand to

```
% newaliases
```

Also,

```
% new [Control]-[D]
```

lists all commands (along `PATH`) that begin with `new`. As an option, if the shell variable `listpathnum` is set, a number indicating the index in `PATH` is printed next to each command on a [Control]-[D] listing.

Autologout

A new shell variable has been added called `autologout`. If the terminal remains idle (no character input) at the shell's top level for a number of minutes greater than the value assigned to `autologout`, you are automatically logged off. The `autologout` feature is temporarily disabled while a command is executing. The initial value of `autologout` is 60. If unset or set to 0, `autologout` is entirely disabled.

Command Line Control

A ^R re-prints the current command line; ^W erases the last word entered on the current command line.

Sanity

C shell restores your terminal to a sane mode if it appears to return from some command in raw, cbreak, or noecho mode.

Saving Your History Buffer

cs has the ability to save your history list between login sessions. If the shell variable `savehist` is set to a number, that number of command events from your history list is saved. For example, placing the line

```
set history=10 savehist=10
```

in your `.cshrc` file maintains a history buffer of length 10 and saves the entire list when you logout. When you log back in, the entire buffer is restored. The commands are saved in the file `.history` in your login directory.

EXTERNAL INFLUENCES

Environment Variables

LC_COLLATE determines the collating sequence used in evaluating pattern matching notation for file name substitution.

LC_CTYPE determines the interpretation of text as single and/or multi-byte characters, the classification of characters as letters, and the characters matched by character class expressions in pattern matching notation.

LANG determines the language in which messages are displayed.

If LC_COLLATE or LC_CTYPE is not specified in the environment or is set to the empty string, the value of LANG is used as a default for each unspecified or empty variable. If LANG is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of LANG. If any internationalization variable contains an invalid setting, cs behaves as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single- and multi-byte character code sets are supported.

WARNINGS

The `.cshrc` file should be structured such that it cannot generate any output on standard output or standard error, including occasions when it is invoked without an affiliated terminal. *rcp(1)* causes `.cshrc` to be sourced, and any output generated by this file, even to standard error causes problems. Commands such as *stty(1)* should be placed in `.login`, not in `.cshrc`, so that their output cannot affect *rcp(1)*.

cs has certain limitations. Words or environment variables can be no longer than 10240 characters. The system limits argument lists to 10240 characters. The number of arguments to a command which involves file name expansion is limited to one-sixth the number of characters allowed in an argument list. Command substitutions may substitute no more characters than are allowed in an 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, cs prints the directory it started in if it is different from the current directory; this can be misleading (i.e., wrong) because 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 interrupt `b`, the shell then immediately executes `c`. This is especially noticeable if this expansion results from an `alias`. It suffices to place the sequence of commands in parentheses to force it into a subshell; i.e., `(a ; b ; c)`.

Because of the signal handling required by cs, interrupts are disabled just before a command is executed, and restored as the command begins execution. There may be a few seconds delay between when a command is given and when interrupts are recognized.

Control over tty output after processes are started is primitive; perhaps this will inspire someone to work on a good virtual terminal interface. In a virtual terminal interface much more interesting things could be done with output control.

Alias substitution is most often used to clumsily simulate shell procedures; shell procedures should be provided rather than aliases.

Commands within loops, prompted for by `?`, are not placed in the *history* list. Control structure 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. All and more than one `:` modifier should be allowed on `$` substitutions.

Terminal type is examined only the first time you attempt recognition.

To list all commands on the system along **PATH**, enter [Space]-[Ctrl]-[D].

The `cs` metasequence `!~` does not work.

In an international environment, character ordering is determined by the setting of `LC_COLLATE`, rather than by the binary ordering of character values in the machine collating sequence. This brings with it certain attendant dangers, particularly when using range expressions in file name generation patterns. For example, the command,

```
rm [a-z]*
```

might be expected to match all file names beginning with a lowercase alphabetic character. However, if dictionary ordering is specified by `LC_COLLATE`, it would also match file names beginning with an uppercase character (as well as those beginning with accented letters). Conversely, it would fail to match letters collated after `z` in languages such as Norwegian.

The correct (and safe) way to match specific character classes in an international environment is to use a pattern of the form:

```
rm [[:lower:]]*
```

This uses `LC_CTYPE` to determine character classes and works predictably for all supported languages and codesets. For shell scripts produced on non-internationalized systems (or without consideration for the above dangers), it is recommended that they be executed in a non-NLS environment. This requires that `LANG`, `LC_COLLATE`, etc., be set to "C" or not set at all.

`cs` implements command substitution by creating a pipe between itself and the command. If the root file system is full, the substituted command cannot write to the pipe. As a result, the shell receives no input from the command, and the result of the substitution is null. In particular, using command substitution for variable assignment under such circumstances results in the variable being silently assigned a NULL value.

Relative path changes (such as `cd ..`), when in a symbolically linked directory, cause `cs`'s knowledge of the working directory to be along the symbolic path instead of the physical path.

Prior to HP-UX Release 9.0, `cs`, when getting its input from a file, would exit immediately if unable to execute a command (such as if it was unable to find the command). Beginning at Release 9.0, `cs` continues on and attempts to execute the remaining commands in the file. However, if the old behavior is desired for compatibility purposes, set the environment variable `EXITONERR` to 1.

AUTHOR

`cs` was developed by the University of California, Berkeley and HP.

FILES

| | |
|-----------------------------|--|
| <code>~/ .cshrc</code> | A <code>cs</code> script sourced (executed) at the beginning of execution by each shell. See WARNINGS |
| <code>~/ .login</code> | A <code>cs</code> script sourced (executed) by login shell, after <code>.cshrc</code> at login. |
| <code>~/ .logout</code> | A <code>cs</code> script sourced (executed) by login shell, at logout. |
| <code>/etc/passwd</code> | Source of home directories for <code>~name</code> . |
| <code>/usr/bin/sh</code> | Standard shell, for shell scripts not starting with a <code>#</code> . |
| <code>/etc/csh.login</code> | A <code>cs</code> script sourced (executed) before <code>~/ .cshrc</code> and <code>~/ .login</code> when starting a <code>cs</code> login (analogous to <code>/etc/profile</code> in the Bourne shell). |
| <code>/tmp/sh*</code> | Temporary file for <code><<</code> . |

SEE ALSO

`cd(1)`, `echo(1)`, `kill(1)`, `nice(1)`, `sh(1)`, `umask(1)`, `access(2)`, `exec(2)`, `fork(2)`, `pipe(2)`, `umask(2)`, `wait(2)`, `tty(7)`, `a.out(4)`, `environ(5)`, `lang(5)`, `regex(5)`.

C Shell tutorial in *Shells Users Guide*.

NAME
csplit - context split

SYNOPSIS
csplit [-s] [-k] [-f *prefix*] [-n *number*] *file* *arg1* [... *argn*]

DESCRIPTION
csplit reads *file*, separates it into *n*+1 sections as defined by the arguments *arg1* ... *argn*, and places the results in separate files. The maximum number of arguments (*arg1* through *argn*) allowed is 99 unless the -n *number* option is used to allow for more output file names. If the -f *prefix* option is specified, the resulting filenames are *prefix*00 through *prefix**NN* where *NN* is the two-digit value of *n* using a leading zero if *n* is less than 10. If the -f *prefix* option is not specified, the default filenames **xx**00 through **xx***NN* are used. *file* is divided as follows:

| Default Filename | Prefixed Filename | Contents |
|---------------------|-------------------------|--|
| xx 00 | <i>prefix</i> 00 | From start of <i>file</i> up to (but not including) the line referenced by <i>arg1</i> . |
| xx 01 | <i>prefix</i> 01 | From the line referenced by <i>arg1</i> up to the line referenced by <i>arg2</i> . |
| | | . |
| | | . |
| xx <i>NN</i> | <i>prefix</i> <i>NN</i> | From the line referenced by <i>argn</i> to end of <i>file</i> . |

If the *file* argument is -, standard input is used.
csplit supports the Basic Regular Expression syntax (see *regex*(5)).

- Options
csplit recognizes the following options:
- s Suppress printing of all character counts (csplit normally prints the character counts for each file created).
 - k Leave previously created files intact (csplit normally removes created files if an error occurs).
 - f *prefix* Name created files *prefix*00 through *prefix**NN* (default is **xx**00 through **xx***NN*).
 - n *number* The output file name suffix will use *number* digits instead of the default 2. This allows creation of more than 100 output files.

- Arguments (*arg1* through *argn*) to csplit can be any combination of the following:
- / *regex* / Create a file containing the section from the current line up to (but not including) the line matching the regular expression *regex*. The new current line becomes the line matching *regex*.
 - / *regex* / + *n*
/ *regex* / - *n* Create a file containing the section from the current line up to (but not including) the *n*th before (-*n*) or after (+*n*) the line matching the regular expression *regex*. (e.g., / **Page** / -5). The new current line becomes the line matching *regex*±*n* lines.
 - % *regex* % equivalent to / *regex* / , except that no file is created for the section.
 - line_number* Create a file from the current line up to (but not including) *line_number*. The new current line becomes *line_number*.
 - { *num* } Repeat argument. This argument can follow any of the above argument forms. If it follows a *regex* argument, that argument is applied *num* more times. If it follows *line_number*, the file is split every *line_number* lines for *num* times from that point until end-of-file is reached or *num* expires.
 - { * } Repeats previous operand as many times as necessary to finish input.

Enclose in appropriate quotes all *regex* arguments containing blanks or other characters meaningful to the shell. Regular expressions must not contain embedded new-lines. csplit does not alter or remove the

original file; it is the user's responsibility to remove it when appropriate.

EXTERNAL INFLUENCES

Environment Variables

LC_COLLATE determines the collating sequence used in evaluating regular expressions.

LC_CTYPE determines the characters matched by character class expressions in regular expressions.

LC_MESSAGES determines the language in which messages are displayed.

If **LC_COLLATE** or **LC_CTYPE** or **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **csplit** behaves as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single- and multi-byte character code sets are supported.

DIAGNOSTICS

Messages are 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. This warning also occurs if the file is exhausted before the repeat count is.

EXAMPLES

Create four files, **cobol100** through **cobol103**. After editing the "split" files, recombine them back into the original file, destroying its previous contents.

```
csplit -f cobol file '/procedure division/' /par5./ /par16./
```

Perform editing operations

```
cat cobol10[0-3] > file
```

Split a file at every 100 lines, up to 10,000 lines (100 files). The **-k** option causes the created files to be retained if there are fewer than 10,000 lines (an error message is still printed).

```
csplit -k file 100 '{99}'
```

Assuming that **prog.c** follows the normal C coding convention of terminating routines with a **}** at the beginning of the line, create a file containing each separate C routine (up to 21) in **prog.c**.

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

SEE ALSO

sh(1), **split(1)**, **environ(5)**, **lang(5)**, **regex(5)**.

STANDARDS CONFORMANCE

csplit: SVID2, SVID3, XPG2, XPG3, XPG4

NAME

ct - spawn getty to a remote terminal (call terminal)

SYNOPSIS

ct [-w *n*] [-x *n*] [-h] [-v] [-s *speed*] *telno*...

DESCRIPTION

ct dials *telno*, the telephone number of a modem that is attached to a terminal, and spawns a *getty*(1M) process to that terminal.

ct tries each line listed in file */etc/uucp/Devices* until it finds an available line with appropriate attributes or runs out of entries. If no lines are free, ct asks whether it should wait for a line, and if so, how many minutes it should wait before giving up. ct searches again for an available line at one-minute intervals until the specified limit is exceeded. Note that normally, ct disconnects the current tty line, so that the line can answer the incoming call. This is because ct assumes that the current tty line is connected to the terminal to spawn the *getty* process.

The *telno* argument specifies the telephone number, which can be composed of characters 0 through 9, -, =, *, and #. Use equal signs to signify secondary dial tones and minus signs for delays at appropriate places. The maximum length of *telno* is 31 characters. If more than one telephone number is specified, ct tries each in succession until one answers; this is useful for specifying alternate dialing paths.

When ct disconnects the current line, *getty* should not be spawned on this line if ct is going to make use of the same line to reconnect. To do this, set the entry for this line in the *inittab* file to *uugetty* instead of *getty* (see *inittab*(4)).

Options

ct recognizes the following options and command-line arguments:

- w*n* Instruct ct to wait for a line a maximum of *n* number of minutes, if lines are busy. If this option is specified, ct does not query the user about whether to wait for a line.
- x*n* Produce detailed output from program execution on the standard error output. This option is used for debugging. The debugging level *n* is a single digit; the most useful value is -x9.
- h Prevent ct from disconnecting ("hanging up") the current tty line. This option is necessary if the user is using a different tty line than the one used by ct to spawn the *getty*.
- v Verbose mode. The -v option is used with the -h option and causes ct to send a running narrative to the standard error output stream.
- s*speed* Set the data rate where *speed* is expressed in baud. The default rate is 1200.

After the user on the destination terminal logs out, ct prompts, **Reconnect?** If the response begins with the letter **n** the line is dropped. Otherwise, *getty* is restarted and the **login:** prompt is printed.

Of course, the destination terminal must be attached to a modem that can automatically answer incoming calls.

FILES

/var/adm/ctlog
/etc/uucp/Devices

SEE ALSO

cu(1), login(1), uucp(1), getty(1M), uugetty(1M).

NAME

ctags - create a tags file

SYNOPSIS

ctags [-xvFBatwu] files ...

DESCRIPTION

ctags makes a tags file for *ex*(1) (or *vi*(1)) from the specified C, Pascal and FORTRAN sources. A **tags** file gives the locations of specified objects (for C, functions, macros with arguments, and typedefs; Pascal, procedures, programs and functions; FORTRAN, subroutines, programs and functions) 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. Output is sorted in ascending collation order (see Environment Variables below). All objects except C *typedefs* are searched with a pattern, *typedefs* with a line number. Specifiers are given in separate fields on the line, separated by spaces or tabs. Using the *tags* file, **ex** can quickly find these objects' definitions.

- x Cause **ctags** to print a simple function index. This is done by assembling a list of function names, file names on which each function is defined, the line numbers where each function name occurs, and the text of each line. The list is then printed on the standard output. No *tags* file is created or changed.
- v Produce a page index on the standard output. This listing contains the function name, file name, and page number within that file (assuming 56-line pages to match *pr*(1)).

Files whose name ends in **.c** or **.h** are assumed to be C source files and are searched for C routine and macro definitions. 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.

Other options are:

- F Use forward searching patterns (/.../) (default).
- B Use backward searching patterns (?...?).
- a Add the information from the files to the *tags* file. Unlike re-building the *tags* file from the original files, this can cause the same symbol to be entered twice in the *tags* file. This option should be used with caution and then only in very special circumstances.
- t Create tags for typedefs.
- w Suppress warning diagnostics.
- u Update the specified files in *tags*; that is, all references to those files are deleted, and the new values are added to the file as in **-a** above. (Beware: this option is implemented in a way which is rather slow; it is usually faster to simply rebuild the *tags* file.)

The tag **main** is treated specially in C programs. The tag formed is created by adding **M** to the beginning of name of the file, with any trailing **.c** removed, and leading pathname components also removed. This makes use of **ctags** practical in directories with more than one program.

EXTERNAL INFLUENCES**Environment Variables**

LC_COLLATE determines the order in which the output is sorted.

LC_CTYPE determines the interpretation of the single- and/or multi-byte characters within comments and string literals.

If **LC_COLLATE** or **LC_CTYPE** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang*(5)) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **ctags** behaves as if all internationalization variables are set to "C". See *environ*(5).

International Code Set Support

Single- and multi-byte character code sets are supported with the exception that multi-byte character file names are not supported.

DIAGNOSTICS

Too many entries to sort.

An attempt to get additional heap space failed; the sort could not be performed.

Unexpected end of function in file *file*, line *line*.

The tags file may be incorrect.

A `}` character was found unexpectedly in the first column. This can lead to incorrect entries in the *tags* file.

Duplicate entry in file *file*, line *line*: *name*. Second entry ignored.

The same name was detected twice in the same file. A *tags* entry was made only for the first name found.

Duplicate entry in files *file1* and *file2*: *name* (Warning only).

The same name was detected in two different files. A *tags* entry was made only for the first name found.

EXAMPLES

Create a tags file named **tags** in the current directory for all C source (***.c**) files and all header (***.h**) files in the current directory:

```
ctags *.c[h]
```

Once the tags file exists in the current directory, it can be used with commands that support tag files (such as **vi** (see *vi(1)*)).

Use the tags file with **vi** to edit a particular function **myfunc()** located in one of the source files:

```
vi -t myfunc
```

While editing a C source file using **vi**, use the **ex**-mode tag command to edit function **myfunc()**:

```
:tag myfunc
```

Use **vi** to find **main()** in file **myprog.c**:

```
vi -t Mmyprog
```

While using **vi**, find **main()** in file **myprog.c** (does not have to be the file currently being edited):

```
:tag Mmyprog
```

WARNINGS

Recognition of **functions**, **subroutines**, and **procedures** for FORTRAN and Pascal is done in a very simple way. No attempt is made to deal with block structure; if there are two Pascal procedures in different blocks with the same name, a warning message is generated.

The method of deciding whether to look for C or Pascal and FORTRAN functions is an approximation, and can be fooled by unusual programs.

ctags does not know about **#ifdefs** and Pascal types.

It relies on the input being well formed to detect *typedefs*.

Use of **-tx** shows only the last line of typedefs.

ex is naive about *tags* files with several identical tags; it simply chooses the first entry its (non-linear) search finds with that tag. Such files can be created with either the **-u** or **-a** options or by editing a *tags* file.

If more than one (function) definition appears on a single line, only the first definition is indexed.

AUTHOR

ctags was developed by the University of California, Berkeley.

FILES

| | |
|--------------|----------------------------------|
| tags | output tags file |
| OTAGS | temporary file used by -u |

ctags(1)

ctags(1)

SEE ALSO

ex(1), vi(1).

STANDARDS CONFORMANCE

ctags: XPG4

C

NAME

cu - call another (UNIX) system; terminal emulator

SYNOPSIS

cu [-s *speed*] [-l *line*] [-h] [-q] [-t] [-d *level*] [-e|-o] [-m] [-n] [*telno* | *systemname* | *dir*]

XPG4 Syntax:

cu [-s *speed*] [-l *line*] [-h] [-q] [-t] [-d] [-e|-o] [-m] [-n] [*telno* | *systemname* | *dir*]

DESCRIPTION

cu calls up another system, which is usually a UNIX operating system, but can be a terminal or a non-UNIX operating system. cu manages all interaction between systems, including possible transfers of ASCII files.

Options

cu recognizes the following options and command-line arguments:

| | |
|-------------------|--|
| -s <i>speed</i> | Specify the transmission speed (110, 150, 300, 600, 1200, 2400, 3600, 4800, 7200, 9600, 19200). The default value is 300. |
| -l <i>line</i> | Specify a device name to use as the communication line. This can be used to override searching 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 obtained from file <code>/etc/uucp/Devices</code> . When the -l and -s options are used simultaneously, cu searches <code>/etc/uucp/Devices</code> to determine whether the requested speed for the requested line is available. If so, the connection is made at the requested speed; otherwise, an error message is printed and the call is not made. The specified device is usually a directly connected asynchronous line (such as <code>/dev/ttyapb</code>). In this case, a telephone number is not required, but the string <i>dir</i> can be used to specify that a dialer is not required. If the specified device is associated with an auto-dialer, a telephone number must be provided. |
| -h | Emulate local echo, supporting calls to other computer systems that expect terminals to be set to half-duplex mode. |
| -q | Use ENQ/ACK handshake (remote system sends ENQ, cu sends ACK.) |
| -t | Used when dialing an ASCII terminal that has been set to auto-answer. Appropriate mapping of carriage-return to carriage-return-line-feed pairs is set. |
| -d <i>level</i> | Print diagnostic traces. <i>level</i> is a number from 0-9, where higher <i>levels</i> produce more detail in the diagnostic messages. |
| -d | (XPG4 only.) Print diagnostic traces. The level is always 9. |
| -e (-o) | Generate even (odd) parity for data sent to the remote. |
| -m | Specify a direct line that has modem controls. Modem controls are ignored by cu. |
| -n | Cause the telephone number that cu dials to be requested interactively from the user rather than taking it from the command line. |
| <i>telno</i> | When using an automatic dialer, <i>telno</i> is the telephone number, with equal signs for secondary dial tone or minus signs for delays appropriately placed in the <i>telno</i> string. |
| <i>systemname</i> | A UUCP system name can be used instead of a telephone number (see <code>uucp(1)</code>); in this case, cu obtains an appropriate direct line or telephone number from <code>/etc/uucp/Systems</code> (including appropriate baud rate). cu dials each telephone number or direct line for <i>systemname</i> in the <code>Systems</code> file until a connection is made or all the entries are tried. |
| <i>dir</i> | Using <i>dir</i> ensures that cu uses the line specified by the -l option. |

After making the connection, cu runs as two processes:

- *transmit* process reads data from the standard input and, except for lines beginning with ~, passes it to the remote system;
- *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 to ensure that the buffer is not overrun. "Prompt handshaking" can be used to control transfer of ASCII files to systems that have no type-ahead capability but require data to be sent only after a prompt is given. This is described in detail below. Lines beginning with `~` have special meanings.

Transmit Process Commands

The *transmit* process interprets the following commands:

- `~.`, `~..` Terminate the conversation. On hard-wired lines, `~.` sends several EOF characters to log out the session, whereas `~..` suppresses the EOF sequence. In general the remote hard-wired machine is unaware of the disconnect if `~..` is used. On dial-up connections, `~.` and `~..` do not differ.
- `~!` Escape to an interactive shell on the local system.
- `~!cmd ...` Run *cmd* on the local system (via `sh -c`).
- `~&` Similar to `~!` but kill the receive process, restarting it upon return from the shell. This is useful for invoking sub-processes that read from the communication line where the receive process would otherwise compete for input.
- `~&cmd ...` Run *cmd* on the local system (via `sh -c`) and kill the receive process, restarting it later.
- `~|cmd` Pipe incoming data from the remote system through the standard input to *cmd* on the local system. To terminate, reset with either a `~&` or `~|` command.
- `~|` Resets the receive process following a `~|cmd` command.
- `~$cmd ...` Run *cmd* locally and send its output to the remote system.
- `~%cd` Change the directory on the local system. *Note:* `~!cd` causes the command to be run by a sub-shell, causing a return to the current directory upon completion.
- `~%take remote_source_file [local_destination_file]`
Copy file *remote_source_file* from the remote system to file *local_destination_file* on the local system. If *local_destination_file* is not specified, the *remote_source_file* argument is used in both places.
- `~%put local_source_file [remote_destination_file]`
Copy file *local_source_file* on local system to file *remote_destination_file* on remote system. If *remote_destination_file* is not specified, the *local_source_file* argument is used in both places.
- `~~ ...` Send the line `~ ...` to the remote system. If you use `cu` on the remote system to access a third remote system, send `~~.` to cause the second remote `cu` to exit.
- `~%break` Transmit a BREAK to the remote system.
- `~%nostop` Toggle between DC3/DC1 input control protocol and no input control. This is useful if the remote system does not respond properly to the DC3 and DC1 characters.
- `~%<file` Send the contents of the local file to the remote system using *prompt handshaking*. The specified file is read one line at a time, and each line is sent to the remote system when the *prompt sequence* is received. If no prompt is received by the time the *prompt timeout* occurs, the line is sent anyway. If the timeout is set to 0 seconds, or if the first character in the prompt sequence is a null character (`@`), the handshake always appears to be satisfied immediately, regardless of whether or not the remote system generates a prompt. This capability is intended mainly to facilitate transfer of ASCII files from HP-UX to an HP 3000 system running MPE. This is usually accomplished by running the MPE `FCOPY` utility and giving the command `from=;to=destfile;new` and then running the `cu` input diversion to send the file to `FCOPY` which saves it in *destfile*. This facility might be useful with other systems also, such as an HP 1000 running RTE.
- `~%setpt n` Specify the number of seconds to wait for a prompt before giving up. The default is 2 seconds. Specifying a timeout of 0 seconds disables handshaking; that is, handshake appears to complete immediately.
- `~%setps xy` Set the handshake prompt to the characters *xy*. The default is DC1. The prompt can be any one or two characters. To specify a control character for *x* or *y*, use the Ctrl-X

form where a circumflex (ASCII 94) precedes the character, as in `^X`. A null character can be specified with `^@`. (A null first character in the prompt implies a "null" prompt, which always appears to be satisfied.) A circumflex is specified by `^^`.

`~%>[>]file`

Divert output from the remote system to the specified file until another `~%>` command is given. When an output diversion is active, typing `~%>` terminates it, whereas `~%> anotherfile` terminates it and begins a new one. The output diversion remains active through a `~&` subshell, but unpredictable results can occur if input/output diversions are intermixed with `~%take` or `~%put`. The `~%>>` command appends to the named file. Note that these commands, which are interpreted by the *transmit* process, are unrelated to the `~>` commands described below, which are interpreted by the receive process.

`~suspend`

Suspend the `cu` session. *suspend* is the suspend character set in the terminal when `cu` was invoked (usually `^Z` — see *stty*(1)). As in all other lines starting with tilde, a `~suspend` line must be terminated by pressing **Return**.

Receive Process

The *receive* process normally copies data from the remote system to its standard output. A line from the remote that begins with `~>` initiates an output diversion to a file. The complete sequence is:

```
~>[>]: file
zero or more lines to be written to file
~>
```

Data from the remote is diverted (or appended, if `>>` is used) to *file*. The trailing `~>` terminates the diversion.

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 the current ones on the local system. Backslashes are inserted at appropriate places.

The use of `~%take` requires that the remote system support the `echo` and `cat` commands (see *echo*(1) and *cat*(1)). Also, `stty tabs` mode should be set on the remote system if tabs are being copied without expansion. When connecting to a machine that uses the eighth bit as a parity bit, `stty istrip` mode should be set on the local system.

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 if `~~` is used. For example, using the keyboard on system X, `uname` can be executed on Z, X, and Y as follows where lines 1, 3, and 5 are keyboard commands, and lines 2, 4, and 6 are system responses:

```
uname
Z
~!uname
X
~~!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.

EXTERNAL INFLUENCES

Environment Variables

LANG determines the language in which messages are displayed.

If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang*(5)) is used instead of **LANG**. If any internationalization variable contains an invalid setting, `cu` behaves as if all internationalization variables are set to "C". See *environ*(5).

International Code Set Support

Single- and multi-byte character code sets are supported.

DIAGNOSTICS

Exit code is zero for normal exit; non-zero (various values) otherwise.

EXAMPLES

To dial a system whose number is 9 201 555 1212 using 1200 baud:

```
cu -s1200 9=2015551212
```

If the speed is not specified, 300 is the default value.

To log in on a system connected by a direct line:

```
cu -l/dev/ttyXpX dir
```

To dial a system with the specific line and a specific speed:

```
cu -s1200 -l/dev/ttyXpX dir
```

To dial a system using a specific line:

```
cu -l/dev/culXpX 2015551212
```

To use a system name (yyyzzz):

```
cu yyyzzz
```

To connect directly to a modem:

```
cu -l/dev/culXX -m dir cu -l/dev/culXX -m dir
```

WARNINGS

cu buffers input internally.

AUTHOR

cu was developed by AT&T and HP.

FILES

```
/etc/uucp/Systems  
/etc/uucp/Devices  
/etc/uucp/Dialers  
/var/spool/locks/LCK..(tty-device)  
/dev/null
```

SEE ALSO

cat(1), ct(1), echo(1), stty(1), uname(1), uucp(1), uuname(1).

STANDARDS CONFORMANCE

cu: SVID2, SVID3, XPG2, XPG3, XPG4

NAME

cue - HP Character-Terminal User Environment (CUE)

SYNOPSIS

`/usr/bin/cue`

DESCRIPTION

CUE provides an easy-to-use, attractive, customizable environment that allows users on Series 800 HP-UX systems to easily identify themselves to the system and begin a work session. See **DEPENDENCIES** for supported terminal types.

A menubar is available for changing the native language of the session, changing the type of session to start upon a successful login, or getting on-line help. To obtain context-sensitive help at any time, press the function key labeled **HELP** (*f1*).

A pulldown menu and function keys (*f1-f8*) are displayed, allowing the user to modify various options or to get help. Before the login is initiated, the user has the option of interactively changing the native language of the session and the type of session to start upon a successful login.

The default native language is **C**, but the language is easily modifiable by entering the Language Menu which is accessible by selecting the Configuration item in the menu bar. The native language can also be specified as a parameter to **cuegetty** (see **cuegetty(1M)**).

The default session type is the POSIX shell, **sh**, but the session type can be easily changed to **tsh**, **keysh**, or **csh** by entering the Session Type Menu which is accessible by selecting the Configuration item in the menu bar.

The following standard **login** features are available:

- password aging
- logging invalid login attempts in `/var/adm/btmp`
- list of valid *ttys* for super-user login

CUE displays a visual screen that prompts for the *username* and corresponding *password*. If your username does not have a password, press the <carriage return> key to skip this field. Terminal echo is turned off (where possible) during typing of the password so that it will not appear on any written record of the session. After three unsuccessful login attempts, a *hangup* signal is issued.

If password aging has been invoked by the super-user on your behalf, your password may have expired. In this case, you will be diverted into **passwd** to change it, after which you can attempt to login again. See **passwd(1)**.

If login is not successfully completed within a certain period of time (e.g., five minutes), the terminal may be silently disconnected.

After a successful login, the accounting files are updated, initializing the user and group ids, group access list, and working directory. If the session type chosen is *tsh*, the **SHELL** to start in each *tsh* session is determined from corresponding user entries in the `/etc/passwd` file. **cue** then forks the appropriate shell by using the last component of the shell pathname preceded by a - (for example, **-sh** or **-ksh**). When the session type is invoked with its name preceded by a minus in this manner, the shell performs its own initialization, including execution of profile, login, or other initialization scripts.

For example, if the user login shell is *sh(1)* or *ksh(1)* the shell executes the profile files `/etc/profile` and `$HOME/.profile` if they exist (and possibly others as well). Depending on the contents of the profile files, messages regarding mail in your mail file or any messages you may have received since your last login may be displayed. At this point, *cuesession* is started to perform accounting procedures, display messages, and start your session.

If `/var/adm/btmp` is present, all unsuccessful login attempts are logged to this file. This feature is disabled if the file is not present. A summary of bad login attempts can be viewed by users with appropriate privileges by using **lastb**, see **last(1M)**.

If `/etc/securetty` is present, login security is in effect, meaning that only users with appropriate privileges are allowed to login successfully on the *ttys* listed in this file. Restricted *ttys* are listed by device name, one per line. Valid *tty* names are dependent on installation. Some examples could be **console**, **tty01**, **ttya1**, etc. Note that this feature does not inhibit a normal user from using **su**.

Starting Cue

There are several methods that can be used to start **cue**.

- An entry for **cuegetty** can be placed in the **/etc/inittab** file. See *cuegetty*(1M). This is the preferred method as the user does not need to do anything further to start **cue**.
- Start **cue** from the command line by typing: **cue**.
- Start **cue** by making it the last entry in the user's **.login** configuration file.

Multiple **cue** logins may run simultaneously on separate terminals attached to the same local host. **cuegetty** can be configured in the **/etc/inittab** file for all users.

Remote users to the CUE system must access CUE by entering the **cue** command at the command-line prompt or as the last item in the user's **.login** configuration file.

EXTERNAL INFLUENCES**Environment Variables**

cue invokes the user's session with the following default environment:

CUESESSION is set to the session type selected. Valid values are:

| | |
|-----------------------|--------------------------------------|
| /usr/bin/sh | POSIX Shell (DEFAULT) |
| /usr/bin/tsm | manages up to 10 sessions at once |
| /usr/bin/keysh | Easy Context-Sensitive Softkey Shell |
| /usr/bin/ksh | Korn Shell |
| /usr/bin/csh | C Shell |

HOME is set to the home directory of the user

LANG is set to the native language selected (C is the default)

LOGNAME is set to the user name

MAIL is set to **/var/mail/\$LOGNAME**

NLSPATH is set to the path applications search for NLS message catalogs, usually **/usr/lib/nls/%L/%N.cat**

PATH is set to the path to be searched for commands : **/usr/bin**

SHELL is set to the user's default shell (from **/etc/passwd**)

Several methods are available to modify or add to this list depending on the desired scope of the resulting environment variable.

Basic environment variables can be set for all CUE users on a system by setting the values in **/etc/profile** and **/etc/csh.login**. Personal environment variables can be set on a per-user basis in the script file **\$HOME/.profile** for **sh** and **ksh** users or **.cshrc** for **csh** users.

Note that alias and function definitions need to be included in the file specified by **ENV** for **ksh**, as this file will be sourced for each invocation of the shell. For **csh** users, the **.cshrc** file should be structured such that it cannot generate any output on standard output or standard error, including occasions when it is invoked without an affiliated terminal. The **rcp** command sources the **.cshrc** file and any output generated by this file, even to standard error, causes problems. Commands such as **stty** should be placed in the **.login** file, not in **.cshrc**, so that their output cannot affect **rcp**.

For users with appropriate privileges, **PATH** is augmented to include **/etc**.

(Series 800 Only)

VT320 Terminal Support

Because the VT320 terminal has predefined local functions for keys labeled as F1, F2, F3 and F4, users should use following mapping when they desire to use function keys:

| HP or Wyse60 | VT320 or HP 700/60 in VT320 mode |
|--------------|--|
| F1 | PF2 ! |
| F2 | PF1 ! |
| F3 | space bar |
| F4 | PF3 ! |
| F5 | F10, [EXIT], F5 * |
| F6 | none |
| F7 | F18, first unlabeled key to right of Pause/Break* |
| F8 | F19, second unlabeled key to right of Pause/Break* |

* When using PC-AT keyboard with HP 700/60 in VT320 mode

! See "Configuration: HP 700/60 in DEC mode, or DEC terminals with PC-AT type keyboard"

Further, since DEC terminals do not support softkey menu, no such menu is displayed on these terminals.

Many applications tend to use TAB for forward navigation (moving from one field to another) and shift-TAB is used for backward navigation. Users having DEC terminals or using terminals in DEC emulation modes such as VT100 or VT320 may note that these terminals/emulators may give out same character for TAB and shift-TAB. As such, it is impossible for an application to distinguish between TAB and shift-TAB, and both of them treated as if a TAB key was pressed. It might present slight overhead to users in case they want to go backwards. Now instead, they should complete rest of the inputs and get back to the desired field later.

VT100 Terminal Support

VT100 does not allow the (f1-f8) function keys to be configured. Therefore, the following keyboard mappings will apply to VT100 terminals:

| HP or Wyse60 | VT100 or HP 700/60 in VT100 mode |
|--------------|------------------------------------|
| F1 | PF2 ! |
| F2 | PF1 ! |
| F3 | space bar |
| F4 | [PF3], [space bar] or [PF3], [=] ! |
| F5 | return |
| F6 | none |
| F7 | none |
| F8 | none |

! See "Configuration: HP 700/60 in DEC mode, or DEC terminals with PC-AT type keyboard"

Further, since DEC terminals do not support softkey menu, no such menu is displayed on these terminals.

Many applications tend to use TAB for forward navigation (moving from one field to another) and shift-TAB is used for backward navigation. Users having DEC terminals or using terminals in DEC emulation modes such as VT100 or VT320 may note that these terminals/emulators may give out same character for TAB and shift-TAB. As such, it is impossible for an application to distinguish between TAB and shift-TAB, and both of them treated as if a TAB key was pressed. It might present slight overhead to users in case they want to go backwards. Now instead, they should complete rest of the inputs and get back to the desired field later.

Configuration: HP 700/60 terminal in DEC mode, or DEC terminal with PC-AT type keyboard

Customers using the following configuration may want to be aware of the following keyboard difference.

(Series 800 Only)

It may be possible for a user with the "HP 700/60 terminal in DEC mode, or DEC terminal with PC-AT type keyboard" configuration to be told to press function key F1 through F4 to achieve some desired result. For HP 700/60 terminal in DEC mode or DEC terminals, these functions keys may be mapped onto PF1-PF4 keys. (see "Keyboard Mappings"). However, the PC-AT type keyboard does not provide PF1, PF2, PF3, or PF4 keys, as does the DEC/ANSI keyboard.

Keyboard Mappings

| | |
|------------|---------------|
| "Num Lock" | maps to "PF1" |
| "/" | maps to "PF2" |
| "*" | maps to "PF3" |
| "_" | maps to "PF4" |

The "Num Lock", "/", "*", and "_" keys are located on the keyboard, in a row, above the number pad on the right side of the keyboard. Please note that although this keyboard is called a PC-AT type keyboard, it is supplied by HP. A PC-AT type keyboard can be recognized by location of ESC key at the left-top of the keyboard.

Wyse60 Terminal Support

On Wyse60, use DEL (located next to Backspace) key to backspace. On an HP 700/60 with a PC-AT type keyboard in Wyse60 mode, the DEL key is located in the bottom row on the number pad.

Wyse60 terminals provide a single line to display softkey labels unlike HP terminals which provide two lines. Sometimes this may result in truncated softkey labels. For example, "Help on Context" label for F1 may appear as "Help on C". Some standard labels for screen-oriented applications such as SAM and swinstall are as follows:

| | |
|----------------------------|-----------------|
| On wyse60 may appear as .. | means |
| Help On C | Help On Context |
| Select/D | Select/Deselect |
| Menubar | Menubar on/off |

Internationalization

All screens, labels, and messages are localizable. The message catalog **cue.cat** contains the localized representations of the default labels and messages. **cue** will read the appropriate message catalog indicated by the **LANG** environment variable and display the localized strings. By selecting a native language in the Language Menu, the language of the CUE screens and the future work session can be specified. If the message catalog exists for **cue** in the language selected, **cue** will be redisplayed in that language. If not, the CUE screens will continue in the current language and the work session that is started after a successful login will be started in the language selected. In either case, the **LANG** environment variable will be set appropriately for the resulting work session.

If **cue** will be started on the command line or as the last item in the **.login** file, the CUE screens will be brought up using the language specified by the **LANG** environment variable. If CUE screens do not exist, then the default native language, C, will be used.

To use **cue** with Asian languages, the **AWTERM** environment variable must be set to **hpterm-asian**. This will allow the text in Asian fonts to be displayed properly by **cue**.

If **cue** will be started by **cuegetty**, it is possible to start up the CUE Login Screens in a language other than the default, C, by invoking **cuegetty** with the **-L nls_language** option. Of course, CUE screens and the **cue.cat** file must exist for the **nls_language** specified.

DEPENDENCIES

CUE is available only on Series 800 systems, and is compatible only with the following terminals:

HP 700/92 HP 700/94 HP 2392 HP 2394 VT 320 VT 100 WYSE 60

WARNINGS

cue is an HP proprietary command, which will be *obsoleted* in a future release, and is not portable to other vendor's platforms.

(Series 800 Only)

FILES

| | |
|---|--|
| <code>/var/adm/btmp</code> | history of bad login attempts |
| <code>/etc/logingroup</code> | group file - defines group access lists |
| <code>/etc/motd</code> | message-of-the-day |
| <code>/etc/passwd</code> | password file - defines users, passwords, and primary groups |
| <code>/etc/profile</code> | system profile (initialization for all users) |
| <code>/etc/securetty</code> | list of valid ttys for root login |
| <code>/etc/utmp</code> | users currently logged in |
| <code>/var/adm/wtmp</code> | history of logins, logouts, and date changes |
| <code>/var/mail/<i>your-name</i></code> | mailbox for user <i>your-name</i> |
| <code>/usr/bin/cue</code> | <i>cue</i> executable |
| <code>/usr/sbin/cuegetty</code> | <i>cuegetty</i> executable |
| <code>/usr/newconfig/etc/cue.inittab</code> | template for <code>/etc/inittab</code> |
| <code>/etc/cue.dm</code> | screen descriptions |
| <code>/usr/sbin/cuesession</code> | starts selected session type |
| <code>/usr/lib/nls/\$LANG/cue.cat</code> | NLS message catalog |
| <code>/usr/share/man/man1.Z/cue.1</code> | man page for <i>cue</i> (1) |
| <code>/usr/share/man/man1m.Z/cuegetty.1m</code> | man page for <i>cuegetty</i> (1M) |

SEE ALSO

csh(1), cuegetty(1M), env(1), keysh(1), ksh(1), login(1), passwd(1), sh(1), tsm(1), btmp(4), environ(5), hpnl(5), lang(5).

NAME

cut - cut out (extract) selected fields of each line of a file

SYNOPSIS

```
cut -c list [file ...]
cut -b list [-n] [file ...]
cut -f list [-d char] [-s] [file ...]
```

DESCRIPTION

cut cuts out (extracts) columns from a table or fields from each line in a file; in data base parlance, it implements the projection of a relation. Fields as specified by *list* can be fixed length (defined in terms of character or byte position in a line when using the **-c** or **-b** option), or the length can vary from line to line and be marked with a field delimiter character such as the tab character (when using the **-f** option). **cut** can be used as a filter; if no files are given, the standard input is used.

When processing single-byte character sets, the **-c** and **-b** options are equivalent and produce identical results. When processing multi-byte character sets, when the **-b** and **-n** options are used together, their combined behavior is very similar, but not identical to the **-c** option.

Options

Options are interpreted as follows:

- | | |
|-----------------------|---|
| <i>list</i> | A comma-separated list of integer <i>byte</i> (-b option), <i>character</i> (-c option), or <i>field</i> (-f option) numbers, in increasing order, with optional - to indicate ranges. For example: <div style="margin-left: 40px;"> 1,4,7 Positions 1, 4, and 7. 1-3,8 Positions 1 through 3 and 8. -5,10 Positions 1 through 5 and 10. 3- Position 3 through last position. </div> |
| -b <i>list</i> | Cut based on a list of bytes. Each selected byte is output unless the -n option is also specified. |
| -c <i>list</i> | Cut based on character positions specified by <i>list</i> (-c 1-72 extracts the first 72 characters of each line). |
| -f <i>list</i> | Where <i>list</i> is a list of fields assumed to be separated in the file by a delimiter character (see -d); for example, -f 1,7 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. |
| -d <i>char</i> | 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. Adjacent field delimiters delimit null fields. <i>char</i> may be an international code set character. |
| -n | Do not split characters. If the high end of a range within a list is not the last byte of a character, that character is not included in the output. However, if the low end of a range within a list is not the first byte of a character, the entire character <i>is</i> included in the output." |
| -s | Suppresses lines with no delimiter characters when using -f option. Unless -s is specified, lines with no delimiters appear in the output without alteration. |

Hints

Use **grep** to extract text from a file based on text pattern recognition (using regular expressions). Use **paste** to merge files line-by-line in columnar format. To rearrange columns in a table in a different sequence, use **cut** and **paste**. See *grep(1)* and *paste(1)* for more information.

EXTERNAL INFLUENCES**Environment Variables**

LC_CTYPE determines the interpretation of text as single and/or multi-byte characters.

If **LC_CTYPE** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a

default of "C" (see *lang(5)*) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **cut** behaves as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

cut supports both single- and multi-byte character code sets. International code set characters may be specified in the *char* given to the **-d** option. **cut** recognizes the international code set characters according to the locale specified in the **LC_CTYPE** environment variable.

EXAMPLES

Password file mapping of user ID to user names:

```
cut -d : -f 1,5 /etc/passwd
```

Set environment variable **name** to current login name:

```
name=`who am i | cut -f 1 -d " "`
```

Convert file **source** containing lines of arbitrary length into two files where **file1** contains the first 500 bytes (unless the 500th byte is within a multi-byte character), and **file2** contains the remainder of each line:

```
cut -b 1-500 -n source > file1
cut -b 500- -n source > file2
```

DIAGNOSTICS

line too long

Line length must not exceed **LINE_MAX** characters or fields, including the new-line character (see *limits(5)*).

bad list for b/c/f option

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

no fields *list* is empty.

WARNINGS

cut does not expand tabs. Pipe text through *expand(1)* if tab expansion is required.

Backspace characters are treated the same as any other character. To eliminate backspace characters before processing by **cut**, use the **fold** or **col** command (see *fold(1)* and *col(1)*).

AUTHOR

cut was developed by OSF and HP.

SEE ALSO

grep(1), *paste(1)*.

STANDARDS CONFORMANCE

cut: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

date - display or set the date and time

SYNOPSIS

```
date [-u]
date [-u] +format
date [-u] [mmddhhmm[[cc]yy]]
date [-a [-]sss[.fff]]
```

DESCRIPTION

The **date** command displays or sets the current HP-UX system clock date and time. Since the HP-UX system operates in Coordinated Universal Time (UTC), **date** automatically converts to and from local standard or daylight/summer time, based on your **TZ** environment variable. See *Environment Variables* in EXTERNAL INFLUENCES below.

Options

date recognizes the following option:

- u Input and output values in Coordinated Universal Time (UTC), functionally equivalent to Greenwich Mean Time (GMT), instead of in local time.
- 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.

Formats

The **date** command has two forms for displaying the date and time and one form for setting them.

date [-u]

Display the current date and time. The output is the same as for the **%c** formatting directive for all languages except the **C** default language. See *Formatting Directives* and **EXAMPLES** below.

date [-u] +format

Display the current date and time according to formatting directives specified in *format*, which is a string of zero or more formatting directives and ordinary characters. If it contains blanks, enclose it in apostrophes or quotation marks.

See *Formatting Directives* below.

All ordinary characters are copied unchanged into the output string.

The output string is always terminated with a newline character.

If **+** is specified and *format* is omitted, only a newline is output.

date [-u] [mmddhhmm[[cc]yy]]

Set the HP-UX system clock to the date and time specified. You require the superuser privilege.

If you include the **-u** option, the specified date and time is assumed to be in Coordinated Universal Time (UTC).

The numeric argument is interpreted left to right in two-digit pairs as follows:

```
mm  Month number [01-12].
dd  Day number in the month [01-31].
hh  Hour number (24-hour system) [00-23].
mm  Minute number [00-59].
cc  Century minus one [19-20].
yy  Last two digits of the year number [70-99, 00-37 (1970-1999, 2000-2037)]. If
    omitted, the current year is used.
```

If you attempt to set the date backwards, **date** generates the warning,

do you really want to run time backwards?[yes/no]

Type **yes** or the equivalent for your locale to set the clock backwards; anything else to cancel the command.

When **date** is used to set the date, a pair of date change records is written to the file `/var/adm/wtmp`.

(XPG4 only.) No warning is generated if date is set backwards.

Formatting Directives

The following formatting directives, shown without the optional field width and precision specification, are replaced by the indicated characters. If a directive is not one of the following, the result is undefined.

The output for digits, characters, and words depends on the language/locale settings. See *Environment Variables* in EXTERNAL INFLUENCES below.

The examples assume that the **date** command was executed on Wednesday, January 12, 1994 at 7:45:58 p.m. Pacific Standard Time, using the **C** default language.

- %a** Abbreviated weekday name. For example, **Wed**.
- %A** Full weekday name. For example, **Wednesday**.
- %b** Abbreviated month name. For example, **Jan**.
- %B** Full month name. For example, **January**.
- %c** Current date and time representation. For example, **Wed Jan 12 19:45:58 1994**.
- %C** Century (the year divided by 100 and truncated to an integer) as a two-digit decimal number [00-99]. For example, **19**.
- %d** Day of the month as a two-digit decimal number [01-31]. For example, **12**.
- %e** Day of the month as a two-character decimal number with leading space fill [" 1"- "31"]. For example, **12**.
- %E** Combined Emperor/Era name and year.
- %H** Hour (24-hour clock) as a two-digit decimal number [00-23]. For example, **19**.
- %I** Hour (12-hour clock) as a two-digit decimal number [01-12]. For example, **07**.
- %j** Day of the year as a three-digit decimal number [001-366]. For example, **012**.
- %m** Month as a decimal two-digit number [01-12]. For example, **01**.
- %M** Minute as a decimal two-digit number [00-59]. For example, **45**.
- %n** Newline character.
- %N** Emperor/Era name.
- %o** Emperor/Era year.
- %p** Equivalent of either AM or PM. For example, **PM**.
- %R** Time as %H:%M
- %S** Second as a two-digit decimal number (allows for possible leap seconds) [00-61]. For example, **58**.
- %t** Tab character.
- %u** Weekday as a one-digit decimal number [1-7 (Monday-Sunday)]. For example, **3**.
- %U** Week number of the year (Sunday as the first day of the week) as a two-digit decimal number [00-53]. All days that precede the first Sunday in the year are considered to be in week 00. For example, **02**.
- %V** Week number of the year (Monday as the first day of the week) as a two-digit decimal number [01-53]. If the week containing January 1 has four or more days in the new year (January 1 is Thursday or sooner), it is designated as week 01; otherwise, (January 1 is Friday or later), it is designated as the last week of the previous year, and the next week is week 01. For example, **02**.

- %w** Weekday as a one-digit decimal number [0-6 (Sunday-Saturday)]. For example, 3.
- %W** Week number of the year (Monday as the first day of the week) as a two-digit decimal number [00-53]. All days that precede the first Monday in the year are considered to be in week 00. For example, 02.
- %x** Current date representation. For example, 01/12/94.
- %X** Current time representation. For example, 19:45:58.
- %y** Year without century as a two-digit decimal number [00-99]. For example, 93.
- %Y** Year with century as a four-digit decimal number [1970-2037]. For example, 1994.
- %Z** Time zone name (or no characters if time zone cannot be determined). For example, PST.
- %%** The % character.

Obsolescent Directives

The following directives are provided for backward compatibility. It is recommended that the preceding directives be used instead.

- %D** Date in usual U.S. format. For example, 01/12/94. Use **%x** or **%m/%d/%y** instead.
- %F** Full month name. For example, **January**. Use **%B** instead.
- %h** Abbreviated month name. For example, **Jan.** Use **%b** instead.
- %r** Time in 12-hour U.S. format. For example, 07:45:58 PM. Use **"%I:%M:%S %p"** instead.
- %T** Time in 24-hour U.S. format. For example, 19:45:58. Use **%X** or **%H:%M:%S** instead.
- %z** Time zone name (or no characters if time zone cannot be determined). For example, PST. Use **%Z** instead.

Modified Formatting Directives

Some Formatting Directives can be modified by the **E** and **O** modifier characters to indicate a different format or specification for the language specified in the **LC_TIME** environment variable.

If the corresponding keyword (**era**, **era_year**, **era_d_fmt**, and **alt_digit**) is not specified or not supported, the unmodified field descriptor value is used. The command

```
LC_ALL=language locale -ck era era_year era_d_fmt alt_digit
```

displays the keywords and their values in the specified *language* (see *locale(1)*).

- %Ec** Alternate appropriate date and time representation.
- %EC** The name of the base year in alternate representation.
- %Ex** Alternate date representation.
- %Ey** Offset from **%EC** (year only) in the alternate representation.
- %EY** Full alternate year representation.
- %Od** Day of month using the alternate numeric symbols.
- %Oe** Day of month using the alternate numeric symbols with leading space-character fill if applicable.
- %OH** Hour (24-hour clock) using the alternate numeric symbols.
- %OI** Hour (12-hour clock) using the alternate numeric symbols.
- %Om** Month using the alternate numeric symbols.
- %OM** Minutes using the alternate numeric symbols.
- %OS** Seconds using the alternate numeric symbols.
- %OU** Week number of the year (Sunday is the first day of the week) using the alternate numeric symbols.
- %Ow** Weekday as number using the alternate numeric symbols (Sunday=0).
- %OW** Weekday number of the year (Monday is the first day of the week) using the alternate numeric symbols.

%Oy Year (offset from %C) in alternate representation.

Field Width and Precision

An optional field width and precision specification can immediately follow the initial % of a formatting directive in the following order:

- [-|0]width** The decimal digit string *width* specifies a *minimum* field width in which the result of the conversion is right- or left-justified. The default is right-justified with space padding on the left. If the string starts with "-", the result is left-justified with space padding on the right. If the string starts with "0", the result is right-justified and padded with zeros on the left.
- .prec** The decimal digit string *prec* specifies the *minimum* number of digits to appear for the **d**, **H**, **I**, **j**, **m**, **M**, **o**, **S**, **U**, **w**, **W**, **y**, and **Y** numeric directives. If a directive supplies fewer digits than specified by the precision, it will be expanded with leading zeros.
prec specifies the *maximum* number of characters to be used from the **a**, **A**, **b**, **B**, **c**, **D**, **E**, **F**, **h**, **n**, **N**, **p**, **r**, **t**, **T**, **x**, **X**, **z**, **Z**, and % text directives. If a directive supplies more characters than specified by the precision, excess characters are truncated on the right.

If no field width or precision is specified for a **d**, **H**, **I**, **m**, **M**, **S**, **U**, **W**, or **y** directive, the default is **.2**; for the **j** directive, the default is **.3**; for **Y**, the default is **.4**; for **w**, the default is **.1**.

EXTERNAL INFLUENCES

Environment Variables

LC_CTYPE determines the interpretation of the bytes within the *format* string as single- and/or multi-byte characters.

LC_NUMERIC determines the characters used to form numbers for those directives that produce numbers in the output. The characters used are those defined by **alt_digit** (see *locale(1)*) and **ALT_DIGIT** in *langinfo(5)*.

LC_TIME determines the content (for example, the weekday names produced by the %**a** directive) and format (for example, the current time representation produced by the %**X** directive) of date and time strings output by the **date** command.

LC_MESSAGES determines the language in which messages (other than the date and time strings) are displayed.

If **LC_CTYPE**, **LC_NUMERIC**, **LC_TIME**, or **LC_MESSAGES** is not specified or is null, it defaults to the value of **LANG**.

If **LANG** is not specified or is null, it defaults to **C** (see *lang(5)*).

If any internationalization variable contains an invalid setting, all internationalization variables default to **C** (see *environ(5)*).

TZ determines the conversion between the system time in UTC and the time in the user's local time zone. See *environ(5)* and *tztab(4)*. **TZ** also determines the content (that is, the time-zone name produced by the %**z** and %**Z** directives) of date and time strings output by the **date** command.

If **TZ** is not set or is set to the empty string, its default value is **EST5EDT**.

International Code Set Support

Single- and multi-byte character code sets are supported.

DIAGNOSTICS

The following messages may be displayed.

bad conversion

The date/time specification is syntactically incorrect. Check it against the usage and for the correct range of each of the digit-pairs.

bad format character - c

The character *c* is not a valid format directive, field width specifier, or precision specifier.

do you really want to run time backwards?[yes/no]

The date/time you specified is earlier than the current clock value. Type **yes** (or the equivalent for your locale) to set the clock backwards; anything else to cancel the command.

no permission

You need the superuser privilege to change the date.

EXAMPLES

Date in Different Languages

Display the date. In this example, the **TZ** environment variable contains **PST8PDT**, and the language environment variables are set as noted.

```
date      → Fri Aug 20 15:03:37 PDT 1993  ← C (default)
date -u   → Fri Aug 20 22:03:37 UTC 1993  ← C (default)
date      → Fri, Aug 20, 1993 03:03:37 PM ← en_US.roman8 (U.S. English)
date      → Fri. 20 Aug, 1993 03:03:37 PM ← en_GB.roman8 (U.K. English)
date      → 20/08/1993 15.47.47          ← pt_PT.roman8 (Portuguese)
```

Set Date

Set the date to Oct 8, 12:45 a.m.

```
date 10080045
```

Display Formatted Date

Display the current date and time using a format. Note the use of quotation marks due to the blanks in the format.

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

The output resembles the following:

```
DATE: 10/08/87
TIME: 12:45:05
```

Display Formatted Date Using Local Language Conversion

With the date as set in the "Set Date" example above and **LC_TIME** set to **de_De.roman8** (German):

```
date +%-%4.4h %2.1d %H:%M'
```

generates output similar to:

```
Okt      8 12:45
```

where the month field is four characters long, flush-left, and space-padded on the right if the month name is shorter than four characters. The day field is two characters long, with leading zeros suppressed.

WARNINGS

The former HP-UX format directive **A** has been changed to **W** for ANSI compatibility.

Changing the date while the system is running in multiuser mode should be avoided to prevent disrupting user-scheduled and time sensitive programs and processes. Also, changing the date can cause *make*(1) and the SCCS and *cron*(1M) subsystems to behave in an unexpected manner. The **cron** daemon should be killed prior to setting the date backwards, then restarted. SCCS files should be checked with the **val** command (see *val*(1)) if deltas have been made while the clock was wrongly set.

The following formatting directives may be deleted from future releases: **%E**, **%F**, **%O**, **%Z**.

Currently, the maximum date supported is December 31, 2037 23:59:00 UTC.

AUTHOR

date was developed by AT&T and HP.

FILES

```
/var/adm/wtmp
```

SEE ALSO

locale(1), stime(2), ctime(3C), strftime(3C), tztab(4), environ(5), lang(5), langinfo(5).

STANDARDS CONFORMANCE

date: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2


d

NAME

dc - desk calculator

SYNOPSISdc [*file*]**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. (See *bc*(1), a preprocessor for **dc** that provides infix notation and a C-like syntax that implements functions. **bc** also provides reasonable control structures for programs.) 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. An end of file on standard input or the **q** command stop dc. 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 or A-F. It can be preceded by an underscore (_) to input a negative number. Numbers can 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 and a warning generated. The remainder is calculated according to the current scale factor; it is not the integer modulus function. 7 % 3 yields .1 (one tenth) if scale is 1 because 7 / 3 is 2.3 with .1 as the remainder. |
| s <i>x</i> | The top of the stack is popped and stored into a register named <i>x</i> , where <i>x</i> can be any character. If the s is capitalized, <i>x</i> is treated as a stack and the value is pushed on it. |
| l <i>x</i> | The value in register <i>x</i> is pushed on the stack. 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. If q is capitalized, 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. Strings can be nested by using nested pairs of brackets. |
| < <i>x</i> > <i>x</i> = <i>x</i> ! <i>x</i> !< <i>x</i> !> <i>x</i> != <i>x</i> | The top two elements of the stack are popped and compared. Register <i>x</i> 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 an HP-UX system command (unless the next character is <, >, or =, in which case appropriate relational operator above is used). |
| 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. See below for notes on output base.
- 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.
- K pushes the scale factor on the top of the stack.
- 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.
- ; and : Used by bc for array operations.
- Y Generates debugging output for dc itself.

The input base may be any number, but only the digits 0-9 and A-F are available for input, thus limiting the usefulness of bases outside the range 1-16. All 16 possible digits may be used in any base; they always take their conventional values.

The output base may be any number. Bases in the range of 2-16 generate the "usual" results, with the letters A-F representing the values from 10 through 16. Bases 0 and 1 generate a string of 1s whose length is the value of the number. Base -1 generates a similar string consisting of ds. Other bases have each "digit" represented as a (multi-digit) decimal number giving the ordinal of that digit. Each "digit" is signed for negative bases. "Digits" are separated by spaces. Given the definition of output base, the command Op always yields "10" (in a representation appropriate to the base); O1-p yields useful information about the output base.

DIAGNOSTICS

- | | |
|---------------------------------|--|
| <code>x is unimplemented</code> | Where <i>x</i> is an octal number. |
| <code>stack empty</code> | There are insufficient elements on the stack to do what was asked. |
| <code>Out of space</code> | The free list is exhausted (too many digits). |
| <code>Out of headers</code> | Too many numbers are being kept around. |
| <code>Out of pushdown</code> | Too many items are on the stack. |
| <code>Nesting Depth</code> | There are too many levels of nested execution. |

EXAMPLES

This example prints the first ten values of n! (n factorial):

```
[1a1+dsa*pla10>y]sy
0sa1
lyx
```

SEE ALSO

bc(1).

DC: *An Interactive Desk Calculator* tutorial in *Number Processing Users Guide*.

NAME

dd - convert, reblock, translate, and copy a (tape) file

SYNOPSIS

dd [*option=value*] ...

DESCRIPTION

dd copies the specified input file to the specified output file with possible conversions. The standard input and output are used by default. Input and output block size can be specified to take advantage of raw physical I/O. Upon completion, dd reports the number of whole and partial input and output records.

Options

dd recognizes the following *option=value* pairs:

| | |
|------------------------|---|
| if= <i>file</i> | Input file name; default is standard input. |
| of= <i>file</i> | Output file name; default is standard output. The output file is created using the same owner and group used by creat (). |
| ibs= <i>n</i> | Input block size is <i>n</i> bytes; default is 512. |
| obs= <i>n</i> | Output block size is <i>n</i> bytes; default is 512. |
| bs= <i>n</i> | Set both input and output block size to the same size, superseding ibs and obs . This option is particularly efficient if no conversion (conv option) is specified, because no in-core copy is necessary. |
| cbs= <i>n</i> | Conversion buffer size is <i>n</i> bytes. |
| skip= <i>n</i> | Skip <i>n</i> input blocks before starting copy. |
| iseek= <i>n</i> | Skip <i>n</i> input blocks before starting copy. (This is an alias for the skip option.) |
| seek= <i>n</i> | Skip <i>n</i> blocks from beginning of output file before copying. |
| oseek= <i>n</i> | Skip <i>n</i> blocks from beginning of output file before copying. (This is an alias for the seek option.) |
| count= <i>n</i> | Copy only <i>n</i> input blocks. |
| files= <i>n</i> | Copy and concatenate <i>n</i> input files. This option should be used only when the input file is a magnetic tape device. |

conv=*value* [, *value* ...]

Where *values* are comma-separated symbols from the following list.

| | |
|----------------|--|
| ascii | Convert EBCDIC to ASCII. |
| ebcdic | Convert ASCII to EBCDIC. |
| ibm | Convert ASCII to EBCDIC using an alternate conversion table. |
| | The ascii , ebcdic , and ibm values are mutually exclusive. |
| block | Convert each newline-terminated or end-of-file-terminated input record to a record with a fixed length specified by cbs . Any newline character is removed, and space characters are used to fill the block to size cbs . Lines that are longer than cbs are truncated; the number of truncated lines (records) is reported (see <i>DIAGNOSTICS</i> below). |
| | The block and unblock values are mutually exclusive. |
| unblock | Convert fixed-length input records to variable-length records. For each input record, cbs bytes are read, trailing space characters are deleted, and a newline character is appended. |
| lcase | Map upper-case input characters to the corresponding lower-case characters. |
| | The lcase and ucase values are mutually exclusive. |
| ucase | Map lower-case input characters to the corresponding upper-case characters. |

| | |
|----------------|--|
| swab | Swap every pair of input bytes. |
| noerror | Do not stop processing on an input error. If the sync conversion symbol is also specified, missing input is replaced with null bytes and processed normally; otherwise, the input block is omitted from the output. |
| notrunc | Do not truncate existing output file. Blocks in the output file not overwritten by this invocation of dd are preserved. |
| sync | Pad every input block to size ibs . If block or unblock is also specified, pad with space characters; otherwise, pad with null bytes. |

Where sizes are required, *n* indicates a numerical value in bytes. Numbers can be specified using the forms:

| | |
|-------------------|--|
| <i>n</i> | for <i>n</i> bytes |
| <i>n</i> k | for <i>n</i> Kbytes ($n \times 1024$), |
| <i>n</i> b | for <i>n</i> blocks ($n \times 512$), or |
| <i>n</i> w | for <i>n</i> words ($n \times 2$). |

To indicate a product, use **x** to separate number pairs.

The **cbs** option is used when **block**, **unblock**, **ascii** or **ebcdic** conversion is specified. In case of **ascii**, *cbs* characters are placed into the conversion buffer, converted to ASCII, trailing blanks are trimmed, and a newline is added before sending the line to the output. In case of **ebcdic**, ASCII characters are read into the conversion buffer, converted to EBCDIC, and blanks are added to make up an output block of size *cbs*.

EXTERNAL INFLUENCES

International Code Set Support

Single- and multi-byte character code sets are supported.

Environment Variables

The following environment variables affect execution of **dd**:

LANG determines the locale when **LC_ALL** and a corresponding variable (beginning with **LC_**) do not specify a locale.

LC_ALL determines the locale used to override any values set by **LANG** or any environment variables beginning with **LC_**.

The **LC_CTYPE** variable determines the locale for the interpretation of sequences of bytes of text data as characters (single-byte/multi-byte characters, upper-case/lower-case characters).

The **LC_MESSAGES** variable determines the language in which messages are written.

RETURN VALUE

Exit values are:

| | |
|----|---------------------------|
| 0 | Successful completion. |
| >0 | Error condition occurred. |

DIAGNOSTICS

Upon completion, **dd** reports the number of input and output records:

| | |
|---------------------------------|--|
| <i>f</i> + <i>p</i> records in | Number of full and partial blocks read. |
| <i>f</i> + <i>p</i> records out | Number of full and partial blocks written. |

When **conv=block** is specified and there is at least one truncated block, the number of truncated records is also reported:

n truncated records

EXAMPLES

Read an EBCDIC tape blocked ten 80-byte EBCDIC card images per block into an ASCII file named **x**:

```
dd if=/dev/rmt/c0t0d0BEST of=x ibs=800 cbs=80 conv=ascii,lcase
```

Note the use of the raw magnetic tape device file. **dd** is especially suited to I/O on raw physical devices because it allows reading and writing in arbitrary block sizes.

WARNINGS

Some devices, such as 1/2-inch magnetic tapes, are incapable of seeking. Such devices may be positioned prior to running **dd** by using *mt(1)* or some other appropriate command. The **skip**, **seek**, **iseek** and **oseek** options do work for such devices. However, skipping blocks using these options is slow on devices that cannot seek, since the blocks must actually be read to get to the desired position on the tape.

ASCII and EBCDIC conversion tables are taken from the 256-character ACM standard, Nov. 1968. The **ibm** conversion, while less widely accepted as a standard, corresponds better to certain IBM print train conventions. There is no universal solution.

Newline characters are inserted only on conversion to ASCII; padding is done only on conversion to EBCDIC. These should be separate options.

If **if** or **of** refers to a raw disk, **bs** should always be a multiple of the sector size of the disk. By default, **bs** is 512 bytes. If the sector size of the disk is different from 512 bytes, **bs** should be specified using a multiple of sector size. The character special (raw) device file should always be used for devices.

It is entirely up to the user to insure there is enough room in the destination file, file system and/or device to contain the output since **dd** cannot pre-determine the required space after conversion.

SEE ALSO

cp(1), *mt(1)*, *tr(1)*, *disk(7)*, *mt(7)*.

STANDARDS CONFORMANCE

dd: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

delta - make a delta (change) to an SCCS file

SYNOPSIS

delta [-r *SID*] [-s] [-n] [-g *list*] [-m *mrlist*] [-y *comment*] [-p] *files*

DESCRIPTION

The **delta** command is used to permanently introduce into the named SCCS file changes that were made to the file retrieved by **get** (called the *g-file*, or generated file). See *get(1)*.

delta makes a delta to each named SCCS file. If a directory is named, **delta** behaves as though each file in the directory was specified as a named file, except that non-SCCS files (last component of the path name does not begin with **.s**) and unreadable files are silently ignored. If a name of - is given, the standard input is read (see WARNINGS). Each line of the standard input is taken to be the name of an SCCS file to be processed.

delta may issue prompts on the standard output, depending upon certain options specified and flags (see *admin(1)*) that may be present in the SCCS file (see the **-m** and **-y** options below).

Options

Option arguments apply independently to each named file.

- r***SID* Uniquely identifies which delta is to be made to the SCCS file. Use of this option is necessary only if two or more outstanding **gets** for editing (**get -e**) on the same SCCS file were done by the same person (login name). The *SID* value specified with the **-r** option can be either the *SID* specified on the **get** command line or the *SID* to be made as reported by the **get** command (see *get(1)*). A diagnostic results if the specified *SID* is ambiguous, or, if necessary and omitted on the command line.
- s** Suppresses issuing, on the standard output, of the created delta's *SID* as well as the number of lines inserted, deleted and unchanged in the SCCS file.
- n** Specifies retention of the edited *g-file* (normally removed at completion of delta processing).
- g***list* Specifies a *list* (see *get(1)* for the definition of *list*) of deltas which are to be *ignored* when the file is accessed at the change level (*SID*) created by this delta.
- m**[*mrlist*] If the SCCS file has the **v** flag set (see *admin(1)*), a Modification Request (MR) number *must* be supplied as the reason for creating the new delta.

If **-m** is not used and the standard input is a terminal, the prompt **MRs?** is issued on the standard output before the standard input is read. If the standard input is not a terminal, no prompt is issued. The **MRs?** prompt always precedes the **comments?** prompt (see the **-y** option).

MRs in a list are separated by blanks and/or tab characters. An unescaped new-line character terminates the MR list.

Note that if the **v** flag has a value (see *admin(1)*), it is assumed to be the name of a program (or shell procedure) that is to validate the correctness of the MR numbers. If a non-zero exit status is returned from the MR number-validation program, **delta** assumes that the MR numbers were not all valid and terminates.
- y**[*comment*] Arbitrary text used to describe the reason for making the delta. A null string is considered a valid *comment*.

If **-y** is not specified and the standard input is a terminal, the prompt **comments?** is issued on the standard output before the standard input is read. If the standard input is not a terminal, no prompt is issued. An unescaped new-line character terminates the comment text.
- p** Causes **delta** to print (on the standard output in a *diff(1)* format) the SCCS file differences before and after the delta is applied.

EXTERNAL INFLUENCES**Environment Variables**

LC_CTYPE determines the interpretation of text as single- and/or multi-byte characters.

LC_MESSAGES determines the language in which messages are displayed.

If **LC_CTYPE** or **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **delta** behaves as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single- and multi-byte character code sets are supported.

DIAGNOSTICS

Use *scsghelp(1)* for explanations.

WARNINGS

SCCS files can be any length, but the number of lines in the text file itself cannot exceed 99 999 lines.

Lines beginning with an ASCII SOH character (octal 001) cannot be placed in the SCCS file unless the SOH is escaped. This character has special meaning to SCCS (see *scsfile(4)*) and will cause an error.

A **get** of many SCCS files, followed by a **delta** of those files, should be avoided when the **get** generates a large amount of data. Instead, multiple **get/delta** sequences should be used.

If the standard input (-) is specified on the **delta** command line, the **-m** (if necessary) and **-y** options *must* also be present. Omission of these options causes an error.

Comments can be of multiple lines. The maximum length of the comment (total length of all comment lines) cannot exceed 1024 bytes. No line in a comment should have a length of more than 1000 bytes.

FILES

All of the auxiliary files listed below, except for the *g-file*, are created in the same directory as the *s-file* (see *get(1)*). The *g-file* is created in the user's working directory.

| | |
|-----------------------|---|
| <i>g-file</i> | Existed before the execution of delta ; removed after completion of delta (unless -n was specified). |
| <i>p-file</i> | Existed before the execution of delta ; may exist after completion of delta . |
| <i>q-file</i> | Created during the execution of delta ; removed after completion of delta . |
| <i>x-file</i> | Created during the execution of delta ; renamed to SCCS file after completion of delta . |
| <i>z-file</i> | Created during the execution of delta ; removed during the execution of delta . |
| <i>d-file</i> | Created during the execution of delta ; removed after completion of delta . |
| /usr/bin/bdiff | Program to compute differences between the file retrieved by get and the <i>g-file</i> . |

SEE ALSO

admin(1), *bdiff(1)*, *cdc(1)*, *get(1)*, *scsghelp(1)*, *prs(1)*, *rmidel(1)*, *scsfile(4)*.

STANDARDS CONFORMANCE

delta: SVID2, SVID3, XPG2, XPG3, XPG4

NAME

deroff - remove nroff, tbl, and neqn constructs

SYNOPSIS

deroff [-mx] [-w] [-i] [*file* ...]

DESCRIPTION

deroff reads each *file* in sequence and removes all **nroff** requests, macro calls, backslash constructs, **neqn** constructs (between **.EQ** and **.EN** lines, and between delimiters — see *neqn(1)*), and **tbl** descriptions (see *tbl(1)*), 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 nroff/troff** formatter 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.

The **-m** option can be followed by an **m**, **s**, or **l**. The **-mm** option causes the macros be interpreted such 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.

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 multi-byte character string or any string that contains at least two letters and is composed of letters, digits, ampersands (&), and apostrophes ('); In a macro call, however, a “word” is a multi-byte character string or a string that begins 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.”

If the **-i** option is specified, **deroff** ignores the **.so** and **.nx nroff/troff** commands.

EXTERNAL INFLUENCES**Environment Variables**

LC_CTYPE determines the interpretation of text and filenames as single and/or multi-byte characters. Note that multi-byte punctuation characters are not recognized when using the **-w** option.

LC_MESSAGES determines the language in which messages are displayed.

If **LC_CTYPE** or **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of “C” (see *lang(5)*) is used instead of **LANG**.

If any internationalization variable contains an invalid setting, **deroff** behaves as if all internationalization variables are set to “C”. See *environ(5)*.

International Code Set Support

Single- and multi-byte character code sets are supported.

WARNINGS

deroff is not a complete **nroff** interpreter; thus 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.

AUTHOR

deroff was developed by the University of California, Berkeley.

SEE ALSO

neqn(1), *nroff(1)*, *tbl(1)*.

NAME

diff - differential file and directory comparator

SYNOPSIS

diff [-C *n*] [-S *name*] [-lrs] [-bcefhintw] *dir1 dir2*

diff [-C *n*] [-S *name*] [-bcefhintw] *file1 file2*

diff [-D *string*] [-biw] *file1 file2*

DESCRIPTION**Comparing Directories**

If both arguments are directories, **diff** sorts the contents of the directories by name, then runs the regular file **diff** algorithm (described below) on text files that have the same name in each directory but are different. Binary files that differ, common subdirectories, and files that appear in only one directory are listed. When comparing directories, the following options are recognized:

- l Long output format; each text file **diff** is piped through **pr** to paginate it (see *pr*(1)). Other differences are remembered and summarized after all text file differences are reported.
- r Applies **diff** recursively to common subdirectories encountered.
- s **diff** reports files that are identical but otherwise not mentioned.
- S *name* Starts a directory **diff** in the middle of the sorted directory, beginning with file *name*.

Comparing Files

When run on regular files, and when comparing text files that differ during directory comparison, **diff** tells what lines must be changed in the files to bring them into agreement. **diff** usually finds a smallest sufficient set of file differences. However, it can be misled by lines containing very few characters or by other situations. If neither *file1* nor *file2* is a directory, either can be specified as -, in which case the standard input is used. If *file1* is a directory, a file in that directory whose filename is the same as the filename of *file2* is used (and vice versa).

There are several options for output format. The default output format contains lines resembling the following:

```
n1 a n3,n4
n1,n2 d n3
n1,n2 c n3,n4
```

These lines resemble **ed** commands to convert *file1* into *file2*. The numbers after the letters pertain to *file2*. In fact, by exchanging **a** for **d** and reading backwards one may ascertain equally how to convert *file2* into *file1*. As in **ed**, identical pairs where *n1* = *n2* or *n3* = *n4* are abbreviated as a single number.

Following each of these lines come all the lines that are affected in the first file flagged by <, then all the lines that are affected in the second file flagged by >.

Except for **-b**, **-w**, **-i**, or **-t** which can be given with any of the others, the following options are mutually exclusive:

- e Produce a script of **a**, **c**, and **d** commands for the **ed** editor suitable for recreating *file2* from *file1*. Extra commands are added to the output when comparing directories with **-e**, so that the result is a shell script for converting text files common to the two directories from their state in *dir1* to their state in *dir2* (see *sh-bourne*(1)).
- f Produce a script similar to that of the **-e** option that is not useful with **ed** but is more readable by humans.
- n Produce a script similar to that of **-e**, but in the opposite order, and with a count of changed lines on each insert or delete command. This is the form used by **rcsdiff** (see *rcsdiff*(1)).
- c Produce a difference list with 3 lines of context. **-c** modifies the output format slightly: the output begins with identification of the files involved, followed by their creation dates, then each change separated by a line containing about twelve asterisks (*)s. Lines removed from *file1* are marked with -, and lines added to *file2* are marked +. Lines that change from one file to the other are marked in both files with with !. Changes that lie within 3 lines of each other in the file are grouped together on output.

- C *n* Output format similar to -c but with *n* lines of context.
- h Do a fast, half-hearted job. This option works only when changed stretches are short and well separated, but can be used on files of unlimited length.
- D *string* Create a merged version of *file1* and *file2* on the standard output, with C preprocessor controls included so that a compilation of the result without defining *string* is equivalent to compiling *file1*, while compiling the result with *string* defined is equivalent to compiling *file2*.
- b Ignore trailing blanks (spaces and tabs) and treat other strings of blanks as equal.
- w Ignore all whitespace (blanks and tabs). For example, `if (a == b)` and `if(a==b)` are treated as equal.
- i Ignores uppercase/lowercase differences. Thus `A` is treated the same as `a`.
- t Expand tabs in output lines. Normal or -c output adds one or more characters to the front of each line. Resulting misalignment of indentation in the original source lines can make the output listing difficult to interpret. This option preserves original source file indentation.

EXTERNAL INFLUENCES

Environment Variables

LANG determines the locale to use for the locale categories when both **LC_ALL** and the corresponding environment variable (beginning with **LC_**) do not specify a locale. If **LANG** is not set or is set to the empty string, a default of "C" (see *lang(5)*) is used.

LC_CTYPE determines the space characters for the **diff** command, and the interpretation of text within file as single- and/or multi-byte characters.

LC_MESSAGES determines the language in which messages are displayed.

If any internationalization variable contains an invalid setting, **diff** and **diffh** behave as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single- and multi-byte character code sets are supported with the exception that **diff** and **diffh** do not recognize multi-byte alternative space characters.

RETURN VALUE

Upon completion, **diff** returns with one of the following exit values:

- 0** No differences were found.
- 1** Differences were found.
- >1** An error occurred.

EXAMPLES

The following command creates a script file **script**:

```
diff -e x1 x2 >script
```

w is added to the end of the script in order to save the file:

```
echo w >> script
```

The script file can then be used to create the file **x2** from the file **x1** using the editor **ed** in the following manner:

```
ed x1 < script
```

The following command produces the difference output with 2 lines of context information before and after the line that was different:

```
diff -C2 x1 x2
```

The following command ignores all blanks and tabs and ignores uppercase-lowercase differences.

```
diff -wi x1 x2
```

WARNINGS

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

When comparing directories with the **-b**, **-w**, or **-i** options specified, **diff** first compares the files in the same manner as **cmp**, then runs the **diff** algorithm if they are not equal. This may cause a small amount of spurious output if the files are identical except for insignificant blank strings or uppercase/lowercase differences.

The default algorithm requires memory allocation of roughly six times the size of the file. If sufficient memory is not available for handling large files, the **-h** option or **bdiff** can be used (see *bdiff(1)*).

With other options if sufficient memory is not available, then either the **swap** or **maxdsiz** values can be increased.

When run on directories with the **-r** option, **diff** recursively descends sub-trees. When comparing deep multi-level directories, more memory may be required than is currently available on the system. The amount of memory required depends on the depth of recursion and the size of the files.

AUTHOR

diff was developed by AT&T, the University of California, Berkeley, and HP.

FILES

/usr/sbin/diffh used by **-h** option

SEE ALSO

bdiff(1), *cmp(1)*, *comm(1)*, *diff3(1)*, *diffmk(1)*, *dircmp(1)*, *ed(1)*, *more(1)*, *nroff(1)*, *rcsdiff(1)*, *sccsdiff(1)*, *sdiff(1)*, *terminfo(4)*.

STANDARDS CONFORMANCE

diff: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

diff3 - 3-way differential file comparison

SYNOPSIS

diff3 [-exEX3] *file1 file2 file3*

DESCRIPTION

diff3 compares three versions of a file, and prints disagreeing ranges of text flagged with these codes:

```
====    all three files differ
====1   file1 is different
====2   file2 is different
====3   file3 is different
```

The type of change required to convert a given range of a given file to some other is indicated in one of these ways:

```
f:n1a      Text is to be appended after line number n1 in file f, where f = 1, 2, or 3.
f:n1,n2c    Text is to be changed in the range line n1 through line n2. If n1 = n2, the range can
               be abbreviated to n1.
```

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.

- e** **diff3** Produces a script for the **ed** editor that can be used to incorporate into *file1* all changes between *file2* and *file3* (see **ed(1)**); i.e., the changes that normally would be flagged **====** and **====3**.
- x** Produces a script to incorporate only changes flagged **====**
- 3** Produces a script to incorporate only changes flagged **====3**
- E** Produces a script that will incorporate all changes between *file2* and *file3*, but treat overlapping changes (that is, changes that would be flagged with **====** in normal listing) differently. The overlapping lines in both files will be inserted by the edit script bracketed by <<<<<< and >>>>>> lines.
- X** Produces a script that will incorporate only changes flagged **====** , but treat these changes in the manner of **-E** option.

The following command applies the resulting script to *file1*.

```
(cat script; echo '1,$p') | ed - file1
```

EXTERNAL INFLUENCES**International Code Set Support**

Single- and multi-byte character code sets are supported.

WARNINGS

Text lines that consist of a single period (.) defeat **-e**.

Files longer than 64K bytes do not work.

FILES

```
/var/tmp/d3*
/usr/sbin/diff3prog
```

SEE ALSO

diff(1).

NAME

diffmk - mark changes between two different versions of a file

SYNOPSIS

diffmk *prevfile currfile markfile*

DESCRIPTION

diffmk compares the previous version of a file with the current version and creates a file that includes **nroff/troff** "change mark" commands. *prevfile* is the name of the previous version of the file and *currfile* is the name of the current version of the file. **diffmk** generates *markfile* which contains all the lines of the *currfile* plus inserted formatter "change mark" (**.mc**) requests. When *markfile* is formatted, changed or inserted text is shown by a | character at the right margin of each line. The position of deleted text is shown by a single *.

If the characters | and * are inappropriate, a copy of **diffmk** can be edited to change them because **diffmk** is a shell script.

EXTERNAL INFLUENCES**International Code Set Support**

Single- and multi-byte character code sets are supported.

EXAMPLES

A typical command line for comparing two versions of an **nroff/troff** file and generating a file with the changes marked is:

```
diffmk prevfile currfile markfile; nroff markfile | pr
```

diffmk can also be used to produce listings of C (or other) programs with changes marked. A typical command line for such use is:

```
diffmk prevfile.c currfile.c markfile.c; nroff mcs markfile.c | pr
```

where the file **mcs** contains:

```
.pl 1
.ll 77
.nf
.eo
```

The **.ll** request can specify a different line length, depending on the nature of the program being printed. The **.eo** request is probably needed only for C programs.

WARNINGS

Aesthetic considerations may dictate manual adjustment of some output.

diffmk does not differentiate between changes in text and changes in formatter request coding. Thus, file differences involving only formatting changes (such as replacing **.sp** with **.sp 2** in a text source file) with no change in actual text can produce change marks.

Although unlikely, certain combinations of formatting requests can cause change marks to either disappear or to mark too much. Manual intervention may be required because the subtleties of various formatting macro packages and preprocessors is beyond the scope of **diffmk**. **tbl** cannot tolerate **.mc** commands in its input (see *tbl(1)*), so any **.mc** request that would appear inside a **.TS** range is silently deleted. The script can be changed if this action is inappropriate, or **diffmk** can be run on two files that have both been run through the **tbl** preprocessor before any comparisons are made.

diffmk uses **diff**, and thus has the same limitations on file size and performance that **diff** may impose (see *diff(1)*). In particular the performance is nonlinear with the size of the file, and very large files (well over 1000 lines) may take extremely long to process. Breaking the file into smaller pieces may be advisable.

diffmk also uses the *ed(1)* editor. If the file is too large for **ed**, **ed** error messages may be embedded in the file. Again, breaking the file into smaller pieces may be advisable.

SEE ALSO

diff(1), *nroff(1)*.

NAME

dircmp - directory comparison

SYNOPSIS

dircmp [-d] [-s] [-wn] *dir1 dir2*

DESCRIPTION

dircmp examines *dir1* and *dir2* and generates various tabulated information about the contents of the directories. Sorted listings of files that are unique to each directory are generated for all the options. If no option is entered, a sorted list is output indicating whether the filenames common to both directories have the same contents.

- 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 *n* characters. The default width is 72.

EXTERNAL INFLUENCES**Environment Variables**

LC_COLLATE determines the order in which the output is sorted.

If LC_COLLATE is not specified in the environment or is set to the empty string, the value of LANG is used as a default. If LANG is not specified or is set to the empty string, a default of "C" (see *lang*(5)) is used instead of LANG. If any internationalization variable contains an invalid setting, **dircmp** behaves as if all internationalization variables are set to "C" (see *environ*(5)).

International Code Set Support

Single- and multi-byte character code sets are supported.

EXAMPLES

Compare the two directories **slate** and **sleet** and produce a list of changes that would make the directories identical:

```
dircmp -d slate sleet
```

WARNINGS

This command is likely to be withdrawn from X/Open standards. Applications using this command might not be portable to other vendors' systems. As an alternative **diff -R** is recommended.

SEE ALSO

cmp(1), **diff**(1).

STANDARDS CONFORMANCE

dircmp: SVID2, SVID3, XPG2, XPG3

NAME

dmpxlt - dump iconv translation tables to a readable format

SYNOPSIS

`/usr/bin/dmpxlt [-f output_filename] [input_filename]`

DESCRIPTION

dmpxlt dumps the compiled version of the **iconv** codeset conversion tables into an ASCII-readable format that can be modified and used as input to **genxlt(1)** to regenerate the table for **iconv(1)**.

Options

dmpxlt recognizes the following options:

-f *output_filename* If this option is not selected, the data will be sent to standard output.

dmpxlt will create an output file in the prescribed format, giving the filecode mapping between the two code sets, which can be edited and reused by **genxlt(1)** to create new tables for **iconv(1)**. The entries are in hexadecimal.

EXTERNAL INFLUENCES**Environment Variables**

LANG provides a default value for the internationalization variables that are unset or null. If **LANG** is unset or null, the default value of "C" (see **lang(5)**) is used. If any of the internationalization variables contains an invalid setting, **dmpxlt** will behave as if all internationalization variables are set to "C". See **environ(5)**.

LC_ALL If set to a non-empty string value, overrides the values of all the other internationalization variables.

LC_MESSAGES determines the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.

NLSPATH determines the location of message catalogues for the processing of **LC_MESSAGES**.

International Code Set Support

Single and multi-byte character code sets are supported.

RETURN VALUE

The following are exit values:

0 Successful completion.
>0 Error condition occurred.

EXAMPLES

This example creates the source file *genxlt_input* from the table **roma8=iso81**:

```
dmpxlt -f genxlt_input /usr/lib/nls/iconv/tables/roma8=iso81
```

FILES

/usr/lib/nls/iconv/tables All tables must be installed in this directory.

SEE ALSO

iconv(1), **genxlt(1)**, **iconv(3C)**, **environ(5)**, **lang(5)**.

NAME

domainname - set or display name of Network Information Service domain

SYNOPSIS

domainname [*name_of_domain*]

DESCRIPTION

Network Information Service (NIS) uses domain names to refer collectively to a group of hosts. Without an argument, *domainname* displays the name of the NIS domain. Only superuser can set the domain name by providing *name_of_domain*. The domain name is usually set in the configuration file `/etc/rc.config.d/namesvrs`, by setting the `NIS_DOMAIN` variable.

DEPENDENCIES

NIS servers use the NIS domain name as the name of a subdirectory of `/var/yp`. Since the NIS domain name can be as long as 64 characters, *name_of_domain* may exceed the maximum file name length allowed on a given file system. If that length is exceeded, the subdirectory name becomes a truncated version of the NIS domain name.

The first 14 characters of all NIS domains on the network must be unique: truncated names should be checked to verify that they meet this requirement.

AUTHOR

domainname was developed by Sun Microsystems, Inc.

SEE ALSO

ypinit(1M), getdomainname(2), setdomainname(2).

NAME

dos2ux, ux2dos - convert ASCII file format

SYNOPSIS

dos2ux *file...*

ux2dos *file...*

DESCRIPTION

dos2ux and **ux2dos** read each specified *file* in sequence and write it to standard output, converting to HP-UX format or to DOS format, respectively. Each *file* can be either DOS format or HP-UX format for either command.

A DOS file name is recognized by the presence of an embedded colon (:) delimiter; see *dosif(4)* for DOS file naming conventions.

If no input file is given or if the argument **-** is encountered, **dos2ux** and **ux2dos** read from standard input. Standard input can be combined with other files.

EXAMPLES

Print file **myfile** on the display:

```
dos2ux myfile
```

Convert **file1** and **file2** to DOS format then concatenate them together, placing them in **file3**.

```
ux2dos file1 file2 > file3
```

RETURN VALUE

dos2ux and **ux2dos** return **0** if successful or **2** if the command failed. The only possible failure is the inability to open a specified file, in which case the commands print a warning.

WARNINGS

Command formats resembling:

```
dos2ux file1 file2 > file1
```

overwrite the data in **file1** before the concatenation begins, causing a loss of the contents of **file1**. Therefore, be careful when using shell special characters.

SEE ALSO

doschmod(1), *doscp(1)*, *dosdf(1)*, *dosls(1)*, *dosmkdir(1)*, *dosrm(1)*, *dosif(4)*.

NAME

doschmod - change attributes of a DOS file

SYNOPSIS

doschmod [-u] *mode device: file ...*

DESCRIPTION

doschmod is the DOS counterpart of **chmod** (see *chmod(1)*).

Options

doschmod recognizes one option:

- u** Disable argument case conversion. In the absence of this option, all DOS file names are converted to uppercase.

A DOS file name is recognized by the presence of an embedded colon (:) delimiter; see *dosif(4)* for DOS file naming conventions.

Metacharacters *, ?, and [...] can be used when specifying DOS file names. These must be quoted when specifying a DOS file name, because file name expansion must be performed by the DOS utilities, not by the shell. DOS utilities expand file names as described in *regex(5)* under *PATTERN MATCHING NOTATION*.

The attributes of each named file are changed according to *mode*, which is an octal number in the range 000 to 0377. *mode* is constructed from the logical OR of the following modes:

| | |
|-----|---|
| 200 | Reserved. Do not use. |
| 100 | Reserved. Do not use. |
| 040 | Archive. Set whenever the file has been written to and closed. |
| 020 | Directory. Do not modify. |
| 010 | Volume Label. Do not modify. |
| 004 | System file. Marks files that are part of the DOS operating system. |
| 002 | Hidden file. Marks files that do not appear in a DOS directory listing using the DOS DIR command. |
| 001 | Read-Only file. Marks files as read-only. |

WARNINGS

Specifying inappropriate *mode* values can make files and/or directories inaccessible, and in certain cases can damage the file system. To prevent such problems, do not change the mode of directories and volume labels.

Normal users should have no need to use *mode* bits other than 001, 002, and 040.

EXAMPLES

Mark file `/dev/rfd9122:memo.txt` as a hidden file:

```
doschmod 002 /dev/rfd9122:memo.txt
```

Mark file `driveC:autoexec.bat` read-only:

```
doschmod 001 driveC:autoexec.bat
```

SEE ALSO

chmod(1), *dos2ux(1)*, *doscpc(1)*, *dosdf(1)*, *dosls(1)*, *dosmkdir(1)*, *dosrm(1)*, *chmod(2)*, *dosif(4)*.

NAME

doscp - copy to or from DOS files

SYNOPSIS

doscp [-fvu] *file1 file2*

doscp [-fvu] *file1 [file2 ...] directory*

DESCRIPTION

doscp is the DOS counterpart of **cp** (see *cp(1)*). **doscp** copies a DOS file to a DOS or HP-UX file, an HP-UX file to an HP-UX or DOS file, or HP-UX or DOS files to an HP-UX or DOS directory. The last name in the argument list is the destination file or directory.

A DOS file name is recognized by the presence of an embedded colon (:) delimiter; see *dosif(4)* for DOS file naming conventions.

Metacharacters *, ?, and [...] can be used when specifying both HP-UX and DOS file names. These must be quoted when specifying a DOS file name, because file name expansion must be performed by the DOS utilities, not by the shell. DOS utilities expand file names as described in *regex(5)* under *PATTERN MATCHING NOTATION*.

The file name - (dash) is interpreted to mean standard input or standard output depending upon its position in the argument list.

Options

doscp recognizes the following options:

- f Unconditionally write over an existing file. In the absence of this option, **doscp** asks permission to overwrite an existing HP-UX file.
- v Verbose mode. **doscp** prints the source name.
- u Disable argument case conversion. In the absence of this option, all DOS file names are converted to upper case.

RETURN VALUE

doscp returns 0 if all files are copied successfully. Otherwise, it prints a message to standard error and returns with a non-zero value.

EXAMPLES

Copy the files in the HP-UX directory **abc** to the DOS volume stored as HP-UX file **hard_disk**:

```
doscp abc/* hard_disk:
```

Copy DOS file **/backup/log** through the HP-UX special file **/dev/rfd9127** to HP-UX file **logcopy** located in the current directory:

```
doscp /dev/rfd9127:/backup/log logcopy
```

Copy DOS file **zulu** on the volume stored as HP-UX file **bb** to standard output:

```
doscp bb:zulu -
```

Copy all files in directory **/dameron** with extension **txt** in the DOS volume **/dev/rdsk/c1t2d0** to the HP-UX directory **abacus** located in the current directory:

```
doscp '/dev/rdsk/c1t2d0:/dameron/*.txt' abacus
```

WARNINGS

doscp works more reliably if you use a raw device special file (**/dev/rdsk/**) than a block device special file.

To use SCSI floppy disk devices, the **sflop** device driver must be configured into the kernel. (You can use the **ioscan** command to verify the configuration.)

SEE ALSO

cp(1), **dos2ux(1)**, **doschmod(1)**, **dosdf(1)**, **dosls(1)**, **dosmkdir(1)**, **dosrm(1)**, **ioscan(1M)** **dosif(4)**.

NAME

dosdf - report number of free disk clusters

SYNOPSIS

dosdf *device*[:]

DESCRIPTION

dosdf is the DOS counterpart of the **df** command (see *df(1)*). It prints the cluster size in bytes and the number of free clusters on the specified DOS volume.

SEE ALSO

df(1), dos2ux(1), doschmod(1), doscp(1), dosls(1), dosmkdir(1), dosrm(1), dosif(4).


d

NAME

dosls, dosll - list contents of DOS directories

SYNOPSIS

dosls [-aAudl] *device*:[*file*] ...

dosll [-aAudl] *device*:[*file*] ...

DESCRIPTION

dosls is the DOS counterpart of **ls** (see *ls(1)*).

For each directory named, **dosls** lists the contents of that directory. For each file named, **dosls** repeats its name and any other information requested. If invoked by the name **dosll**, the **-l** (ell) option is implied.

Options

dosls and **dosll** recognizes the following options:

- a** List all directory entries. In the absence of this option, hidden files, system files, and files whose names begin with a dot (.) are not listed.
- A** Same as **-a**, except the current directory and the parent directory are not listed. For the superuser, this option defaults to being set, and is disabled by **-A**.
- u** Disable argument case conversion. In the absence of this option, all DOS file names are converted to uppercase.
- d** If an argument is a directory, list only its name. Often used with **-l** to get the status of a directory.
- l** List in long format, giving file attribute, size in bytes, and the date and time of last modification for each file, as well as listing the DOS volume label. Long listing is disabled if this option is used with the **dosll** command.

A DOS file name is recognized by the presence of an embedded colon (:) delimiter; see *dosif(4)* for DOS file naming conventions.

Metacharacters *, ?, and [...] can be used when specifying DOS file names. These must be quoted when specifying a DOS file name, because file name expansion must be performed by the DOS utilities, not by the shell. DOS utilities expand file names as described in *regexp(5)* under *PATTERN MATCHING NOTATION*.

EXAMPLES

These examples assume that a DOS directory structure exists on the device accessed through HP-UX special file **/dev/rdisk/c2t1d0**.

The following example lists all of the files in the root directory of the DOS directory structure:

```
dosls -a /dev/rdisk/c2t1d0:
```

The following example lists all of the files with extension **bat** in the root directory of the DOS directory structure:

```
dosls -a '/dev/rdisk/c2t1d0:*.bat'
```

The following example produces a long-format listing of all the information about the DOS directory **/dos/math**, but does not list the files in the directory:

```
dosls -ld /dev/rdisk/c2t1d0:/dos/math
```

SEE ALSO

dos2ux(1), **doschmod(1)**, **dosecp(1)**, **dosdf(1)**, **dosmkdir(1)**, **dosrm(1)**, **ls(1)**, **dosif(4)**.

NAME

dosmkdir - make a DOS directory

SYNOPSIS

dosmkdir [-u] *device:directory* ...

DESCRIPTION

dosmkdir is the DOS counterpart of the **mkdir** command (see *mkdir(1)*). It creates specified directories. The standard entries, **.** for the directory itself and **..** for its parent, are made automatically.

There is one option:

-u Disable argument case conversion. In the absence of this option, all DOS file names are converted to uppercase.

A DOS file name is recognized by the presence of an embedded colon (**:**) delimiter; see *dosif(4)* for DOS file naming conventions.

DIAGNOSTICS

dosmkdir returns 0 if all directories were successfully created. Otherwise, it prints a message to standard error and returns non-zero.

EXAMPLES

Create an empty subdirectory named **numbers** under the directory **/math/lib** on the device accessed through HP-UX special file **/dev/rfd9122**:

```
dosmkdir /dev/rfd9122:/math/lib/numbers
```

SEE ALSO

dos2ux(1), *doschmod(1)*, *doscp(1)*, *dosdf(1)*, *dosls(1)*, *dosrm(1)*, *mkdir(1)*, *dosif(4)*.

NAME

dosrm, dosrmdir - remove DOS files or directories

SYNOPSIS

dosrm [-friu] *device:file* ...

dosrmdir [-u] *device:file* ...

DESCRIPTION

dosrm and **dosrmdir** are DOS counterparts of **rm** and **rmdir** (see *rm(1)* and *rmdir(1)*, respectively).

dosrm removes the entries for one or more files from a directory. If a specified file is a directory, an error message is printed unless the optional argument **-r** is specified (see below).

dosrmdir removes entries for the named directories, provided they are empty.

Options

dosrm and **dosrmdir** recognize the following options:

- f** (force) Unconditionally remove the specified file, even if the file is marked read-only.
- r** Cause **dosrm** to recursively delete the entire contents of a directory, followed by the directory itself. **dosrm** can recursively delete up to 17 levels of directories.
- i** (interactive) Cause **dosrm** to ask whether or not to delete each file. If **-r** is also specified, **dosrm** asks whether to examine each directory encountered.
- u** Disable argument case conversion. In the absence of this option, all DOS file names are converted to uppercase.

A DOS file name is recognized by the presence of an embedded colon (:) delimiter; see *dosif(4)* for DOS file naming conventions.

Metacharacters *, ?, and [...] can be used when specifying DOS file names. These must be quoted when specifying a DOS file name, because file name expansion must be performed by the DOS utilities, not by the shell. DOS utilities expand file names as described in *regexp(5)* under *PATTERN MATCHING NOTATION*.

EXAMPLES

These examples assume that a DOS directory structure exists on the device accessed through the HP-UX special file */dev/rfd9122*.

Recursively comb through the DOS directory */tmp* and ask if each DOS file should be removed forcibly (that is, with no file mode checks):

```
dosrm -irf /dev/rfd9122:/tmp
```

Remove the DOS directory *doug* from the DOS volume stored as HP-UX file *hard_disk*:

```
dosrmdir hard_disk:doug
```

SEE ALSO

dos2ux(1), *doschmod(1)*, *doscp(1)*, *dosdf(1)*, *dosls(1)*, *dosmkdir(1)*, *rm(1)*, *rmdir(1)*, *dosif(4)*.

NAME

du - summarize disk usage

SYNOPSIS

du [-a] [-s] [-bkrx] [-t *type*] [*name* ...]

DESCRIPTION

The **du** command gives the number of 512-byte blocks allocated for all files and (recursively) directories within each directory and file specified by the *name* operands. The block count includes the indirect blocks of the file. A file with two or more links is counted only once. If *name* is missing, the current working directory is used.

By default, **du** generates an entry only for the *name* operands and each directory contained within those hierarchies.

Options

The **du** command recognizes the following options:

- a Print entries for each file encountered in the directory hierarchies in addition to the normal output.
- b For each *name* operand that is a directory for which file system swap has been enabled, print the number of blocks the swap system is currently using.
- k Gives the block count in 1024-byte blocks.
- r Print messages about directories that cannot be read, files that cannot be accessed, etc. **du** is normally silent about such conditions.
- s Print only the grand total of disk usage for each of the specified *name* operands.
- x Restrict reporting to only those files that have the same device as the file specified by the *name* operand. Disk usage is normally reported for the entire directory hierarchy below each of the given *name* operands.
- t *type* Restrict reporting to file systems of the specified *type*. (Example values for *type* are **hfs**, **cdfs**, **nfs**, etc.) Multiple -t *type* options can be specified. Disk usage is normally reported for the entire directory hierarchy below each of the given *name* operands.

EXAMPLES

Display disk usage for the current working directory and all directories below it, generating error messages for unreadable directories:

```
du -r
```

Display disk usage for the entire file system except for any **cdfs** or **nfs** mounted file systems:

```
du -t hfs /
```

Display disk usage for files on the root volume (/) only. No usage statistics are collected for any other mounted file systems:

```
du -x /
```

WARNINGS

Block counts are incorrect for files that contain holes.

SEE ALSO

df(1M), bdf(1M), quot(1M).

STANDARDS CONFORMANCE

du: SVID2, SVID3, XPG2, XPG3, XPG4

NAME

echo - echo (print) arguments

SYNOPSIS

echo [*arg*] ...

DESCRIPTION

echo writes its arguments separated by blanks and terminated by a new-line on the standard output. It also understands C-like escape conventions; beware of conflicts with the shell's use of \:

| | |
|---------------|---|
| \a | write an alert character |
| \b | backspace |
| \c | print line without appending a new-line |
| \f | form-feed |
| \n | new-line |
| \r | carriage return |
| \t | tab |
| \v | vertical tab |
| \\ | backslash |
| \n | the 8-bit character whose ASCII code is the 1-, 2-, 3- or 4-digit octal number <i>n</i> , whose first character must be a zero. |
| \0 <i>num</i> | write an 8-bit value that is the zero-, one-, two- or three-digit octal number <i>num</i> |

echo is useful for producing diagnostics in command files and for sending known data into a pipe.

Notes

Berkeley **echo** differs from this implementation. The former does not implement the backslash escapes. However, the semantics of the \c escape can be obtained by using the -n option. The echo command implemented as a built-in function of **cs**h follows the Berkeley semantics (see **cs**h(1)).

EXTERNAL INFLUENCES**Environment Variables**

LC_CTYPE determines the interpretation of *arg* as single and/or multi-byte characters.

If **LC_CTYPE** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang*(5)) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **echo** behaves as if all internationalization variables are set to "C". See *environ*(5).

International Code Set Support

Single- and multi-byte character code sets are supported.

AUTHOR

echo was developed by OSF and HP.

SEE ALSO

sh(1).

BUGS

No characters are printed after the first \c. This is not normally a problem.

STANDARDS CONFORMANCE

echo: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

ed, red - line-oriented text editor

SYNOPSIS

ed [-p *string*] [-s | -] [-x] [*file*]

red [-p *string*] [-s | -] [-x] [*file*]

DESCRIPTION

The **ed** command executes a line-oriented text editor. It is most commonly used in scripts and noninteractive editing applications because, even though it can be used interactively, other editors such as **vi** and **ex** are typically easier to use in an interactive environment.

If *file* is specified, **ed** performs 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.

Options

The following options are recognized:

- p *string* Use *string* as the prompt string when in command mode. By default, there is no prompt string.
- s | - Suppress printing of byte counts by **e**, **E**, **r**, and **w** commands, and suppress the **!** prompt after a **!** *command*. The - option is obsolescent and will be removed in a future release.
- x Perform an **X** command first to handle an encrypted file.

File Handling

ed operates on a copy of the file it is editing; changes made to the copy have no effect on the original file until a **w** (write) command is given. The copy of the text being edited resides in a temporary file called the *buffer*. There is only one buffer.

red is a restricted version of **ed** that only allows editing of files in the current directory and prohibits executing shell commands via **!***shell-command*. Attempts to bypass these restrictions result in the error message **restricted shell**.

Both **ed** and **red** support the *fspec*(4) formatting capability. After including a format specification as the first line of *file* and invoking **ed** with the controlling terminal in **stty -tabs** or **stty tab3** mode (see *stty*(1)), the specified tab stops are automatically used when scanning *file*. For example, if the first line of a file contained

```
<:t5,10,15 s72:>
```

the tab stops would be set at columns 5, 10, and 15, and a maximum line length of 72 would be imposed.

Note: When you input text, **ed** expands tab characters as they are typed to every eighth column as a default.

Editor Commands Structure

Commands to **ed** have a simple and regular structure: zero, one, or two *addresses* followed by a single-character *command*, 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.

In general, only one command is allowed on a line. Append, change, and insert commands accept text input which is then placed in the buffer as appropriate. While **ed** is accepting text following an append, change, or insert command, it is said to be in *input mode*. While in input mode, *no* editor commands are recognized; all input is merely collected. To terminate input mode, type a period (.) alone at the beginning of a line.

Regular Expressions

ed supports the Basic Regular Expression (RE) syntax (see *regex*(5)), with the following additions:

- The null RE (for example, **/**) is equivalent to the last RE encountered.
- If the closing delimiter of an RE or of a replacement string (for example, **/**) would be the last character before a newline, that delimiter can be omitted, in which case the addressed line is printed. The following pairs of commands are equivalent:

| | | |
|------------------------|---------------------|-------------------|
| <code>s/s1/s2</code> | <code>g/s1</code> | <code>?s1</code> |
| <code>s/s1/s2/p</code> | <code>g/s1/p</code> | <code>?s1?</code> |

Line Addresses

To understand line addressing, remember that **ed** maintains a pointer to the *current line*. Generally speaking, the current line is the last line affected by a command. The exact effect of a given command on the current line is discussed under the description of each command. Addresses are interpreted according to the following rules:

1. The character `.` refers to the current line.
2. The character `$` refers to the last line of the buffer.
3. A decimal number *n* refers to the *n*th line of the buffer.
4. A `'x` refers to the line marked with the mark name character *x*, which must be a lower-case letter. Lines are marked with the **k** command described below.
5. An RE enclosed by slashes (*/RE/*) refers to the first line found by searching forward from the line following the current line toward the end of the buffer and stopping at the first line containing a string matching the RE. If necessary, the search wraps around to the beginning of the buffer and continues up to and including the current line, so that the entire buffer is searched. (Also see WARNINGS below.)
6. An RE enclosed by question marks (*?RE?*) addresses the first line found by searching backward from the line preceding the current line toward the beginning of the buffer and stopping at the first line containing a string matching the RE. If necessary, the search wraps around to the end of the buffer and continues up to and including the current line. (Also see WARNINGS below.)
7. An address followed by a plus (+) or minus (-) sign followed by a decimal number specifies that address plus or minus the indicated number of lines. The plus sign can be omitted.
8. If an address begins with + or -, the addition or subtraction is calculated with respect to the current line. For example, `-5` is interpreted as `.-5`.
9. If an address ends with + or -, 1 is added to or subtracted from the address, respectively. As a consequence of this and rule 8 above, the address `-` refers to the line preceding the current line. (To maintain compatibility with earlier versions of the editor, the circumflex (^) and - characters are interpreted identically when encountered in addresses.) Moreover, multiple trailing + and - characters have a cumulative effect, so `--` refers to the second line preceding the current line.
10. For convenience, a comma (,) represents the address pair `1,$`, while a semicolon (;) represents the pair `.,$`.

Commands require zero, one, or two addresses. Commands that do not use addresses treat the presence of an address as an error. Commands that accept one or two addresses assume default addresses when the number of addresses specified is insufficient. If more addresses are specified than a given command requires, the last one or two are used as appropriate.

Addresses are usually separated from each other by a comma (,). They can also be separated by a semicolon (;), in which case the current line (.) is set to the first address, after which 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.

Editor Commands

In the following list of **ed** commands, the default addresses are shown in parentheses (parentheses are not part of the address and should not be placed in an actual command except for other purposes).

It is generally illegal for more than one command to appear on a line. However, any command (except **e**, **f**, **r**, or **w**) can be suffixed by **l**, **n**, or **p** in which case the current line is respectively either listed, numbered, or printed, as discussed below under the **l**, **n**, and **p** commands.

| | |
|--|--|
| <p>(.)a <i>text</i> .</p> | <p>The a (append) command reads <i>text</i> and appends it after the addressed line. Upon completion, the new current line is the last inserted line, or, if no text was added, at the addressed line. Address 0 is legal for this command, causing the appended <i>text</i> to be placed at the beginning of the buffer.</p> |
|--|--|

- (.,.)**c**
text
.
The **c** (change) command deletes the addressed lines then accepts input text to replace the deleted lines. Upon completion, the new current line is the last line in *text* or, if no text was provided, at the first line after the deleted line or lines.
- (.,.)**d**
The **d** (delete) command deletes the addressed lines from the buffer. Upon completion, the new current line is the first line following the deleted text, or the last line in the file if the deleted line or lines were at the end of the buffer.
- e** *file*
The **e** (edit) command deletes the entire contents of the buffer, then reads in the named *file*. Upon completion, the new current line is the last line in the buffer. If no file name is given, the remembered file name, if any, is used (see the **f** command). The number of characters read is displayed, and *file* is remembered for possible use as a default file name in subsequent **e**, **r**, or **w** commands.
- If the *file* name starts with **!**, the rest of the line is interpreted as a shell command whose standard output is to be read. Such a shell command is not remembered as the current file name.
- Also see DIAGNOSTICS below.
- E** *file*
The **E** (forced edit) command is identical to **e** except that no check is made to ensure that the current buffer has not been altered since the last **w** command.
- f** *file*
If *file* is specified, the **f** (file name) command changes the remembered file name to *file*. Otherwise, it prints the remembered file name.
- (1,\$)**g** / *RE* / *command-list*
The **g** (global) command first marks 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 multiple-line list except the last line must end with a backslash (****). **a**, **i**, and **c** commands and associated input are permitted. The **.** that normally terminates input mode can 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*. (Also see WARNINGS below.)
- (1,\$)**G** / *RE* /
The interactive **G** (Global) command first marks every line that matches the given *RE*. Then, for every such line, the line is printed, then the current line is changed to that line and one command (other than **a**, **c**, **i**, **g**, **G**, **v**, or **V**) can be input and executed. After executing that command, the next marked line is printed, and so on. A newline character acts as a null command, and an **&** causes the re-execution of the most recent command executed within the current invocation of **G**. Note that 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 **h** (help) command gives a short error message explaining the reason for the most recent **?** diagnostic.
- H**
The **H** (Help) command causes **ed** to enter a mode in which error messages are printed for all subsequent **?** diagnostics. It also explains the previous **?** if there was one. The **H** command alternately turns this mode on and off. Initially, it is off.
- (.)**i**
text
.
The **i** (insert) command inserts the given *text* before the addressed line. Upon completion, the current line is the last inserted line, or, if there were none, 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.
- (.,.+1)**j**
The **j** (join) command joins contiguous lines by removing the appropriate newline characters. If exactly one address is given, this command does nothing.
- (.)**k***x*
The **k** (mark) command marks the addressed line with the name *x*, which must be a lower-case letter. The address '*x*' then addresses this line. Upon completion, the new current line remains unchanged from before.
- (.,.)**l**
The **l** (list) command writes the addressed lines to standard output in a visually unambiguous form. Characters listed in the following table are written as the corresponding escape sequence. Nonprintable characters not in the table are written as a three-digit octal number (with a preceding backslash character) for each byte in the character (most significant byte first).

Long lines are folded with the point of folding indicated by writing a backslash character followed by a newline. The end of each line is marked with a \$. An **l** (ell) command can be appended to any command other than **e**, **E**, **f**, **q**, **Q**, **r**, **w**, or **!**. The current line number is set to the address of the last line written.

| Escape Sequence | Represents | ASCII Name | Escape Sequence | Represents | ASCII Name |
|-----------------|------------|------------|-----------------|-----------------|------------|
| \\ | backslash | \ | \r | carriage return | CR |
| \a | alert | BEL | \t | horizontal tab | HT |
| \b | backspace | BS | \v | vertical tab | VT |
| \f | formfeed | FF | | | |

(.,.)ma The **m** (move) command repositions the addressed lines after the line addressed by *a*. Address 0 is legal for *a*, causing the addressed lines to be moved to the beginning of the file. It is an error if address *a* falls within the range of moved lines; Upon completion, the new current line is the last line moved.

(.,.)n The **n** (number) command prints the addressed lines, preceding each line by its line number and a tab character. Upon completion, the new current line is the last line printed. The **n** command can be appended to any command other than **e**, **f**, **r**, or **w**.

(.,.)p The **p** (print) command prints the addressed lines. Upon completion, the new current line is the last line printed. The **p** command may be appended to any other command other than **e**, **E**, **f**, **q**, **Q**, **r**, **w**, or **!**. For example, **dp** deletes the current line and prints the new current line.

P The **P** (prompt) command causes **ed** to prompt with an asterisk (*) (or with *string* if the **-p** option was specified in the command line) for all subsequent commands. The **P** command alternately turns this mode on and off. It is initially on if the **-p** option was specified; otherwise, off. The current line number is unchanged.

q The **q** (quit) command causes **ed** to exit. No automatic write of a file is done (but see **DIAGNOSTICS** below).

Q The editor exits unconditionally without checking for changes in the buffer since the last **w** command.

(\$)r file The **r** (read) command reads the specified *file* into the buffer after the addressed line. If no file name is given, the remembered file name, if any, is used (see the **e** and **f** commands). The remembered file name is not changed unless *file* is the very first file name mentioned since **ed** was invoked. Address 0 is legal for **r** and places the contents of *file* at the beginning of the buffer. If the read is successful, the number of characters read is displayed. Upon completion, the new current line is the last line read into the buffer. If the *file* name starts with **!**, the rest of the line is interpreted as a shell command whose standard output is to be read. For example, **\$r !ls** appends a listing of files in the current directory to the end of the file being edited. A shell command is not remembered as the current file name.

(.,.)s / RE/ replacement/ flags

The **s** (substitute) command searches each addressed line for an occurrence of the specified **RE**. In each line in which a match is found, all (nonoverlapped) matched strings are replaced by *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 for the substitution to fail on all addressed lines. Any character other than space or newline can be used instead of **/** to delimit the **RE** and *replacement*. Upon completion, the new current line is the last line on which a substitution occurred. (Also see **WARNINGS** below.)

If an ampersand (&) appears in *replacement*, it is replaced by the string matching the **RE** on the current line. The special meaning of & in this context can be suppressed by preceding it with ****.

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 *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 containing more than one character or when preceded by a \.

A line can be split by substituting a newline character into it. The newline in *replacement* must be escaped by preceding it by \. Such substitution cannot be done as part of a **g** or **v** command list.

The value of *flags* is zero or more of:

- n** Substitute for the *n*th occurrence only of the RE found on each addressed line.
- g** Substitute for all nonoverlapped occurrences of the RE on each addressed line.
- l** Write to standard output the final line in which a substitution was made. The line is written in the format specified for the **l** command.
- n** Write to standard output the final line in which a substitution was made. The line is written in the format specified for the **n** command.
- p** Write to standard output the final line in which a substitution was made. The line is written in the format specified for the **p** command.

(.,.)t a Same as **m** command, except that a copy of the addressed lines is placed after address *a* (which can be 0). Upon completion, the new current line is the last line of the copy.

u The **u** (undo) command nullifies the effect of the most recent command that modified anything in the buffer, that is, the most recent **a**, **c**, **d**, **g**, **G**, **i**, **j**, **m**, **r**, **s**, **t**, **v**, or **V** command. All changes made to the buffer by a **g**, **G**, **v**, or **V** global command are "undone" as a single change; if no changes were made by the global command (such as with **g/RE/p**), the **u** command has no effect. The current line number is set to the value it had immediately before the command started.

(1,\$)v/RE/command-list

The complement of the global command **g** in that the lines marked during the first step are those that do *not* match the RE.

(1,\$)V/RE/

The complement of the interactive global command **G** in that the lines marked during the first step are those that do *not* match the RE.

(1,\$)w file

The **w** (write) command writes the addressed lines into the named file. If the file does not exist, it is created with mode 666 (readable and writable by everyone), unless the current **umask** setting dictates otherwise (see **umask**(1)). The remembered file name is not changed unless *file* is the very first file name encountered since **ed** was invoked. If no file name is given, the remembered file name, if any, is used (see the **e** and **f** commands). Upon completion, the current line address is unchanged. If the command is successful, the number of characters written is displayed.

If the *file* name starts with **!**, the rest of the line is interpreted as a shell command whose standard input is the addressed lines. Such a shell command is not remembered as the current file name.

x

A key string is demanded from the standard input. Subsequent **e**, **r**, and **w** commands will encrypt and decrypt the text with this key, using the algorithm of **crypt**(1). An explicitly empty key turns off encryption.

(\$)=

The line number of the addressed line is displayed. The current line address is unchanged by this command.

! shell-command

The remainder of the line after the **!** is sent to the shell to be interpreted and executed 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, **!!** repeats the last shell command. If any expansion is performed, the expanded line is echoed. Upon completion, the current line address is unchanged.

(.+1) newline An address alone on a line causes the addressed line to be printed. A newline alone is equivalent to **+.1p**. This technique 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.

The following size limitations apply: 256 characters per global command list, 64 characters per file name, and 32 MB characters in the buffer. The limit on the number of lines depends on the amount of user memory: each line takes 1 word.

EXTERNAL INFLUENCES

Environment Variables

SHELL determines the preferred command-line interpreter for use in all !-style commands. If this variable is null or not set, the POSIX shell, **/usr/bin/sh**, is used (see *sh-posix(1)*).

When set, **TMPDIR** specifies a directory to be used for temporary files, overriding the default directory, **/tmp**.

LANG provides a default value for internationalization variables that are unset or null. If **LANG** is unset or null, the default value is "C" (see *lang(5)*). If any internationalization variable contains an invalid setting, all internationalization variables default to "C". See *environ(5)*.

If **LC_ALL** is set to a nonempty string value, it overrides the values of all the other internationalization variables, including **LANG**.

LC_CTYPE determines the interpretation of text as single- and/or multibyte characters, the classification of characters as printable, and the characters matched by character class expressions in regular expressions.

LC_MESSAGES determines the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.

NLSPATH determines the location of message catalogues for the processing of **LC_MESSAGES**.

International Code Set Support

Single- and multibyte character code sets are supported.

DIAGNOSTICS

? Command error. Use **h** or **H** to get a detailed explanation.

?file Inaccessible file. Use **h** or **H** to get a detailed explanation.

If changes have been made in the buffer since the last **w** command that wrote the entire buffer, **ed** warns you if you attempt to destroy the buffer with an **e** or **q** command. **ed** displays ? or **warning: expecting 'w'**, then continues normal editing unless you enter a second **e** or **q** command, in which case the second command is executed. The **-s** or **-** command-line option inhibits this feature.

EXAMPLES

Make a simple substitution in **file-1** from a shell script, changing the first occurrence of **abc** in any line to **xyz**, and save the changes in **file-2**.

```
cat - << EOF | ed -s file-1
1,$ s/abc/xyz/
w file-2
q
EOF
```

Note that, if a command fails, the editor exits immediately.

WARNINGS

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 regular expression does not match a newline character.

The **l** command does not handle DEL correctly.

Files encrypted directly with the **crypt** command with the null key cannot be edited (see *crypt(1)*).

If the editor input is coming from a command file (e.g., **ed file < ed-cmd-file**), the editor exits at the first failure of a command in the command file.

When reading a file, **ed** discards ASCII NUL characters and all characters after the last newline. This can cause unexpected behavior when using regular expressions to search for character sequences containing NUL characters or text near end-of-file.

AUTHOR

ed was developed by HP and OSF.

FILES

| | |
|----------------|---|
| /tmp/ep | Temporary buffer file where <i>p</i> is the process number. |
| ed.hup | Work is saved here if the terminal is hung up. |

SEE ALSO

awk(1), csh(1), crypt(1), ex(1), grep(1), ksh(1), sed(1), sh(1), sh-bourne(1), sh-posix(1), stty(1), vi(1), fspec(4), environ(5), lang(5), regexp(5).

The **ed** section in *Text Processing: User's Guide*.

STANDARDS CONFORMANCE

ed: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

red: SVID2, SVID3, XPG2, XPG3


e

(64-Bit Applications Only)

NAME

elfdump - dump information contained in object files

SYNOPSIS

elfdump [-acCfghHjkLopqrsStuUv] [-D *num*] [+D *num2*] [-n *name*] [+s *section*] [-T *num*] [+T *num2*] *files...*

DESCRIPTION

elfdump takes one or more object files or libraries and dumps information about them. The following options are supported:

- a Dumps archive headers from an archive library.
- c Dumps the string table(s).
- C (Modifier) Demangles C++ symbol names before printing them. This modifier is valid with -c, -r, -s, and -t. If specified with -H, this modifier is ignored. If specified with -n *name*, the symbol whose unmangled name matches *name* will be printed, and its symbol name will be printed as a demangled name.
- D *num* (Modifier) Prints the section whose index is *num*.
- +D *num2* (Modifier) Prints the sections in the range 1 to *num2*. If used with -D, the sections in the range *num* to *num2* are printed. Valid with -h, -r, -s. If used with -r, only the relocations which apply to the section(s) in the range are printed.
- f Dumps the file header.
- g Dumps global symbols from an archive.
- h Dumps the section headers.
- H (Modifier) Dumps output information in hexadecimal, octal, or decimal format, with all options.
- j Prints the object dictionary for one or more executable files, if the source file was compiled with +objdebug or linked with +tools. The object dictionary entry contains the name of the object file that contributed to a particular section, the relative offset within the section, size of the object file's contribution, and attributes of the entry.
- k Prints the CTTI section headers according to the directory member relationship.
- L Dumps the .dynamic section in shared libraries and dynamically linked program files.
- n *name* (Modifier) Dumps information about the specified section or symbol *name*. This option is valid with -h, -r, -s, and -t. If used with -t, *name* pertains to a symbol name and **elfdump** will only dump the symbol entry whose name matches *name*. If used with the other options, *name* pertains to a section name and **elfdump** will only dump the section whose name matches it.
- o Dumps the optional headers (program headers).
- p (Modifier) Do not print titles, with all options.
- q (Modifier) Suppress printing CTTI section headers. Valid with -k option.
- r Dumps the relocations.
- s Dumps the section contents.
- +s *name* (Modifier) Dumps the section specified by *name*. Valid with -c and -t only.
- S (Modifier) Dumps output information in short format. Valid with the -h and -o options.
- t Dumps the symbol table entries.
- T *num* Prints the symbol whose index is *num*.
- +T *num2* (Modifier) Prints the symbols in the range 0 to *num2*. If used with -T, print the symbols in the range *num* to *num2*. Valid with -t.
- u Prints the usage menu.
- U Prints the unwind table.

-v (Modifier) Verifies the CTTI section headers before printing. Valid with the **-k** option.

SEE ALSO**System Tools**

ld(1) Invoke the link editor

Miscellaneous

a.out(4) Assembler, compiler, and linker output
elf(3E) Executable and Linking Format

Texts and Tutorials:

HP-UX Linker and Libraries Online User Guide

(See the **+help** option)

HP-UX Linker and Libraries User's Guide

(See *manuals*(5) for ordering information)

HP-UX Software Transition Toolkit (STK) -- ELF Object Formats

<http://www.software.hp.com/STK/>

NAME

elm - process electronic mail through a screen-oriented interface

SYNOPSIS

```
elm [-aKkmtVz] [-f folder]
elm [-s subject] address-list
elm -c [alias-list]
elm -h
elm -v
```

DESCRIPTION

The **elm** program is a screen-oriented electronic mail processing system. It supports the industry-wide MIME standard for nontext mail, a special forms message and forms reply mechanism, and an easy-to-use alias system for individuals and groups. **elm** operates in three principal modes:

- Interactive mode, running as an interactive mail interface program. (First syntax.)
- Message mode, sending a single interactive message to a list of mail addresses from a shell command line. (Second syntax.)
- File mode, sending a file or command output to a list of mail addresses via a command-line pipe or redirection. (Second syntax.)

In all three cases, **elm** honors the values that are set in your **elmrc** initialization file, in your **elm** alias database, and in the system **elm** alias database.

The modes are described below in inverse order (shortest description to longest).

Options

The following options are recognized:

- a Set **arrow=ON**. Use the arrow (->) instead of the inverse bar to mark the current item in the various indexes. This overrides the setting of the **arrow** boolean variable (see the ELM CONFIGURATION section).
- c Check alias. Check the aliases in *alias-list* against your personal **elm** alias database and the system **elm** alias database. The results are written to standard output. Errors are reported first, in the form:

(alias "*alias*" is unknown)

Successes are reported in a header-entry format, with group aliases replaced by their members, in the form:

Expands to: *alias-address* (*fullname*),
 alias-address (*fullname*),
 ...
 alias-address (*fullname*)

If there is no *fullname*, the " (*fullname*)" portion is omitted.
- f *folder* Folder file. Read mail from the *folder* file rather than from the incoming mailbox. A folder file is in the standard mail file format, as created by the mail system or saved by **elm** itself.
- h Help. Display an annotated list of command-line options.
- k Set **softkeys=OFF**. Disable the use of softkeys (HP 2622 function keys). This overrides the setting of the **softkeys** boolean variable (see the ELM CONFIGURATION section).
- K Set **keypad=OFF** and **softkeys=OFF**. Disable the use of softkeys and arrow cursor keys. If your terminal does not have the HP 2622 function key protocols, this option is required. This overrides the settings of the **keypad** and **softkeys** boolean variables (see the ELM CONFIGURATION section).
- m Set **menu=OFF**. Do not display the command menus on several Interactive Mode screens. This overrides the setting of the **menu** boolean variable (see the ELM

- CONFIGURATION section).
- s** *subject* Subject. Specify the subject for a File Mode or Message Mode message.
 - t** Set **usetite=OFF**. Do not use the **termcap ti/te** and **terminfo cup** cursor-positioning entries. This overrides the setting of the **usetite** boolean variable (see the ELM CONFIGURATION section).
 - V** Verbose transmission. Pass outbound messages to the **sendmail** mail transport agent using the **-v** option (see *sendmail(1M)*).
 - v** Version. Print out the **elm** version information. This displays the revision number and build date as well as the compilation features that were specified or omitted.
 - z** Zero. Do not enter **elm** if there is no mail in the incoming mailbox.

Operands

The following operands are recognized:

- address-list* A blank-separated list of one or more mail addresses, your **elm** user aliases, or **elm** system aliases.
- alias-list* A blank-separated list of one or more of your **elm** user aliases or **elm** system aliases.

Terminology

The following terms are used throughout this manpage.

blank

A space or a tab character, sometimes known as linear white space.

body

The body of a message. See **message**.

boolean variable

See **configuration variable**.

configuration variable

A boolean, numeric, or string variable that defines default behavior in the **elm** mail system. See the ELM CONFIGURATION section.

elm system alias text file

The source file, **/var/mail/.elm/aliases.text**, for the **elm** system alias database.

elm user alias text file

The source file, **\$HOME/.elm/aliases.text**, for a user's own **elm** alias database.

elm user headers file

A file, **\$HOME/.elm/elmheaders**, where a user can specify special header entries that are included in all outbound messages.

elmrc configuration file

A file, **\$HOME/.elm/elmrc**, that defines the initial values for **elm** configuration variables.

environment variable

A global variable set in the shell that called **elm**. See the EXTERNAL INFLUENCES section.

folder

A file that contains mail messages in the format created by **sendmail** or **elm**.

full name

The first and last name of a user, as extracted from an alias text file or from the **/etc/passwd** file.

header

The header of a message. See **message**.

header entry

An entry in the header portion of a message, sometimes called a header field.

incoming mailbox

The mailbox where you receive your mail, usually `/var/mail/loginname`.

mail directory

The directory, defined by the `maildir` string variable, where a user normally stores mail messages in folders.

mail transport agent (MTA)

The program that sends and receives mail messages to and from other systems. On HP-UX systems, the MTA is `sendmail` (see `sendmail(1M)`).

mailcap

A file that contains information on how to compose and display mail messages that are not just seven- and eight-bit ASCII characters.

metamail

A system program that processes nontext mail messages.

message

In a folder, a sequence of text lines comprised of a message delimiter, a header, and a body. The message delimiter is a line in the form:

From *sender date*

The **header** starts after the message delimiter and ends with the first null line. The **body** begins at the null line and ends at the next message delimiter. A body can have subsections, called **attachments** or **body parts**, which have are comprised of a boundary delimiter, a header, and a body. This process can be recursive. See the METAMAIL CONFIGURATION section for more details.

numeric variable

See **configuration variable**.

sendmail alias database

The alias database, `/etc/mail/aliases`, that is used by the `sendmail` MTA to direct local mail.

signature file

A file that is appended to your outbound messages, usually containing information about yourself. You can have two signature files, one for messages to your local machine and one for other messages. See the `localsignature` and `remotesignature` string variables.

string variable

See **configuration variable**.

user name

Usually the login or mailbox name of someone you send mail to.

variable

See **configuration variable** and **environment variable**.

FILE MODE

If standard input is connected to a pipe or to a file, and an *address-list* is specified, `elm` operates in File Mode.

The output of the previous command in the pipe, or the content of the file, is mailed to the members of the *address-list*. The *address-list* is expanded, based on your `elm` alias database and the system `elm` alias database, and placed in the **To:** header entry.

If `-s` is omitted or *subject* is null, *subject* defaults to:

`no subject (file transmission)`

The expressed or default value of *subject* is placed in the **Subject:** header entry.

See the EXAMPLES section.

MESSAGE MODE

If standard input is connected to your terminal, and an *address-list* is specified, **elm** operates in Message Mode.

The *address-list* is expanded, based on your **elm** alias database and the system **elm** alias database, and placed in the **To:** header entry. The **To:** header entry is displayed, in the same form as for the Message Menu **m** (mail) command in Interactive Mode.

The value of *subject*, if nonnull, or a null string, is placed in the **Subject:** header entry and the **Subject:** line is displayed for modification.

If **askcc** is **ON** in your **elmr**c file, you are prompted for **Copies to:**.

Then the editor defined by the **editor** string variable (if a signature file is not added) or the **altditor** string variable (if a signature file is added) is started so that you can write your message.

When you leave your editor, you enter the Send Menu, as described for Interactive Mode.

If you choose the Send Menu **s** (send) command, the message is sent and the program terminates. If you select the Send Menu **f** (forget) command, the message is stored in **\$HOME/Canceled.mail** and the program terminates. If you select other commands, the appropriate action occurs.

See the EXAMPLES section.

INTERACTIVE MODE

If standard input is connected to your terminal, and there is no *address-list*, **elm** operates in a screen-oriented Interactive Mode.

If you do not have a **\$HOME/.elm** directory, or if you do not have a mail directory, defined by the **mail-dir** string variable, you are asked in turn if they should be created. You can answer **y** for *yes*, **n** for *no*, or **q** for *quit*. For **y** or **n**, the directories are created or not, as appropriate, and the program continues. For **q**, the program terminates.

Overview

When invoked, **elm** reads customized variables from file **\$HOME/.elm/elmr**c (if it exists) to initialize parameters. This file can be saved from within **elm** and some of these variables can also be modified with the Message Menu **o** (option) command.

elm first displays the Main or Message Menu, which shows index entries for the messages in your incoming mailbox or selected mail folder. Among other options, you can read, print, reply to, and forward these messages, as well as initiate new mail messages to other users.

You can also move to the Alias Menu, where you can create, modify, and delete your personal aliases. From the Alias Menu, you can select one or more of your aliases and send a message to the corresponding users.

When you send a message, you can include attachments in a number of formats, such as PostScript, images, audio, and video, as well as plain text. The attachments are managed separately, which can be convenient both for you and your correspondents.

Sending Messages

When you send a message, you use the editor defined by the **editor** or **altditor** string variable. If **builtin** is your editor, a set of commands described in the Built-In Editor subsection is available while composing your message.

If the **elmheaders** file exists (see the HEADER FILE section), all nonblank lines in the file are copied to the headers of all outbound mail. This is useful for adding special information headers such as **X-Organization:**, **X-Phone:**, and so forth.

MIME Support

elm supports the MIME protocols for headers and messages (RFC 1521 and RFC 1522) enabling it to view and send mail containing other than normal ASCII text. For example, the mail contents can be audio, video, images, etc., or a combination of these.

This also enables conformance with SMTP (RFC 821), which allows only 7-bit characters in the message, by using MIME-encoding (**base64** and **quoted-printable**) to convert 8-bit data to 7-bit.

elm also provides a facility to view multipart MIME messages. If **elm** receives a message whose type is not **text/plain**, it invokes **metamail**, which invokes the appropriate utility (for example, **ghost-view**, **xv**, an audio editor, **mpeg**) to display the different mail parts according to the content type (for example, **application/postscript**, **image**, **audio**, **video**).

Aliases

elm has its own alias system that supports both personal and system-wide aliases. Personal aliases are specific to a single user; system aliases are available to everyone on the system where the system aliases reside (see **newaliases(1)**). You can access the Alias Menu by executing the Message Menu **a** (alias) command. You can then create and save an alias for the current message, create and check other aliases, and send messages to one or more aliases.

Aliases are limited to 2500 bytes. If you wish to create a group alias that is longer than 2500 bytes, please ask your system administrator to create it for you in the **sendmail** system alias file, **/etc/mail/aliases** (see **sendmail(1M)**).

INTERACTIVE MODE MENUS AND COMMANDS

This section begins with the Message Menu, which is the main screen for Interactive Mode. The rest of the menus are presented alphabetically.

Message Menu

The Message Index is displayed on the Message Menu. You can use the following commands to manipulate and send messages. Some commands use a series of prompts to complete their action. You can use **Ctrl-D** to cancel their operations.

The commands are:

| | |
|--------------------------|---|
| ! <i>command</i> | Shell Escape. Send <i>command</i> to the shell defined by the shell string variable without leaving elm . |
| # | Display all known information about the current message. |
| \$ | Resynchronize the messages without leaving elm . If there are any messages marked for deletion, you are asked if you want to delete them. If any messages are deleted or any status flags have changed, the messages are written back to the mailbox file. All tags are removed. |
| % | Display the computed return address of the current message. |
| * | Set the current message pointer to the last message. |
| + | Display the next message index page, when applicable. |
| - | Display the previous message index page, when applicable. |
| / <i>pattern</i> | Pattern match. Search for <i>pattern</i> in the <i>from</i> and <i>subject</i> fields of the current message index. The search starts at the current message and wraps around to the beginning of the index. The current message pointer is set to the first message that matches. Uppercase and lowercase are treated as equivalent. |
| // <i>pattern</i> | Pattern match. Search for <i>pattern</i> through all the lines of the current folder. The search starts at the current message and wraps around to the beginning of the folder. The current message pointer is set to the first message that matches. Uppercase and lowercase are treated as equivalent. |
| < | Calendar. Scan message for calendar entries and add them to your calendar file. A calendar entry is defined as a line whose first nonblank characters are -> , as in: <div style="text-align: center;">->calendar-entry</div> <p>The delimiter -> and surrounding blanks are removed before the entry is added to the calendar file. Resultant blank lines are ignored. You can define the calendar file name in your elmrc file or with the Options Menu.</p> |
| = | Set the current message pointer to the first message. |
| > | Save in folder. Same as the Message Menu s (save) command. |

| | |
|-----------------|---|
| ?key ... | Help on key. Display a one-line description of what each <i>key</i> does. ? displays a summary listing for each command available. A period (.) returns you to the Message Menu. |
| @ | Display a summary of the messages indexed on the current screen. |
| | Pipe the current message or the set of tagged messages through other filters as desired. Use the shell defined by the shell string variable. |
| n | New current message. Change the current message pointer to the one indexed as <i>n</i> . If the message is not on the current page of headers, the appropriate page displayed. |
| Return | Read current message. The screen is cleared and the current message is displayed by the pager defined by the pager string variable. |
| a | Alias. Switch to the Alias Menu. |
| b | Bounce mail. This is similar to forwarding a message, except that you do not edit the message and the return address is set to the original sender's address, rather than to your address. |
| c | Change folder. This command is used to change the file whose messages are displayed on the Message Menu. You are asked for a file name. The file must be in message format; otherwise, elm aborts. You can use the customary wildcards for your shell, as well as the following special names: <ul style="list-style-type: none"> ! Your incoming mail folder. > Your received folder, defined by the receivedmail string variable. < Your sent folder, defined by the sentmail string variable. . The previously used folder. @alias The default folder for the login name associated with the <i>alias</i> alias. =filename A file in the directory defined by the maildir string variable. |
| C | Copy message. Save the current message or the set of tagged messages to a folder. You are prompted for a file name with a default value. The default value is a file in the maildir directory with the user name of the sender of the first message in the set being saved. Any tags are cleared. Unlike the > and s commands, the messages are not marked for deletion and the current message pointer is not moved. |
| d | Delete. Mark the current message for deletion. See also Ctrl-D , u , and Ctrl-U . |
| Ctrl-D | Delete. Mark all messages for deletion that contain a specified pattern in the From: and Subject: header entries. See also d , u , and Ctrl-U . |
| e | Edit. Allows you to physically edit the current mail folder using the editor defined by the editor string variable. When you exit from your editor, elm resynchronizes your mail folder (see the \$ command). |
| f | Forward the current message. You are asked if you want to edit the outbound message. If you answer y , the characters defined by the prefix string variable are prefixed to each line of the message and the editor defined by the editor string variable will be invoked to allow you to edit the message. If you answer n , the characters are not prefixed and the editor will not be invoked. In either case, you are prompted for To: recipients, allowed to edit the Subject: header entry, and, if the askcc boolean variable is ON, you are prompted for Cc: recipients. <p>If the userlevel numeric variable is 1 (intermediate) or 2 (expert), and there was a previous sent or forgotten message in this session, you are asked if you would like to</p> <p style="text-align: center;">Recall last kept message instead? (y/n)</p> <p>If you answer y, the previous message is returned to the send buffer. If you answer n, the current message is copied into the send buffer and your signature file (if any) is appended.</p> <p>Then the editor is invoked if you chose to edit the outbound message (above). When you leave the editor, or if it was not invoked, the Send Menu is displayed.</p> |

- g** Group reply. The reply is automatically sent **To:** the sender of the message, with **Cc:** to all the original **To:** and **Cc:** recipients. Otherwise, the action is the same as for the **r** command.
- h** Same as **Return**, except that the message is displayed with all headers.
- j** Move down. Move the current message pointer down to the next message.
- J** Move down. Move the current message pointer down to the next undeleted message.
- k** Move up. Move the current message pointer up to the previous message.
- K** Move up. Move the current message pointer up to the previous undeleted message.
- l (ell)** Limit the displayed messages to those that contain certain string values. You are prompted with **Enter criteria:.** To set, add to, or clear the limiting criteria, type one of:
- all** Clear all the criteria and restore the normal display.
 - from *string*** Restrict to entries that contain *string* in the **From:** header.
 - subject *string*** Restrict to entries that contain *string* in the **Subject:** header.
 - to *string*** Restrict to entries that contain *string* in the **To:** header.
- You can add limiting criteria by repeating the **l** command.
- Ctrl-L** Redraw the screen.
- m** Mail. Send mail to one or more addresses. You are prompted for **To:** recipients, a **Subject:** and, if the **askcc** boolean variable is **ON**, **Cc:** recipients.
- If the **userlevel** numeric variable is **1** (intermediate) or **2** (expert), and there was a previous sent or forgotten message in this session, you are asked if you would like to
- Recall last kept message instead? (y/n)**
- If you answer **y**, the previous message is returned to the send buffer. If you answer **n**, the signature file (if any) is copied into the send buffer.
- Then, the editor defined by the **editor** string variable is invoked. After you exit from your editor, the Send Menu is displayed.
- n** Next message. Advances the current message pointer to the next message, and displays that message as for the **Return** command.
- o** Options. Invokes the Options Menu, permitting you to change certain configuration options. The changeable options are defined by the **configoptions** string variable.
- p** Print. Print the current message or the set of tagged messages using the command defined by the **print** string variable. The current message pointer does not move. Tagged messages remain tagged.
- q** Quit. Gracefully terminate, performing message cleanup according to defined personal preferences. You can choose to actually delete messages marked for deletion. For your incoming mailbox, you can choose to keep undeleted mail in the mailbox or move it to the received folder defined by the **receivedmail** string variable.
- If the **ask** boolean variable is **ON**, you may be asked the following questions. The actions described are all performed after you have answered all the relevant questions.
- Delete messages? (y/n)**
- This question is asked if you have messages marked for deletion. The default answer is provided by the **alwaysdelete** boolean variable (**ON** means **y** (yes) and **OFF** means **n** (no)).
- If you answer **y**, all messages marked for deletion will be deleted.
- If you answer **n**, all messages marked for deletion will be restored to their former read, unread, or new state.

Move read messages to "received" folder? (y/n)

This question is asked if you are reading your incoming mailbox and if you have messages that have been read. The default answer is provided by the **always-store** boolean variable (ON means **y** (yes) and OFF means **n** (no)).

If you answer **y**, undeleted messages that have been read will be moved to the folder defined by the **receivedmail** string variable and the next question will also be asked.

If you answer **n**, all undeleted messages are returned to your incoming mailbox and the next question is not asked.

Keep unread messages in incoming mailbox? (y/n)

This question is asked if you are reading your incoming mailbox, if you answered **y** to the **Move read messages...** question (or it was not asked), and if you have messages that have not been read. The default answer is provided by the **alwayskeep** boolean variable (ON means **y** (yes) and OFF means **n** (no)).

If you answer **y**, all undeleted unread (new and old) messages are returned to your incoming mailbox.

If you answer **n**, all undeleted unread messages will be moved to the folder defined by the **receivedmail** string variable.

If the **ask** boolean variable is OFF, the answers to the questions (which are not displayed) are taken automatically from the values of the **alwaysdelete**, **always-store**, and **alwayskeep** boolean variables, respectively.

- Q** Quick quit. This is equivalent to executing the **q** command with the **ask** boolean variable set to OFF.
- r** Reply to the sender of the current message. If the **autocopy** boolean variable is OFF, you are asked if the source message should be copied into the edit buffer. If it is ON, the message is copied automatically. If copied in, all lines from the message are preceded by the prefix string defined by the **prefix** string variable. The **To:** header is set to the sender of the message (or the address in the **Reply-To:** header, if one was set), the **Subject:** is set to the subject of the message, preceded by **Re:**, and presented for you to edit. If the **askcc** boolean variable is ON, you are prompted for **Cc:** recipients. Then, the editor defined by the **editor** string variable is invoked. After you exit from your editor, the Send Menu is displayed.
- s** Save in folder (same as >). Save the current message or the set of tagged messages to a folder. You are prompted for a file name with a default value. The default value is a file in the **maildir** directory with the login name of the sender of the first message in the set being saved. Any tags are cleared and the messages are marked for deletion. The current message pointer is moved to the first undeleted message after the last saved message.
- t** Tag toggle. Tag the current message for a later operation and move the current message pointer to the next undeleted message. The operation can be one of **|**, **C**, **p**, and **s**.
- Or, remove the tag from a tagged message. See also the **Ctrl-T** command.
- T** Tag toggle. Tag the current message for a later operation and remain at the current message. The operation can be one of **|**, **C**, **p**, and **s**.
- Or, remove the tag from a tagged message. See also the **Ctrl-T** command.
- Ctrl-T** Tag all messages containing the specified pattern. Or remove the tags from all tagged messages.
- If any messages are currently tagged, you are asked if the tags should be removed. Answer **y** to remove the old tags; answer **n** to keep them. In either case, you are prompted for a string to match in either the **From:** or **Subject:** line of each message. All messages that match the criterion are tagged. If you enter a null string (carriage-return alone), no more messages are tagged.

| | |
|---------------|--|
| u | Undelete. Remove the deletion mark from the current message. See also d , Ctrl-D , and Ctrl-U . |
| Ctrl-U | Undelete. Remove any deletion mark from all messages that contain a specified pattern in the From: and Subject: header entries. See also d , Ctrl-D , and u . |
| v | View attachments. Invoke the Attachment View Menu for the current message. |
| x | Exit. Exit without changing the mailbox. If changes are pending, such as deletions, you are asked if they can be abandoned. If you answer y , the changes are abandoned and the program terminates. If you answer n the exit is abandoned and you return to the Message Menu command prompt. |
| X | Exit immediately without changing the mailbox. All pending changes are abandoned. |

e

Message Index

The messages in the current folder are indexed on the Message Menu, one per line, in the format:

```
sssnum mmm d from (lines) subject
```

defined as:

| | |
|----------------|--|
| <i>sss</i> | A three-character status field, described in the Message Status subsection. |
| <i>num</i> | The ordinal message index number. |
| <i>mmm</i> | The month from the last Date: header entry, or from the From message header. |
| <i>d</i> | The day from the last Date: header entry, or from the From message header. |
| <i>from</i> | Either the sender name from the last From: header entry or from the From message header. |
| <i>lines</i> | The number of lines in the message. |
| <i>subject</i> | The subject description from the first Subject: header entry, truncated to fit your screen. |

The current message index entry is either highlighted in inverse video or marked in the left margin with an arrow (->). See the **-a** option in the Options subsection and the **arrow** string variable in the ELM CONFIGURATION section.

Message Status

The first three characters of each message index entry describe the message status. Each can be blank or one of the values described below in descending order of precedence.

When a message has more than one status flag of a particular type set, the highest-precedence indicator is displayed on the index line. For example, if a forms message (**F**) is also marked as company confidential (**C**), the **C** rather than the **F** status character is displayed.

Column One: Variable Status

| | |
|----------|--|
| D | Deleted. The message is marked for deletion. |
| E | Expired. The date specified in the Expires: header entry is older than today. elm accepts the following date formats: <pre>Mon, 11 Jun 90 (format produced by elm in the Header Menu) Jun 11, 90 11 Jun, 90 9006111324GMT (ISO X.400 format: YYMMDDhhmmzzz)</pre> |
| N | New. The message was received after your last elm session or during the current session. The message has not been read. |
| O | Old. The message was received before or during your last elm session. It was marked N in your last session and it was not read. |
| | Blank. The message has been read. |

Column Two: Permanent Status

- C** Confidential. The **Sensitivity: 3** header entry is present. The message is considered company confidential, as specified by the ISO X.400 standard. You can set this value for outbound mail with the user-defined option of the Header Menu.
- U** Urgent. The message contains a **Priority:** header entry.
- P** Private. The **Sensitivity: 2** header entry is present. The message is considered private, as specified by the ISO X.400 standard. You can set this value for outbound mail with the user-defined option of the Header Menu.
- A** Action. The message contains an **Action:** header entry.
- F** Forms. The message is an **elm** forms message. The message contains a **Content-Type: mailform** header entry.
- M** MIME. The message or its attachments is in a MIME format that can be displayed using **metamail**.
- ?** MIME. The message or its attachments is in a MIME format whose version is not supported.
- Blank. Normal status.

Column Three: Tagged Status

- +** Tagged. Tagged messages are handled as a group by some commands. See **t** and other commands in the Message Menu subsection.
- Blank. The message is not tagged.

Built-In Editor

When you are composing an outbound message with the **builtin** built-in editor, it prompts you for text lines with an empty line. Enter a period (.) to end the message and continue with the Send Menu.

Built-in editor commands are lines that begin with an escape character, defined by the **escape** string variable. The default escape character is tilde (~).

Note: Some remote login programs use tilde as their default escape character when it is the first character on a line. (You can tell, because the tilde does not print.) Usually, the tilde is transmitted when you enter a second character that is not recognized by the program or when you enter a second tilde. See the program documentation for further information.

The built-in editor commands are:

- ~!** [*command*] Execute the shell *command*, if one is given (as in **~!ls**), or start an interactive shell, using the shell defined by the **shell** string variable.
- ~<** *command* Execute the shell *command* and place the output of the command into the editor buffer. For example, "**~< who**" inserts the output of the **who** command in your message.
- ~?** Print a brief help menu.
- ~~** Start a line with a single tilde (~) character.
- ~b** Prompt for changes to the Blind-Carbon-Copy (**Bcc:**) list.
- ~c** Prompt for changes to the Carbon-Copy (**Cc:**) list.
- ~e** Invoke the editor defined for the **easyeditor** string variable on the message, if possible.
- ~f** [*options*] Add the specified list of messages or the current message. This uses **readmail** which means that all **readmail** options are available (see **readmail(1)**).
- ~h** Prompt for changes to all the available headers (**To:**, **Cc:**, **Bcc:**, and **Subject:**).
- ~m** [*options*] Same as **~f**, but each line is prefixed with the current prefix. See the **prefix** string variable.
- ~o** Prompt for the name of an editor to use on the message.
- ~p** Print out the message as typed in so far.

| | |
|---------------------------|--|
| ~r <i>filename</i> | Include (read in) the contents of the specified file. |
| ~s | Prompt for changes to the Subject: line. |
| ~t | Prompt for changes to the To: list. |
| ~v | Invoke the editor defined for the visualeditor string variable on the message, if possible. |

Alias Menu

The Alias Menu is invoked with the Message Menu **a** command. The source text for your alias file is stored in the file `$HOME/.elm/aliases.text`. You can edit this file directly or with the following commands.

The aliases currently compiled into your database and the system database are displayed in an indexed list similar to the Message Menu. The entry format is described in the Alias Index subsection. The index is sorted in the order defined by the **aliassortby** string variable.

The commands are:

| | |
|------------------|--|
| \$ | Resynchronize your alias text file and your alias database by rebuilding the database from the text file by running newalias . Aliases marked for deletion are removed, tagged aliases are untagged, and new and changed aliases are recognized. The alias screen is updated to reflect these changes. |
| + | Display the next alias index page, when applicable. |
| - | Display the previous alias index page, when applicable. |
| /pattern | Pattern match. Search for <i>pattern</i> in the alias and user name fields of the alias list. The search starts at the current alias and wraps around to the beginning of the alias list. The current alias pointer is set to the first alias that matches. Uppercase and lowercase are treated as equivalent. |
| //pattern | Pattern match. Search for <i>pattern</i> through all the fields of the alias list (alias, user name, comment, and address). The search starts at the current alias and wraps around to the beginning of the alias list. The current alias pointer is set to the first alias that matches. Uppercase and lowercase are treated as equivalent. /pattern Pattern match. This command allows you to search through all the alias and username entries in the alias list, starting at the current alias and continuing through the end. If the first character of the pattern is a / , then the comment and the fully expanded address fields are also included in the search. The search is case-insensitive. This allows you to find a specific alias in a situation where there are a large number of aliases. |
| ?key ... | Help on key. Display a one-line description of what each <i>key</i> does. ? displays a summary listing for each command available. A period (.) returns you to the Alias Menu. |
| a | Alias current message. This allows you to create an alias that has the return address of the current message as the address field of the alias. It prompts for a unique alias name and allows you to edit the comment and address fields. |
| c | Change the current user alias. The old values of the alias fields are used as the defaults in the prompts for the new values. You cannot change the alias name. If the alias name is one of a multiple-alias record, it is removed from that record and stored as a separate record. The old alias is marked N . Changes are effective after the next alias resynchronization. |
| d | Mark the current user alias for deletion. The deletions are made when you exit from the Alias Menu with an q , r , or i command or you resynchronize your alias database with the \$ command. (You cannot delete a system alias in this way.) |
| Ctrl-D | Delete user aliases with a specified search pattern. |
| e | Edit your aliases.text file, using the editor defined in the editor string variable. Your aliases are resynchronized when you finish editing (see the \$ command). |
| f | Display a fully expanded alias. The currently selected alias is fully expanded and displayed. |

| | | | | | | | | | | | | | | | |
|----------------------------|---|------------|--|----------------------------|--|---------------------------|--|--------------|--|---------------|---|---------------|--|-------------|--|
| i,I | See the Alias Menu q and Q commands. | | | | | | | | | | | | | | |
| j | Move down. Move the current alias pointer down to the next alias. | | | | | | | | | | | | | | |
| J | Move down. Move the current alias pointer down to the next undeleted alias. | | | | | | | | | | | | | | |
| k | Move up. Move the current alias pointer up to the previous alias. | | | | | | | | | | | | | | |
| K | Move up. Move the current alias pointer up to the previous undeleted alias. | | | | | | | | | | | | | | |
| l (ell) | Limit the displayed aliases to certain types or those that contain certain string values. You are prompted with Enter criteria:. To set, add to, or clear the limiting criteria, type one of: <table> <tr> <td>all</td><td>Clear all the criteria and restore the normal display.</td></tr> <tr> <td>alias <i>string</i></td><td>Restrict to alias names containing <i>string</i>.</td></tr> <tr> <td>name <i>string</i></td><td>Restrict to full names (first name and last name) containing <i>string</i>.</td></tr> <tr> <td>group</td><td>Restrict to group aliases (can include system and user aliases).</td></tr> <tr> <td>person</td><td>Restrict to person aliases (can include system and user aliases).</td></tr> <tr> <td>system</td><td>Restrict to system aliases (can include group and person aliases).</td></tr> <tr> <td>user</td><td>Restrict to system aliases (can include group and person aliases).</td></tr> </table> <p>You can add limiting criteria by repeating the l command.</p> | all | Clear all the criteria and restore the normal display. | alias <i>string</i> | Restrict to alias names containing <i>string</i> . | name <i>string</i> | Restrict to full names (first name and last name) containing <i>string</i> . | group | Restrict to group aliases (can include system and user aliases). | person | Restrict to person aliases (can include system and user aliases). | system | Restrict to system aliases (can include group and person aliases). | user | Restrict to system aliases (can include group and person aliases). |
| all | Clear all the criteria and restore the normal display. | | | | | | | | | | | | | | |
| alias <i>string</i> | Restrict to alias names containing <i>string</i> . | | | | | | | | | | | | | | |
| name <i>string</i> | Restrict to full names (first name and last name) containing <i>string</i> . | | | | | | | | | | | | | | |
| group | Restrict to group aliases (can include system and user aliases). | | | | | | | | | | | | | | |
| person | Restrict to person aliases (can include system and user aliases). | | | | | | | | | | | | | | |
| system | Restrict to system aliases (can include group and person aliases). | | | | | | | | | | | | | | |
| user | Restrict to system aliases (can include group and person aliases). | | | | | | | | | | | | | | |
| Ctrl-L | Redraw the screen. | | | | | | | | | | | | | | |
| m | Mail to the current alias or to the set of tagged aliases. The corresponding expanded addresses are placed in the To: header entry, and processing continues as for the Message Menu m (mail) command. The tags are cleared. | | | | | | | | | | | | | | |
| n | Make a user alias. elm prompts for a unique alias name, then for an address. The information provided is added to your individual alias_text file (<code>\$HOME/.elm/aliases.text</code>), then added to the database. | | | | | | | | | | | | | | |
| q | Exit. Return to the Message Menu. If aliases are marked for deletion, you are asked if you want to delete them. The alias index pointer is retained. If the alias text file was changed, the database is resynchronized. | | | | | | | | | | | | | | |
| Q | Exit. Return to the Message Menu. If aliases are marked for deletion, the mark is retained and the alias is not deleted. The alias index pointer is retained. If the alias text file was changed, the database is resynchronized. | | | | | | | | | | | | | | |
| r,R | See the Alias Menu q and Q commands. | | | | | | | | | | | | | | |
| t | Tag the current alias for a later operation and move the current alias pointer to the next undeleted alias. The operation can be one of c , m , or n . Or, remove the tag from a tagged alias. See also the Ctrl-T command. | | | | | | | | | | | | | | |
| T | Tag. Tag the current alias for a later operation and remain at the current alias. The operation can be one of c , m , or n . Or, remove the tag from a tagged alias. See also the Ctrl-T command. | | | | | | | | | | | | | | |
| Ctrl-T | Tag all aliases containing a specified pattern for a later operation. The operation can be one of c , m , or n . Or remove the tags from all tagged aliases. If any aliases are currently tagged, you are asked if the tags should be removed. Answer y to remove the old tags; answer n to keep them. In either case, you are prompted for a string to match in either the alias name or user name fields of each alias. All aliases that match the criterion are tagged. If you enter a null string (carriage-return alone) no more aliases are tagged. | | | | | | | | | | | | | | |
| u | Undelete. Remove the deletion mark from the current alias. See also d , Ctrl-D , and Ctrl-U . | | | | | | | | | | | | | | |
| Ctrl-U | Undelete. Remove any deletion mark from all messages that contain a specified pattern in the From: and Subject: header entries. See also d , Ctrl-D , and u . | | | | | | | | | | | | | | |

- v** View. Display the *address-list* for the current alias.
- x** Exit from the alias menu without processing any deletions. Aliases marked for deletion are unmarked and **newalias** is not run, even if alias additions have been made.

Alias Index

The aliases in the current database are indexed on the Alias Menu, one per line. The database values are defined in *newalias(1)*.

ssnum fullname[, comment] type [(S)] alias

defined as:

- ss** A two-character status field. The first character can be:
 - D** Delete. The alias is marked for deletion.
 - N** New. The alias is new or changed in the alias text file but is not included in the current database. Resynchronization is needed.
 - Blank. The alias is in the current database.
 The second character can be:
 - +** Tag. The alias is tagged.
 - Blank. The alias is not tagged.
- num** The index number of the alias.
- fullname** The full name for the alias, as it will be used in an expanded address. It has the form:
 - firstname lastname*
 - firstname* The first name, from the alias database.
 - lastname* The last name, from the alias database.
- comment** Comment, from the alias database.
- type** Type of alias. This is **Person** for an alias with a single address or **Group** for an alias with two or more addresses.
- (S)** If present, the entry is from the **elm** system alias database. If absent, the entry is from your personal alias database.
- alias** The alias name. If the record has multiple alias names, there is one index entry per name.

Attachment Configuration Menu

The Attachment Configuration Menu is invoked with the Attachment Send Menu **a** (add) or **m** (modify) command. The menu displays the default or current specification for an attachment. If it is called with the **a** command, it automatically prompts for a file name. The commands are:

- d** Description. The value is placed in a **Content-Description:** body-part header entry. The default is the file name.
- e** Content-Transfer-Encoding. This is the method by which the file is encoded to allow it to pass through various Mail Transport Agents. The choices are:
 - 7bit** Unencoded, normal **US-ASCII** text.
 - 8bit** Unencoded 8-bit characters with the high-order bit set.
 - quoted-printable** Text with control characters and high-order-bit characters converted to a string in the form **=hh**, where **hh** is the hexadecimal representation of the character. An **=** at the end of a line indicates that the source line was broken into two lines.
 - base64** Any file type with bits encoded in 6-bit groups and rendered in numeric order as the **US-ASCII** characters **A-Z**, **a-z**, **0-9**, **+**, and **/**. The last line may be padded to a multiple of 4 characters with **=** characters.
 - binary** Unencoded binary data.
 The value is placed in a **Content-Transfer-Encoding:** body-part header entry. The default is **7bit**.

- f** File name. The name of the file to be attached. **elm** examines the file and sets the values of Content-Transfer-Encoding, Content-Disposition, and Content-Type accordingly.
- p** Content-Disposition. The value is placed in a **Content-Disposition:** body-part header entry. The default is **attachment; filename=filename**.
- t** Content-Type. The type of the file and supporting parameters, in the form:

type/subtype [**;** *parameter*]...

The *type* can be one of **application**, **audio**, **image**, **message**, **text**, or **video**, as defined in RFC 1521. Although **multipart** is also a valid type, you cannot specify it directly; **elm** provides it as necessary and handles messages that contain it. The value is placed in a **Content-Type:** body-part header entry. The default is:

text/plain; charset=US-ASCII

Some common entries are described below. See the METAMAIL CONFIGURATION section for additional information.

text/ subtype [**;** **charset=charset**]

This is relatively readable text that may be formatted with embedded text characters, as for possible subtypes **richtext** or **html**. The default *subtype* is **plain** (unformatted in any way). The default *charset* is **US-ASCII**.

application/octet-stream

This is a catch-all for files such as program binary, or files that contain control characters or characters with high-order bits set.

application/postscript

The file can be displayed with a PostScript-equipped printer or viewer.

message/rfc822

This specifies that the file is in message format, as described in the Terminology subsection.

image/jpeg, **image/gif**

These are picture formats that require a display program.

audio/basic

This is an audio format that requires a reproduction program.

video/mpeg

This is an audio/video format that requires a reproduction program.

Attachment Send Menu

The Attachment Send Menu is invoked with the Send Menu **a** command. This menu displays a list of the attachments that will be sent in a message, one per line, as described in the Attachment Index subsection.

The commands are:

- a** Add attachments. Call the Attachment Configuration Menu and prompt for a file name.
- d** Delete an attachment.
- e** Edit an attachment. Call the editor associated with the attachment if it is editable.
- j** Move the current attachment pointer down to the next attachment.
- k** Move the current attachment pointer up to the previous attachment.
- Ctrl-L** Redraw the screen.
- m** Modify the attributes of an attachment. Call the Attachment Configuration Menu.
- p** Print an attachment. See the Message Menu **p** (print) command.
- q** Quit. Return to the Send Menu.
- s** Save an attachment. See the Message Menu **C** (copy) command.

Attachment View Menu

The Attachment View Menu is invoked with the Send Menu **v** command. This menu displays a list of the attachments in a folder message, one per line, as described in the Attachment Index subsection.

The commands are:

| | |
|---------------|---|
| Return | Display the current attachment. |
| j | Move the current attachment pointer down to the next attachment. |
| k | Move the current attachment pointer up to the previous attachment. |
| Ctrl-L | Redraw the screen. |
| p | Print the attachment. See the Message Menu p (print) command. |
| q | Quit. Return to the previous attachment level or the Message Menu. |
| s | Save the attachment. The attachment is saved in the form it was received, as with the Message Menu s (save) command. |
| v | View the subattachment list, if any. |

Attachment Index

Attachments are listed on the Attachment Send Menu and the Attachment View Menu in the following format:

```
num filename (size) format [encoding]
```

defined as:

| | |
|---------------------|---|
| <i>num</i> | The index number of the attachment. |
| <i>filename</i> | The name of the attached file. |
| <i>size</i> | The size of the attachment in bytes, computed from the file or the message. |
| <i>type/subtype</i> | The type and subtype of the attachment. This value is placed or found in a Content-Type: header. |
| <i>encoding</i> | The encoding type. This value is placed or found in a Content-Transfer-Encoding: header. |

Header Menu

The Header Menu is invoked with the Send Menu **h** command. It allows you to add, change, and delete a set of common header entries in your message. In general, if an item is empty, it is not included in the message.

The commands are:

| | |
|-----------------|--|
| Return | Return to Send Menu. |
| !command | Shell. Execute <i>command</i> with the shell defined by the shell string variable. |
| a | Action: header. Enter any string. If this entry is present in a received message, elm displays an A in the Permanent Status column of the Message Index. |
| b | Bcc: header. Enter a list of aliases and actual addresses. Aliases are expanded and shown as addresses and user names. |
| c | Cc: header. Enter a list of aliases and actual addresses. Aliases are expanded and shown as addresses and user names. |
| d | Domainize. Convert non-Internet addresses to Internet format. The UUCP ! format (<i>host.domain!user</i>) becomes the Internet @ format (<i>user@host.domain</i>). If <i>.domain</i> is omitted, it defaults to .uucp . |
| e | Expires: header. Enter any numeric value from 1 to 56 (8 weeks). If this entry is present in a received message, elm displays an E in the Variable Status column of the Message Index when the computed date has passed. |
| i | In-Reply-To: header. Enter a string. |
| n | Precedence: header. Enter a precedence name. If the precedences string variable is set and nonnull, the name must be one defined by the variable. If the |

name is associated with a priority, and the **Priority:** header is null, the priority value is inserted in the **Priority:** header. If **precedences** is null or not set, you can enter any value.

If the precedence name matches one defined in the **sendmail** configuration file **/etc/mail/sendmail.cf**, the transmission priority is modified accordingly. If there is no match, the priority is not changed.

- p** **Priority:** header. Enter a string. If this entry is present in a received message, **elm** displays a **U** in the Permanent Status column of the Message Index
- r** **Reply-To:** header. Enter a personal alias or a single address. If it is present, **elm** and other mailers use this header instead of the **From:** header when choosing the address for a reply (Message Menu **r** (reply) command).
- s** **Subject:** header. Enter a string.
- t** **To:** header. Enter a list of aliases and actual addresses. Aliases are expanded and shown as addresses and user names.
- u** User-defined header. Define your own header entry in the form:

header-name: header-string

header-name: must not contain blanks. You can use this command to create a **Sensitivity:** header entry, as described in the Message Status subsection, or a different header, but only one. See the **HEADER FILE** section for another way to include user-defined header entries.

Options Menu

The Options Menu is invoked with the Message Menu **o** command. It displays a list of the options, defined by the **configoptions** string variable, that you can modify while **elm** is running. Enter the appropriate letter (in upper- or lowercase) that is followed with a right parenthesis (**)** and follow the directions on the screen. The full set of option prompts and the corresponding variables is listed below. The default options are marked with an *****.

| | |
|------------------------------|--|
| A)rrow cursor | The arrow string variable. * |
| B)order on copy | The prefix string variable. |
| C)alendar file | The calendar string variable. * |
| D)isplay mail using | The pager string variable. * |
| E)ditor (primary) | The editor string variable. * |
| F)older directory | The maildir string variable. * |
| H)old sent message | The copy boolean variable. |
| J) reply editor | The altditor string variable. |
| K) pause after pager | The promptafter boolean variable. |
| A (l)ias Sorting | The aliassortby string variable. |
| M)enu display | The menu boolean variable. * |
| N)ames only | The names boolean variable. * |
| O utbound mail saved | The sentmail string variable. * |
| P)rint mail using | The print string variable. * |
| R)epley copies msg | The autocopy boolean variable. |
| S)orting criteria | The sortby string variable. * |
| T)ext editor (~e) | The easyeditor string variable. |
| U)ser level | The userlevel numeric variable. * |
| V)isual Editor (~v) | The visualeditor string variable. * |
| W)ant Cc: prompt | The askcc boolean variable. |
| Y)our full name | The fullname string variable. * |
| Z) signature dashes | The sigdashes boolean variable. |

Note: The menu displays the first *screen-height-6* lines from the defined set. *screen-height* is the number of text lines on the screen. If an option is not displayed, it cannot be modified.

When you are done, enter one of the following values:

- > Save the current configuration values in your configuration file, **\$HOME/.elm/elmrc**. If the file does not exist, it is created. This is a convenient way to make an configuration file that you can edit directly, as well as with the Options Menu.

i,I Return to the Message Menu.

q,Q Return to the Message Menu.

x,X Exit immediately from **elm** without changing the mailbox. All pending changes are abandoned.

Send Menu

The Send Menu is invoked when your outbound message has been prepared to be mailed after a Message Menu **f**, **g**, **m**, or **r** command or the Alias Menu **m** command.

The commands are:

- !** *command* Shell. Execute a shell command. See the Message Menu **!** (shell) command.
- a** Attachments. Invoke the Attachments Send Menu.
- c** Copy. Copy to a file. See the Message Menu **C** (copy) command.
- e** Edit. Invoke your editor, as defined by the **altditor** string variable, to revise the message.
- f** Forget. Do not send the message. At user levels **1** and **2**, the message may be returned to the send buffer when you execute a subsequent Message Menu **f**, **g**, **m**, or **r** command or the Alias Menu **m** command.
- h** Edit the header entries. Invoke the Header Menu.
- m** Make form. Convert the message to the forms message format. See the FORMS MESSAGES section. This command is only available if the **forms** boolean variable is **ON** and the **userlevel** numeric variable is either **1** or **2**.
- s** Send. Send the message.

FORMS MESSAGES

A feature that is unique to **elm** is the ability to compose and reply to forms messages.

Creating a Forms Message

- In your **elmrc** file, set **forms=ON**.
- Set your **userlevel** numeric variable to **1** (moderately experienced) or **2** (expert). You can do this in your **elmrc** file or on the default Options Menu.
- As you compose the message, define the fields to be filled in by the recipient with a colon (:), followed by either the number of spaces allowed for the field value, or a newline to indicate that the field may fill the remainder of the line.

A colon on a line by itself indicates that the recipient will be prompted for multiline input. There can be no blanks before the colon.

Every line containing a colon is a prompt line. During the response process, all text starting at the first nonblank character after the last colon on each line is deleted and the line is evaluated for response fields.

- After you have created the message, enter the Send Menu **m** (make form) command to set up the special format. Then enter the Send Menu **s** (send) command to send the message.

Here is an example of a simple forms message:

```
On-Line Phone and Address Database
Please fill out and return as soon as possible.
Name:
Manager:
Department:                               Division:
Your home address:
Home phone number:
Thank you for your cooperation.
```

Replying to a Forms Message

When you receive a forms message, the message index entry is flagged with an **F** status letter. You can view it in the normal way with the **Return** or **h** commands.

To reply, use the Message Menu **r** (reply) command (you cannot use the Message Menu **g** (group reply) command). **elm** prompts you for each field, with any text present between the fields displayed as appropriate. The example above is presented line-by-line; user input is in italic type:

```
On-Line Phone and Address Database
Please fill out and return as soon as possible.
Name: my name
Manager: my manager
Department: my department
Division: my division
Your home address: home address
Home phone number: phone number
Thank you for your cooperation.
```

The received message would look like this:

```
On-Line Phone and Address Database
Please fill out and return as soon as possible.
Name: my name
Manager: my manager
Department: my department
Division: my division
Your home address: home address
Home phone number: phone number
Thank you for your cooperation.
```

HEADER FILE

The `$HOME/.elm/elmheaders` file provides you with a way to specify special information headers such as **X-Organization:**, **X-Phone:**, and so forth. The nonblank lines from this file are added to the headers of all outbound mail.

Entries in the **elmheaders** file should have the following format:

```
header-name: header-string
```

header-name: must not contain blanks. *header-string* can be continued over several lines by preceding each continuation line with blanks, as indicated by the output below.

Within the **elmheaders** file, you can use backquotes (left apostrophes) to execute shell commands when the file is read, so that an entry of the form:

```
X-Operating-System: `uname -a`
```

would produce a header entry like:

```
X-Operating-System:
  HP-UX hpulpc17 B.10.10 A 9000/710 2012505939 two-user license
```

According to RFC 822, user-defined header names should begin with **X-** or **x-**. Otherwise, they risk having their usage overridden if the name is later standardized with some other meaning.

Defined Headers

The following header names are defined for the message header in RFC 822 and RFC 1521.

| | | | |
|-------------------------|--------|-----------------------------------|--------|
| Bcc: | (822) | Cc: | (822) |
| Comments: | (822) | Content-Description: | (1521) |
| Content-ID: | (1521) | Content-Transfer-Encoding: | (1521) |
| Content-Type: | (1521) | Date: | (822) |
| Encrypted: | (822) | From: | (822) |
| In-Reply-To: | (822) | Keywords: | (822) |
| MIME-Version: | (1521) | Message-ID: | (822) |
| Received: | (822) | References: | (822) |
| Reply-To: | (822) | Resent-Bcc: | (822) |
| Resent-Cc: | (822) | Resent-Date: | (822) |
| Resent-From: | (822) | Resent-Message-ID: | (822) |
| Resent-Reply-To: | (822) | Resent-Sender: | (822) |
| Resent-To: | (822) | Return-Path: | (822) |
| Sender: | (822) | Subject: | (822) |
| To: | (822) | X-user-defined: | (822) |

Other commonly used headers:

| | |
|-----------------------------|------------------------|
| Action: | Apparently-To: |
| Content-Disposition: | Content-Length: |
| Expires: | Mailer: |
| Newsgroups: | Precedence: |
| Priority: | Sensitivity: |
| Status: | X-Mailer: |

ELM CONFIGURATION

elm supports user configuration by means of the `$HOME/.elm/elmrc` configuration file. You can create the configuration file with the Options Menu `>` command. It can contain any combination of the string, numeric, and boolean variables described below.

String Variables

String variables have the form

string-name = *string-value*

The following string variables are defined.

| | |
|-----------------------|---|
| aliassortby | The sort order for the alias index in the Alias Menu. The recognized values are: alias Sort by alias name. name Sort by the full name of the alias, last name first. text Sort by the order of the aliases in the alias text file. Prefix the value with reverse- to reverse the sort order. The default is name . |
| altditor | The name of the editor to use for messages that have initial text (a copied message in a reply, a signature in any outbound message, etc.). when the editor string variable is set to none or builtin . The default is the value of the EDITOR environment variable, if set and nonnull, or <code>/usr/bin/vi</code> otherwise. See also the editor string variable. |
| alternatives | A list of other machine and user name combinations that you receive forwarded mail from. elm uses this information when a group reply is being processed to ensure that a reply message is not sent to a user and/or machine address that would simply forward the reply message back to you. The default is none. |
| attribution | Attribution string for replies. When you reply to a message and include the message in the reply, this string is placed at the top of the message. The characters %s are replaced by the full name of the author of the original message. The default is none. For example: attribution = %s wrote: |
| calendar | The name of your calendar file. This is used in conjunction with the Message Menu <code><</code> (calendar) command, which scans messages for calendar entries. The default is <code>\$HOME//calendar</code> . |
| charset | The name of the character set used with the MIME Content-Type: header for the text/plain type. It can be any Internet-defined character set name that is a superset of US-ASCII . The default is US-ASCII . For example, Content-Type: text/plain; charset=US-ASCII |
| compatcharsets | A list of Internet-defined character sets that are supersets of US-ASCII , so that messages with charset=US-ASCII can be displayed without the help of metamail . The default is a string containing the following values: Extended_UNIX_Code_Packed_Format_for_Japanese ISO-2022-JP ISO-8859-1 ISO-8859-2 ISO-8859-3 |

| | |
|-----------------------|---|
| | ISO-8859-4 ISO-8859-5 ISO-8859-7 ISO-8859-8 ISO-8859-9 KOI8-R Shift_JIS |
| configoptions | <p>A string of options that can be configured on the Options Menu. Specify the options as a single letter each, in the order they should be displayed. The default is "^_cdefsopyv_am_un". The defaults are marked below with an *.</p> <p>The option characters include:</p> <ul style="list-style-type: none"> ^ The menu title. _ A blank line. a The arrow string variable. * b The prefix string variable. c The calendar string variable. * d The pager string variable. * e The editor string variable. * f The maildir string variable. * h The copy boolean variable. j The altditor string variable. k The promptafter boolean variable. l The aliassortby string variable. m The menu boolean variable. * n The names boolean variable. * o The sentmail string variable. * p The print string variable. * r The autocopy boolean variable. s The sortby string variable. * t The easyeditor string variable. u The userlevel numeric variable. * v The visualeditor string variable. * w The askcc boolean variable. y The fullname string variable. * z The sigdashes boolean variable. |
| displaycharset | <p>The name of the character set supported by the display. This is independent of the charset string variable. This is also copied to the MM_CHARSET environment variable when metamail is called. The default is US-ASCII.</p> |
| easyeditor | <p>The name of the editor for the ~e command of the built-in editor. See also the editor string variable. The default is none.</p> |
| editor | <p>The name of the editor to use when creating new mail. The default is the value of the EDITOR environment variable, if set and nonnull, or /usr/bin/vi otherwise.</p> <p>You can use none or builtin to specify the built-in editor. The built-in editor is available for all outbound mail that does not already have text in the send buffer (no forwarded message, no copied message in a reply, no signature in any outbound message, etc.). If there is text in the send buffer and builtin is specified, the editor defined by the altditor variable is used instead.</p> <p>See also the altditor, easyeditor, and visualeditor string variables.</p> |
| escape | <p>The escape character used in the built-in editor. The default is tilde (~).</p> |
| fullname | <p>The name the mailer will use when sending mail from you. The default is the full name portion (everything up to the first comma) of the pw_gecos field from your entry in the /etc/passwd file. This field can be set with</p> |

the **chfn** command (see *chfn(1)*, *finger(1)*, and *passwd(4)*).

localsignature

A signature file that is automatically appended to outbound mail to the local host before the editor is invoked. This usually contains personal data about the sender. See also the **remotesignature** string variable. The default is none.

All the addresses in the **To:** header must be apparently for the local host. Local addresses are those that, after any **elm** alias conversion, do not contain a domain name. That is, they have only a user name (for example, **santaclaus**) or a user name and the local host name (for example, **santaclaus@northpole**).

santaclaus@northpole.arcticsea.org is considered to be a remote address, even if it points to the local host. A user name that is readdressed by the **sendmail** system alias list is treated as local if it matches the preceding criteria.

maildir

Your mail directory, where you usually store your folders for received and outbound mail. The default is **\$HOME//Mail**.

In **elm**, you can use the **=** metacharacter to specify this directory. For example, if you save a message to file **=/archive**, the **=** is expanded to the current value of **maildir**. (The slash (/) is optional.)

When you start **elm**, if the directory specified by **maildir** does not exist, you are asked if you want to create it. If you answer **y** (yes), the directory is created, with access permissions set to **700**.

pager

The program to display each message. The default is the value of the **PAGER** environment variable, if set and nonnull, or the built-in pager, **builtin+**, otherwise.

The built-in pager, **builtin+**, also allows you to execute some Message Menu commands while you are viewing the message and it has some simple forward and backward scrolling commands. While it is active, enter **?** for a list of commands. An alternative is the **more** utility.

precedences

A list of precedence values that you can place in a **Precedence:** header entry in outbound mail, using the Header Menu. Each precedence value can be optionally paired with a priority value that is automatically placed in a **Priority:** header entry, causing the received message to be marked as urgent. The default is none.

The HP-UX mail transport agent, **sendmail**, recognizes this header. If the precedence value is defined by a **P** control line in the **sendmail** configuration file, **/etc/mail/sendmail.cf**, the transmission priority of the message is adjusted accordingly. See *sendmail(1M)*.

The format of the entry is

```
precedences = precedence[:priority] [precedence[:priority]] ...
```

precedence is a precedence name. The default list defined in **/etc/mail/sendmail.cf** is:

| | |
|-------------------------|--------------------------------------|
| first-class | Transmission priority 0, the default |
| special-delivery | Transmission priority 100 |
| list | Transmission priority -30 |
| bulk | Transmission priority -60 |
| junk | Transmission priority -100 |

priority is an arbitrary string that is placed in a **Priority:** header entry.

prefix

The prefix for an included line in an outbound message. When you reply to a message or forward a message to another person, you can optionally include the original message. This prefix marks the included line. The default is **>_** (the **_** is interpreted as a space character).

print

The command to run when the **p** (print) command is executed from various menus. There are two possible formats for this string: If the string

contains the special variable **%s**, the variable is replaced by the name of a temporary file that contains the messages, and the command is executed by the shell defined by the **shell** string variable. If the string does not contain **%s**, the temporary file name is appended to it, and the command is executed. The default is

```
cat %s | lp
```

receivedmail The file where the received messages will be saved. The default is **=received**, the file **received** in the directory defined by **maildir**.

remotesignature A signature file that is automatically appended to all outbound mail to remote hosts before the editor is invoked. This usually contains personal data about the sender. See also the **localsignature** string variable. The default is none.

If any of the addresses in the **To:** header entry are not local, as described for the **localsignature** string variable, the remote signature file is attached.

savecharset The character set to be used to save a message in a folder. Possible values are **JIS**, **SJIS**, and **EUC**. If a value is not specified, the message will be saved according to your locale (given by the **LC_TYPE** and/or **LANG** environmental variables). This option is applicable only for the Japanese locale. The default is none. See also the **jisconversion** boolean variable.

sentmail The file where copies of outbound mail can be saved. One possibility is your incoming mailbox, **/var/mail/loginname**. The default is **=sent**, the file **sent** in the directory defined by **maildir**.

See the **copy** boolean variable for further details.

shell The shell to use for **!** escapes and other such operations. The default is the value of the **SHELL** environment variable, if set and nonnull, or **/usr/bin/ksh** otherwise.

sortby The way to sort the index of the current folder. The choices are:

| | |
|-----------------|--|
| from | The name of the sender. |
| sent | The date the message was sent. |
| received | The date the message was received. |
| subject | The subject of the message. A leading Re: (and some others) is ignored, so replies sort with original messages. |
| lines | The number of lines in the message. |
| status | The read status: blank, O, and N. |

You can prefix these values with **reverse-** to reverse the order of the sort. The value can be modified on the Options Menu. The default is **reverse-sent**.

textencoding Type of encoding to put into the **MIME Content-Transfer-Encoding:** header entry. The choices are **7bit** or **8bit**. The default is **7bit**.

tmpdir Where to create temporary files. The default is the value of the **TMPDIR** environment variable, if set and nonnull, or to **/tmp/** otherwise.

visualeditor Name of the editor to use for the **~v** command of the built-in editor. The default is the value of the **VISUAL** environment variable, if set and nonnull, or **/usr/bin/vi** otherwise.

weedout A list of header-entry initial strings that you don't want to see when you are reading mail. This list is made effective by setting the **weed** boolean variable to **ON**.

The list can continue for as many lines as desired, as long as the continued lines all have leading blanks. To include blanks in a string, enclose it in

quotation marks ("). The strings you specify are normally appended to the default list, which is:

```
>From
Apparently-To:
Content-Length
Content-Transfer-Encoding
Content-Type:
From
In-Reply-To:
MIME-Version
Mailer:
Message-Id:
Newsgroups:
Received:
References:
Status:
X-Mailer:
```

There are two special values:

clear-weed-list

Clear the default list. The default headers are removed from the **weedout** list, allowing you to completely define your own list.

end-of-user-headers

Mark the end of the **weedout** list, in case any following lines could be mistaken for headers in the list.

The default value of **weedout** is ***end-of-user-headers***.

The underscore (_) character can be used to specify a space.

Note that **From** weeds out both **From** and **From:.** If, for example, you want to weed out **From** but not **From:.**, specify ***clear-weed-list*** followed by **From_** and any other headers that you don't want to see.

Numeric Variables

Numeric variables have the form

numeric-name = *numeric-value*

The following numeric variables are defined.

| | |
|---------------------|--|
| bounceback | Threshold for returning copies of remote UUCP messages. If the destination host is greater than the specified number of hops (gateways) from your local host, the destination host sends you a copy of the message when it is received. If the value is 0, this feature is disabled. The default is 0. |
| builtinlines | <p>Determines if the built-in pager should be used on some messages, even if you usually use an external pager, defined in the pager string variable. There are two ways of defining whether the built-in pager should be used.</p> <ul style="list-style-type: none"> • If you want to use the built-in pager on any message that is shorter than <i>n</i> lines, set the value to <i>n</i>. • If you want to use the built-in pager on any message that is <i>m</i> lines shorter than the number of lines on your screen, set the value to <i>-m</i>. <p>If you set the value to 0, the message will always be sent through the external pager. The default is -3.</p> |
| noencoding | <p>This enables you to send raw 8-bit or binary data when the mail transfer agent doesn't support the 8BITMIME and the -B8BITMIME options. The default is 0.</p> <p>The possible values are:</p> <ul style="list-style-type: none"> 0 Always convert 8-bit and binary messages to 7-bit before sending them. |

| | | |
|-------------------|---|---|
| | 1 | Convert 8-bit messages to 7-bit, but depend on sendmail to handle binary messages. |
| | 2 | Depend on sendmail to handle both 8-bit and binary messages. |
| readmsginc | | The value by which the Reading in folder, message: counter is incremented while reading a new folder. If you set this value to a number larger than one, it will speed up the time it takes to read a large folder when you are using a slow terminal. The default is 1. |
| sleepmsg | | The time in seconds that elm will wait after displaying a diagnostic message before erasing it. The value can be 0 or a positive integer. The default is 2. |
| timeout | | The interval, in seconds, after which elm rechecks the incoming mailbox for new mail. The default is 600 (10 minutes). |
| userlevel | | The relative level of your user sophistication. Acceptable values are: |
| | 0 | Novice user (the default). Command menus are a small verbose subset of the available commands. |
| | 1 | Moderately experienced user. Command menus are a larger terse subset of the available commands. Outbound message commands allow you to recover previously unsent messages as the text of the current outbound message. |
| | 2 | Expert. The features are the same as for 1. |
| | | Level 1 or 2 is required if you want to send a forms message. |

Boolean Variables

Boolean variables have the forms

boolean-name = **ON** -and- *boolean-name* = **OFF**

The following boolean variables are defined.

| | |
|---------------------|---|
| alwaysdelete | If ON , the default answer of the Message Menu q (quit) prompt Delete messages? (y/n) is set to y (yes). If OFF , the default answer is set to n (no). The default is OFF . See the Message Menu q command. |
| alwayskeep | If ON , the default answer of the Message Menu q (quit) prompt Keep unread messages in incoming mailbox? (y/n) is set to y (yes). If OFF , the default answer is set to n (no). The default is ON . See the Message Menu q command. |
| alwaysstore | If ON , the default answer of the Message Menu q (quit) prompt Move read messages to "received" folder? (y/n) is set to y (yes). If OFF , the default answer is set to n (no). The default is OFF . See the Message Menu q command. |
| arrow | If ON , use an arrow (->) to mark the current item in a menu index. If OFF , use an inverse bar. If the program is invoked with the -a command line option, arrow is set to ON . The default is OFF . |
| ask | If ON , you are asked the questions Delete messages? (y/n) Move read messages to "received" folder? (y/n) Keep unread messages in incoming mailbox? (y/n) (as appropriate) each time you leave the program with the Message Menu q (quit) command. See that command for details of the process. If OFF , or if you use the Message Menu Q command, elm uses the values defined by the alwaysdelete , alwaysstore , and alwayskeep boolean variables, respectively, without prompting. The default is ON . |

| | |
|-----------------------|--|
| askcc | If ON , elm prompts you for "carbon copies" with the prompt Copies To: each time you send, forward, or reply to a message. If OFF , the prompt is omitted. In either case, you can still explicitly include Cc: addresses with the ~c command in the built-in editor, or with the Header Menu commands. The default is ON . |
| autocopy | If ON , elm automatically copies the text of the message you are replying to into the edit buffer. If OFF , elm prompts you with Copy message?. The default is OFF . |
| confirmappend | If ON , you are asked to confirm before messages are appended to an existing file, whether it is a file in your mail directory or a file in another directory. If OFF , see confirmappend and confirmfiles Operation below. The default is OFF . |
| confirmcreate | If ON , you are asked to confirm before a new file is created. whether it is a file in your mail directory or a file in another directory. If OFF , see confirmcreate and confirmfolders Operation below. The default is OFF . |
| confirmfiles | If ON , you are asked to confirm before messages are appended to an existing file that is not in your mail directory. This does not affect files in your mail directory. If OFF , see confirmappend and confirmfiles Operation below. The default is OFF . |
| confirmfolders | If ON , you are asked to confirm before a new file is created in your mail directory. This does not affect files in other directories. If OFF , see confirmcreate and confirmfolders Operation below. The default is OFF . |

confirmcreate and confirmfolders Operation

confirmcreate=ON and confirmfolders=ON

Confirm before creating a file in your mail directory.
Confirm before creating a file another directory.

ON and OFF Confirm before creating a file in your mail directory.
Confirm before creating a file another directory.

OFF and ON Confirm before creating a file in your mail directory. Do
not confirm before creating a file another directory.

OFF and OFF Do not confirm before creating a file in your mail directory.
Do not confirm before creating a file another directory.

confirmappend and confirmfiles Operation

confirmappend=ON and confirmfiles=ON

Confirm before appending to a file in your mail directory.
Confirm before appending to a file in another directory.

ON and OFF Confirm before appending to a file in your mail directory.
Confirm before appending to a file in another directory.

OFF and ON Confirm before appending to a file in your mail directory.
Do not confirm before appending to a file in another directory.

OFF and OFF Do not confirm before appending to a file in your mail directory.
Do not confirm before appending to a file in another directory.

copy If **ON**, save silent copies of all outbound mail on the outbound step. If **OFF**, do not save copies. The default is **OFF**.

If **ON**, and the **savename** boolean variable is **ON**, **elm** first tries to save it to a file named as defined by **savename**. If the file exists, the message is saved. If the file does not exist, but the **forcename** boolean variable is **ON**, the file is created and the message is saved. If **forcename=OFF**, the message is saved to the file defined by the **sentmail** string variable. If **savename=OFF**, the message is saved to the file defined by the

| | |
|----------------------|--|
| | sentmail string variable. |
| forcename | If ON , create the folder when saving outbound messages by the login name of the recipient, even if the folder doesn't already exist. If OFF , do not create the folder. The default is OFF . See the copy boolean variable for further details. |
| forms | If ON , and the userlevel numeric variable is 1 or 2, you can create a forms message. The Send Menu m (make form) command converts your message into a forms message. If OFF , you cannot. The default is ON . |
| jisconversion | If ON , convert outbound mail to JIS (Japanese Industrial Standard) before sending it. If OFF , do not convert it. This option is applicable only to the Japanese locales, ja_JP.SJIS and ja_JP.eucJP . The default is OFF . savecharset string variable. |
| keepempty | If ON , keep folders from which all the messages are deleted. If OFF , delete empty folders. The default is OFF . |
| keypad | If ON , enable the HP 2622 terminal cursor keys. If OFF , disable the cursor keys. If the program is invoked with the -K command line option, keypad is set to OFF . See also the softkeys boolean variable. The default is ON . |
| menu | If OFF , this inhibits the menu display on all program screen displays. If ON , the menus are displayed. If the program is invoked with the -m command line option, menu is set to OFF . The default is ON . |
| metoo | If ON , you are sent a copy of the message that you send to an alias that includes your name also. If OFF , the copy is not sent. The default is OFF . |
| mimeforward | If ON , a forwarded message is sent as an attachment. If OFF , a forwarded message is sent as part of the main message. The default is ON . |
| movepage | If ON , commands that move through the mailbox by pages (the + and - commands) also move the current index pointer to the top of the new page of messages. If OFF , moving through the pages does not alter the current message pointer location. The default is ON . |
| names | If ON , show only the user names when expanding the To: aliases for an outbound message. If OFF , show the entire expanded addresses. The default is ON . |
| nohdrencoding | If ON , don't do RFC 1522 encoding for header values that contain 8-bit or multibyte characters. If OFF , do the encoding. The default is OFF . |
| noheader | If ON , do not include the headers of messages when copying a message into a file buffer for replying to or forwarding. If OFF , copy all headers. The default is ON . |
| noheaderfwd | If ON , do not include headers when copying a message into a file buffer for forwarding. If OFF , copy all headers. For forwarding, this option overrides the setting of noheader . The default is OFF . |
| pagemultipart | If ON , use the value of the pager variable to display MIME multipart messages with unknown subparts or with unknown subtypes. If OFF , call metamail to view multipart messages. The default is OFF . |
| pointnew | If ON , automatically point to the first new message in your message index at start-up. If OFF , point to the first message. In either case, if the start-up folder is not your incoming mailbox, or if there are no new messages, point to the first message. The default is ON . |
| promptafter | If ON , prompt for a command after the external pager exits. If OFF , return to the calling menu. The default is ON . |
| resolve | If ON , move the pointer to the next message in the index, after deleting, undeleting, saving, or forwarding a message. If OFF , keep the pointer at the current message. The default is ON . |

| | |
|------------------|---|
| savename | <p>If ON, and you are saving a message, elm constructs a suggested file name in your maildir directory from the user name of the person who sent the message, in the form =username. If OFF, no file name is suggested.</p> <p>If ON, and you are sending a message that will be saved, elm constructs a file name based on the user name of the first entry in the To: list, in the same form as above. If OFF, no file name is constructed. See the copy boolean variable for further details.</p> <p>The default is ON.</p> |
| sigdashes | <p>If ON, insert two dashes above the signature text, included from a local or remote signature file. This is a common convention. If OFF, omit the dashes. The default is ON.</p> |
| softkeys | <p>If ON, enable the HP 2622 terminal function-key protocol. If OFF, disable the function-key protocol. If the program is invoked with the -k or -K command line option, softkeys is set to OFF. See also the keypad boolean variable. The default is OFF.</p> |
| titles | <p>If ON, title a displayed message with a line in the form:</p> <p style="text-align: center;">Message <i>number/total</i> <i>sendername</i> <i>date</i> <i>time</i></p> <p><i>sendername</i>, <i>date</i>, and <i>time</i> are extracted from the message headers in the manner described in Message Index. This is useful if you have suppressed the relevant header entries with the weedout list. If OFF, the message is not titled. The default is ON.</p> |
| usetite | <p>If ON, use the termcap ti/te and terminfo cup cursor-positioning entries (see terminfo(4)). If OFF, do not use those entries. If the program is invoked with the -t command line option, usetite is set to OFF. The default is ON.</p> |
| weed | <p>If ON, do not display the headers defined by weedout variable when displaying a message for reading. If OFF, display all headers. The default is ON.</p> |

METAMAIL CONFIGURATION

MIME (Multipurpose Internet Mail Extensions) encoding classifies the message and its attachments according to a Content-Transfer-Encoding, which is the encoding, if any, that is used to make the message mailable, and a Content-Type, which is the type and form of the message part after it has been decoded. The encoding and types are described in more detail in the Attachment Configuration Menu subsection and in RFC 1521.

elm provides built-in support for the following Content-Types:

text/plain [**;** **charset=charset**]

The text is all in the displayable character set *charset* which defaults to **US-ASCII**.

multipart/mixed **;** **boundary=boundary-string**

The message is composed of a number of individual "body parts", separated by **--boundary-string**, each having optional headers defining Content-Type and Content-Transfer-Encoding. The default Content-Type is **text/plain**.

multipart/digest **;** **boundary=boundary-string**

This is similar to **multipart/mixed**, except that the default Content-Type is **message/rfc822**.

multipart/report **;** **boundary=boundary-string**

message/rfc822

The message consists of another message in standard message format.

metamail is a system program that is invoked by **elm** to manage the display of messages and attachments that are not displayable in ordinary ASCII text.

metamail provides external support for other Content-Types, as defined in one or more **mailcap** files. The system **mailcap** file is **/etc/mail/mailcap**. You can define your own default **mailcap** file in **\$HOME/.mailcap**. You can also specify your own list of **mailcap** files by setting the **MAILCAPS** environment variable. The **mailcap** files are searched in order until an entry is found that matches the Content-Type and any qualifications.

A minimum **mailcap** entry consists of a line in the form:

```
content-type ; command
```

The *command* is the command that you would type to view a file of the indicated Content-Type, with the string **%s** replaced by a file name. For example, to view body part that was HTML source text and had the Content-Type **text/html**, you could have the entry

```
text/html; netscape %s
```

Similarly, for a GIF image file, you could have the entry

```
image/gif; xv %s
```

RFC 1521 defines a number of Content-Types that **elm** leaves for **metamail** to handle:

```
text/richtext
multipart/alternative
multipart/parallel
multipart/digest
message/partial
message/external-body
image/jpeg
image/gif
audio/basic
video/mpeg
application/octet-stream
application/postscript
```

Check the system **mailcap** file for entries that handle many of them.

EXTERNAL INFLUENCES

Environment Variables

| | |
|-----------------|---|
| HOME | Your home (login) directory. |
| EDITOR | If set and nonnull, provides a default value for the altditor and editor string variables. |
| LANG | If set and nonnull, determines the language in which messages are displayed. The default is C . See <i>environ</i> (5). |
| MAILCAPS | If set, defines the search path for mailcap files used by metamail . The default is \$HOME/.mailcap:/etc/mail/mailcap |
| PAGER | If set and nonnull, provides a default value for the pager string variable. |
| SHELL | If set and nonnull, provides a default value for the shell string variable. |
| TMPDIR | If set and nonnull, provides a default value for the tmpdir string variable. |
| VISUAL | If set and nonnull, provides a default value for the visualeditor string variable. |

International Code Set Support

Single- and multibyte character code sets are supported.

EXAMPLES

Message Mode Example

To send a message without loading the main **elm** mail-processing program, use the simple command form consisting of the name of the program followed by the recipient's login name and optional address. **elm** prompts for subject and copies, then starts an editor so you can compose the message (user responses are in italic type):

```
$ elm j_doe
To: doe (John Doe)
```

```

Subject:  this is a test
Copies To: ...
...invokes editor, you compose message, then...
Your options now are:
a)ttachments e)dit message edit h)header s)end it f)orget it.

```

```

What is your choice? s
mail sent!

```

If you "forget" the message, it is saved in `$HOME/Canceled.mail`.

File Mode with Redirection

To send a file by use of command-line redirection, use a command like:

```
$ elm j_doe < help.c
```

which reads file `help.c` and sends it with the default subject.

File Mode with a Pipe

To mail the output of a command and include a subject line:

```
$ ls -a | elm -s "Directory Listing" j_doe
```

WARNINGS

Using two separate mail programs to access the same mail file simultaneously (usually inadvertently from two separate windows) can cause unpredictable results.

AUTHOR

elm was developed by Hewlett-Packard Company.

FILES

| | |
|--|---|
| <code>\$HOME/.elm</code> | Directory for the user's elm alias, configuration, header, and other files |
| <code>\$HOME/.elm/aliases</code> | User alias database data table |
| <code>\$HOME/.elm/aliases.dir</code> | User alias database directory table |
| <code>\$HOME/.elm/aliases.pag</code> | User alias database hash table |
| <code>\$HOME/.elm/aliases.text</code> | User alias source text |
| <code>\$HOME/.elm/elheaders</code> | User-defined additional headers |
| <code>\$HOME/.elm/elmrc</code> | User configuration file |
| <code>\$HOME/Canceled.mail</code> | Canceled message in noninteractive use. |
| <code>/tmp/alias.pid</code> | Temporary file for deleting alias |
| <code>/tmp/form.pid</code> | Editor buffer for forms message |
| <code>/tmp/mbox. loginname</code> | Temporary mailbox for user <i>loginname</i> |
| <code>/tmp/print.pid</code> | Temporary file for printing message |
| <code>/tmp/snd.pid</code> | Outgoing mail message edit buffer |
| <code>/tmp/sndh.pid</code> | Outgoing mail header edit buffer |
| <code>/usr/lib/nls/msg/C/elm.cat</code> | Location of the message catalog |
| <code>/usr/share/lib/elm/elmrc-info</code> | Comment file for <code>\$HOME/.elm/elmrc</code> file |
| <code>/var/mail</code> | Directory for incoming mail; it must have mode 755 and group ID mail |
| <code>/var/mail/.elm</code> | Directory for elm mailer system aliases |
| <code>/var/mail/.elm/aliases</code> | System alias database data table |
| <code>/var/mail/.elm/aliases.dir</code> | System alias database directory table |
| <code>/var/mail/.elm/aliases.pag</code> | System alias database hash table |
| <code>/var/mail/.elm/aliases.text</code> | System alias source text |
| <code>/var/mail/ loginname</code> | Incoming mailbox for user; it must have mode 660 and group ID mail |
| <code>/var/mail/ loginname.lock</code> | Lock for mail directory |

SEE ALSO

answer(1), chfn(1), elmalias(1), fastmail(1), finger(1), mailfrom(1), newalias(1), newmail(1), readmail(1), vi(1), sendmail(1M), passwd(4), terminfo(4), environ(5).

RFC 821 "Simple Mail Transfer Protocol (SMTP)"

RFC 822 "Standard for the Format of Internet Text Messages"

RFC 1521 "MIME (Multipurpose Internet Mail Extensions) Part One: Mechanisms for Specifying and Describing the Format of Internet Message Bodies"

RFC 1522 "MIME (Multipurpose Internet Mail Extensions) Part Two: Message Header Extensions for Non-ASCII Text"


e

NAME

elmalias - display and verify elm user and system aliases

SYNOPSIS

elmalias [-dersu] [-a|-f *format*|-n|-v|-V] [*alias-name-list*]

Remarks

The former functionality of the **elmalias** command has been taken over by the **newalias** command (see *newalias(1)*).

DESCRIPTION

The **elmalias** command displays and verifies user and system **elm** aliases.

The system database must have been created by the **newalias** command (see *newalias(1)*). The user database must have been created by either the **newalias** command or the **elm** mail system (see *elm(1)*). If the same alias is in both databases, the user version is used. Missing database files are silently ignored.

Each database entry can have the following fields, which are described in detail in *newalias(1)*:

| | |
|---------------------|--|
| <i>alias-list</i> | A list of one or more aliases for the entry. |
| <i>address-list</i> | A list of one or more addresses for the entry. An address can be an alias from another entry's <i>alias-list</i> . |
| <i>comment</i> | An optional field containing information about the entry. This field is not included in outbound mail. |
| <i>firstname</i> | An optional field interpreted as the first name of the person or group. It is used in <i>fullname</i> . |
| <i>lastname</i> | An optional field interpreted as the last name of the person or group. It is used in <i>fullname</i> . |
| <i>fullname</i> | A combination value made up from the <i>firstname</i> and <i>lastname</i> fields, in the form: <i>firstname lastname</i> . |

elmalias recognizes three types of alias names:

| | |
|----------------|--|
| person | A database entry that has one address in <i>address-list</i> . elmalias assumes this address is a valid mailing address. |
| group | A database entry that has two or more addresses in <i>address-list</i> . elmalias assumes initially that these addresses are aliases to person or group entries. |
| unknown | An address in a group entry or an alias in <i>alias-name-list</i> that is not an alias in the database. In both cases, the item is reported as both the alias name and its address. |

With no options or operands, **elmalias** displays the *address-list* field for each alias in the two databases. If an entry has more than one alias, the *address-list* field is displayed multiple times.

With an *alias-name-list* and no options, **elmalias** displays the *address-list* field of each alias name in the list. If an alias name is not found in the databases, it is treated as **unknown**, without comment.

Options

elmalias recognizes the following options:

| | |
|-------------------------|---|
| -a | Change the display to alias name followed by <i>address-list</i> field. |
| -d | Turn debugging on. |
| -e | Fully expand group aliases. This option can be used only when an <i>alias-name-list</i> is given. If an address in a group address list is an alias name, it is replaced by that alias entry. The process is recursive until the results are either person or unknown types. If a group address is not an alias name, it is reported as both the alias name and the address, with type unknown . Duplicate alias names are reported only once. person entries are not expanded, even if their addresses are actually aliases. |
| -f <i>format</i> | Display the <i>format</i> string for each alias name in the file or in the <i>alias-name-list</i> . The following character pairs are replaced in the <i>format</i> by the corresponding value for each alias. |

%a The alias name.
%c The *comment* field.
%l The *lastname* field.
%n The *fullname* value.
%t The alias type: **person**, **group**, or **unknown**.
%v The *address-list* field.

-n Change the display to *address-list* followed by *fullname*, if any, in parentheses.
-r Report an error if a name in *alias-name-list* does not correspond to an alias in the database. Display a message for each unknown name and exit with a nonzero status.
-s Use the system alias database only, unless **-u** is also specified.
-u Use the user alias database only, unless **-s** is also specified.
-v Use a verbose output format. Change the display to alias name, followed by *address-list*, followed by *fullname*, if any, in parentheses.
-V Use a very verbose, multiline output format with the following titles, corresponding to the format codes of the **-f** option. If a field is empty, the title is omitted.

Alias:
Address:
Type:
Name:
Last Name:
Comment:

EXIT STATUS

elmalias sets the following exit status values:

- 0 Normal completion.
- <>0 An error occurred. You may have specified:
 - An invalid option.
 - The **-e** option without an *alias-name-list*.
 - The **-f** option without a *format*.
 - The **-r** option with an unknown alias name in *alias-name-list*.

EXAMPLES

Consider a user database that contains the following entries:

```
# sample alias file
mom = My Mommy, Work: x2468 = my_mother@a.computer
dad,father,pop = Father; Dear, Work: x1357 = host!otherhost!dad
parents = The Folks = mom dad parent@host
siblings = The Kids = brother1
              brother2
              sister
brother1 = Son; First = bro1@kid.computer
```

Since **brother2** and **sister** do not refer to alias entries, they are typed as **unknown**.

elmalias with no options or operands produces the following output.

```
my_mother@a.computer
host!otherhost!dad
host!otherhost!dad
host!otherhost!dad
mom,dad,parent@host
brother1,brother2,sister
bro1@kid.computer
```

elmalias -v produces the alias names, address, and full name, as follows:

```

mom                my_mother@a.computer (My Mommy)
dad                host!otherhost!dad (Dear Father)
father             host!otherhost!dad (Dear Father)
pop               host!otherhost!dad (Dear Father)
parents            mom,dad,parent@host (The Folks)
siblings           brother1,brother2,sister (The Kids)
brother1           bro1@kid.computer (First Son)

```

To expand a set of aliases and format them with field titles, use the **-e** and **-f** options, as in the following command:

```
elmalias -ef "Alias: %a Address: %v Type: %t" parents siblings
```

producing:

```

Alias: mom Address: my_mother@a.computer Type: Person
Alias: dad Address: host!otherhost!dad Type: Person
Alias: parent@host Address: parent@host Type: Unknown
Alias: brother1 Address: bro1@kid.computer Type: Person
Alias: brother2 Address: brother2 Type: Unknown
Alias: sister Address: sister Type: Unknown

```

AUTHOR

elmalias was developed by HP.

FILES

| | |
|-----------------------------|---------------------------------------|
| \$HOME/.elm/aliases | User alias database data table |
| \$HOME/.elm/aliases.dir | User alias database directory table |
| \$HOME/.elm/aliases.pag | User alias database hash table |
| \$HOME/.elm/aliases.text | User alias source text |
| /var/mail/.elm/aliases | System alias database data table |
| /var/mail/.elm/aliases.dir | System alias database directory table |
| /var/mail/.elm/aliases.pag | System alias database hash table |
| /var/mail/.elm/aliases.text | System alias source text |

SEE ALSO

elm(1), newalias(1).

NAME

enable, disable - enable/disable LP printers

SYNOPSIS

enable *printers*

disable [-c] [-r[*reason*]] *printers*

DESCRIPTION

The **enable** command activates the named *printers*, enabling them to print requests taken by **lp**. Use **lpstat** to find the status of printers (see *lp(1)* and *lpstat(1)*).

disable deactivates the named *printers*, disabling them from printing requests taken by **lp**. By default, any requests that are currently printing on the designated printers are reprinted in their entirety either on the same printer or on another member of the same class. Use **lpstat** to find the status of printers.

Options

disable recognizes the following options:

- c Cancel any requests that are currently printing on any of the designated printers.
- r[*reason*] Associate a *reason* with the deactivation of the printers. This reason applies to all printers mentioned up to the next -r option. If the -r option is not present or the -r option is given without a reason, a default *reason* is used. *reason* is reported by **lpstat**. The maximum length of the *reason* message is 80 bytes.

EXTERNAL INFLUENCES**Environment Variables**

LANG determines the language in which messages are displayed.

If **LANG** is not specified or is null, it defaults to **C** (see *lang(5)*).

If any internationalization variable contains an invalid setting, all internationalization variables default to **C** (see *environ(5)*).

International Code Set Support

Single- and multibyte character code sets are supported.

EXAMPLES

Enable printer **snowwhite** to accept requests:

```
enable snowwhite
```

Deactivate printer **snowwhite** and cancel any logged jobs:

```
disable -c snowwhite
```

WARNINGS

If the restrict cancel feature (selected by the **lpadmin -orc** option — see *lpadmin(1M)*) is enabled, **disable** ignores the -c option.

enable and **disable** perform their operation on the local system only.

FILES

```
/etc/lp/*
/usr/lib/lp/*
/var/adm/lp/*
/var/spool/lp/*
```

SEE ALSO

lp(1), *lpstat(1)*, *accept(1M)*, *lpadmin(1M)*, *lpsched(1M)*, *rcancel(1M)*, *rlp(1M)*, *rlpdaemon(1M)*, *rlpstat(1M)*.

NAME

env - set environment for command execution

SYNOPSIS

env [-] [-i] [*name* = *value*] ... [*command* [*arguments* ...]]

DESCRIPTION

env obtains the current *environment*, modifies it according to its arguments, then executes the command with the modified environment. Arguments of the form *name=value* are merged into the inherited environment before the command is executed. The **-i** option causes the inherited environment to be ignored completely so that the command is executed with exactly the environment specified by the arguments. The **-** option is obsolete and has the same effect as the **-i** option.

If no command is specified, the resulting environment is printed, one name-value pair per line.

RETURN VALUE

If *command* is invoked, the exit status of **env** is the exit status of *command*; otherwise, **env** exits with one of the following values:

| | |
|--------------|--|
| 0 | env completed successfully. |
| 1-125 | env encountered an error. |
| 126 | <i>command</i> was found but could not be invoked. |
| 127 | <i>command</i> could not be found. |

EXTERNAL INFLUENCES**Environment Variables**

LC_MESSAGES determines the language in which messages are displayed.

If **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang*(5)) is used instead of **LANG**.

If any internationalization variable contains an invalid setting, **env** behaves as if all internationalization variables are set to "C". See *environ*(5).

International Code Set Support

Single- and multi-byte character code sets are supported.

WARNING

The **-** option is obsolete. Use **-i** instead.

SEE ALSO

sh(1), *exec*(2), *profile*(4), *environ*(5).

STANDARDS CONFORMANCE

env: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

eucset - sets and gets code widths for ldterm

SYNOPSIS

eucset [-p]

eucset [[-c *HP15-codeset*] or [-c *UTF8*] or [*cswidth*]]

DESCRIPTION

The **eucset** command sets or gets (reports) the encoding and display widths of the Extended UNIX Code (EUC) and UCS Transformation Format (UTF8) characters processed by the current input terminal. EUC is an encoding method for codesets composed of single or multiple bytes. It permits applications and the terminal hardware to use the 7-bit US ASCII code and up to three single byte or multibyte code sets simultaneously.

The **eucset** command without any options, first tries to set the codeset to one of the four HP15 codesets. If unsuccessful, 7-bit US ASCII is used as the default codeset. This command must be used to specify any other EUC codesets, whether they are single byte or multibyte. See the WARNINGS section, for special warnings on the values of the *cswidth* argument.

To utilize this command for UTF8, specify UTF8 with the -c option.

Options

The **eucset** command recognizes the following options and arguments:

- p Displays the current settings of the EUC character widths for the terminal.
- c Sets the codeset to one of the four HP15 codesets or the **UTF8** codeset. The HP15 codesets supported are **SJIS**, **CCDC**, **GB**, and **BIG5**.

EUC Code Set Classes

EUC divides codesets into four classes. Each codeset has two characteristics: the number of bytes for encoding the characters in the codeset, and the number of display columns to display the characters in the codeset. All characters within a codeset possess the same characteristics.

- Codeset 0 consists of all 7-bit, single byte ASCII characters. The most significant bit of each of these characters is 0 (zero). Characters in codeset 0 require one byte for encoding, and occupy one display column. These values are fixed for codeset 0 (zero). The 7-bit US ASCII code is the primary EUC codeset, which is available to users without direct specification.
- Codeset 1 is a supplementary EUC codeset. Codeset 1 characters have an initial byte whose most significant bit is 1. Characters in codeset 1 may require more than one byte for encoding, and may require more than one display column. The **eucset** command must be used to set the characteristics for codeset 1.
- Codesets 2 and 3 are supplementary EUC codesets. Characters in these codesets have an initial byte of SS2 or SS3, respectively. They require more than one byte for encoding, and may require more than one display column. The **eucset** command must be used to set the characteristics for codesets 2 and 3.

The *cswidth* argument in the **eucset** command line is a character string that describes the character widths for codesets 1 through 3. This command does not allow the user to modify the settings for codeset 0. The character string is of the following format:

X1[: Y1], X2[: Y2], X3[: Y3]

The value *X1* is the number of bytes required to encode a character in codeset class 1. *Y1* is the number of display columns needed to display characters in this class. *X2* is the number of bytes required to encode a character in codeset 2, not counting the SS2 byte, and *Y2* is the number of display columns for codeset 2 characters. *X3* is the number of bytes needed to encode characters in codeset 3, not counting the SS3 byte, and *Y3* is the number of display columns required for these characters. The values for the column widths may be omitted if they are equal to the number of encoding bytes. If the encoding value of any of the EUC codesets is set to 0 (zero), this indicates that the codeset does not exist. See the WARNINGS section for special warnings on the values of the *cswidth* argument.

If no *cswidth* argument is supplied, the **eucset** command uses the value of the **CSWIDTH** environment variable. If this variable is not present, the following default string is substituted:

```
1:1,0:0,0:0
```

This default string designates that the environment uses a single byte EUC codeset that has characters in the EUC codeset 1 format. If the environment uses a multibyte EUC codeset in the codeset 1 format, single byte or multibyte EUC codesets in the codeset 2 or 3 format, or both, the default setting cannot be used.

EXTERNAL INFLUENCES

Environment Variables

| | |
|--------------------|--|
| LANG | Provide a default value for the internationalization variables that are unset or null. If LANG is not specified or is set to the empty string, a default of C (see <i>lang(5)</i>) is used instead of LANG . If any of the internationalization variables contain an invalid setting, eucset behaves as if all internationalization variables are set to C . See <i>environ(5)</i> . |
| LC_ALL | If set to a nonempty string value, override the values of all other internationalization variables. |
| LC_MESSAGES | Determines the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output. |
| NLSPATH | Determines the location of message catalogs for the processing of LC_MESSAGES . |

EXAMPLES

To display the encoding and display widths for the EUC codesets 1 to 3 in your environment, enter:

```
eucset -p
```

Assuming **eucset** has been previously used to set for **ja_JP.eucJP**, the entry generates the following:

```
cswidth 2:2,1:1,2:2
```

To change the current settings of the encoding and display widths for the EUC characters in codesets 1 and 2 to two bytes each, enter one of the following:

```
eucset 2:2,2:2,0:0
```

```
eucset 2,2,0
```

To set the encoding and display widths for the EUC characters in the locale **ja_JP.eucJP**, enter:

```
eucset 2:2,1:1,2:2
```

For **zh_TW.eucTW**, enter:

```
eucset 2:2,3:2
```

For **ko_KR.eucKR**, enter:

```
eucset 2:2
```

To set the codewidth to that of **UTF8**, enter:

```
eucset -c UTF8
```

WARNINGS

The *cswidth* argument does not include the SS2 or SS3 bytes in the byte width values.

This command is not specified by standards, may not be available on other vendor's systems, and may be subject to change or obsolescence in a future release.

AUTHOR

eucset was developed by OSF and HP.

SEE ALSO

dtterm(1), *ldterm(1)*.

NAME

ex, edit - extended line-oriented text editor

SYNOPSIS

ex [-] [-l] [-r] [-R] [-t *tag*] [-v] [-wsize] [-x] [-C] [+*command*] [*file* ...]

XPG4 Synopsis

ex [-rR] [-s | -v] [-c *command*] [-t *tag*] [-w *size*] [*file* ...]

Obsolescent Options

ex [-rR] [- | -v] [+*command*] [-t *tag*] [-w *size*] [*file* ...]

edit [-] [-l] [-r] [-R] [-t *tag*] [-v] [-wsize] [-x] [-C] [+*command*] [*file* ...]

Remarks

The program names **ex**, **edit**, **vi**, **view**, and **vedit** are separate personalities of the same program. This manual entry describes the behavior of the **ex/edit** personality. On many HP-UX and other similar systems, **e** is a synonym for **ex**.

DESCRIPTION

The **ex** program is the line-oriented personality of a text editor that also supports screen-oriented editing (see *vi*(1)).

(XPG4 only.) Certain block-mode terminals do not have all the capabilities necessary to support the complete **ex** definition, such as the full-screen editing commands (**visual** mode or **openmode**). When these commands cannot be supported on such terminals, this condition shall neither produce an error message such as "not an editor command" nor report a syntax error.

The **edit** program is identical to **ex**, except that some editor option defaults are altered to make the editor somewhat friendlier for beginning and casual users (see Editor Options below).

Options and Arguments

ex recognizes the following command-line options and arguments:

- (Obsolescent) Suppress all interactive-user feedback. This is useful when editor commands are taken from scripts.
- s (XPG4 only.)
Suppress all interactive-user feedback. This is useful when editor commands are taken from scripts.
Ignore the value of the **TERM** and any implementation terminal type and assume the terminal is a type incapable of supporting visual mode.
Suppress the use of the **EXINIT** environment variable and the reading of the **.exrc** file.
- l Set the **lisp** editor option (see Editor Options below).
- r Recover the specified *files* after an editor or system crash. If no *file* is specified, a list of all saved files is printed. You must be the owner of the saved file in order to recover it (superuser cannot recover files owned by other users).
- R Set the **readonly** editor option to prevent overwriting a file inadvertently (see Editor Options below).
- t *tag* (XPG4 only.) Edit the file containing the specified *tag* and proceed as if the first command were **:tag tag**. The tags represented by the **-t tag** and the **ta** command is optional. It shall be provided on any system that also provides a confirming implementation of **ctags**. Otherwise, the use of the **-t** produces undefined results.
Execute the **tag tag** command to load and position a predefined file. See the **tag** command in Command Descriptions and the **tags** editor option in Editor Options below.
- v Invoke visual mode (**vi**).
- w *size* Set the value of the **window** editor option to *size* (see Editor Options below). If *size* is omitted, it defaults to 3.

- x** Set encryption mode. You are prompted for a key to initiate the creation or editing of an encrypted file (see the **crypt** command in Command Descriptions below).
 - C** Encryption option. Same as the **-x** option, except that all text read in is assumed to have been encrypted.
 - c *command*** (XPG4 only.)
 - +*command*** (Obsolescent) Begin editing by executing the specified **ex** search or positioning *command*.
 - file*** Specify the file or files to be edited. If more than one *file* is specified, they are processed in the order given. If the **-r** option is also specified, the files are read from the recovery area.
- (XPG4 only.) If both the **-t *tag*** and **-c *command*** options are given, the **-t *tag*** shall be processed first; i.e., the file containing the tag is selected by the **-t** and then the command is executed.

e

Definitions

Current file. The name of the file being edited by **ex** is called the current file. Text from the current file is read into a work area, and all editing changes are performed on this work area. Changes do not affect the original file until the work area is explicitly written back to the file. If the % character is used as a file name, it is replaced by the current file name.

Alternate file. The alternate file is the name of the last file mentioned in an editor command, or the previous current file name if the last file mentioned becomes the current file. If the # character is used as a file name, it is replaced by the alternate file name.

Buffers. Twenty-six buffers named **a** through **z** can be used for saving blocks of text during the edit. If the buffer name is specified in uppercase, text is appended to the existing buffer contents rather than overwriting it.

Readonly flag. The **readonly** flag can be cleared from within the editor by setting the **noreadonly** editor option (see Editor Options below). Writing to a different file is allowed even when the **readonly** flag is set. Also, a write can be forced to a **readonly** file by using **!** after the write command (see the **write** command in Command Descriptions below).

Interrupt. If an interrupt signal is received, and commands are being supplied from a keyboard, **ex** returns to command mode. If editor commands are coming from a file, an interrupt signal causes **ex** to abort.

System crash. If the system crashes or **ex** aborts due to an internal error or unexpected signal, **ex** attempts to preserve the work area if any unwritten changes were made. Use the **-r** command-line option to retrieve the saved changes.

Command mode/input mode. **ex** starts up in command mode, as indicated by the colon (:) prompt. **ex** switches to input mode whenever an **append**, **change**, or **insert** command is encountered. To terminate input mode and return to command mode, type a period (.) alone at the beginning of a line.

Comments. Command lines beginning with a quotation mark (") are ignored (this is useful for placing comments in an editor script).

Multiple commands can be combined on a single line by separating them with a vertical bar character (|). However, global commands, comments, and the shell escape command must be the last command on a line because they cannot be terminated by a | character.

Addressing

(XPG4 only.) Addressing in **ex** relates to the current line. In general, the current line shall be the last line affected by the command; the exact effect on the current line is discussed under the description of each command. When the buffer contains no lines, the current line shall be set to zero.

ex recognizes the following line address forms:

- .** Dot or period (.) refers to the current line. There is always a current line whose position can be the result of an explicit movement command or the result of a command that affects multiple lines (in which case it is usually the last line affected).
- n*** The *n*th line in the work area. Lines are numbered sequentially, starting at line 1.
- \$** The last line in the work area.
- %** Abbreviation for **1,\$**, meaning the entire work area.

| | |
|--------------------------|--|
| <code>+n, +[+]...</code> | An offset relative to the current line or the preceding line specification. <code>+</code> means forward; <code>-</code> means backward. For example, the forms <code>.+3</code> , <code>+3</code> , and <code>+++</code> are equivalent. |
| <code>-n, -[-]...</code> | |
| <code>/re/</code> | The line containing the pattern <i>re</i> , scanning forward (<code>/</code>) or backward (<code>?</code>). The trailing <code>/</code> or <code>?</code> can be omitted if the line is only being displayed. If <i>re</i> is omitted, ex uses the more recently set of either the scanning string or the substitution string (see Regular Expressions below). |
| <code>?re?</code> | |
| <code>'x</code> | Lines can be marked using single lowercase letters (see the mark command in Command Descriptions below). <code>'x</code> refers to the line marked with <i>x</i> . In addition, the previous current line is marked before each nonrelative motion. This line can be referred to by using <code>'</code> for <i>x</i> (thus <code>''</code> refers to the previous current line). |
| | (XPG4 only.) Commands require zero, one or two addresses. Commands that require zero addresses shall regard the presence of an address as an error. |

(XPG4 only.) Adjacent address in a **range** shall be separated from each other by a comma (,) or a semicolon(;). In the latter case, the current line(.) shall be set to the first address, and only then is the second address calculated. This feature can be used to determine the starting line for forwards and backwards searches. The second address of any two-address sequence shall correspond to the first address. The first address shall be less than or equal to the second address. The first address shall be greater than or equal to the first line of the editing buffer, and the last address shall be less than or equal to the last line of the editing buffer. Any other case shall be an error.

Addresses for commands consist of a series of line addresses (specified as above), separated by a comma (,) or semicolon (;). Such address lists are evaluated left-to-right. When the separator is a semicolon, the current line is set to the value of the previous address before the next address is interpreted. If more addresses are given than the command requires, then all but the last one or two are ignored. Where a command requires two addresses, the first line addressed must precede the second one in the work area. A null (missing) address in a list defaults to the current line.

Regular Expression

The editor maintains copies of two regular expression strings at all times: the substitution string, and the scanning string. The substitute command sets the substitution string to the regular expression used. Both the global-command and the regular-expression form of line addressing (see Addressing above) for all commands set the scanning string to the regular expression used. These strings are used as default regular expressions as described under Addressing, the **global** command, and the **substitute** command.

The editor supports Basic Regular Expressions (see *regex*(5)) with the following modifications:

| | |
|-----------------------|--|
| <code>\<</code> | The <code>\<</code> matches the beginning of a "word"; that is, the matched string must begin in a letter, digit, or underline, and must be preceded by the beginning of the line or a character other than the above. This construct can only be used at the beginning of a regular expression (as in <code>\<word</code>), but not in the middle (<code>word1 \<word2</code>). |
| <code>\></code> | The <code>\></code> matches the end of a "word" (see previous paragraph). This construct can only be used at the end of a regular expression (as in <code>word1\></code>), but not in the middle (<code>word1\> word2</code>). |
| <code>~</code> | Match the replacement part of the last substitute command. |
| <code>[string]</code> | The positional quoting within bracket expressions defined by Basic Regular Expressions is replaced by the use of the backslash (\) to quote bracket-expression special characters. |
| <code>nomagic</code> | When the editor option nomagic is set, the only characters with special meanings are <code>^</code> at the beginning of a pattern, <code>\$</code> at the end of a pattern, and <code>\</code> . The characters <code>.</code> , <code>*</code> , <code>[</code> , and <code>~</code> lose their special meanings unless escaped by a <code>\</code> . |

Replacement Strings

The character `&` in the replacement string stands for the text matched by the pattern to be replaced. Use `\&` if the **nomagic** editor option is set.

The character `~` is replaced by the replacement part of the previous **substitute** command. Use `\~` if the **nomagic** editor option is set.

The sequence `\n`, where *n* is an integer, is replaced by the text matched by the subpattern enclosed in the *n*th set of parentheses `(` and `)`.

The sequence \u (\l) causes the immediately following character in the replacement to be converted to uppercase (lowercase), if the character is a letter. The sequence \U (\L) turns case conversion on, until the sequence \E or \e is encountered, or the end of the replacement string is reached.

Command Names and Abbreviations

The following table summarizes the line-mode commands. The commands whose names are enclosed in parentheses are available only in their abbreviated forms.

| Command | Abbr. | Command | Abbr. | Command | Abbr. |
|------------|--------|------------|----------|------------------|-------|
| abbreviate | ab | next | n | tag | ta |
| append | a | number | nu # | unabbreviate | una |
| args | ar | open | o | undo | u |
| change | c | pop | | unmap | unm |
| chdir | chd cd | preserve | pre | version | ve |
| copy | co t | print | p | visual | vi |
| crypt | cr X | put | pu | write | w wq |
| delete | d | quit | q | xit | x |
| edit | e ex | read | r | yank | ya |
| file | f | recover | rec | | |
| global | g v | rewind | rew | (execute buffer) | * @ |
| insert | i | set | se | (line number) | = |
| join | j | shell | sh | (left shift) | < |
| list | l | source | so | (right shift) | > |
| map | | stop | st ^Z | (scroll) | ^D |
| mark | ma k | substitute | s sr & ~ | (shell escape) | ! |
| move | m | suspend | su ^Z | (window) | z |

Command Descriptions

In the following command descriptions, some arguments appear frequently. They are described below.

- line* A single line address, in any of the forms described in Addressing above. The default is the current line.
- range* A pair of line addresses separated by a comma or semicolon, as described in Addressing above. The default is the current line (.,.).
- count* A positive integer specifying the number of lines to be affected by the command. The default is 1 or the number of lines in *range*.
When *count* is specified, *range* is ineffective. Instead, only a line number should be specified to indicate the first line affected by the command. (If a range is given, the last line of the range is interpreted as the starting line for the command.)
- flags* One or more of the characters #, p, and l. The corresponding command to print the line is executed after the command completes. Any number of + or - characters can also be given with these flags. The default is no flags.

These modifiers are all optional.

When only a *line* or a *range* is specified (with a null command), the implied command is **print**. If a null line is entered, the next line is printed (equivalent to .+lp)

- buffer* XPG4 Feature. One of a number of named areas for saving text. The named buffers are specified by the lowercase letters of the POSIX locale. Specifying **buffer** shall cause the area of the text affected by the command to be stored into the buffer as it was before the command took effect. This argument is also used on the **put** command and the visual mode "put" commands (p and P), to specify the buffer that shall provide the text to insert.
If the buffer name is specified in uppercase, and the buffer is to be modified, the buffer shall be appended to rather than being overwritten. If the buffer is not to be modified, the buffer name can be specified in lowercase or uppercase with the same results. There shall be also one unnamed buffer, which is the repository for all text deleted or yanked when no buffer is specified.

There are also numbered buffers, 1 through 9, which shall be accessible only from visual mode. These buffers are special in that, in the visual mode, when deleted text is placed in the unnamed buffer, it also shall be placed in buffer 1, the previous contents buffer 1 shall be placed in buffer 2 and so on. Any text in the buffer 9 shall be lost. Text that is yanked into the unnamed buffer shall not modify the numbered buffers. Text cannot be placed directly into the numbered buffered, although it can be retrieved from them by using a visual mode "put" command with the buffer name given as *s* number. When the buffer modifier is not used in the commands below, the unnamed buffer shall be the default.

word **XPG4 Feature.** In the POSIX Locale, a **word** consists of a maximal sequence of letters, digits and underscores, delimited at both ends by characters other than letters, digits, or underscores, or by the beginning or end of a word or the file. **!** A character that can be appended to the command to modify its operation, as detailed in the individual command descriptions.

If both a **count** and **range** is specified for a command that uses them, the number of lines affected shall be taken from the **count** value rather than the **range**. The starting line for the command shall be taken to be the first line addressed by the range.

When only a **line** or **range** is specified with no command, the implied command shall be either **print**, **list**, or **number** (**p**, **l**, or **#**). The command selected shall be the last of these three commands to be used. When no range or count is specified and the command line is a blank line, the current line shall be written, and the current line shall be set to **.+1**.

Zero or mode <blank> characters can precede or follow the addresses, count or command name. Any object following a command name (such as buffer, file etc) that begins with an alphabetic character shall be separated from the command name with at least one <blank>.

For each of the commands listed below, the command can be entered as the abbreviation (those characters in the Synopsis command word preceding the **]**), the full command (all characters shown for the command word, omitting the **[** and **]**), or any subset of the characters of the full command down to the abbreviation.

abbreviate **ab[breviate]** *word replacement*

Add the named abbreviation to the current list. In visual mode, if *word* is typed as a complete word during input, it is replaced by the string *replacement*.

append **line a[ppend][!]**

Enter input mode; the input text is placed after the specified line. If line 0 is specified, the text is placed at the beginning of the work area. The last input line becomes the current line, or the target line if no lines are input.

Appending **!** to the command toggles the **autoindent** editor option setting for this insert only.

args **ar[gs]**

Prints the argument, placing the current argument between **[** and **]**.

change **range c[hange][!]** *count*

Enter input mode; the input text replaces the specified lines. The last input line becomes the current line; if no lines are input, the effect is the same as a delete.

Appending **!** to the command toggles the **autoindent** editor option setting for this insert only.

chdir **chd[ir][!]** [*directory*]
cd[!] [*directory*]

Change the working directory to *directory*. If *directory* is omitted, the value of the **HOME** environment variable is used. If the work area has been modified since the last write and the name of the file being edited does not begin with a slash (/), a warning is issued and the working directory is not changed. To force a change of directory in this case, append the character **!** to the command.

copy **range co[py]** *line flags*
 range t *line flags*

A copy of the specified lines (*range*) is placed after the specified destination *line*; line 0 specifies that the lines are to be placed at the beginning of the work area. (The letter **t** is an alternative abbreviation for the **copy** command.)

crypt **cr[yp]t**
X

The user is prompted for a key with which to enter encryption mode. This command can also be used to change the key entered from a previous **crypt** command or the **-x** command line option. If no key is supplied in response to the prompt (that is, only carriage return is pressed), encryption mode is canceled and the work area is written out in plain-text form by subsequent write commands.

While in encryption mode, all file input is decrypted using the current key. However, while an input file is being processed, if a block of text (approximately 1024 bytes) is encountered that contains only 7-bit ASCII characters, that block of text is assumed to be plain-text and is not decrypted. All file output, except that piped via a **!** shell escape to another command, is encrypted using the current key.

The temporary file used by the editor to manage the work area is not encrypted until the current work area is discarded (or written out) and editing begins on a new file. When creating a new file that requires encryption protection, ensure that the work area file is also encrypted by specifying the **-x** option when invoking the editor.

cr[yp]t
C

Encryption option. Same as the **X** command, except that all text read in is assumed to have been encrypted.

delete *range* **d[ele]te** *buffer count*

The specified lines are deleted from the work area. If a named *buffer* is specified, the deleted text is saved in it. If no buffer is specified, the unnamed buffer is used (that is, the buffer where the most recently deleted or yanked text is placed by default). The new current line is the line after the deleted lines or the last line of the file if the deleted lines were at the end of the file.

edit **e[dit][!]** *[+ line] file*
ex[!] *[+ line] file*

Begin editing a new file (**ex** is an alternative name for the **edit** command). If the current work area has been modified since the last write, a warning is printed and the command is aborted. This action can be overridden by appending the character **!** to the command (**e! file**). The current line is the last line of the work area unless it is executed from within *vi*, in which case the current line is the first line of the work area. If the *+line* option is specified, the current line is set to the specified position, where *line* can be a number (or **\$**) or specified as */re* or *?re*.

file **f[ile]**

Print the current file name and other information, including the number of lines and the current position.

global *range* **g[lobal][!]** */re/ command...*
range **v** */re/ command...*

Perform *command* on lines within *range* (or on the entire work area if no *range* is given) that contain *re*. First mark the lines within the given *range* that match the pattern *re*. If the pattern is omitted, the more recently set of either the substitution string or the scanning string is used (see Regular Expressions above). Then the given *commands* are executed with **.** set to each marked line. Any character other than a letter or a digit can be used to delimit the pattern instead of the **/**.

command can be specified on multiple lines by hiding new-lines with a backslash. If *command* is omitted, each line is printed. **append**, **change**, and **insert** commands are allowed; the terminating dot can be omitted if it ends *command* or *commands*. The **visual** command is also permitted (unless the **global** command itself has been issued from visual mode), and takes input from the terminal. (If *command* contains a visual-mode command (that is, **open** or **visual**), the visual-mode command must be

terminated by the visual-mode **Q** command in order to proceed to the next marked line.) The **global** command itself and the **undo** command are not allowed in *command*. The editor options **autoprint**, **autoindent**, and **report** are inhibited.

Appending a **!** to the **global** command (that is, **g!** ...) or using the alternate name **v** causes *command* to be run on the lines within *range* that do not match the pattern.

insert *line i***nsert****[!]**

Enter input mode; the input text is placed before the specified line. The last line input becomes the current line, or the line before the target line, if no lines are input.

Appending **!** to the command toggles the **autoindent** editor option setting for this insert only.

join *range j***oin****[!]** *count flags*

Join together the text from the specified lines into one line. White space is adjusted to provide at least one blank character (two if a period appears at the end of a line, or none if the first character of a line is a closing parenthesis ())). Extra white space at the beginning of a line is discarded.

Appending a **!** to the command causes a simpler join with no white-space processing.

list *range l***ist** *ount flags*

Print the specified lines with tabs displayed as **^I** and the end of each line marked with a trailing **\$**. (The only useful flag is **#** for line numbers.) The last line printed becomes the current line.

map **map** *key* | **#***n* *action*
map! *key* | **#***n* *action*

The **map** and **map!** commands define macros for use in visual mode. The first argument, *key*, can be a single character or a multicharacter sequence. In the special sequence, **#***n*, *n* is a digit referring to the function key *n*. Special characters, whitespace, and newline must be escaped with a **^V** to be entered in the arguments. The *key* argument cannot contain a colon (**:**) as its first character, nor can a multicharacter sequence begin with an alphabetic character.

Macros defined by **map** are effective in visual command mode. Macros defined by **map!** are effective in visual input mode. When *key* or the function key corresponding to **#***n* is entered, the editor interprets the operation as though *action* were typed.

The **map** or **map!** command without options displays the corresponding current list of macros.

See also the editor options **keyboardeedit**, **keyboardeedit!**, **timeout**, and **timeoutlen** in Editor Options below.

mark *line m***a****[rk]** *x*
line k *x*

The specified line is given the specified mark *x*, which must be a single lowercase letter (**a-z**). *x* must be preceded by a space or tab. The current line position is not affected. **k** is an alternate name for **mark**.

move *range m***[ove]** *line*

Move the specified lines (*range*) to follow the target *line*. The first line moved becomes the current line.

next **n****[ext]****[!]** [*file ...*]

The next file from the command line argument list is edited. Appending a **!** to the command overrides the warning about the work area having been modified since the last write (and discards any changes unless the **autowrite** editor option is set). The argument list can be replaced by specifying a new one on this command line.

number *range nu***m****ber** *count flags*
range # *count flags*

(The **#** character is an alternative abbreviation for the **number** command.) Print the lines, each preceded by its line number (the only useful flag is **1**). The last line printed becomes the current line.

open *line* **o[pen]** / *re/* *flags*

Enter open mode, which is similar to visual mode with a one-line window. All the visual-mode commands are available. If a match is found in *line* for the optional regular expression, the cursor is placed at the start of the matching pattern. Use the visual mode command **Q** to exit from open mode. For more information, see *vi(1)*.

pop **pop[!]**

Load the file whose name is stored at the top of the tag stack and set the current line to the stored location. The top entry of the tag stack is deleted. (The current file name is placed on the stack when you execute the line mode **tag** command or the visual mode **^]** command.)

! overrides the warning about the work area having been modified since the last write; any changes are discarded unless the **autowrite** editor option is set).

preserve **pre[serve]**

The current editor work area is saved as if the system had just crashed. Use this command in emergencies, for example when a write does not work and the work area cannot be saved in any other way. Use the **-r** command-line option to recover the file. After the file has been preserved, a mail message shall be sent to the user. The message shall contain the name of the file, the time of preservation and an **ex** command that could be used to recover the file. Additional information may be included in the mail message.

print **range** **p[rint]** *count*

Print the specified lines, with non-printing characters printed as control characters in the form **^x**; DEL is represented as **^?**. The last line printed becomes the current line.

put *line* **pu[t]** *buffer*

Place deleted or "yanked" lines after *line*. A buffer can be specified; otherwise, the text in the unnamed buffer (that is, the buffer in which deleted or yanked text is placed by default) is restored. The current line indicator shall be set to the first line put back.

quit **q[uit][!]**

Terminate the edit. If the work area has been modified since the last write, a warning is printed and the command fails. To force termination without preserving changes, append **!** to the command.

read *line* **r[ead]** *file*

Place a copy of the specified *file* in the work area after the target line (which can be line 0 to place text at the beginning). If no *file* is named, the current file is the default. If no current file exists, *file* becomes the current file. The last line read becomes the current line except in visual mode where the first line read becomes the current line.

If *file* is given as **!string**, *string* is interpreted as a system command and passed to the command interpreter; the resultant output is read into the work area. A blank or tab must precede the **!**.

recover **rec[over][!]** *file*

Recover *file* from the save area, after an accidental hangup or a system crash. If the current work area has been modified since the last write, a warning is printed and the command is aborted. This action can be overridden by appending the character **!** to the command (**rec! file**).

rewind **rew[ind][!]**

The argument list is rewound, and the first file in the list is edited. This shall be equivalent to a **next** command with the current argument list as its operands. If the current buffer has been modified since the last write, a warning shall be written and the command shall be aborted. Any warnings can be overridden by appending a **!**. The

current indicator line shall be affected by the editor options, **autowrite** and **wri-teany**.

```
set      se[t] [all]
        se[t] [no]boolean-option?
        se[t] value-option[?]
        se[t] boolean-option
        se[t] noboolean-option
        se[t] value-option=value
```

Set and display the values of the editor options (see Editor Options below).

With no arguments, the command prints those editor options whose values have been changed from the default settings. If **all** is specified, it prints all current option values.

The second and third forms display the current value of the specified option. The **?** is necessary only for Boolean options.

The fourth form turns a Boolean option on. The fifth form turns a Boolean option off.

The sixth form assigns values to string and numeric options. Spaces and tabs in strings must be escaped with a leading backslash (****).

The last five forms can be combined; interpretation is left-to-right.

```
shell    sh[ell]
```

Execute the command interpreter specified by the **shell** editor option (see Editor Options below). Editing is resumed when you exit from the command interpreter.

```
source    so[urce] file
```

Read and execute commands from the specified *file*. **so** commands can be nested. The maximum supported nesting depths is implementation defined, but shall be at least one.

```
substitute range s[substitute] / re / repl / options count flags
           range s options count flags
           range & options count flags
           range sr options count flags
           range ~ options count flags
           range s\? repl
           range s\&repl
```

On each specified line, the first instance of the pattern *re* is replaced by the string *repl*. (See Regular Expressions and Replacement Strings above.) Any character other than a letter or a digit can be used to delimit the pattern instead of the **/**.

If you include the **g** (global) option, all instances of the pattern in the line are substituted.

If you include the **c** (confirm) option, you are queried about whether to perform each individual substitution, as follows: Before each substitution the line is displayed with the pattern to be replaced marked underneath with carets (**^**). Type **y** to cause the substitution to be performed; any other input to abort it. The last line substituted becomes the current line.

If the substitution pattern *re* is omitted (**s//repl/**), the more recently set of either the substitution string or the scanning string is used (see Regular Expressions above).

If the **s** or **&** forms of the command are used, the substitution pattern defaults to the previous substitution string and the replacement string defaults to the previous replacement string used.

If the **sr** or **~** forms of the command are used, the substitution pattern defaults to the more recently set of either the substitution string or the scanning string and the replacement string defaults to the previous replacement string used.

The form **s\?repl** is equivalent to **s/scan-re/repl/**, where *scan-re* is the previous scanning string.

The form **s\&repl** is equivalent to **s/subs-re/repl/**, where *subs-re* is the previous substitution string.

suspend **su[suspend][!]**
stop **st[op][!]**
 susp

Suspend the editor job and return to the calling shell. **stop** and *susp* are equivalent to **suspend**. *susp* is the user process control suspend character, which is typically the character **^Z** (ASCII SUB) (see *stty*(1)). This command is disabled if the calling shell does not support job control or has disabled it.

The work area is written to the current file before the editor is suspended if the **autowrite** editor option is set, the **readonly** editor option is not set, and the work area has been modified since the last write. To override this action, append the **!** character to the **suspend** or **stop** command.

tag **ta[g][!] tag**

Search the files specified by the **tags** editor option (see Editor Options below) sequentially until a tag definition for *tag* is found. If *tag* is found, load the associated file into the work area and set the current position to the address specified in the tag definition.

The work area is written to the current file before the new file is loaded if the new file is different from the current file, the **autowrite** editor option is set, the **readonly** editor option is not set, and the work area has been modified since the last write. To override this action, append the **!** character to the command.

If the **tagstack** editor option is set, the current file name and line number is pushed onto the tag stack for later recall with the line mode **pop** command or the visual mode **^]** command.

unabbreviate **una[bbreviate] word**

Delete *word* from the list of abbreviations (see the **abbreviate** command above).

undo **u[ndo]**

Reverse the changes made by the previous editing command. For this purpose, **global** and **visual** are considered single commands. Commands that affect the external environment, such as **write**, **edit**, and **next**, cannot be undone. An **undo** can itself be reversed.

unmap **unm[ap][!] key**

The macro definition for *key* is removed (see the **map** command above).

version **ve[rsion]**

Print the current version information for the editor.

visual **line vi[sual] type count flags**

Enter visual mode at the specified *line*.

The *type* can be one of the characters **+**, **-**, **.**, or **^**, as in the **z** (window) command, to specify the position of the specified line on the screen window. The default is to place the line at the top of the screen window.

A *count* specifies an initial window size; the default is the value of the editor option **window**.

The flags **#** and **l** (ell) cause the lines in the visual window to be displayed in the corresponding mode (see the **number** and **list** commands).

Use the **Q** command to exit visual mode. For more information, see *vi*(1).

write **[range] w[rite][!]>> file**
 [range] wq[!]>> file

Write the specified lines (or the entire work area, if no *range* is given) out to *file*, printing the number of lines and characters written. If *file* is not specified, the default is the current file (the command fails with an error message if there is no current file and no file is specified).

If an alternate file is specified and the file exists, the write fails, but can be forced by appending **!** to the command. To append to an existing file, append **>>** to the command. If the file does not exist, an error is reported.

If the file is specified as **!string**, *string* is interpreted as a system command, the command interpreter is invoked, and the specified lines are passed as standard input to the command.

The command **wq** is equivalent to a **w** followed by a **q**. **wq!** is equivalent to **w!** followed by **q**. **wq>>** is equivalent to **w>>** followed by **q**.

xit

x[it][!][>>] file

If changes have been made to the work area, a **write** command is executed with any options (such as **!**, **>>**, or *file*) used by the **write** command. Then (in any case) the **quit** command is executed.

yank

range ya[nk] buffer count

Place the specified lines in the named *buffer*. If no buffer is specified, the unnamed buffer is used (that is, the buffer where the most recently deleted or yanked text is placed by default).

(execute buffer)

*** [buffer]**
@ [buffer]

Execute the contents of *buffer* as an editor command. *buffer* can be the letter of a named buffer (**a-z**) or ***** or **@**. The ***** and the **@** forms of this command are equivalent. If a buffer is not specified or *buffer* is ***** or **@**, the buffer last named in a ***** or **@** command is executed.

(line number)

line = flags

Print the line number of the specified *line*. The default is the last line. The current line position is not affected.

(scroll)

^D

Print the next *n* lines, where *n* is the value of the **scroll** editor option.

(shell escape)

! command
range ! command

Pass the remainder of the line after the **!** to the system command interpreter for execution. A warning is issued if the work area has been changed since the last write. A single **!** is printed when the command completes. The current line position is not affected.

Within the text of *command*, **%** and **#** are expanded as file names, and **!** is replaced with the text of the previous **!** command. Thus, **!!** repeats the previous **!** command. When such expansion is performed, the expanded line is echoed.

If you specify *range*, the specified lines are passed to the command interpreter as standard input. The output from the *command* replaces the specified lines.

(shift left)

range < count flags

Shift the specified lines to the left. The number of spaces to be deleted is determined by the editor option **shiftwidth**. Only whitespace (blanks and tabs) is lost in shifting; other characters are not affected. The last line changed becomes the current line.

(shift right)

range > count flags

Shift the specified lines to the right by inserting whitespace. The number of spaces inserted is determined by the editor option **shiftwidth**. The last line changed becomes the current line.

(window)

line z type count flags

The number of lines specified by *count* are displayed. The default for *count* is the value of the editor option **window**.

If *type* is omitted, *count* lines following the specified *line* are printed.

If *type* is specified, it must be one of the following characters:

- + Display a window of lines following the addressed line.
- Place the addressed line at the bottom of the window of displayed lines.
- . Place the addressed line at the center of the window.
- ^ Display a window of lines that is two windows prior to the addressed line.
- = Display the addressed line at the center of the window with a line of dashes above and below the addressed line.

The last line printed becomes the current line, except for the =, where the addressed line becomes the current line.

Editor Options

The command **ex** has a number of options that modify its behavior. These options have default settings, which can be changed using the **set** command (see above). Options can also be set at startup by putting a **set** command string in the environment variable **EXINIT**, or in the file **.exrc** in the **HOME** directory, or in **.exrc** in the current directory. If **EXINIT** exists, the **.exrc** file in the **HOME** directory is not executed. If the current directory is not the **HOME** directory and the **exrc** editor option is set (see below), the **.exrc** file in the current directory is executed after **EXINIT** or the **HOME** directory **.exrc**.

The editor obtains the horizontal and vertical size of the terminal screen from the **terminfo** database (see *terminfo*(4)). These values can be overridden by setting the **UNIX95** environment variable, which specifies to use the XPG4 behavior for this command. **COLUMNS** and **LINES** environment variables. See the **window** editor option below for more information.

The following table shows the defaults that differ for the various editor personalities:

| Name | Default Editor Options | | | | |
|-------|------------------------|--------|------------|----------|------------|
| edit | nomagic | novice | noreadonly | report=1 | showmode |
| ex | magic | novice | noreadonly | report=5 | noshowmode |
| vedit | nomagic | novice | noreadonly | report=1 | showmode |
| vi | magic | novice | noreadonly | report=5 | noshowmode |
| view | magic | novice | readonly | report=5 | noshowmode |

Editor options are Boolean unless otherwise specified. Abbreviations are shown in parentheses.

autoindent (ai) Indent each line in input mode (using blanks and tabs) to align with the previous line. Indentation begins after the line appended, or before the line inserted or the first line changed. Additional indentation can be provided as usual. Succeeding lines are automatically indented to the new alignment.

Reducing the indent is achieved by typing **^D** one or more times: the cursor is moved back to the next multiple of **shiftwidth** spaces for each **^D**. A **^** followed by a **^D** removes all indentation temporarily for the current line. A **0** followed by a **^D** removes all indentation.

Reversed by **noautoindent (noai)**. The default is **noautoindent**.

autoprint (ap) The current line is printed after each command that changes work area text. Auto-print is suppressed in **global** commands. Reversed by **noautoprint (noap)**. The default is **autoprint**.

autowrite (aw) The work area is written out to the current file if the work area has been modified and a **next**, **rewind**, or **!** command is given. Reversed by **noautowrite (noaw)**. The default is **noautowrite**.

beautify (bf) Cause all control characters other than tab, newline, and formfeed to be discarded from the input text. Reversed by **nobeautify (nobf)**. The default is **nobeautify**.

directory=dirname (dir)

Specify the directory in which the editor work area should be placed. This option only takes effect when a new work area is created. It should be set in **EXINIT** or **.exrc** to affect the location of the work area file for the edit file specified on the command line. The default is **/var/tmp**.

If the specified directory is set from **EXINIT** or a **.exrc** file and is not writable by the user, the editor quits; if set interactively by the user, the editor issues an error

message.

doubleescape When set, two consecutive ESC (escape) characters are required to leave input mode. In input mode, a single ESC character followed by a different character causes **vi** to issue an audible or visual warning (see the **flash** editor option) and insert both characters into the work area. Reversed by **nodoubleescape**. The default is **nodoubleescape**.

The character sequences transmitted by the keyboard editing keys of some terminals are identical to some sequences of **vi** user commands. If the mapping of these keys is enabled (see the **keyboardeit** and **keyboardeit!** options), **vi** might not be able to reliably distinguish between the character sequence transmitted by an editing key and the same character sequence typed by a user. This problem is most likely to occur when the user types ESC to terminate input mode immediately followed by another **vi** command. If you set the **doubleescape** option, the ambiguity of this case is removed.

edcompatible (**ed**)

Cause the presence of **g** and **c** suffixes on substitute commands to be remembered, and toggled by repeating the suffixes. Reversed by **noedcompatible** (**noed**). The default is **noedcompatible**.

errorbells (**eb**) When set, error messages are preceded with a bell only on terminals that do not support a standout or highlighting mode such as inverse video. If the terminal supports highlighting, the bell is never used prior to error messages and this option has no effect. Note that visual-mode errors are signaled by the bell (regardless of the setting of this option) without an accompanying error message.

Reversed by **noerrorbells** (**noeb**). The default is **noerrorbells**.

exrc

When set, the **.exrc** file in the current directory is processed during editor initialization if the current directory is not the **HOME** directory. This option is not set by default and must be set in the **EXINIT** environment variable or the **HOME** directory **.exrc** file to have any effect. See the Editor Options introductory text above. Reversed by **noexrc**. The default is **noexrc**.

flash (**fl**)

When set, the screen flashes instead of beeping, provided an appropriate **flash_screen** entry is present in the **/usr/share/lib/terminfo** database for the terminal being used. Reversed by **noflash** (**nofl**). The default is **flash**.

hardtabs=number (**ht**)

Define the spacing between hardware tab settings and the number of spaces used by the system when expanding tab characters. Tab stops are placed in each column number (starting at the left edge of the screen) that corresponds to an integer multiple of **number**. The default is **hardtabs=8**.

ignorecase (**ic**) All uppercase characters in the text are mapped to lowercase in regular expression matching. Also, all uppercase characters in regular expressions are mapped to lowercase, except in character class specifications. Reversed by **noignorecase** (**noic**). The default is **noignorecase**.

keyboardeit

When set, any keyboard editing key mappings that are loaded automatically at initialization for command-mode use are enabled. If not set, these mappings are disabled (but not deleted). Use the **map** command to get a list of the currently enabled command-mode mappings. Reversed by **nokeyboardeit**. The default is **keyboardeit**.

keyboardeit!

When set, the keyboard editing key mappings automatically loaded at initialization for input mode use are enabled. If not set, these mappings are disabled (but not deleted). Use the **map!** command to list the currently enabled input-mode mappings. Reversed by **nokeyboardeit!**. The default is **nokeyboardeit!** for terminals whose keyboard editing keys send HP-style escape sequences (an ESC followed by a single letter). The default is **keyboardeit!** for all other terminals.

lisp

Modify **autoindent** mode and the **(,)**, **[,]**, **{ , }** commands in visual mode for **lisp** source code. Reversed by **nolisp**. The default is **nolisp**.

list

Display all printed lines with tabs shown as **^I**, and the end of line marked by a **\$**. Reversed by **nolist**. The default is **nolist**.

| | |
|--------------------------------------|--|
| magic | Affect the interpretation of characters in regular expressions and substitution replacement strings (see Regular Expressions and Replacement Strings above). Reversed by nomagic . The ex , vi , and view default is magic . The edit and vedit default is nomagic . |
| mesg | Allows other users to use the write command (see <i>write(1)</i>) to send messages to your terminal, possibly disrupting the screen display. Unsetting this option (nomesg) blocks write permission to your terminal from other system users while you are using the editor. Reversed by nomesg . The default is mesg . |
| modelines (ml) | If set when the editor reads in a file, any ex commands embedded in the first five and last five lines of the file are executed after .exrc and EXINIT commands are processed but before editing control is given to the user. The ex commands must be prefixed by ex: or vi: and terminated by : in a single line. Any number of other characters with the exception of the colon (:) can precede or follow the embedded command. Reversed by nomodelines (noml) . The default is nomodelines . |
| novice | Use the version of the editor available for novices, known as edit or vedit . Reversed by nonovice . The ex , vi , and view default is nonovice . The edit , and vedit default is novice . |
| number (nu) | Cause lines to be printed with line numbers. Reversed by nonumber (nonu) . The default is nonumber . |
| optimize (opt) | Suppress automatic carriage returns on terminals that do not support direct cursor addressing. This streamlines text output in certain situations such as when printing multiple lines that contain leading whitespace. Reversed by nooptimize (noopt) . The default is nooptimize . |
| paragraphs=pair-string (para) | <p>The value of this option is a string whose successive pairs of characters specify the names of text-processing macros that begin paragraphs. (A macro appears in the text in the form .xx, where the . is the first character in the line.)</p> <p>If any macros have a single-character name, use a space character to substitute for the missing second character in the name. To type a space character in such situations, precede the space with a backslash (\) to prevent the editor from interpreting it as a delimiter.</p> <p>The default is paragraphs=IPLPPPQPP\ LIpplpipnbp.</p> |
| prompt | When set, command mode input is prompted for with a colon (:); when unset, no prompt is displayed. Reversed by noprompt . The default is prompt . |
| readonly (ro) | Set the readonly flag for the file being edited, thus preventing accidental overwriting at the end of the session. This option is equivalent to invoking ex , edit , vi , or vedit with the -R option or using the view command. Reversed by noreadonly (nor) . The ex , edit , vi , and vedit default is noreadonly . The view default is readonly . |
| redraw | Simulate an intelligent terminal on a dumb terminal. During input mode, lines are continuously reprinted as text is entered. Since this is likely to require a large amount of output to the terminal, it is useful only at high transmission speeds. If noredraw is set, lines are reprinted only when input mode is terminated and deleted lines are marked with an @ in the left margin. Reversed by noredraw . The default is redraw . |
| remap | If set, macro translation allows for macros defined in terms of other macros; translation continues until the final product is obtained. If unset, a one-step translation only is done. Reversed by noremap . The default is remap . |
| report=n | The value of n gives the number of lines that must be changed by a command before a report is displayed on the number of lines affected. If n is 5, then changes are reported for 6 or more lines. The ex , vi , and view default is report=5 . The edit , and vedit default is report=1 . |
| scroll=n | The value of n determines the number of lines scrolled by a ^D command and the number of lines displayed by the z command (twice the value of scroll). The default is half the value of the window option. |

sections=*pair-string*

The value of this option is a string, in that successive pairs of characters specify the names of text-processing macros that begin sections. See the **paragraphs** editor option above. The default is **sections=NHSHH\ HUuhsh+c**.

shell=*filename* (**sh**)

Set the file name of the shell to be used for the **!** shell escape and the **shell** command. It defaults to the value of your **SHELL** environment variable, if set, and otherwise to **/usr/bin/sh**.

shiftwidth=*n* (**sw**)

Sets the indentation step value used by **autoindent** and the shift (**<** and **>**) commands. The default is **shiftwidth=8**.

showmatch (**sm**)

In visual mode, jump momentarily to the matching **(** or **{** when you type a **)** or **}**, if the match is still on the screen. Reversed by **noshowmatch** (**nosm**). The default is **noshowmatch**.

showmode (**smd**)

Display the current editor mode (such as **INPUT MODE**, **REPLACE 1 CHAR**, **REPLACE MODE**) in the lower right-hand corner of the screen during visual and open mode. Reversed by **noshowmode** (**nosmd**). The **ex**, **vi**, and **view** default is **noshowmode**. The **edit**, and **vedit** default is **showmode**.

slowopen (**slow**)

In visual mode, **slowopen** prevents screen updates during input to improve throughput on unintelligent terminals. Reversed by **noslowopen** (**noslow**). The default is **noslowopen**.

tabstop=*n* (**ts**)

Sets the spacing of the software tab stops used by the editor to expand tabs in the input file. The default is **tabstop=8**.

taglength=*n* (**tl**)

Set the maximum number of characters that should be treated as significant in a tag. Characters beyond the limit are ignored. A value of zero means that all characters in the tag are significant. The default is **taglength=0**.

tags=[*filename*] ...

Specify the tags files to be used by the **tag** command and the **-t** command-line option. The default is **tags=tags /usr/lib/tags**, specifying the file **tags** in the current directory and the file **/usr/lib/tags**. File names are separated by whitespace.

Each line of a tags file contains the following three fields separated by whitespace: the tag name, the name of the file to be edited, and an address specification (see Addressing above). A **tags** file must be sorted in order by tag name.

The **ctags** command (see **ctags(1)**) creates tags files from C, Pascal and FORTRAN source files.

tagstack (**tgst**)

Enable the pushdown stack of activated tags. Reversed by **notagstack** (**notgst**). The default is **tagstack**.

When you enter a line mode **tag** command or visual mode **^]** command, the current line number and file name are stored on the tag stack. A future line mode **pop** command or visual mode **^T** command will return to the stored file name at the stored line number.

If the tag stack is disabled and then reenabled again, the stack continues where it left off. The **pop** command does not work when the tag stack is disabled.

term=*termtype*

Define the type of terminal being used with the editor. The default value is obtained from the **TERM** environment variable. If **TERM** is unset or null, **term** is set to **unknown**. There is no difference between the **term** and **ttytype** editor options. Setting either one results in both being changed.

terse

Use shorter error messages. Reversed by **noterse**. The default is **noterse**.

timeout (**to**)

If set, require that all the characters of a multicharacter macro name (the first argument in a **map** command) must be received within the amount of time specified by the **timeoutlen** option in order to be accepted as a match for the macro name. If not set, no limit is placed on how long to wait for the completion of a macro name. Reversed by **notimeout** (**noto**). The default is **timeout**.

| | |
|---|--|
| timeoutlen = <i>n</i> | Set, in milliseconds (ms), the length of the macro timeout period (see the timeout editor option). This option has no effect unless timeout is set. The value of <i>n</i> must be at least 1. The default is timeoutlen =500 (half a second). |
| ttytype = <i>termtype</i> (tty) | Define the type of terminal being used with the editor. See the term editor option for details. There is no difference between the term and ttytype editor options. Setting either one results in both being changed. |
| warn | Before executing a ! or shell command escape, display the message [No write since last change] if the work area has been modified since it was last loaded or fully written to a file. Reversed by nowarn . The default is warn . |
| window = <i>lines</i> (wi) | Set the number of lines in a text window in visual mode. The default value is one less than the size of your terminal screen (as defined by the LINES environment variable, if set, or the entry for your terminal in the <i>terminfo</i> (4) data base otherwise). However, if the terminal baud rate (see <i>stty</i> (1)) is set to less than 1200 or 2400, the default value is reduced to a maximum of 8 or 16 lines, respectively. The startup value can be specified with the -w command-line option. |
| w300 = <i>lines</i> | If the terminal baud rate is less than 1200, set the window editor option to the value specified. |
| w1200 = <i>lines</i> | If the terminal baud rate is greater than or equal to 1200 but less than 2400, set the window editor option to the value specified. |
| w9600 = <i>lines</i> | If the terminal baud rate is greater than or equal to 2400, set the window editor option to the value specified. |
| wrapmargin = <i>n</i> (wm) | In visual mode only, if <i>n</i> is greater than zero, a newline is automatically inserted in an input line at a word boundary, so that lines end at least <i>n</i> spaces from the right margin of the terminal screen. The default is wrapmargin =0. |
| wrapscan (ws) | When set, editor searches using /re/ (or ?re?) continue silently from the beginning (or end) of the file upon reaching the end (or beginning) of the file (that is, the scan "wraps around"). When unset, editor searches stop at the beginning or the end of the file, as appropriate. Reversed by nowrapscan (nows). The default is wrapscan . |
| writeany (wa) | Inhibits the checks otherwise made before write commands, allowing a write to any file (provided the system allows it). Reversed by nowriteany (nowa). The default is nowriteany . |

EXTERNAL INFLUENCES

Environment Variables

COLUMNS This variable shall override the system-selected horizontal screen size.

LINES overrides the system-selected vertical screen size, used as the number of lines in a screenful and as the vertical screen size in visual mode.

PATH determines the search path for the shell command specified in the editor commands, **shell**, **read**, and **write**.

SHELL is variable that shall be interpreted as the preferred command-line interpreter for use in **!**, **shell**, **read**, and other commands with an operand of the form **!string**. For the **shell** command, the program shall be invoked with the single argument **-i**. For all others, it shall be invoked with the two arguments **-c** and **string**. If no **SHELL** environment variable is set, or it is set to a null string, the **sh** utility shall be used.

TERM is a variable that shall be interpreted as the name of the terminal type. If this variable is unset or null, an unspecified default terminal type shall be used.

EXINIT is a variable that shall be interpreted to contain a list of **ex** commands that are executed on editor startup, before reading the first file. The list can contain multiple commands by separating them using a vertical line (|) character.

HOME shall be interpreted as a pathname of a directory that shall be searched for an editor startup file name **.exrc**.

LC_COLLATE determines the collating sequence used in evaluating regular expressions and in processing the **tags** file. If it is not specified or is null, it defaults to the value of **LANG**.

LC_CTYPE determines the interpretation of text as single and/or multibyte characters, the classification of characters as uppercase or lowercase letters, the shifting of the case of letters, and the characters matched by character class expressions in regular expressions. If it is not specified or is null, it defaults to the value of **LANG**.

LANG determines the language in which messages are displayed. If it is not specified or is null, it defaults to "C" (see *lang(5)*).

LC_ALL determines the locale to be used to override any values for locale categories specified by the setting of **LANG** or any environment variable (beginning with **LC_**).

LC_MESSAGES determines the processing of affirmative responses and the language in which messages should be written.

If any internationalization variable contains an invalid setting, all internationalization variables default to "C" (see *environ(5)*).

When set, the **TMPDIR** environment variable specifies a directory to be used for temporary files, overriding the default directory **/var/tmp**.

International Code Set Support

Single- and multibyte character code sets are supported.

ASYNCHRONOUS EVENTS (XPG4 Only)

The following actions shall be taken upon receipt of signals:

SIGINT

When an interrupt occurs, **ex** shall alert the terminal and write a message. The current editor command shall be aborted, and **ex** shall return to the command level and prompt for another command. If the standard input is not a terminal device, **ex** shall exit at the interrupt and return a nonzero exit status.

SIGCONT

The screen shall be refreshed.

SHIGHUP

If the current buffer has changed since the last **e** or **w** command, **ex** shall attempt to save the current file in a state such that it can be recovered later by an **ex -r** command.

The action taken for all other signals is unspecified.

EXTENDED DESCRIPTION (XPG4 Only)

The pathname of the file being edited by **ex** is referred to as the **current** file. The text of the file shall be read into a working version of the file (called **buffer** in this clause), and all editing changes shall be performed on that version; the changes shall have no effect on the original file until an **ex** command causes the file to be written out. Lines in the buffer may be limited to { **LINE_MAX** } bytes, and an error message may be written if the limit is exceeded during editing.

The **alternate** pathname is the name of the last file mentioned in an editor command, or the previous current pathname if the last file mentioned became the current file. When the **%** appears in a pathname entered as part of a command argument, it shall be replaced by the alternate pathname. Any character, including **%** and **#** shall retain its literal value when preceded by a backslash.

When an error occurs, **ex** shall alert the terminal and write a message.

If the system crashes, **ex** shall attempt to preserve the buffer if any unwritten changes were made. The command-line option **-r** can be used to retrieve the saved changes.

During initialization (before the first file is read or any user commands from the terminal are processed), if the environment variable **EXINIT** is set, the editor shall execute **ex** commands contained in that variable. If the variable is not set, **ex** shall attempt to read commands from the **\$HOME/.exrc**. If and only if **EXINIT** or **\$HOME/.exrc** sets the editor option **exrc**, **ex** finally shall attempt to read commands from a file **.exrc** in the current directory. In the event that **EXINIT** is not set and the current directory is the home directory of the user, any **.exrc** file shall only be processed once. No **.exrc** shall be read unless it is owned by the same user ID as the effective user ID of the process. After any **.exrc** files are processed, any commands specified by the **-c** option shall be processed.

By default, **ex** shall start in the command mode, which shall be indicated by the ":" prompt. The input mode can be entered by **append**, **insert**, or **change** commands. There is one other mode, visual mode, in which full screen editing is available. This is described more fully under the **visual** command. The command line can consist of multiple **ex** commands separated by vertical-line characters(|). The use of commands that enter input or visual modes in this manner, unless they are the final command on the line, produces undefined results.

Command lines beginning with the double-quote character (") shall be ignored. This can be used for comments in an editor script.

WARNINGS

The **undo** command causes all marks to be lost on lines that are changed and then restored.

The **z** command prints a number of logical rather than physical lines. More than a screenful of output can result if long lines are present.

Null characters are discarded in input files and cannot appear in resultant files.

On some systems, the recovery of an edit file with the **-r** option is possible only if certain system-dependent actions are taken when the system is restarted.

Edit preserve files can only be recovered on systems running the same HP-UX release in which they were preserved. Preserve files are not recoverable across different releases.

On HP terminals, the attribute field of any function key specified by a **map #n ...** command should be set to **normal** rather than to the default of **transmit**.

Do not use the **-C** option to edit unencrypted files. The **-C** option is meant to be used only on files that are already encrypted. If the **-C** option is used on files which are not yet encrypted, a write in the edit session is likely to corrupt the file.

For information about line length limits, file size limits, etc., see the WARNINGS section of *vi(1)*.

EXIT STATUS (XPG4 Only)

The **ex** utility shall exit with one of the following values:

- 0 Successful completion.
- >0 An error occurred.

AUTHOR

ex was developed by the University of California, Berkeley. The 16-bit extensions to **ex** are based in part on software of the Toshiba Corporation.

FILES

| | |
|------------------------------------|--------------------------------------|
| \$HOME/.exrc | Primary editor initialization file |
| ./ .exrc | Secondary editor initialization file |
| /usr/sbin/expreserve | Preserve command |
| /usr/sbin/exrecover | Recover command |
| /usr/share/lib/terminfo/*/* | Description of terminal capabilities |
| /var/preserve | Preservation directory |
| /var/tmp/Ex nnnnn | Editor temporary file |
| /var/tmp/Rx nnnnn | Named buffer temporary file |

SEE ALSO

ctags(1), *ed(1)*, *stty(1)*, *vi(1)*, *write(1)*, *terminfo(4)*, *environ(5)*, *lang(5)*, *regex(5)*.

The Ultimate Guide to the vi and ex Text Editors, Benjamin/Cummings Publishing Company, Inc., ISBN 0-8053-4460-8, HP part number 97005-90015.

STANDARDS COMPLIANCE

ex: SVID2, SVID3, XPG2, XPG3, XPG4

NAME

expand, unexpand - expand tabs to spaces, and vice versa

SYNOPSIS

expand [-t *tablist*] [*file* ...]

unexpand [-a] [-t *tablist*] [*file* ...]

Obsolescent:

expand [-*tabstop*] [-*tab1*,*tab2*,..., *tabn*] [*file* ...]

DESCRIPTION

expand processes the named files or the standard input and writes to the standard output with tabs changed into spaces. Backspace characters are preserved in the output, and the column count is decreased by one column for tab calculations. For proper tab calculation, if a multi-column character is to be "backspace'd", it should be followed by multiple backspace characters which equal to it's column width. If a tab character is found after the last tab position, it is replaced by a single space. **expand** is useful for preprocessing character files that contain tabs (before sorting, looking at specific columns, etc).

expand recognizes the following command-line options and arguments:

[-t *tablist*] *tablist* specifies where to set the tab positions instead of the default 8. *tablist* can take two forms. If it is a single number, tabs are set *tablist* spaces apart. *tablist* can also be a blank- or comma-separated list of increasing positions where tabs are to be set.

[-*tabstop*] This option is obsolescent and is equivalent to using -t *tabstop*.

[-*tab1*,*tab2*,..., *tabn*] This option is obsolescent and is equivalent to using -t *tab1*,*tab2*, ... ,*tabn*.

unexpand processes the named files or the standard input and writes to the standard output with spaces changed into tabs where possible. By default, only leading spaces and tabs are converted to maximal strings of tabs. The default tab position is every 8 characters. Backspace characters are preserved into the output, and the column count is decreased by one column for tab calculations. For proper tab calculation, if a multi-column character is to be "backspace'd", it should be followed by multiple backspace characters which equal to it's column width.

unexpand recognizes the following command-line options and arguments:

-a Tabs are inserted whenever they would compress the resultant file by replacing two or more spaces before a tab position.

-t *tablist* *tablist* specifies the tab positions. *tablist* can take two forms. If it is a single number, tabs are set every *tablist* spaces apart. If *tablist* is a blank- or comma-separated list of increasing positions, tabs are set at those locations. The -t option implies the -a option. If the -t option is not specified, the default is equivalent to specifying -t 8 except that -a is not implied for this case.

EXTERNAL INFLUENCES**Environment Variables**

LC_CTYPE determines the interpretation of text as single and/or multi-byte characters.

LC_MESSAGES determines the language in which messages are displayed.

If **LC_CTYPE** or **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**.

If any internationalization variable contains an invalid setting, **expand** and **unexpand** behave as if all internationalization variables are set to "C". See *environ(5)*.

If **LC_ALL** is set to a non-empty string value, it overrides the values of all the other internationalization variables.

International Code Set Support

Single- and multi-byte character code sets are supported with the exception that **unexpand** do not recognize multi-byte alternative space characters.

expand(1)

expand(1)

STANDARDS CONFORMANCE

expand: XPG4, POSIX.2

unexpand: XPG4, POSIX.2



e

NAME

`expand_alias` - recursively expands the sendmail aliases

SYNOPSIS

`expand_alias [-rmax_recursion] [-t] [-tt] alias`

DESCRIPTION

`Expand_alias` is a shell script that recursively expands the `sendmail` aliases. Through use of `telnet host 25` and the `expn` command, each alias is recursively expanded into its destination(s). Indentation is used to show each level of recursion. Because of the recursive use of `telnet`, `expand_alias` is slow. If the local `telnet` cannot directly connect to a remote system, due to a firewall configuration, `expand_alias` will not be able to succeed. If the local `telnet` is to transparently connect across the firewall, `expand_alias` will be able to contact sendmail daemons outside the firewall, allowing the alias to be more fully expanded. (For example, some local telnet clients use a `socksd` located on the firewall to permit the local `telnet` client to transparently connect to Internet hosts. If the local default `telnet` uses a `socksd` in such a manner, `expand_alias` will use that `telnet` functionality to more fully expand an alias.)

`max_recursion` defaults to 10. After `max_recursion` expansions, no further expansion is attempted.

If `-t` is specified, only the terminal aliases will be displayed.

`-tt` is similar to `-t` except that if a terminal line has a pipe, its printing is suppressed and the previous level of expansion is printed instead.

EXAMPLES

```
expand_alias root
expand_alias root@cat
expand_alias root@cat.cup.hp.com
expand_alias root@cup.hp.com
```

AUTHOR

`expand_alias` was developed by the Hewlett-Packard Company.

NAME

expr - evaluate arguments as an expression

SYNOPSIS

expr *arguments*

DESCRIPTION

expr takes *arguments* as an expression, evaluates, then writes the result on the standard output. Terms in the expression must be separated by blanks. Characters special to the shell must be escaped. Note that 0, rather than the null string, is returned to indicate a zero value. Strings containing blanks or other special characters should be quoted. Integer-valued arguments can be preceded by a unary minus sign. Internally, integers are treated as 32-bit, 2's complement numbers.

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.

expr \ | expr Returns the first *expr* if it is neither null nor 0, otherwise returns the second *expr*.

expr \& expr Returns the first *expr* if neither *expr* is null or 0, otherwise returns 0.

expr { =, >, >=, <, <=, != } expr

If both arguments are integers, and if the comparison is satisfied, *expr* returns 1 otherwise it returns 0. *expr* returns the result of an integer comparison if both arguments are integers; otherwise returns the result of a lexical comparison (note that = and == are identical, in that both test for equality).

expr { +, - } expr

Addition or subtraction of decimal integer-valued arguments.

expr { *, /, % } expr

Multiplication, division or remainder of decimal integer-valued arguments producing an integer result.

expr : expr

The matching operator **:** compares the first argument with the second argument which must be a regular expression. *expr* supports the Basic Regular Expression syntax (see *regex(5)*), except that all patterns are "anchored" (i.e., 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.

length expr The length of *expr*.

substr expr expr expr

Takes the substring of the first *expr*, starting at the character specified by the second *expr* for the length given by the third *expr*.

index expr expr Returns the position in the first *expr* which contains a character found in the second *expr*.

match Match is a prefix operator equivalent to the infix operator **:**.

\(... \)

Grouping symbols. Any expression can be placed within parentheses. Parentheses can be nested to a depth of **EXPR_NEST_MAX** as specified in the header file **<limits.h>**.

EXTERNAL INFLUENCES**Environment Variables**

LC_COLLATE determines the collating sequence used in evaluating regular expressions and the behavior of the relational operators when comparing string values.

LC_CTYPE determines the interpretation of text as single- and/or multi-byte characters, and the characters matched by character class expressions in regular expressions.

LANG determines the language in which messages are displayed.

If **LC_COLLATE** or **LC_CTYPE** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **expr** behaves as if all internationalization variables are set to "C" (see *environ(5)*).

International Code Set Support

Single- and multi-byte character code sets are supported.

RETURN VALUE

As a side effect of expression evaluation, *expr* returns the following exit values:

- 0 Expression is neither null nor zero.
- 1 Expression is null or zero.
- 2 Invalid expression.
- >2 An error occurred while evaluating the expression.

DIAGNOSTICS

| | |
|-----------------------------|----------------------------------|
| syntax error | Operator or operand errors |
| non-numeric argument | Arithmetic attempted on a string |

EXAMPLES

Add 1 to the shell variable **a**:

```
a=`expr $a + 1`
```

For **\$a** equal to either **/usr/abc/file** or just **file**, return the last segment of a path name (i.e., **file**). Beware of **/** alone as an argument because *expr* interprets it as the division operator (see WARNINGS below):

```
expr $a : '.*\/(.*\)' \| $a
```

A better representation 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 : '.*'
```

WARNINGS

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 = '='
```

resembles:

```
expr = = =
```

as the arguments are passed to *expr* (and they will all be taken as the **=** operator). The following works:

```
expr X$a = X=
```

AUTHOR

expr was developed by OSF and HP.

SEE ALSO

sh(1), test(1), environ(5), lang(5), regexp(5).

STANDARDS CONFORMANCE

expr: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

factor, **primes** - factor a number, generate large primes

SYNOPSIS

factor [*number*]

primes [*start* [*stop*]]

DESCRIPTION

If no arguments are provided on the command line, **factor** waits for a number to be typed in. If a positive number is typed, it factors the number and print its prime factors; each one is printed the proper number of times. It then waits for another number. **factor** exits if it encounters a zero or any non-numeric character.

If an argument is provided on the command line, **factor** factors the number as above, then exits.

Maximum time to factor is proportional to \sqrt{n} and occurs when n is prime or the square of a prime.

The largest number that can be dealt with by **factor** is 1.0e14.

primes prints prime numbers between a lower and upper bound. If no arguments are provided on the command line, **primes** waits for two numbers to be typed in. The first number is interpreted as the lower bound; the second as the upper bound. All prime numbers in the resulting inclusive range are printed.

If *start* is specified, all primes greater than or equal to *start* are printed. If both *start* and *stop* are given, all primes occurring in the inclusive range *start* through *stop* are printed.

start and *stop* values must be integers represented as long integers.

If the stop value is omitted in either case, **primes** runs either until overflow occurs or until it is stopped by typing the interrupt character.

The largest number that can be dealt with by **primes** is 2,147,483,647.

DIAGNOSTICS

Both commands print **Ouch** when the input is out of range, illegal characters are encountered, or when *start* is greater than *stop*.

EXAMPLES

Print the prime factorization for the number 12:

```
factor 12
```

Print all prime numbers between 0 and 20:

```
primes 0 20
```

NAME

fastbind - prepare an incomplete executable for faster program start-up

SYNOPSIS

fastbind [-nu] *incomplete-executable...*

DESCRIPTION

fastbind is a tool that can improve the start-up time of programs that use shared libraries (incomplete executables) by storing information about needed shared library symbols in the executable file.

fastbind performs analysis on the symbols used to bind an executable and all of its dependent shared libraries, and stores this information in the executable file. The next time the executable is run, the dynamic loader (`/usr/lib/dld.sl` for 32-bit PARISC or `/usr/lib/pa20_64/dld.sl` for 64-bit PARISC) will notice that this information is available, and it will use this fastbind information to bind the executable instead of the standard search method for binding the symbols.

Since **fastbind** writes the fastbind information in the executable file, you must have write permission on the executable file. Also, if the executable file being analyzed is being run as another process, the file will be locked against modifications by the kernel, and **fastbind** will fail.

If the shared libraries that an executable is dependent on are modified after the fastbind information is created, the dynamic loader will silently revert to standard search method for binding the symbols. The fastbind information can be re-created by running **fastbind** on the executable again. **fastbind** will automatically erase the old fastbind information and generate the new one.

The **ld** option **+fb** can be used to instruct the linker to run the fastbind tool on an incomplete executable it has produced.

Environment Variables

If **dld** determines that the fastbind information is out of date, it will silently revert to standard search method for binding the symbols. If the environment variable **_HP_DLDOPTS** is set to **-fbverbose** the dynamic loader will emit a warning message when the fastbind information is out of date.

The environment variable **_HP_DLDOPTS** can be set to **-nofastbind** to make the dynamic loader ignore the fastbind information and revert to the standard search method for binding the symbols.

Options

fastbind recognizes the following options:

- n** Remove the fastbind information from the executable, returning it to the same state it was in before **fastbind** was originally run on it.
- u** Normally, if **fastbind** detects any unsatisfied symbols while building the fastbind information, it will generate an error message and not modify the executable file. When **fastbind** is invoked with **-u** option however, unresolved symbols are allowed.

EXTERNAL INFLUENCES**Environment Variables**

The following internationalization variables affect the execution of **fastbind**:

LANG

Determines the locale category for native language, local customs and coded character set in the absence of **LC_ALL** and other **LC_*** environment variables. If **LANG** is not specified or is set to the empty string, a default of **C** (see *lang(5)*) is used instead of **LANG**.

LC_ALL

Determines the values for all locale categories and has precedence over **LANG** and other **LC_*** environment variables.

LC_MESSAGES

Determines the locale that should be used to affect the format and contents of diagnostic messages written to standard error.

LC_NUMERIC

Determines the locale category for numeric formatting.

LC_CTYPE

Determines the locale category for character handling functions.

NLSPATH

Determines the location of message catalogs for the processing of `LC_MESSAGES`.

If any internationalization variable contains an invalid setting, **fastbind** behaves as if all internationalization variables are set to `C`. See *environ(5)*.

In addition, the following environment variable affects **fastbind**:

TMPDIR

Specifies a directory for temporary files (see *tmpnam(3S)*).

DIAGNOSTICS

fastbind returns zero when the operation is successful. A non-zero return code indicates that an error occurred.

EXAMPLES

To run **fastbind** on the executable file **a.out** enter:

```
fastbind a.out
```

To later remove the fastbind information from the executable file **a.out** enter:

```
fastbind -n a.out
```

WARNINGS

32-bit PARISC **fastbind** does not work with EXEC_MAGIC executables.

fastbind effectively enforces bind restricted and bind immediate. For example, consider an executable linked bind deferred, which calls a function `foo()` defined in an implicitly loaded library. Before the actual call is made, if it explicitly loads a shared library (using *shl_load(3X)* with **BIND_FIRST**) having a definition for `foo()`, when `foo()` is finally called, it will be resolved from the explicitly loaded library. But after running **fastbind**, the symbol `foo()` will be resolved from the implicitly loaded library.

AUTHOR

fastbind was developed by Hewlett-Packard.

FILES

| | |
|---|------------------------------|
| a.out | output file |
| <code>/usr/lib/dld.sl</code> | 32-bit PARISC dynamic loader |
| <code>/usr/lib/pa20_64/dld.sl</code> | 64-bit PARISC dynamic loader |
| <code>/usr/lib/nls/\$LANG/fastbind.cat</code> | message catalog |
| <code>/var/tmp/___FB*</code> | temporary files |

SEE ALSO**System Tools:**

ld(1) invoke the link editor

Miscellaneous:

a.out(4) assembler, compiler, and linker output
dld.sl(5) dynamic loader

Texts and Tutorials:

HP-UX Linker and Libraries User's Guide

NAME

fastmail - quick batch mail interface

SYNOPSIS

```
fastmail [-b bcc-list] [-c cc-list] [-C comments] [-f from-name] [-F from-addr] [-i in-reply-to]
[-r reply-to] [-R references] [-s subject] filename address-list
```

DESCRIPTION

The **fastmail** command is a simple interface to the mail system that allows you to send a message without the overhead of an interactive mailer. It is particularly efficient in batch-processing mail to very large groups of people.

All addresses should be full e-mail addresses, **sendmail** aliases in the `/etc/mail/aliases` file, or local login names.

Options

fastmail recognizes the following options:

- b *bcc-list*** Include a **Bcc:** header entry. Send blind carbon copies to the comma-separated list of addresses in *bcc-list*.
- c *cc-list*** Include a **Cc:** header entry. Send carbon copies to the comma-separated list of addresses in *cc-list*.
- C *comments*** Include a **Comments:** header entry with the string value *comments*.
- d** Debug. Display information on processing steps.
- f *from-name*** Replace the user name in the **From:** header entry with *from-name*.
If the user is **x@y**, and the user name is **MrX**, then the default **From:** line is:
From: x@y (MrX) .
The option **-f Joe** changes it to:
From: x@y (Joe)
- F *from-addr*** Replace the address in the **From:** header entry with *from-addr*. In the **-f** example above, **-F a@b** changes the original entry to
From: a@b (MrX)
- i *in-reply-to*** Include the **In-Reply-To:** header entry with the string value *in-reply-to*. This is usually used to identify a message that you are replying to.
- r *replyto*** Include the **Reply-To:** header entry with the single address given in *replyto*. This is the address where replies will usually be sent, instead of to the address given in the **From:** header entry, very common with mailing lists.
- R *references*** Include a **References:** header entry containing the string value *references*.
- s *subject*** Include a **Subject:** header entry containing the value *subject*. If this option is omitted, the message is sent without a subject entry.

Operands

fastmail recognizes the following operands:

- address-list* A list of one or more blank-separated addresses for the **To:** header line. These are the principal recipients of the message.
- filename* Either the name of a file containing the message, or a dash (-) to read from standard input.

EXAMPLES**A Fully Specified Command**

This command has every option specified.

```
fastmail \
-b "bcc1,bcc2,bcc3,bcc4" \
-C "Just a Comment" \
-c "cc1,cc2,cc3,cc4" \
```

```

-d \
-F me@anotherhost.com \
-f My Name \
-i "Your recent message" \
-R REF:13579 \
-r oscar \
-s "Testing fastmail" \
message-file \
addr1 addr2 addr3 addr4

```

The online execution displays the following debug messages:

```

Mailing to addr1,addr2,addr3,addr4 cc1,cc2,cc3,cc4 bcc1,bcc2,bcc
3,bcc4 [via sendmail]
cat /tmp/fastmail.5578 message-file | /usr/sbin/sendmail addr1,a
ddr2,addr3,addr4 cc1,cc2,cc3,cc4 bcc1,bcc2,bcc3,bcc4

```

The received message has the following relevant header entries:

```

From realsender@mycomputer.myhost.com Tue Oct 22 21:14:04 EDT 1996
Subject: Testing fastmail
From: me@anotherhost.com (My Name)
Reply-To: oscar@mycomputer.myhost.com
To: addr1@mycomputer.myhost.com, addr2@mycomputer.myhost.com,
    addr3@mycomputer.myhost.com, addr4@mycomputer.myhost.com
Cc: cc1@mycomputer.myhost.com, cc2@mycomputer.myhost.com,
    cc3@mycomputer.myhost.com, cc4@mycomputer.myhost.com
References: REF:13579
In-Reply-To: Your recent message
Comments: Just a Comment

```

The Bcc: header entry is not transmitted.

A Batch Process

Suppose you are user **big** on machine **big-machine** and you have a shell script named **batch-mail** that contains the following lines:

```

#
# Batch Mail - batch mailing of a file to a LOT of users
#
# Usage: batch-mail "<from>" "<subject>" <filename>

sender_copy=$LOGIN
replyto=The-Mr-Big-list

fastmail -b $sender_copy -r $replyto -f "$1" -s "$2" $3 person1
sleep 10
fastmail -r $replyto -f "$1" -s "$2" $3 person2
sleep 10
fastmail -r $replyto -f "$1" -s "$2" $3 person3
sleep 10
fastmail -r $replyto -f "$1" -s "$2" $3 person4

```

The command:

```
batch-mail "Mr. Big" "Warning to all" warning.text
```

would mail a copy of the **warning.text** file to **person1**, **person2**, **person3**, and **person4**, staggered ten seconds apart.

\$LOGIN would also silently receive a copy of the first message in the mail. Each resultant message would include the header lines:

```

From: big@big-machine (Mr. Big)
Subject: Warning to all
Reply-To: The-Mr-Big-list

```


FILES

| | |
|---------------------------------|-------------------------------------|
| <code>/etc/mail/aliases</code> | <code>sendmail</code> aliases file. |
| <code>/usr/sbin/sendmail</code> | Mail transport agent. |
| <code>/tmp/fastmail.pid</code> | Temporary file. |

AUTHOR

`fastmail` was developed by HP.

SEE ALSO

`elm(1)`, `sendmail(1M)`.

RFC 822 "Standard for the Format of Internet Text Messages"


f

NAME

file - determine file type

SYNOPSIS

file [-m *mfile*] [-c] [-f *ffile*] [-h] *file* ...

DESCRIPTION

file performs a series of tests on each *file* in an attempt to classify it. If *file* appears to be an ASCII file, **file** examines the first 512 bytes and tries to guess its language. If *file* is an executable **a.out** file, **file** prints the version stamp, provided it is greater than 0 (see the description of the **-V** option in *ld(1)*).

file uses the file **/etc/magic** to identify files that have some sort of **magic number**, that is, any file containing a numeric or string constant that indicates its type. Commentary at the beginning of **/etc/magic** explains the format.

Options

file recognizes the following command-line options:

- m *mfile*** Use alternate magic file *mfile*.
- c** Check the magic file for format errors. This validation is not normally carried out for reasons of efficiency. No file classification is done when this option is specified.
- f *ffile*** Obtain the list of files to be examined from file *ffile*. **file** classifies each file whose name appears in *ffile*.
- h** Do not follow symbolic links.

EXTERNAL INFLUENCES**Environment Variables**

LC_MESSAGES determines the language in which messages are displayed.

If **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**.

If any internationalization variable contains an invalid setting, **file** behaves as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single- and multi-byte character code sets are supported. However, all non-ASCII text files are identified as "data".

WARNINGS

The **file** command for a release interprets the core files for that particular release correctly. Using the **file** command on a core file generated on a different release will report incorrect results.

SEE ALSO

ld(1).

STANDARDS CONFORMANCE

file: SVID2, SVID3, XPG2, XPG4

NAME

find - find files

SYNOPSIS

find *pathname_list* [*expression*]

DESCRIPTION

The **find** command recursively descends the directory hierarchy for each path name in *pathname_list* (that is, one or more path names) seeking files that match a Boolean *expression* written in the primaries given below. By default, **find** does not follow symbolic links.

The Boolean expression is evaluated using short-circuit evaluation. This means that whenever the result of a Boolean operation (AND or OR) is known from evaluating the left-hand argument, the right-hand argument is not evaluated.

In the descriptions of the primaries, the argument *n* represents a decimal integer; **+n** means more than *n*, **-n** means less than *n*, and *n* means exactly *n*.

The following primaries are recognized:

- depth** A position-independent term which 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. It is also useful when using *cpio*(1) and the modification dates of directories must be preserved. Always true.
- follow** A position-independent term which causes **find** to follow symbolic links. 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. Always true.
- fsonly** *FStype* A position-independent term which causes **find** to stop descending any directory whose file system is not of the type specified by *FStype*, where *FStype* is one of **cdfs**, **hfs**, **vxfs**, or **nfs**, representing the CDFS, HFS, JFS (VXFS) or NFS file system type, respectively.

In this context, mount points inherit the *FStype* of their parent directory. This means that when **-fsonly hfs** has been specified and **find** encounters an NFS mount point that is mounted on an HFS file system, the mount point will be visited but entries below that mount point will not. It is important to note that when **-fsonly nfs** has been specified, any HFS file systems that are beneath the mount point of an NFS file system are not traversed. Always true.
- local** True if the file physically resides on the local system. This does not restrict the search to only files which physically reside on the local system, it merely matches such files. See *EXAMPLES*.
- xdev** A position-independent term that causes **find** to avoid crossing any file system mount points that exist below starting points enumerated in *pathname_list*. The mount point itself is visited, but entries below the mount point are not. Always true.
- mountstop** Identical to **-xdev**. This primary is provided for backward compatibility only. **-xdev** is preferred over **-mountstop**.
- name** *file* True if *file* matches the last component of the current file name. The matching is performed according to Pattern Matching Notation (see *regex*(5)). Pattern may contain supplementary code set characters.
- path** *file* Same as **-name** except the full path (as would be output by **-print**) is used instead of just the base name. Note that / characters are not treated as a special case. For example, ***./profile** matches **./home/fred/.profile**.
- perm** [**-**]*mode* In this primary, the argument *mode* is used to represent file mode bits. The argument is identical in format to the *mode* operand as described in *chmod*(1), with the exception that the first character must not be the **-** operator. When using the symbolic form of *mode*, the starting template is assumed to have all file mode bits cleared.

If the leading minus is omitted, this primary is true when the file permission bits exactly match the value of *mode*. Bits associated with the symbolic attributes **s** (set-user-ID, set-group-ID) and **t** (sticky bit) are ignored when the minus is omitted.

If *mode* is preceded by a minus, this primary is true if all of the bits that are set in *mode* are also set in the file permission bits. In this case, the bits associated with the symbolic attributes **s** and **t** are significant.

-fstype *FStype*

True if the file system to which the file belongs is of type *FStype*, where *FStype* is one of **cdfs**, **hfs**, or **nfs**, corresponding to the CDFS, HFS, or NFS file system type, respectively.

-type *c*

True if the type of the file is *c*, where *c* is one of:

- f** Regular file
- d** Directory
- b** Block special file
- c** Character special file
- p** FIFO (named pipe)
- l** Symbolic link
- s** Socket
- n** Network special file
- M** Mount point

-links *n*

True if the file has *n* links.

-user *uname*

True if the file belongs to the user *uname*. If *uname* is numeric and does not appear as a login name in the **/etc/passwd** file, it is taken as a user ID. The *uname* operand can be preceded by a **+** or **-** to modify the comparison of the primaries. If the argument *n* represents a decimal integer; **+n** means more than *n*, **-n** means less than *n*, and *n* means exactly *n*.

-group *gname*

True if the file belongs to the group *gname*. If *gname* is numeric and does not appear in the **/etc/group** file, it is taken as a group ID. The *gname* operand can be preceded by a **+** or **-** to modify the comparison of the primaries. If the argument *n* represents a decimal integer; **+n** means more than *n*, **-n** means less than *n*, and *n* means exactly *n*.

-nouser

True if the file belongs to a user ID that is not listed in the password database. See **passwd(4)**.

-nogroup

True if the file belongs to a group ID that is not listed in the group database. See **group(4)**.

-size *n*[**c**]

True if the file is *n* blocks long (512 bytes per block). If *n* is followed by a **c**, the size is in bytes.

-atime *n*

True if the file access time subtracted from the initialized time is *n*-1 to *n* multiples of 24 h. The initialization time shall be a time between the invocation of the **find** utility and the first access by that invocation of the **find** utility to any file specified by its **path** operands. The access time of directories in *pathname_list* is changed by **find** itself.

-mtime *n*

True if the file modification time subtracted from the initialization time is *n*-1 to *n* multiples of 24 h. The initialization time shall be a time between the invocation of the **find** utility and the first access by that invocation of the **find** utility to any file specified in its **path** operands.

-ctime *n*

True if the time of last change of file status information subtracted from the initialization time is *n*-1 to *n* multiples of 24 h. The initialization time shall be a time between the invocation of the **find** utility and the first access by that invocation of the **find** utility to any file specified by its **path** operands.

-newer *file*

True if the current file has been modified more recently than the argument *file*.

-newer[*tv1*[*tv2*]] *file*

True if the indicated time value (*tv1*) of the current file is newer than the indicated time value (*tv2*) of *file*. The time values *tv1* and *tv2* are each selected from the set of characters:

- a** The time the file was last accessed
- c** The time the inode of the file was last modified
- m** The time the file was last modified

If the *tv2* character is omitted, it defaults to **m**. Note that the **-newer** option is equivalent to **-newermm**.

Syntax examples;

- newera file**
- newermc file**
- inum *n*** True if the file serial number (inode number) is *n*. Note that file serial numbers are unique only within a given file system. Therefore, matching file serial numbers does not guarantee that the referenced files are the same unless you restrict the search to a single file system.
- linkedto *path*** True if the file is the same physical file as the file specified by *path* (i.e., linked to *path*). This primary is similar to **-inum**, but correctly detects when a file is hard-linked to *path*, even when multiple file systems are searched.
- print** Causes the current path name to be printed. Always true.
- exec *cmd*** True if the executed *cmd* returns a zero value as exit status. The end of *cmd* must be punctuated by a semicolon (;) or a plus sign (+) (semicolon and plus are special to the shell and must be escaped). When a plus sign is used, *cmd* aggregates a set of pathnames and executes on the set. The reason for preferring + to a semicolon is vastly improved performance. Any command argument { } is replaced by the current path name. *cmd* may contain supplementary code set characters.
- ok *cmd*** Same as **-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**. The form of the affirmative response is locale dependent: y in the C locale, see LANG on environ(5). *cmd* may contain supplementary code set characters.
- cpio *device*** Write the current file on *device* in *cpio*(4) format (5120-byte records). The use of **-cpio** implies **-depth**. Always true.
- ncpio** Same as **-cpio** but adds the **-c** option to *cpio*. The use of **-ncpio** implies **-depth**. Always true.
- prune** If the current entry is a directory, cause **find** to skip that directory. This can be useful to avoid walking certain directories, or to avoid recursive loops when using **cpio -p**. Note, however, that **-prune** is useless if the **-depth** option has also been given. See the description of **-only** and the *EXAMPLES* section, below, for more information. Always true.
- only** This is a positive-logic version of **-prune**. A **-prune** is performed after every directory, unless **-only** is successfully evaluated for that directory. As an example, the following three commands are equivalent:


```
find . -fsonly hfs -print
find . -print -fstype hfs -only
find . -print ! -fstype hfs -prune
```

Note, however, that **-only** is useless if the **-depth** option has also been given. Always true.
- (*expression*)** True if the parenthesized expression is true. The spaces are required. Parentheses are special to the shell and must be escaped, as in \ (and \).

Primaries can be combined by using the following operators (in order of decreasing precedence):

- ! *expression*** Logical NOT operator. True if *expression* is not true.
- expression* [-a] *expression*** Logical AND operator. True if both of the *expressions* are true.
- expression* -o *expression*** Logical OR operator. True if either or both of the *expressions* are true.

If *expression* is omitted, or if none of **-print**, **-ok**, **-exec**, **-cpio**, or **-ncpio** is specified, **-print** is assumed. The **-user**, **-group**, and **-newer** primaries each evaluate their respective arguments once.

HFS Access Control Lists

The **-acl** primary enables the user to search for HFS access control list entries. It is true if the file's access control list matches an access control list pattern or contains optional access control list entries (see *acl(5)*). It has three forms:

- acl** *aclpatt* Match all files whose access control list includes all (zero or more) pattern entries specified by the *aclpatt* pattern.
- acl** **=***aclpatt* Match a file only if its access control list includes all (zero or more) pattern entries specified by the *aclpatt* pattern, and every entry in its access control list is matched by at least one pattern entry specified in the *aclpatt* pattern.
- acl** **opt** Match all files containing optional access control list entries.

The *aclpatt* string can be given as an operator or short form pattern; see *acl(5)*.

By default, **-acl** is true for files whose access control lists include all the (zero or more) access control list patterns in *aclpatt*. A file's access control list can also contain unmatched entries.

If *aclpatt* begins with **=**, the remainder of the string must match all entries in a file's access control list.

The *aclpatt* string (by default, or the part following **=**) can be either an access control list or an access control list pattern. However, if it is an access control list, *aclpatt* must include at least the three base entries (*user.%*, *mode*), (*%group*, *mode*), and (*%.%*, *mode*)).

As a special case, if *aclpatt* is the word **opt**, the primary is true for files with access control list entries.

JFS Access Control Lists

The **-aclv** primary enables the user to search for JFS access control list entries. It is true if the file's access control list matches an access control list pattern or contains optional access control list entries (see *aclv(5)*). It has three forms:

- aclv** *aclpatt* Match all files whose access control list includes all (zero or more) pattern entries specified by the *aclpatt* pattern.
- aclv** **=***aclpatt* Match a file only if its access control list includes all (zero or more) pattern entries specified by the *aclpatt* pattern, and every entry in its access control list is matched by at least one pattern entry specified in the *aclpatt* pattern.
- aclv** **opt** Match all files containing optional access control list entries.

By default, **-aclv** is true for files whose access control lists include all the (zero or more) access control list patterns in *aclpatt*. A file's access control list can also contain unmatched entries.

If *aclpatt* begins with **=**, the remainder of the string must match all entries in a file's access control list.

An *aclpatt* consists of a *type* field, an *ID* field, and a *mode* field, separated by colons. Multiple comma-separated *aclpatts* may be specified.

The *type* field is one of **user**, **group**, **class**, **other** or *****, optionally preceded by **default:**. **user**, **group**, **class**, **other** and **default** can be abbreviated to **u**, **g**, **c**, **o** and **d**, respectively. A *type* field of ***** matches any of the above types.

The *ID* field is either a numeric user or group ID, a user or group ID string from */etc/passwd* or */etc/group* respectively, or *****, which matches any ID.

The *mode* field consists of a string of three characters. The first character is either **r**, indicating that read permission is granted; **-**, indicating that read permission is denied; or **?**, which matches either state of read permission. The second character is either **w**, **-**, or **?**, similarly indicating the state of write permission; and the third character is either **x**, **-**, or **?**, indicating the state of execute permission.

As a special case, if *aclpatt* is the word **opt**, the primary is true for files with optional access control list entries.

EXTERNAL INFLUENCES

Environment Variables

If an internationalization variable is not specified or is null, it defaults to the value of **LANG**.

If **LANG** is not specified or is null, it defaults to **C** (see *lang(5)*).

If **LC_ALL** is set to a nonempty string value, it overrides the values of all the other internationalization variables.

If any internationalization variable contains an invalid setting, all internationalization variables default to C (see *environ*(5)).

LC_CTYPE determines the interpretation of text as single and/or multibyte characters, the classification of characters as printable, and the characters matched by character class expressions in regular expressions.

LC_MESSAGES determines the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.

NLSPATH determines the location of message catalogues for the processing of **LC_MESSAGES**.

International Code Set Support

Single- and multibyte character code sets are supported.

EXAMPLES

Search the two directories **/example** and **/new/example** for files containing the string **Where are you** and print the names of the files:

```
find /example /new/example -exec grep -l 'Where are you' {} \;
```

Remove all files named **a.out** or ***.o** that have not been accessed for a week:

```
find / \( -name a.out -o -name '*.o' \) -atime +7 -exec rm {} \;
```

Note that the spaces delimiting the escaped parentheses are required.

Print the names of all files on this machine. Avoid walking **nfs** directories while still printing the **nfs** mount points:

```
find / -fsonly hfs -print
```

Match only local files, and do not examine the contents of any directory found to be remotely mounted:

```
find / ! -local -prune -o -size +50 -print
```

This only works correctly if there are no local file systems mounted on top of remote directories. This example will print all local files on the system larger than 50 blocks, without wasting time accessing remote files.

To get the same effect, but to check for files in local file systems mounted on remote directories, use:

```
find / -local -size +50 -print
```

Copy the entire file system to a disk mounted on **/Disk**, avoiding the recursive copy problem. Both commands are equivalent (note the use of **-path** instead of **-name**):

```
cd /; find . ! -path ./Disk -only -print | cpio -pdxm /Disk
cd /; find . -path ./Disk -prune -o -print | cpio -pdxm /Disk
```

Copy the root disk to a disk mounted on **/Disk**, skipping all mounted file systems below **/**. Note that **-xdev** does not cause **/** to be skipped, even though it is a mount point. This is because **/** is the starting point and **-xdev** only affects entries below starting points.

```
cd /; find . -xdev -print | cpio -pdm /Disk
```

Change permissions on all regular files in a directory subtree to mode 444, and permissions on all directories to 555:

```
find pathname -type f -print | xargs chmod 444
find pathname -type d -print | xargs chmod 555
```

Note that output from **find** was piped to **xargs(1)** instead of using the **-exec** primary. This is because when a large number of files or directories is to be processed by a single command, the **-exec** primary spawns a separate process for each file or directory, whereas **xargs** collects file names or directory names into multiple arguments to a single **chmod** command, resulting in fewer processes and greater system efficiency. The **+** delimiter for the **-exec** primary can be used to achieve the same efficiency.

Access Control List Examples

Find all files not owned by user **karl** that have access control lists with at least one entry associated with **karl**, and one entry for no specific user in group **bin** with the read bit on and the write bit off:

```
find / ! -user karl -acl 'karl.*, %.bin+r-w' -print
```

Find all files that have a read bit set in any access control list entry:

```
find / -acl '*.+r' -print
```

Find all files that have the write bit unset and execute bit set in every access control list entry:

```
find / -acl '=*.*-w+x' -print
```

Find all files that have optional access control list entries:

```
find / -acl opt -print
```

DEPENDENCIES

NFS

The **-acl** primary is always false for NFS files.

WARNINGS

Because of interoperability goals, **cpio** does not support archiving files larger than 2GB or files that have user/group IDs larger than 60,000 (60K). Files with user/group IDs greater than 60K are archived and restored under the user/group ID of the current process.

AUTHOR

find was developed by AT&T and HP.

FILES

| | |
|--------------------|--------------|
| /etc/group | Group names |
| /etc/mnttab | Mount points |
| /etc/passwd | User names |

SEE ALSO

chacl(1), chmod(1), cpio(1), setacl(1), sh(1), test(1), xargs(1), mknod(2), stat(2), cpio(4), fs(4), group(4), passwd(4), acl(5), aclv(5), environ(5), lang(5), regexp(5).

STANDARDS CONFORMANCE

find: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

findmsg, dumpmsg - create message catalog file for modification

SYNOPSIS

```
findmsg [-aiv] [[-D sym] [-U sym]] file ...
```

```
dumpmsg file ...
```

DESCRIPTION

The **findmsg** command extracts messages from a C program source *file* and writes them to the standard output in a format suitable for input to **gencat** (see *gencat*(1)). The input file will be preprocessed using **cpp** (see *cpp*(1)) in order to select print specifiers and handle **ifdef**, **ifndef**... conditional **cpp** primitives. If multiple input files are specified and the **-a** option is not used, the files are processed sequentially such that message catalog comment lines identifying the input *file* are written before the output for each input *file*.

The **findmsg** command scans the source files for uncommented lines with one of the following three formats embedded within it:

```
catgets( any_var, NL_SETN, n, <message> )
<message>                               /* catgets n */
/* catgets n */                         <message>
```

or any combination of these formats wholly contained on a single physical line. *<message>* could be a string constant or a combination of string constants and print specifiers (PRI*). Any number of spaces or tabs can separate the **catgets** comment from the *message*. The digit *n*, which can be any valid message number (see *gencat*(1)), is combined with the *message* string to produce a message catalog source line. The message source line is assigned to the set whose number is the current value of **NL_SETN** as set by the last **#define** directive encountered. If **NL_SETN** has not yet been defined when a message line is found, the message is output without a set number specification. If more than one message is found belonging to the same set and message number, the last message found is output; any others are silently discarded. Conditional compilation and **#include** instructions in the C source files are ignored.

Options

findmsg recognizes the following command-line options:

- a** Merge identically numbered sets from multiple input files so that **gencat** can process the **findmsg** output.
- Dsym** Define symbol *sym*.
- Usym** Cause symbol *sym* to be undefined.
- i** Consider all **#ifdefs** to extract messages from the input file. Options **-D** and **-U** will be used to select print specifiers if this option is not used.
- v** Outputs all error messages issued by **cpp**. By default, **findmsg** does not display the error messages issued by **cpp**.

The **dumpmsg** command extracts messages from a message catalog *file* created by **gencat**. Messages are written to standard output in a format suitable for editing and re-input to **gencat**.

EXTERNAL INFLUENCES**Environment Variables**

LC_CTYPE determines the interpretation of messages as single-byte and/or multi-byte characters.

LC_MESSAGES determines the language in which messages are displayed.

If **LC_CTYPE** or **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of **C** (see *lang*(5)) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **findmsg** and **dumpmsg** behave as if all internationalization variables are set to **C**. See *environ*(5).

International Code Set Support

Single-byte and multi-byte character code sets are supported.

WARNINGS

The **findmsg** and **dumpmsg** commands are HP proprietary, not portable to other vendors' systems, and will not be provided in future HP-UX releases.

AUTHOR

findmsg and **dumpmsg** were developed by HP.

SEE ALSO

findstr(1), **gencat(1)**, **insertmsg(1)**, **catgets(3C)**.


f

NAME

findstr - find strings for inclusion in message catalogs

SYNOPSIS

findstr *file* ...

DESCRIPTION

findstr examines files of C source code for uncommented string constants which it places, along with the surrounding quotes, on the standard output, preceding each by the file name, start position, and length. This information is used by **insertmsg** (see *insertmsg(1)*). **findstr** does not output strings that are parameters of the **catgets**() routine (see *catgets(3C)*).

EXTERNAL INFLUENCES**Environment Variables**

LC_CTYPE determines the interpretation of comments and string literals as single- and/or multi-byte characters.

LC_MESSAGES determines the language in which messages are displayed.

If **LC_CTYPE** or **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **findstr** behaves as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single- and multi-byte character code sets are supported.

WARNINGS

findstr outputs initialization strings of static string variables. Calling **insertmsg** with these strings causes their replacement with a call to **catgets**() (see *catgets(3C)*). Since the initializer must be a string, this assignment results in an invalid C declaration. For example, the following line:

```
static char *x[] = "message"
```

is modified by **insertmsg** (see *insertmsg(1)*) to:

```
static char *x[] = (catgets(catd,NL_SETN,1,"message"))
```

These strings should be manually removed from **findstr** output before being input to **insertmsg**.

findstr will not be provided in future HP-UX releases.

SEE ALSO

insertmsg(1).

NAME

finger - user information lookup program

SYNOPSIS

finger [*options*] *user_name* ...

DESCRIPTION

By default, **finger** lists for each *user_name* on the system:

- Login name,
- Full given name,
- Terminal write status (if write permission is denied),
- Idle time,
- Login time,
- User's home directory and login shell,
- Any plan the user has placed in file **.plan** in their home directory,
- Project on which they are working from the file **.project**, also in the home directory,
- office location and phone number (if known),
- last time the user received the mail, and last time the user read the mail.

Idle time is in minutes if listed as a single integer, hours and minutes if a **:** is present, or days and hours if a **d** is present. Account names as well as first and last names of users are accepted.

finger can also be used to list users on a remote machine. The format for *user_name* is *user_name@host*. If *user_name* is not specified, the remote system (HP-UX or non-HP-UX) uses its default standard format for listing user information.

Options

finger recognizes the following options:

- b** Suppress printing the user's home directory and shell.
- f** Suppress printing the header that is normally printed in a short-format printout.
- h** Suppress printing the **.project** file in a long-format printout.
- i** Force "idle" output format. Similar to short format except that only the login name, terminal, login time, and idle time are printed.
- l** Force long output format.
- m** Match arguments only on user name.
- p** Suppress printing of the **.plan** files
- q** Force quick output format. Similar to short format except that only the login name, terminal, and login time are printed.
- R** Print the user's host name.
- s** Force short output format.
- w** Suppress printing the full name in a short-format printout.

WARNINGS

Only the first line of the **.project** file is printed.

AUTHOR

finger was developed by the University of California, Berkeley.

FILES

| | |
|---------------|-------------------------------|
| /etc/utmp | who file |
| /var/adm/wtmp | last login file |
| /etc/passwd | for users names, offices, ... |
| ~/.plan | plans |
| ~/.project | projects |
| /var/mail | mail directory |

finger(1)

finger(1)

SEE ALSO
chfn(1), who(1).

f

NAME

fmt - format text

SYNOPSIS

fmt [-**cs**] [-**w** *width*] [*file...*]

DESCRIPTION

The **fmt** command 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 *width* is 72. **fmt** concatenates the **file** arguments. If none are given, **fmt** formats text from the standard input.

Blank lines are preserved in the output, as is the spacing between words. **fmt** does not fill lines beginning with a period (.), for compatibility with **nroff**. Nor does it fill lines starting with **From:**.

Indentation is preserved in the output and input lines with differing indentation are not joined (unless **-c** is used).

fmt can also be used as an in-line text filter for **vi**; the **vi** command:

```
!}fmt
```

reformats the text between the cursor location and the end of the paragraph.

Options

fmt recognizes the following options:

- 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.
- 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.
- w** Fill output lines to up to *width* columns.

WARNINGS

The **-w** *width* option is acceptable for BSD compatibility, but it may go away in future releases.

SEE ALSO

nroff(1), **vi**(1).

NAME

fold - fold long lines for finite width output device

SYNOPSIS

fold [-b] [-s] [-w *width*] [*file* ...]

Obsolete form:

fold [-s] [-width] [*file* ...]

DESCRIPTION

The **fold** command is a filter that folds the contents of the specified files, breaking the lines to have a maximum of *width* column positions (or bytes, if the **-b** option is specified). The **fold** command breaks lines by inserting a newline character so that each output line is the maximum width possible that does not exceed the specified number of column positions (or bytes). A line cannot be broken in the middle of a character. If no files are specified or if a *file* name of - is specified, the standard input is used.

The **fold** command is often used to send text files to line printers that truncate, rather than fold, lines wider than the printer is able to print.

If the backspace, tab, or carriage-return characters are encountered in the input, and the **-b** option is not specified, they are treated specially as follows:

| | |
|-----------------|--|
| Backspace | The current count of line width is decremented by one, although the count never becomes negative. Thus, the character sequence <i>character-backspace-character</i> counts as using one column position, assuming both characters each occupy a single column position. fold does not insert a newline character immediately before or after any backspace character. |
| Tab | Each tab character encountered advances the column position pointer to the next tab stop. Tab stops are set 8 columns apart at column positions 1, 9, 17, 25, 33, etc. |
| Carriage-return | The current count of line width is set to zero. fold does not insert a newline character immediately before or after any carriage-return character. |

Note that **fold** may affect any underlining that is present.

Options

The **fold** command recognizes the following options and command-line arguments:

| | |
|------------------------|---|
| -b | Count <i>width</i> in bytes rather than in column positions. |
| -s | Break the line on the last blank character found before the specified number of column positions (or bytes). If none are found, break the line at the specified line length. |
| -w <i>width</i> | Specify the maximum line length, in column positions (or bytes if -b is specified). |
| -width | The default value is 80. <i>width</i> should be a multiple of 8 if tabs are present, or the tabs should be expanded using expand before processing by fold (see <i>expand(1)</i>). The -width option is obsolescent and may be removed in a future release. |

EXTERNAL INFLUENCES**Environment Variables**

LC_CTYPE determines the interpretation of text as single- and/or multi-byte characters.

LC_MESSAGES determines the language in which messages are displayed.

If **LC_CTYPE** or **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**.

If any internationalization variable contains an invalid setting, **fold** behaves as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single- and multi-byte character code sets are supported.

fold(1)

fold(1)

SEE ALSO

expand(1).

STANDARDS CONFORMANCE

fold: XPG4, POSIX.2



f

NAME

forder - convert file data order

SYNOPSIS

forder [-a] [-l] [-n] [*file* ...]

DESCRIPTION

The text orientation (mode) of a file can be right-to-left (non-Latin) or left-to-right (Latin). This text orientation can affect the way data is arranged in the file. The data arrangements that result are called screen order and keyboard order. **forder** converts the order of characters in the file from screen order to keyboard order or vice versa.

forder reads the concatenation of input files (or standard input if none are given) and produces on standard output a converted version of its input. If **-** appears as an input file name, **forder** reads standard input at that point (use **--** to delimit the end of options in such instances).

forder converts input files for all languages that are read from right-to-left. Unless the **-a** option is used, the command merely copies input files to standard output for languages that are read from left-to-right.

Options

forder recognizes the following options:

- a** Convert file data order for languages read from left-to-right.
- l** Identify the file as having been created in Latin mode.
- n** Identify the file as having been created in non-Latin mode.

EXTERNAL INFLUENCES**Environment Variables**

The **LANGOPTS** environment variable determines the mode and order of the file. The syntax of **LANGOPTS** is:

[*mode*] [*_order*]

where *mode* describes the mode of a file: **l** represents Latin mode, and **n** represents non-Latin mode. Non-Latin mode is assumed for values other than **l** and **n**. The *order* describes the data order of a file: **k** is keyboard, and **s** is screen. Keyboard order is assumed for values other than **k** and **s**. Mode information in **LANGOPTS** can be overridden from the command line.

The **LC_ALL** environment variable determines the direction of a language (left-to-right or right-to-left).

The **LC_NUMERIC** environment variable determines whether a language has alternative numbers.

The **LANG** environment variable determines the language in which messages are displayed.

International Code Set Support

Single-byte character code sets are supported.

EXAMPLES

The following command begins with *file1*, which exists in screen order, converts it to keyboard order, sorts the keyboard-ordered output, converts it back to screen order, and redirects the output to *file2*. Note that **-n** is given to inform **forder** that *file1* was created in non-Latin mode.

```
forder -n file1 | sort | forder -n > file2
```

WARNINGS

It is the user's responsibility to ensure that the **LANGOPTS** environment variable accurately reflects the status of the file.

If present, alternative numbers always have a left-to-right orientation.

The **forder** command is HP proprietary, not portable to other vendors' systems, and will not be provided in future HP-UX releases.

AUTHOR

forder was developed by HP.

SEE ALSO

`environ(5)`, `strord(3C)`, `nljust(1)`.


f

NAME

from - who is my mail from?

SYNOPSIS

from [-**s** *sender*] [*user*]

DESCRIPTION

from prints the mail header lines in your mailbox file to show who sent you mail. If *user* is specified, *user's* mailbox is examined instead of your own. If the **-s** option is given, only headers of mail from *sender* are printed.

EXAMPLES

List header lines for all current mail in your mailbox that was sent by **ken**.

from -s ken

FILES

/var/mail/*

AUTHOR

from was developed by the University of California, Berkeley.

SEE ALSO

biff(1), mail(1), prmail(1).


f

NAME

fruled - turn on/off attention LEDs (cell, cabinet and I/O chassis attention LEDs)

SYNOPSIS

```
fruled [-f|-o] [-B] -c cell [-c...]
fruled [-f|-o] [-B] -i I/Ochassis [-i...]
fruled [-f|-o] -b cabinet [-b...]
fruled [-f] -C [-l cabinet] [-l...]
fruled [-f] -I [-l cabinet] [-l...]
```

DESCRIPTION

The **fruled** command turns on/off attention LEDs of cells or I/O chassis. The command can also be used to start or stop flashing cabinet number LEDs.

If a cell attention LED is being illuminated, the cabinet number LED of the cabinet that contains the component can be made to flash by using the **-B** option. Likewise, if an internal component's LED is being turned off, the cabinet number LED can also be made to stop flashing using the **-B** option.

Options and Arguments

fruled recognizes the following command line options and arguments:

- f** Turn off specified attention LED(s). This is the default.
The **-f** and **-o** options are mutually exclusive.
- o** Turn on specified attention LED(s).
The **-o** option is unavailable with **-C** or **-I**.
- B** Start or stop flashing the cabinet number LED of the cabinet that contains the cell or I/O chassis.
The **-B** option is only available with **-c** and **-i**.
- c *cell*** Turn on/off the specified *cell* attention LED.
cell can be specified either in the local (*cabinet#/slot#*) or global (*cell_ID*) format. For example, the cell located in cabinet 2, slot 4 is locally identified as 2/4 or globally as simply 20.
- i *I/Ochassis*** Turn on/off the specified *I/Ochassis* attention LED.
An *I/Ochassis* can be specified in the form of *cabinet#/enclosure#/chassis#*. For example, the I/O chassis located in cabinet 1, enclosure 4 and I/O chassis slot 1 is identified as 1/4/1.
- b *cabinet*** Start or stop flashing the cabinet number LED of the specified *cabinet*.
- C** Turn off all cell attention LEDs.
By default the scope will be the whole complex if **-l** option is not specified.
- I** Turn off all I/O chassis LEDs.
By default the scope will be the entire complex if the **-l** option is not specified.
- l *cabinet*** Limit the scope of the **-C** or **-I** option to a given *cabinet*.

EXIT STATUS

The **fruled** utility exits with one of the following values:

- 0 Successful completion.
- 1 Error condition occurred.
- 2 No LED associated with specified object.

EXAMPLES

Turn on the attention LED of the cell located in cabinet 0 slot 4 and also turn on the attention LED of the cabinet in which it is contained.

```
fruled -o -B -c 0/4
```

Turn off the attention LEDs of 2 cells located in cabinet 0, slot 4 and cabinet 0, slot 6.

```
fruled -f -c 0/4 -c 0/6
```

WARNINGS

The presence of ? in the command output indicates a problem in address translation. If you see ? (example: "I/O chassis 0/?/3") in the command output, please contact your Hewlett Packard representative.

AUTHOR

fruled was developed by the Hewlett-Packard Company.

SEE ALSO

parstatus(1), partition(1), frupower(1M), parcreate(1M), parmodify(1M), parremove(1M), parunlock(1M).


f

NAME

ftio - faster tape I/O

SYNOPSIS

```
ftio -o|-O [achpvxAELM] [-B blksize] [-D type] [-e extarg] [-K comment] [-L filelist]
      [-N datefile] [-S script] [-T tty] [-Z nobufs] tapedev [pathnames] [-F ignorenames]

ftio -i|-I [cdfmptuvxAEMPR] [-B blksize] [-S script] [-T tty] [-Z nobufs] tapedev [patterns]

ftio -g [v] tapedev [patterns]
```

DESCRIPTION

ftio is a tool designed specifically for copying files to tape drives. It performs faster than either **cpio** or **tar** in comparable situations (see *cpio(1)* and *tar(1)*). **ftio** uses multiple processes (to read/write the file system and to write/read the tape device), with large amounts of memory sharing between processes as well as a large block size for reading and writing to the tape.

ftio is compatible with **cpio** in that output from **cpio** is always readable by **ftio**, and output from **ftio** is readable by **cpio**, except as explained in the "cpio Compatibility" section, later in the manpage.

ftio must be invoked with exactly one of the following options: **-o**, **-O**, **-i**, **-I**, or **-g**. The **-o** and **-O** options specify that **ftio** is writing "out" from file system to tape; the **-i** and **-I** options specify that **ftio** is writing "in" from tape to file system. The **-o**, **-O**, **-i**, and **-I** options can be followed by modifiers that must appear immediately after the option with no spaces between the option and the modifier, as in **ftio -idxE** (see Modifiers section below).

tapedev specifies the name of a device special file for the tape device to which the output is written. A device on a remote machine can be specified in the form

machine:device_special_file

ftio creates a server process from **/usr/sbin/rmt** on the remote machine to access the tape device. If **/usr/sbin/rmt** does not exist on the remote system, **ftio** creates a server process from **/etc/rmt**, on the remote machine to access the tape device.

Options

ftio recognizes the following options:

- o** Copy (out) files from the file system to *tapedev*, including path name and status information. If *pathnames* are specified, **ftio** recursively descends *pathnames* looking for files, and copies those files to *tapedev*. If *pathnames* are not specified, **ftio** reads the standard input to obtain a list of path names to copy. **ftio** can copy to multiple tapes if required. For every tape used, **ftio** generates a tape header containing the current tape volume number, machine node name and type, operating system name, release and version numbers (all from the **uname()** system call; see *uname(2)*), username of the person issuing the **ftio** command, the time and date the command was executed, the number of consecutive times the current media has been used, a comment field, and other items used internally by **ftio**. The tape header is separated from the main body of the tape archive by an end-of-file mark. The tape header can be read by invoking **cat** with the device file name as the first argument (see *cat(1)*). Note, character and block device special files written with the **-o** option are not transportable to other HP-UX implementations.
- O** Copy out files in the same way as **ftio -ocva**, when no modifiers are used with the **-O**. However, if the **.ftiorc** file exists in the user's home directory, **ftio** opens this file and scans for lines preceded by **O=**. Options defined on matching lines are passed to **ftio** as if they had been specified on the command line. See **EXAMPLES** section.
- i** Extract (copy into the file system) files from *tapedev*, which is assumed to be a tape and the product of a previous **ftio -o** operation. Only files with names that match *patterns*, according to the rules of Pattern Matching Notation (see *regex(5)*), are selected. In addition, a leading **!** within a pattern indicates that only those names that do *not* match the remainder of the pattern should be selected. Multiple *patterns* can be specified. If no *patterns* are specified, the default for *patterns* is ***** (that is, select all files). The extracted files are conditionally created and copied into the current directory tree, based upon the options described below. The permissions of

the files are those of the previous **-o** operation.

- I** Extract (copy into the file system) files in the same way as for **ftio -icdmv**, when no modifiers are used with the **-I**. However, if the **.ftiorc** file exists in the user's home directory, **ftio** opens this file, and scans for lines preceded by **I=**. Options defined on matching lines are passed to **ftio** as if they had been specified on the command line. See **EXAMPLES** section.
- g** Read the file list in *tapedev*. If *patterns* is specified, only file names that match are printed. Note that file names are always preceded by the volume that **ftio** expected the file to be on when the file list was created; thus only the last volume is valid in this respect.
- e extarg** Specifies the handling of any extent attributes of the file[s] to be archived. Extent attributes cannot be preserved when archiving files with **ftio**. *extarg* takes one of the following values:
 - warn** Issue a warning message and archive the file without extent attributes.
 - ignore** A file with extent attributes will be archived, without preserving the extent attributes and without issuing a warning message.
 - force** A file with extent attributes will not be archived and a warning message will be issued.

If **-e** is not specified, the default value for *extarg* is **warn**.
- B blksize** Specify the size (in bytes) of blocks written to tape. This number can end with **k**, which specifies multiplication by 1024. The use of larger blocks generally improves performance and tape usage. The maximum allowable block size is limited by the tape drive used. A default of 16384 bytes is set because this is the maximum block size on most Hewlett-Packard tape drives.
- D type** Descend a directory recursively, only if the file system to which it belongs is *type*, where *type* can be **hfs**, **vxfs**, or **nfs**.
- F ignorenames** Arguments following **-F** specify *patterns* that should not be copied to the tape. The same rules apply to *ignorenames* as to *patterns*; see the earlier description for **ftio -i**.
- K comment** Specify a comment to be placed in the **ftio** tape header.
- L filelist** Create a list of the files being backed up. *filelist* specifies the output file. If *pathnames* is specified, perform the file search and generate a list of files prior to actually commencing the backup. This list is then appended to the tape header of each tape in the backup as a list of files that **ftio** attempted to fit onto this tape. The last tape in the backup contains a catalog identifying where the files are in the archive set. If *pathnames* is not also specified, the file list is taken from standard input before the backup begins. In addition to generating file lists, the **-L** option implements tape checkpointing, allowing the backup to restart from a write failure on bad media.
- M** Make fully compatible with **cpio**. That is, do not generate or expect tape headers and change the default block size to 5120 bytes. (See the **cpio Compatibility** section below.)
- N datefile** Only files newer than the file specified in *datefile* are copied to tape.
- R** Resynchronize automatically, when **ftio** goes out of phase. This is useful when restoring from a multi-tape backup from tapes other than the first. By default, **ftio** asks the user if resynchronization is required.
- S script** Specify a command to be invoked every time a tape is completed in a multi-tape backup. The command is invoked by the Bourne shell (see *sh-bourne(1)*) with the following arguments: *script tape_no user_name*. *script* is the string argument *script* specified with the **-S** option. *tape_no* is the number of the tape required, and *user_name* is the user who invoked **ftio**. Typically, the string *script* specifies a shell script which is used to notify the user that a tape change is required.
- T tty** Specify alternative to **/dev/tty**. Normally **/dev/tty** is opened by **ftio** when terminal interaction is required.

- Z *nobufs* Specify the number of *blksize* chunks of memory to use as buffer space between the two processes, where *blksize* is the size of blocks written to the tape. More chunks is usually better, but a point is reached where no improvement is gained, and performance might deteriorate as buffer space is swapped out of main memory. A default value of 16 is set for *nobufs*, but using 32 or 64 might improve performance if your system is not heavily loaded. Best results are obtained when backups are performed with the system in single-user mode (see *shutdown(1M)*).

Modifiers

The following modifiers can be used with certain options as indicated in the SYNOPSIS:

- a After files are copied to tape, reset their access time to appear as though the files were not accessed by **ftio**.
- c Write header information in ASCII character form, for portability.
- d When restoring files, create directories as needed.
- f Copy in all files except those that match *patterns*.
- h Archive the files to which symbolic links point, as if they were normal files or directories. By default, **ftio** archives the link itself.
- m Retain previous file modification time and ownership of file. Restoring modification time does not apply to directories that are being restored.
- p At the end of the backup, print the number of blocks transferred, the total time taken (excluding tape rewind and reel-change time), and the effective transfer rate calculated from these figures. These values are printed at the end of each tape if **p** is specified twice.
- t Print only a table of contents of the input. No files are created, read, or copied.
- u Copy unconditionally (by default, **ftio** does not replace a newer file with a older file of the same name).
- v Be verbose. Print a list of file names and tape headers. When used with the **t** modifier, the table of contents looks the same as the output of the **ls -l** (ell) command (see *ls(1)*).
- x Save or restore device special files. **ftio** uses *mknod(2)* to recreate these files during a restore operation. Thus, this modifier is restricted to users with appropriate privileges. This is intended for intrasystem (backup) use. Restoring device files onto a different system can be very dangerous.
- A If copying from tape (**-i** or **-I** option), print all file names found on the tape archive, noting which files have been restored. This is useful when the user restores selected files, but wants to know which (if any) files are on the tape.

If copying to tape (**-o** or **-O** option), the **A** modifier suppresses warning messages regarding optional access control list entries. **ftio(1)** does not back up optional access control list entries in a file's access control list (see *acl(5)*). Normally, a warning message is printed for each file that has optional access control list entries.
- E When archiving, store all files having absolute path names (that is, path names beginning with **/**) with path names relative to the root directory (in other words, remove the leading **/**). On restoration, any files in the archive that had an absolute path name before archiving are restored relative to the current directory.
- L Same as the **-L** option, except that the file list is left in the current directory as the file **ftio.list**, instead of the file named in *filelist*.
- P On restoration, use **prealloc()** to allocate disk space beforehand for the file (see *prealloc(2)*). This vastly improves the localization of file fragments.

When end-of-tape is reached, **ftio** invokes *script* if the **-S** option was specified, rewinds the current tape, then asks the user to mount the next tape.

To pass one or more metacharacters to **ftio** without having the shell expand them, protect them either by preceding each of them with a backslash (as in **/usr***), or enclosing them in protective single quotes (as in **'/usr*'**).

cpio Compatibility

ftio uses the same archive format as **cpio**. However, by default **ftio** creates tape headers and uses a tape block size of 16KB. **cpio** by default uses 512-byte blocks. When used with the **-B** option, **cpio** uses 5120 byte blocks. To achieve full compatibility with **cpio** in either input or output mode, the user should specify the **M** modifier. **ftio -oM** creates a single- or multi-tape archive that has no tape headers, and, by default, the same block size as **cpio -[o|i]B**. An archive created by a **cpio -oB** command can be restored using **ftio -iM**. If the **M** modifier of **ftio** is combined with a **-B 512** block-size specification, full compatibility with **cpio -[o|i]** (no **-B**) is achieved.

EXTERNAL INFLUENCES**Environment Variables**

LC_COLLATE determines the collating sequence used in evaluating pattern matching notation for file name generation.

LC_CTYPE determines the characters matched by character class expressions in pattern matching notation.

LC_TIME determines the format and contents of date and time strings.

LANG determines the language in which messages are displayed.

If **LC_COLLATE**, **LC_CTYPE**, or **LC_TIME** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of C (see *lang(5)*) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **ftio** behaves as if all internationalization variables are set to C. See *environ(5)*.

International Code Set Support

Single-byte character code sets are supported.

EXAMPLES

Copy the entire contents of the file system (including special files) onto tape drive **/dev/rmt/c0t0d0BEST**:

```
ftio -ox /dev/rmt/c0t0d0BEST /
```

Restore all the files on **/dev/rmt/c0t0d0BEST**, relative to the current directory:

```
ftio -idxE /dev/rmt/c0t0d0BEST
```

List the contents of a backup set created using **ftio -o**. Note that use of the **v** modifier gives a more detailed listing, and displays the contents of tape headers.

```
ftio -itv /dev/rmt/c0t0d0BEST
```

Show how to use the **.ftiorc** file:

Assume a **.ftiorc** file exists in the user's home directory and contains the following:

```
# Sample .ftiorc file.
I= cdmuvEpp -B 16k -S /usr/local/bin/ftio.change
O= cavEpp -Z 8 -B 16k -S /usr/local/bin/ftio.change
```

Invoke **ftio** with the following command line to back up the user's home directory and the operating system commands directory:

```
ftio -O /dev/rmt/c0t0d0BEST /home/my_home /usr/sbin
```

Specifying the **-O** option causes **ftio** to check the **.ftiorc** file for additional options. In this case, character headers are generated, access times are reset, a listing of the files copied are printed to standard output, all file names are copied to **/dev/rmt/c0t0d0BEST** with path names relative to **/**, performance data is printed when the backup is complete (and at every tape change), and, if the backup goes beyond one media the script, **/usr/local/bin/ftio.change** is invoked by **ftio** after each media is completed.

WARNINGS

Because of industry standards and interoperability goals, **ftio** does not support the archival of files larger than 2GB or files that have user/group IDs greater than 60K. Files with user/group IDs greater than 60K are archived and restored under the user/group ID of the current process.

ftio operates using System V shared memory and semaphores. The resources committed to these functions are not freed automatically by the system when the process terminates. **ftio** does this only when it terminates normally, or when it terminates after receiving one the following signals: **SIGHUP**, **SIGINT**, **SIGTERM**. Any other signal is handled in the default manner described by *signal(2)*. Note that the behavior for **SIGKILL** is to terminate the process without delay. Thus, if **ftio** receives a **SIGKILL** signal (as might be produced by the indiscriminate use of **kill -9** (see *kill(1)*), system resources used for shared memory and semaphores are not returned to the system. If it becomes necessary to terminate an invocation of **ftio**, use **kill -15** instead. Current system usage of shared memory and semaphores can be checked using the **ipcs** command (see *ipcs(1)*). Committed resources can be removed using **ipcrm** (see *ipcrm(1)*).

AUTHOR

ftio was developed by HP.

SEE ALSO

cpio(1), *find(1)*, *ipcs(1)*, *ipcrm(1)*, *kill(1)*, *ls(1)*, *rmt(1M)*, *mknod(2)*, *prealloc(2)*, *signal(2)*, *uname(2)*, *acl(5)*, *environ(5)*, *lang(5)*, *regexp(5)*, *mt(7)*.

f

NAME

ftp - file transfer program

SYNOPSIS

ftp [-g] [-i] [-n] [-c] [-v] [-B *size*] [*server-host*]

DESCRIPTION

ftp is a user interface to the File Transfer Protocol. **ftp** copies files over a network connection between the local “client” host and a remote “server” host. **ftp** runs on the client host.

Options

The **ftp** command supports the following options:

- g Disable file name “globbing”; see the **glob** command, below. By default, when this option is not specified, globbing is enabled.
- i Disable interactive prompting by multiple-file commands; see the **prompt** command, below. By default, when this option is not specified, prompting is enabled.
- n Disable “auto-login”; see the **open** command, below. By default, when this option is not specified, auto-login is enabled.
- c When this option is set, the **SYST** and **TYPE** calls are not made by the **ftp** client to the server upon establishing a connection. The **-c** option takes effect only when auto-login is disabled i.e. when it is invoked along with the **-n** option. This option does not disable the **SYST** and **TYPE** commands, but only refrains from invoking these commands upon establishing a connection.
- v Enable verbose output; see the **verbose** command, below. If this option is not specified, **ftp** displays verbose output only if the standard input is associated with a terminal.
- B Set the buffer size of the data socket to *size* blocks of 1024 bytes. The valid range for *size* is an integer from 1 to 64 (default is 56).

Note: A large buffer size will improve the performance of **ftp** on fast links (e.g., FDDI), but may cause long connection times on slow links (e.g., X.25).

The name of the server host that **ftp** communicates with can be specified on the command line. If the server host is specified, **ftp** immediately opens a connection to the server host; see the **open** command, below. Otherwise, **ftp** waits for commands from the user.

File Transfer Protocol specifies file transfer parameters for *type*, *mode*, *form*, and *struct*. **ftp** supports the **ASCII**, **binary**, and **tenex** File Transfer Protocol *types*. **ASCII** is the default FTP *type*. (It should be noted though that, whenever **ftp** establishes a connection between two similar systems, it switches automatically to the more efficient **binary** type.) **ftp** supports only the default values for the file transfer parameters *mode* which defaults to **stream**, *form* which defaults to **non-print**, and *struct* which defaults to **file**.

COMMANDS

ftp supports the following commands. Command arguments with embedded spaces must be enclosed in quotes (for example, “argument with embedded spaces”).

!*command* [*args*]

Invoke a shell on the local host. The **SHELL** environment variable specifies which shell program to invoke. **ftp** invokes **/usr/bin/sh** if **SHELL** is undefined. If *command* is specified, the shell executes it and returns to **ftp**. Otherwise, an interactive shell is invoked. When the shell terminates, it returns to **ftp**.

\$ *macro-name* [*args*]

Execute the macro *macro-name* that was defined with the **macdef** command. Arguments are passed to the macro unglobbed.

account [*passwd*]

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 is prompted for an account password in a non-echoing input mode.

append *local-file* [*remote-file*]

Copy *local-file* to the end of *remote-file*. If *remote-file* is left unspecified, the local file name is used in

naming the remote file after being altered by any *ntrans* or *nmap* setting.

ascii

Set the file transfer *type* to network ASCII. This is the default type.

bell

Sound a bell after each file transfer completes.

binary

Set the file transfer *type* to **binary**.

bye Close the connection to the server host if a connection was open, and exit. Typing an end-of-file (EOF) character also terminates and exits the session.

case

Toggle remote computer file name case mapping during **mget** commands. When **case** is on (the default is off), remote computer file names with all letters in uppercase are written in the local directory with the letters mapped to lowercase.

cd remote-directory

Set the working directory on the server host to *remote-directory*.

cdup

Set the working directory on the server host to the parent of the current remote working directory.

chmod mode file-name

Change the permission modes of the file *file-name* on the remote system to *mode*.

close

Terminate the connection to the server host. The **close** command does not exit **ftp**. Any defined macros are erased.

cr Toggle carriage return stripping during **ascii** type file retrieval. Records are denoted by a carriage-return/line-feed sequence during **ascii** type file transfer. When **cr** is on (the default), carriage returns are stripped from this sequence to conform with the UNIX single line-feed record delimiter. Records on non-UNIX remote systems may contain single line-feeds; when an **ascii** type transfer is made, these line-feeds can be distinguished from a record delimiter only when **cr** is off.

delete remote-file

Delete *remote-file*. The *remote-file* can be an empty directory. No globbing is done.

dir [remote-directory] [local-file]

Write a *remote-directory* listing to standard output or optionally to *local-file*. If neither *remote-directory* nor *local-file* is specified, list the remote working directory to standard output. If interactive prompting is on, **ftp** prompts the user to verify that the last argument is indeed the target file for **dir** output. Globbing characters are always expanded.

disconnect

A synonym for **close**.

form format

Set the file transfer *form* to *format*. The only supported format is **non-print**

get remote-file [local-file]

Copy *remote-file* to *local-file*. If *local-file* is unspecified, **ftp** uses the specified *remote-file* name as the *local-file* name, subject to alteration by the current *case*, *ntrans*, and *nmap* settings.

glob

Toggle file name globbing. When file name globbing is enabled, **ftp** expands *cs*(1) metacharacters in file and directory names. These characters are *, ?, [,], ~, {, and }. The server host expands remote file and directory names. Globbing metacharacters are always expanded for the **ls** and **dir** commands. If globbing is enabled, metacharacters are also expanded for the multiple-file commands **mdelete**, **mdir**, **mget**, **mls**, and **mput**.

hash

Toggle printing of a hash-sign (#) for each 1024 bytes transferred. Note that the use of this feature may cause performance degradation.

help [command]

Print an informative message about the **ftp** command called *ftp-command*. If *ftp-command* is unspecified, print a list of all **ftp** commands.

- idle** [*seconds*]
Set the inactivity timer on the remote server to *seconds* seconds. If *seconds* is omitted, **ftp** prints the current inactivity timer.
- lcd** [*local-directory*]
Set the local working directory to *local-directory*. If *local-directory* is unspecified, set the local working directory to the user's local home directory.
- ls** [*remote-directory*] [*local-file*]
Write a listing of *remote-directory* to *local-file*. The listing includes any system-dependent information that the server chooses to include; for example, most UNIX systems produce output from the command **ls -l** (see also **nlist**). If neither *remote-directory* nor *local-file* is specified, list the remote working directory. If globbing is enabled, globbing metacharacters are expanded.
- macrodef** *macro-name*
Define a macro. Subsequent lines are stored as the macro *macro-name*; an empty input line 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 to the 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 *remote-files*. If globbing is enabled, globbing metacharacters are expanded.
- mdir** *remote-files local-file*
Write a listing of *remote-files* to *local-file*. If globbing is enabled, globbing metacharacters are expanded. If interactive prompting is on, **ftp** prompts the user to verify that the last argument is indeed the target local file for **mdir** output.
- mget** *remote-files*
Copy *remote-files* to the local system. If globbing is enabled, globbing metacharacters are expanded. The resulting local file names are processed according to *case*, *ntrans*, and *nmap* settings.
- mkdir** *directory-name*
Create remote *directory-name*.
- mls** *remote-files local-file*
Write an abbreviated listing of *remote-files* to *local-file*. If globbing is enabled, globbing metacharacters are expanded. If interactive prompting is on, **ftp** prompts the user to verify that the last argument is indeed the target local file for **mls** output.
- mode** [*mode-name*]
Set the FTP file transfer *mode* to *mode-name*. The only supported mode is **stream**.
- modtime** *remote-file*
Show the last modification time of *remote-file*.
- mput** *local-files*
Copy *local-files* from the local system to the remote system. The remote files have the same name as the local files processed according to *ntrans* and *nmap* settings. If globbing is enabled, globbing characters are expanded.
- newer** *file-name*
Get the file only if the modification time of the remote file is more recent than the file on the current system. If the file does not exist on the current system, the remote file is considered *newer*. Otherwise, this command is identical to **get**.
- nlist** [*remote-directory*] [*local-file*]
Write an abbreviated listing of *remote-directory* to *local-file*. If *remote-directory* is left unspecified, the current working directory is used. If interactive prompting is on, **ftp** prompts the user to verify that the last argument is indeed the target local file for **nlist** output.
- 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 remote computer 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**. Spaces can 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. If arguments are specified, 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 remote computer 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.

open *server-host* [*port-number*]

Establish a connection to *server-host*, using *port-number* (if specified). If *auto-login* is enabled, **ftp** attempts to log into the server host.

passive

Toggle passive mode of transfer. By default, the passive mode of transfer is disabled. This command enables the server to specify the data port for the ftp transfer.

prompt

Toggle interactive prompting. By default, **ftp** prompts the user for a yes or no response for each output file during multiple-file commands. If interactive prompting is disabled, **ftp** performs the command for all specified 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** does not define new macros during the auto-login process, **close** does 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**, **mput**, 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 FTP protocol **PASV** command by the server on the secondary control connection.

put *local-file* [*remote-file*]

Copy *local-file* to *remote-file*. If *remote-file* is unspecified, **ftp** assigns the *local-file* name, processed according to any **ntrans** or **nmap** settings, to the *remote-file* name.

pwd Write the name of the remote working directory to *stdout*.

quit

A synonym for **bye**.

quote *arguments*

Send *arguments*, verbatim, to the server host. See *ftpd(1M)*.

recv *remote-file* [*local-file*]
A synonym for **get**.

reget *remote-file* [*local-file*]
reget acts like **get**, except that if *local-file* exists and is smaller than *remote-file*, *local-file* is presumed to be a partially transferred copy of *remote-file* and the transfer is continued from the apparent point of failure. This command is useful when transferring very large files over networks that tend to drop connections.

rhel [*command-name*]
Request help from the server host. If *command-name* is specified, supply it to the server. See *ftpd(1M)*.

rstatus [*file-name*]
With no arguments, show status of remote machine. If *file-name* is specified, show status of *file-name* on remote machine.

rename *remote-from* *remote-to*
Rename *remote-from*, which can be either a file or a directory, to *remote-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.

restart *marker*
Restart the immediately following **get** or **put** at the indicated *marker*. On UNIX systems, marker is usually a byte offset into the file.

rmdir *remote-directory*
Delete *remote-directory*. *remote-directory* must be an empty directory.

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. **ftp** reports the unique filename. Note that **runique** does not affect local files generated from a shell command (see below). The default value is **off**.

send *local-file* [*remote-file*]
A synonym for **put**.

sendport
Toggle the use of **PORT** commands. By default, **ftp** attempts to use a **PORT** command when establishing a connection for each data transfer. If the **PORT** command fails, **ftp** uses the default data port. When the use of **PORT** commands is disabled, **ftp** makes no attempt to use **PORT** commands for each data transfer. This is useful for certain FTP implementations that ignore **PORT** commands but (incorrectly) indicate that they've been accepted. See *ftpd(1M)*. Turning **sendport** off may cause delays in the execution of commands.

site *arguments*
Send *arguments*, verbatim, to the server host as a **SITE** command. See *ftpd(1M)*.

size *remote-file*
Show the size of *remote-file*.

status
Show the current status of **ftp**.

struct [*struct-name*]
Set the FTP file transfer *struct* to *struct-name*. The only supported *struct* is **file**.

sunique
Toggle storing of files on remote machine under unique file names. The remote server reports the unique name. By default, **sunique** is **off**.

system
Show the type of operating system running on the remote machine.

tenex

Set the FTP file transfer *type* to **tenex**.

type [*type-name*]

Set the FTP file transfer *type* to *type-name*. If *type-name* is unspecified, write the current *type* to *stdout*. **Ascii**, **binary**, and **tenex** are the *types* currently supported.

umask [*newmask*]

Set the default umask on the remote server to *newmask*. If *newmask* is omitted, the current umask is printed.

user *user-name* [*password*] [*account*]

Log into the server host on the current connection, which must already be open. A **.netrc** file in the user's local home directory can provide the *user-name*, *password*, and optionally the *account*; see *netrc*(4). Otherwise **ftp** prompts the user for this information. The HP-UX FTP server does not require an *account*. For security reasons, **ftp** always requires a password. It does not log into remote accounts that do not have a password.

verbose

Toggle verbose output. If verbose output is enabled, **ftp** displays responses from the server host, and when a file transfer completes it reports statistics regarding the efficiency of the transfer.

? [*command*]

A synonym for the **help** command. Prints the **help** information for the specified *command*.

Aborting A File Transfer

To abort a file transfer, use the terminal interrupt key (usually **Ctrl-C**). Sending transfers are halted immediately. **ftp** halts incoming (receive) transfers by first sending a FTP protocol **ABOR** command to the remote server, then discarding any further received data. 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 does not appear until the remote server completes sending the requested file.

The terminal interrupt key sequence is ignored while **ftp** awaits 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 manually.

File Naming Conventions

Files specified as arguments to **ftp** commands are processed according to the following rules.

- If the file name **-** is specified, **ftp** uses the standard input (for reading) or standard output (for writing).
- If the first character of the file name is **|**, **ftp** interprets the remainder of the argument as a shell command. **ftp** forks a shell, using **popen()** (see *popen*(3S)) with the supplied argument, and reads (writes) from standard output (standard input). If the shell command includes spaces, the argument must be quoted, as in:

```
"| ls -lt".
```

Some useful examples of this mechanism are:

```
ls . "| more".
```

The above command lists the files in the current directory page by page.

```
put "| tail -20 loc_file" rem_file.
```

This command copies the last twenty lines of the local file "loc_file" to the remote system as "rem_file".

- Otherwise, if globbing is enabled, **ftp** expands local file names according to the rules used by the C shell (see *cs*(1)); see the **glob** command, below. If the **ftp** command expects a single local file (e.g., **put**), only the first filename generated by the globbing operation is used.
- For **mget** commands and **get** commands with unspecified local file names, the local filename is named the same as 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.
- For **mput** commands and **put** commands with unspecified remote file names, the remote filename is named the same as 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.

WARNINGS

Correct execution of many commands depends upon proper behavior by the remote server.

DIAGNOSTICS

Error! could not retrieve authentication type.

Please notify sys admin.

There are two authentication mechanisms used by **ftp**. One authentication mechanism is based on Kerberos and the other is not. The type of authentication mechanism is obtained from a system file which is updated by **inetsvcs_sec** (see *inetsvcs_sec*(1M)). If the system file does not contain known authentication types, the above error is displayed.

AUTHOR

ftp was developed by the University of California, Berkeley.

SEE ALSO

csh(1), *rcp*(1), *ftpd*(1M), *inetsvcs_sec*(1M), *netrc*(4), *ftpusers*(4), *hosts*(4).

NAME

ftp - file transfer program

SYNOPSIS

ftp [-g] [-i] [-n] [-c] [-P] [-v] [-B *size*] [*server-host*]

DESCRIPTION

ftp is a user interface to the File Transfer Protocol. **ftp** copies files over a network connection between the local “client” host and a remote “server” host. **ftp** runs on the client host.

Options

The **ftp** command supports the following options:

- g Disable file name “globbing”; see the **glob** command, below. By default, when this option is not specified, globbing is enabled.
- i Disable interactive prompting by multiple-file commands; see the **prompt** command, below. By default, when this option is not specified, prompting is enabled.
- P Disables Kerberos authentication and authorization. Only applicable in a secure environment based on Kerberos V5. When this option is specified, a password is required and the password is sent across the network in a readable form. By default, if this option is not specified, a password is not required and Kerberos authentication and authorization takes place instead. See *sis*(5).
- n Disable “auto-login”; see the **open** command, below. By default, when this option is not specified, auto-login is enabled.
- c When this option is set, the **SYST** and **TYPE** calls are not made by the **ftp** client to the server upon establishing a connection. The **-c** option takes effect only when auto-login is disabled i.e. when it is invoked along with the **-n** option. This option does not disable the **SYST** and **TYPE** commands, but only refrains from invoking these commands upon establishing a connection.
- v Enable verbose output; see the **verbose** command, below. If this option is not specified, **ftp** displays verbose output only if the standard input is associated with a terminal.
- B Set the buffer size of the data socket to *size* blocks of 1024 bytes. The valid range for *size* is an integer from 1 to 64 (default is 56).

Note: A large buffer size will improve the performance of **ftp** on fast links (e.g., FDDI), but may cause long connection times on slow links (e.g., X.25).

The name of the server host that **ftp** communicates with can be specified on the command line. If the server host is specified, **ftp** immediately opens a connection to the server host; see the **open** command, below. Otherwise, **ftp** waits for commands from the user.

File Transfer Protocol specifies file transfer parameters for *type*, *mode*, *form*, and *struct*. **ftp** supports the **ASCII**, **binary**, and **tenex** File Transfer Protocol *types*. **ASCII** is the default FTP *type*. (It should be noted though that, whenever **ftp** establishes a connection between two similar systems, it switches automatically to the more efficient **binary** type.) **ftp** supports only the default values for the file transfer parameters *mode* which defaults to **stream**, *form* which defaults to **non-print**, and *struct* which defaults to **file**.

COMMANDS

ftp supports the following commands. Command arguments with embedded spaces must be enclosed in quotes (for example, “argument with embedded spaces”).

!*command* [*args*]

Invoke a shell on the local host. The **SHELL** environment variable specifies which shell program to invoke. **ftp** invokes **/usr/bin/sh** if **SHELL** is undefined. If *command* is specified, the shell executes it and returns to **ftp**. Otherwise, an interactive shell is invoked. When the shell terminates, it returns to **ftp**.

\$ *macro-name* [*args*]

Execute the macro *macro-name* that was defined with the **macdef** command. Arguments are passed to the macro unglobbed.

account [*passwd*]

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 is prompted for an account password in a non-echoing input mode.

append *local-file* [*remote-file*]

Copy *local-file* to the end of *remote-file*. If *remote-file* is left unspecified, the local file name is used in naming the remote file after being altered by any *ntrans* or *nmap* setting.

ascii

Set the file transfer *type* to network ASCII. This is the default type.

bell

Sound a bell after each file transfer completes.

binary

Set the file transfer *type* to **binary**.

bye Close the connection to the server host if a connection was open, and exit. Typing an end-of-file (EOF) character also terminates and exits the session.

case

Toggle remote computer file name case mapping during **mget** commands. When **case** is on (the default is off), remote computer file names with all letters in uppercase are written in the local directory with the letters mapped to lowercase.

cd *remote-directory*

Set the working directory on the server host to *remote-directory*.

cdup

Set the working directory on the server host to the parent of the current remote working directory.

chmod *mode file-name*

Change the permission modes of the file *file-name* on the remote system to *mode*.

close

Terminate the connection to the server host. The **close** command does not exit **ftp**. Any defined macros are erased.

cr Toggle carriage return stripping during **ascii** type file retrieval. Records are denoted by a carriage-return/line-feed sequence during **ascii** type file transfer. When **cr** is on (the default), carriage returns are stripped from this sequence to conform with the UNIX single line-feed record delimiter. Records on non-UNIX remote systems may contain single line-feeds; when an **ascii** type transfer is made, these line-feeds can be distinguished from a record delimiter only when **cr** is off.

delete *remote-file*

Delete *remote-file*. The *remote-file* can be an empty directory. No globbing is done.

dir [*remote-directory*] [*local-file*]

Write a *remote-directory* listing to standard output or optionally to *local-file*. If neither *remote-directory* nor *local-file* is specified, list the remote working directory to standard output. If interactive prompting is on, **ftp** prompts the user to verify that the last argument is indeed the target file for **dir** output. Globbing characters are always expanded.

disconnect

A synonym for **close**.

form *format*

Set the file transfer *form* to *format*. The only supported format is **non-print**

get *remote-file* [*local-file*]

Copy *remote-file* to *local-file*. If *local-file* is unspecified, **ftp** uses the specified *remote-file* name as the *local-file* name, subject to alteration by the current *case*, *ntrans*, and *nmap* settings.

glob

Toggle file name globbing. When file name globbing is enabled, **ftp** expands *cs*(1) metacharacters in file and directory names. These characters are *, ?, [,], ~, {, and }. The server host expands remote file and directory names. Globbing metacharacters are always expanded for the **ls** and **dir** commands. If globbing is enabled, metacharacters are also expanded for the multiple-file commands **mdelete**, **mdir**, **mget**, **mls**, and **mput**.

hash

Toggle printing of a hash-sign (#) for each 1024 bytes transferred. Note that the use of this feature

may cause performance degradation.

help [*command*]

Print an informative message about the **ftp** command called *ftp-command*. If *ftp-command* is unspecified, print a list of all **ftp** commands.

idle [*seconds*]

Set the inactivity timer on the remote server to *seconds* seconds. If *seconds* is omitted, **ftp** prints the current inactivity timer.

lcd [*local-directory*]

Set the local working directory to *local-directory*. If *local-directory* is unspecified, set the local working directory to the user's local home directory.

ls [*remote-directory*] [*local-file*]

Write a listing of *remote-directory* to *local-file*. The listing includes any system-dependent information that the server chooses to include; for example, most UNIX systems produce output from the command **ls -l** (see also **nlist**). If neither *remote-directory* nor *local-file* is specified, list the remote working directory. If globbing is enabled, globbing metacharacters are expanded.

macrodef *macro-name*

Define a macro. Subsequent lines are stored as the macro *macro-name*; an empty input line 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 to the 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 *remote-files*. If globbing is enabled, globbing metacharacters are expanded.

mdir *remote-files local-file*

Write a listing of *remote-files* to *local-file*. If globbing is enabled, globbing metacharacters are expanded. If interactive prompting is on, **ftp** prompts the user to verify that the last argument is indeed the target local file for **mdir** output.

mget *remote-files*

Copy *remote-files* to the local system. If globbing is enabled, globbing metacharacters are expanded. The resulting local file names are processed according to *case*, *ntrans*, and *nmap* settings.

mkdir *directory-name*

Create remote *directory-name*.

mls *remote-files local-file*

Write an abbreviated listing of *remote-files* to *local-file*. If globbing is enabled, globbing metacharacters are expanded. If interactive prompting is on, **ftp** prompts the user to verify that the last argument is indeed the target local file for **mls** output.

mode [*mode-name*]

Set the FTP file transfer *mode* to *mode-name*. The only supported mode is **stream**.

modtime *remote-file*

Show the last modification time of *remote-file*.

mput *local-files*

Copy *local-files* from the local system to the remote system. The remote files have the same name as the local files processed according to *ntrans* and *nmap* settings. If globbing is enabled, globbing characters are expanded.

newer *file-name*

Get the file only if the modification time of the remote file is more recent than the file on the current system. If the file does not exist on the current system, the remote file is considered *newer*. Otherwise, this command is identical to **get**.

nlist [*remote-directory*] [*local-file*]

Write an abbreviated listing of *remote-directory* to *local-file*. If *remote-directory* is left unspecified, the current working directory is used. If interactive prompting is on, **ftp** prompts the user to verify

that the last argument is indeed the target local file for **nlist** output.

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 remote computer 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**. Spaces can 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. If arguments are specified, 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 remote computer 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.

open *server-host* [*port-number*]

Establish a connection to *server-host*, using *port-number* (if specified). If **auto-login** is enabled, **ftp** attempts to log into the server host.

prompt

Toggle interactive prompting. By default, **ftp** prompts the user for a yes or no response for each output file during multiple-file commands. If interactive prompting is disabled, **ftp** performs the command for all specified files.

put *local-file* [*remote-file*]

Copy *local-file* to *remote-file*. If *remote-file* is unspecified, **ftp** assigns the *local-file* name, processed according to any **ntrans** or **nmap** settings, to the *remote-file* name.

pwd Write the name of the remote working directory to *stdout*.

quit

A synonym for **bye**.

quote *arguments*

Send *arguments*, verbatim, to the server host. See **ftpd(1M)**.

recv *remote-file* [*local-file*]

A synonym for **get**.

reget *remote-file* [*local-file*]

reget acts like **get**, except that if *local-file* exists and is smaller than *remote-file*, *local-file* is presumed to be a partially transferred copy of *remote-file* and the transfer is continued from the apparent point of failure. This command is useful when transferring very large files over networks that tend to drop connections.

rhel [*command-name*]

Request help from the server host. If *command-name* is specified, supply it to the server. See

ftpd(1M).

rstatus [*file-name*]

With no arguments, show status of remote machine. If *file-name* is specified, show status of *file-name* on remote machine.

rename *remote-from remote-to*

Rename *remote-from*, which can be either a file or a directory, to *remote-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.

restart *marker*

Restart the immediately following **get** or **put** at the indicated *marker*. On UNIX systems, marker is usually a byte offset into the file.

rmdir *remote-directory*

Delete *remote-directory*. *remote-directory* must be an empty directory.

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. **ftp** reports the unique filename. Note that **runique** does not affect local files generated from a shell command (see below). The default value is **off**.

send *local-file* [*remote-file*]

A synonym for **put**.

sendport

Toggle the use of **PORT** commands. By default, **ftp** attempts to use a **PORT** command when establishing a connection for each data transfer. If the **PORT** command fails, **ftp** uses the default data port. When the use of **PORT** commands is disabled, **ftp** makes no attempt to use **PORT** commands for each data transfer. This is useful for certain FTP implementations that ignore **PORT** commands but (incorrectly) indicate that they've been accepted. See *ftpd(1M)*. Turning **sendport** off may cause delays in the execution of commands.

site *arguments*

Send *arguments*, verbatim, to the server host as a **SITE** command. See *ftpd(1M)*.

size *remote-file*

Show the size of *remote-file*.

status

Show the current status of **ftp**.

struct [*struct-name*]

Set the FTP file transfer *struct* to *struct-name*. The only supported *struct* is **file**.

sunique

Toggle storing of files on remote machine under unique file names. The remote server reports the unique name. By default, **sunique** is **off**.

system

Show the type of operating system running on the remote machine.

tenex

Set the FTP file transfer *type* to **tenex**.

type [*type-name*]

Set the FTP file transfer *type* to *type-name*. If *type-name* is unspecified, write the current *type* to *stdout*. **Ascii**, **binary**, and **tenex** are the *types* currently supported.

umask [*newmask*]

Set the default umask on the remote server to *newmask*. If *newmask* is omitted, the current umask is printed.

user *user-name* [*password*] [*account*]

Log into the server host on the current connection, which must already be open. A `.netrc` file in the user's local home directory can provide the *user-name*, *password*, and optionally the *account*; see *netrc*(4). Otherwise **ftp** prompts the user for this information. In a secure environment based on Kerberos V5, **ftp** will not require a password. Instead, Kerberos authentication and authorization will be performed as described in *sis*(5). In all other environments, users are considered authenticated if they have a password and that password is correct, and authorized if an *account* exists for them on the remote system.

verbose

Toggle verbose output. If verbose output is enabled, **ftp** displays responses from the server host, and when a file transfer completes it reports statistics regarding the efficiency of the transfer.

? [*command*]

A synonym for the **help** command. Prints the **help** information for the specified *command*.

Aborting A File Transfer

To abort a file transfer, use the terminal interrupt key (usually **Ctrl-C**). Sending transfers are halted immediately. **ftp** halts incoming (receive) transfers by first sending a FTP protocol **ABOR** command to the remote server, then discarding any further received data. 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 does not appear until the remote server completes sending the requested file.

The terminal interrupt key sequence is ignored while **ftp** awaits 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 manually.

File Naming Conventions

Files specified as arguments to **ftp** commands are processed according to the following rules.

- If the file name - is specified, **ftp** uses the standard input (for reading) or standard output (for writing).
- If the first character of the file name is |, **ftp** interprets the remainder of the argument as a shell command. **ftp** forks a shell, using **popen()** (see *popen*(3S)) with the supplied argument, and reads (writes) from standard output (standard input). If the shell command includes spaces, the argument must be quoted, as in:

```
"| ls -lt"
```

Some useful examples of this mechanism are:

```
ls . "| more"
```

The above command lists the files in the current directory page by page.

```
put "| tail -20 loc_file" rem_file
```

This command copies the last twenty lines of the local file "loc_file" to the remote system as "rem_file".

- Otherwise, if globbing is enabled, **ftp** expands local file names according to the rules used by the C shell (see *cs*(1)); see the **glob** command, below. If the **ftp** command expects a single local file (e.g. **put**), only the first filename generated by the globbing operation is used.
- For **mget** commands and **get** commands with unspecified local file names, the local filename is named the same as 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.
- For **mput** commands and **put** commands with unspecified remote file names, the remote filename is named the same as 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.

WARNINGS

Correct execution of many commands depends upon proper behavior by the remote server.

DIAGNOSTICS

Error! could not retrieve authentication type.

Please notify sys admin.

There are two authentication mechanisms used by **ftp**. One authentication mechanism is based on Kerberos and the other is not. The type of authentication mechanism is obtained from a system file which is updated by `inetsvcs_sec(1M)`. If the system file does not contain known authentication types, the above error is displayed.

AUTHOR

ftp was developed by the University of California, Berkeley.

SEE ALSO

`ssh(1)`, `rcp(1)`, `ftpd(1M)`, `inetsvcs_sec(1M)`, `netrc(4)`, `ftpusers(4)`, `hosts(4)`, `sis(5)`.

f

NAME

ftpcount - show current number of users for each class

SYNOPSIS

`/usr/bin/ftpcount`

DESCRIPTION

The **ftpcount** command shows the current number of users (and the limit) for each class defined in the **ftpaccess** file. If the **ftpaccess** file does not exist, the **ftpcount** command will not display anything. However, if the **ftpaccess** file exists and it is of zero bytes, **ftpcount** will display an error message:

ftpcount: no service classes defined, no usage count kept.

EXIT STATUS

ftpcount returns:

0 successful

1 failure

AUTHOR

ftpcount was developed by the Washington University, St. Louis, Missouri.

SEE ALSO

ftpwho(1), ftpaccess(4).

NAME

ftprestart - remove the shutdown message file created by ftpshut utility.

SYNOPSIS

`/usr/bin/ftprestart`

DESCRIPTION

The **ftprestart** command removes all the shutdown message files from the real, anonymous, and virtual user accounts. The message files are created by the **ftpshut** utility in the path as specified by the 'shutdown' directive in the `/etc/ftpd/ftpaccess` file (see *ftpshut(1)* for more details). This command is always used after the **ftpshut** command is executed.

Note: For *guest* user accounts, the message files have to be removed manually. The removal will not be done by the **ftprestart** command.

EXIT STATUS

ftprestart returns:

0 successful

1 failure

AUTHOR

ftprestart was developed by the Washington University, St. Louis, Missouri.

SEE ALSO

ftpshut(1), *ftpaccess(4)*.

NAME

ftpshtut - create shutdown message file to shut down the ftp servers at a given time

SYNOPSIS

```
/usr/bin/ftpshtut [ -l min ] [ -d min ] time [ warning-message ... ]
```

DESCRIPTION

The **ftpshtut** command provides an automated shutdown procedure that a superuser can use to notify ftp users when the ftp server is shutting down. This command will create a shutdown message file in the path specified by the 'shutdown' directive in the `/etc/ftpd/ftppaccess` file in the real, anonymous and virtual user accounts. For guest accounts the system administrator must copy the message file created in the real user account to the guest accounts manually. The server will check this file regularly to see if the server is going to be shut down.

-l min This option is used as `deny_offset`. New FTP access is disabled 'min' minutes before shutdown. The default value of 'min' is 10 minutes. This value can be reset by the user.

-d min This option is used as `disc_offset`. All current FTP connections will be dropped 'min' minutes before shutdown. The default value of 'min' is 5 minutes. This value can be reset by the user.

time *time*, is the time at which the ftp server will be shutdown. If **time** is set to the word 'now' the shutdown will be immediate. *time* can also be set to a future time. Future time can be specified in one of the two formats: *+ number* or *HHMM*. The first format brings the ftp servers down in *number* minutes. The second format brings the ftp servers down at the time of day indicated, using a 24-hour clock format.

warning-message

The warning-message is the message the server will flash to its clients on shut down. The user can use a message of his choice or use the 'macros' or 'magic cookies' that are available. The server will replace the macro with the specified text string. The *warning-message* will be formatted to be 75 characters long, including the length of any expanded macros ("magic cookies"). The default warning message is "System shutdown at %s". The following magic cookies are available:

| | |
|----|--|
| %s | time system is going to shut down |
| %r | time new connections will be denied |
| %d | time current connections will be dropped |
| %C | current working directory |
| %E | the maintainer's email address as defined in <code>ftppaccess</code> |
| %L | local host name |
| %M | maximum allowed number of users in this class |
| %N | current number of users in this class |
| %R | remote host name |
| %T | local time (form Thu Nov 15 17:12:42 1990) |
| %U | username given at login time |

WARNINGS

You can kill the servers only between now and 23:59, if you use the absolute time for **ftpshtut**.

EXIT STATUS

ftpshtut returns:

0 successful

1 failure

-1 the wrong parameter was passed to **ftpshtut**.

AUTHOR

ftpshut was developed by the Washington University, St. Louis, Missouri.

SEE ALSO

ftpaccess(4).


f

NAME

ftpwho - show current process information for each ftp user.

SYNOPSIS

`/usr/bin/ftpwho`

DESCRIPTION

The `ftpwho` command shows the current process information for each user logged into the ftp server. If the `ftpaccess` file does not exist, this command will not display anything. However, if the `ftpaccess` file exists and it is of zero bytes then this command will display an error message:

`ftpwho: no service classes defined, no usage count kept.`

EXIT STATUS

`ftpwho` returns:

0 successful

1 failure

AUTHOR

`ftpwho` was developed by the Washington University, St. Louis, Missouri.

SEE ALSO

`ftpcount(1)`, `ftpaccess(4)`.

NAME

gencat - generate a formatted message catalog file

SYNOPSIS

gencat [-1] *catfile* *msgfile* ...

DESCRIPTION

Message catalogs allow a program to process input and produce output according to local customs and languages. For details, see *Native Language Support Users Guide*.

The **gencat** command merges each message source *msgfile* into a formatted message catalog *catfile* that can be accessed by **catgets()** (see *catgets(3C)*). If *catfile* does not exist, it is created. If *catfile* exists, its messages are included in the new *catfile*. If set and message numbers collide, the new message text in *file* replaces the old message text in *catfile*. A *msgfile* consists of message, directive, and comment lines (all without leading spaces or tabs) described below. Except as noted, fields are separated by one or more space or tab characters.

If - is specified as catalog file, standard output is used.

If - is specified for an instance of message file, standard input is used.

\$set *s* [*comment*] A **\$set** directive specifies the set *s*, of the messages that follow until the next **\$set** or end-of-file appears. The set number *s* is an unsigned integer in the range 1 through **NL_SETMAX**. Any string following the set number is treated as a comment. If a **\$set** directive is not specified, messages are put in the default set **NL_SETD**.

Set numbers must be in ascending order within a *msgfile* but need not be contiguous.

\$delset *s* [*comment*] A **\$delset** directive deletes the message set identified by the set number *s*, from an existing message catalog. Any string following the set number is treated as a comment.

m *message_text* A message line specifies a message number *m*, and associated message text. The message number *m* is an unsigned integer in the range 1 through **NL_MSGMAX**. The *message_text* is a C string, including spaces, tabs and \ (backslash) escapes, but by default without surrounding quotes (see **\$quote** directive below). The message number *m* is separated from the *message_text* by a single space or tab character. The *message_text* begins with the first character following the separator and ends at new-line. Extra spaces or tabs (including any trailing spaces or tabs) are considered part of the *message_text*.

The *message_text* of a message line is stored in *catfile* with message number *m* and set number *s* specified by the most recent **\$set** directive.

Message numbers must be in ascending order within a set, but need not be contiguous.

Note that the space or tab separator distinguishes insertion of a null message from deletion of a message. If a message line has a number and separator but no text, the message number and an associated null message string are stored in *catfile*. If a message line has a number but neither separator nor text, the message number and its associated message text are deleted from *catfile*.

-1 If the -1 option is specified, the length of *message_text* must be no more than **MAX_BUFLEN** - 1 bytes. If the -1 option is not specified, the length of *message_text* must be no more than **NL_TEXTMAX** bytes. See *catgets(3C)*, for message length limits imposed by these routines.

\$quote [*q* *comment*] A **\$quote** directive specifies a quote character *q*, used to surround *message_text* and make leading and trailing space visible in a message line. Any string following the specified quote character *q* is treated as a comment. By default, or if a quote character *q* not is supplied, quoting of *message_text* is not recognized.

\$ *comment*

A **\$** followed by a space or tab is treated as a comment and can appear anywhere in a file. A line consisting of zero or more spaces or tabs is treated as a comment line.

NL_TEXTMAX, **NL_SETMAX**, and **NL_MSGMAX** are defined in `<limits.h>`. **NL_SETD** is defined in `<nl_types.h>`. **MAX_BUFLEN** is defined in `<msgcat.h>`.

EXTERNAL INFLUENCES

Environment Variables

LANG provides a default value for the internationalization variables that are unset or null. If **LANG** is unset or null, the default value of "C" (see *lang(5)*) is used. If any of the internationalization variables contains an invalid setting, **gencat** will behave as if all internationalization variables are set to "C". See *environ(5)*.

LC_ALL, if set to a non-empty string value, overrides the values of all of the other internationalization variables.

LC_CTYPE determines the interpretation of text as single and/or multi-byte characters, the classification of characters as printable, and the characters matched by character class expressions in regular expressions.

LC_MESSAGES determines the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.

NLSPATH determines the location of message catalogs for the processing of **LC_MESSAGES**.

International Code Set Support

Single- and multi-byte character code sets are supported.

WARNINGS

The **\$quote** directive is not currently supported on the HP MPE and RTE operating systems.

AUTHOR

gencat was developed by HP and the X/Open Company, Ltd.

SEE ALSO

dumpmsg(1), *findmsg(1)*, *insertmsg(1)*, *catgets(3C)*, *catopen(3C)*.

Native Language Support Users Guide.

STANDARDS CONFORMANCE

gencat: XPG2, XPG3

NAME

genxlt - generate iconv translation tables

SYNOPSIS

```
genxlt [-f output_filename] [input_filename]
```

DESCRIPTION

genxlt generates a compiled, non-readable binary version of the iconv table that is suitable for use by *iconv(1)* and *iconv(3C)*. If *input_filename* or *output_filename* is not supplied, standard input and/or standard output will be used.

Since the output of **genxlt** is a binary, non-readable file, if the **-f** option is not used, the redirection symbol **>** maybe used to redirect the standard output to a file.

Options

genxlt recognizes the following options:

-f output_filename

If this option is not selected, the data will be sent to standard output, from where it could be redirected to a file.

genxlt creates tables that are in a prescribed format and which can be interpreted by the default conversion routines of *iconv(3C)*. The input file has two columns, giving the filecode mapping between the two code sets. The entries are in hexadecimal.

The input file must be formatted as two columns of hexadecimal digits. Characters in the first column are translated into the characters in the second column. Lines preceded with **#** in the first column are ignored as comments on all lines except in the case of the following keywords: *#Galley:* and *#What:*

In addition to the data, which defines the filecode mapping, a *Galley* character (see *iconv(3C)*) may also be defined for that particular conversion. This is done by adding the line **#Galley: 0xxxxxx**, to the beginning of the input file. The **xxxxxx** is any multi-byte character (see *EXAMPLES*). A *What* string (see *what(1)*), may also be defined in the input file using the entry **#What: <any_string>**. This string may contain information like version number, type of conversion, etc., which are not used in any way for the conversions. Note that if the *What* string is defined, it should appear before the *Galley* definition.

EXTERNAL INFLUENCES**Environment Variables**

LANG provides a default value for the internationalization variables that are unset or null. If **LANG** is unset or null, the default value of "C" (see *lang(5)*) is used. If any of the internationalization variables contains an invalid setting, **genxlt** will behave as if all internationalization variables are set to "C" (see *environ(5)*).

If **LC_ALL** is set to a non-empty string value, it overrides the values of all the other internationalization variables.

LC_MESSAGES determines the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.

NLSPATH determines the location of message catalogues for the processing of **LC_MESSAGES**.

International Code Set Support

Single and multi-byte character code sets are supported.

RETURN VALUE

The exit values are:

```
0      Successful completion.
>0     Error condition occurred.
```

EXAMPLES

This example compiles the *iconv_input* and puts the output binary in */usr/lib/nls/iconv/tables/roma8=iso81*. The following iconv statement uses the *roma8=iso81* table to convert the *data_file* from code set *roman8* to code set *iso8859-1*.

```
% genxlt iconv_input > /usr/lib/nls/iconv/tables/roma8=iso81
% iconv -f roma8 -t iso81 data_file
```


This is an example of the input_file:

```
#What: CodesetA to CodesetB: version 1.0
#Galley: 0xffff
# the conversion data is as follows:
0x01      0x01
0x02      0x42
...
0xff87    0x4567
...
etc.
```

WARNINGS

Because **genxlt** will write over the existing table, it is wise to save the existing table into another file before using **genxlt**.

Warnings are not given for incorrect data in the input_file.

You must have super-user privileges to install files in **/usr/lib/nls/iconv/tables**.

FILES

/usr/lib/nls/iconv/tables All tables must be installed in this directory.

SEE ALSO

dmpxlt(1), **iconv(1)**, **iconv(3C)**.

STANDARDS COMPLIANCE

genxlt: XPG4 tables

NAME

get - get a version of an SCCS file

SYNOPSIS

```
get [-r SID] [-c cutoff] [-e] [-b] [-i list] [-x list] [-k] [-l[p]] [-p] [-s] [-m] [-n] [-g] [-t]
    [-w string] [-a seq-number] file ...
```

DESCRIPTION

The **get** command generates an ASCII text file from each named SCCS file according to the specifications given by its option arguments, which begin with **-**. The arguments can be specified in any order, but all option arguments apply to all named SCCS files. If a directory is named, **get** behaves as if each file in the directory was specified as a named file, except that non-SCCS files (last component of the path name does not begin with **s.**) and unreadable files are silently ignored. If a file name of **-** is given, the standard input is read and each line of the standard input is assumed to be the name of an SCCS file to be processed. Again, non-SCCS files and unreadable files are silently ignored.

The generated text is normally written into a file called the *g-file* whose name is derived from the SCCS file name by simply removing the **s.** prefix (see FILES below).

Options

Explanation of the option arguments below is based on processing only one SCCS file. When processing multiple SCCS files, the effects of any option argument applies independently to each named file.

-rSID The SCCS IDentification string (SID) of the version (delta) of an SCCS file to be retrieved. Table 1 shows, for the most useful cases, which version of an SCCS file is retrieved (as well as the SID of the version to be eventually created by **delta** if the **-e** option is also used), as a function of the SID specified (see *delta(1)*).

-c cutoff *cutoff* date-time, in the form:

```
YY[MM[DD[HH[MM[SS]]]]]
```

No changes (deltas) to the SCCS file which were created after the specified *cutoff* date-time are included in the generated ASCII text file. Units omitted from the date-time default to their maximum possible values; that is, **-c7502** is equivalent to **-c750228235959**. Any number of non-numeric characters can separate the various 2-digit pieces of the *cutoff* date-time. This feature allows one to specify a *cutoff* date in the form: **-c77/2/2 9:22:25**. Note that this implies that one can use the **%E%** and **%U%** identification keywords (see below) for nested **gets** within a command:

```
~!get "-c%E% %U%" s.file
```

-e Indicates that the **get** is for the purpose of editing or making a change (delta) to the SCCS file via a subsequent use of **delta**. The **-e** option used in a **get** for a particular version (SID) of the SCCS file prevents further **gets** for editing on the same SID until **delta** is executed or the **j** (joint edit) flag is set in the SCCS file (see *admin(1)*). Concurrent use of **get -e** for different SIDs is always allowed. Note, however, that only one user is permitted to do a concurrent **get -e** (see *admin(1)*).

If the *g-file* generated by **get** with an **-e** option is accidentally ruined in the process of editing it, it can be regenerated by re-executing the **get** command with the **-k** option in place of the **-e** option.

SCCS file protection specified via the ceiling, floor, and authorized user list stored in the SCCS file (see *admin(1)*) are enforced when the **-e** option is used.

-b Used with the **-e** option to indicate that the new delta should have an SID in a new branch as shown in Table 1. This option is ignored if the **b** flag is not present in the file (see *admin(1)*) or if the retrieved *delta* is not a leaf *delta*. (A leaf *delta* is one that has no successors on the SCCS file tree.)

Note: A branch *delta* can always be created from a non-leaf *delta*.

-i list A *list* of deltas to be included (forced to be applied) in the creation of the generated file. The *list* has the following syntax:

```
list ::= range | list, range
range ::= SID | SID - SID
```

- SID, the SCCS Identification of a delta, can be in any form shown in the "SID Specified" column of Table 1. Partial SIDs are interpreted as shown in the "SID Retrieved" column of Table 1. See WARNINGS.
- xlist** A *list* of deltas to be excluded (forced not to be applied) in the creation of the generated file. See the **-i** option for the *list* format.
 - k** Suppresses replacement of identification keywords (see below) in the retrieved text by their value. The **-k** option is implied by the **-e** option.
 - l[p]** Causes a delta summary to be written into an *l-file*. If **-lp** is used, an *l-file* is not created; the delta summary is written on the standard output instead. See FILES for the format of the *l-file*. The user must have *s-file* read permission in order to use the **-l** option.
 - p** Causes the text retrieved from the SCCS file to be written on the standard output. No *g-file* is created. All output that normally goes to the standard output goes to file descriptor 2 (standard error) instead, unless the **-s** option is used, in which case it disappears.
 - s** Suppresses all output normally written on the standard output. However, fatal error messages (which always go to file descriptor 2) remain unaffected.
 - m** Causes each text line retrieved from the SCCS file to be preceded by the SID of the delta that inserted the text line in the SCCS file. The format is: SID, followed by a horizontal tab, followed by the text line.
 - n** Causes each generated text line to be preceded with the %M% identification keyword value (see below). The format is: %M% value, followed by a horizontal tab, followed by the text line. When both the **-m** and **-n** options are used, the format is: %M% value, followed by a horizontal tab, followed by the **-m** option-generated format.
 - g** Suppresses the actual retrieval of text from the SCCS file. It is primarily used to generate an *l-file*, or to verify the existence of a particular SID.
 - t** Used to access the most recently created ("top") delta in a given release (e.g., **-r1**), or release and level (e.g., **-r1.2**).
 - w string** Substitute *string* for all occurrences of @%M% when **getting** the file.
 - aseq-number** The delta sequence number of the SCCS file delta (version) to be retrieved (see *sccsfile(4)*). This option is used by the **comb** command (see *comb(1)*); it is not a generally useful option, and should be avoided. If both the **-r** and **-a** options are specified, the **-a** option is used. Care should be taken when using the **-a** option in conjunction with the **-e** option, because the SID of the delta to be created may not be what one expects. The **-r** option can be used with the **-a** and **-e** options to control the naming of the SID of the delta to be created.

For each file processed, **get** responds (on the standard output) with the SID being accessed and with the number of lines retrieved from the SCCS file.

If the **-e** option is used, the SID of the delta to be made appears after the SID accessed and before the number of lines generated. If there is more than one named file, or if a directory or standard input is named, each file name is printed (preceded by a new-line) before it is processed. If the **-i** option is used included deltas are listed following the notation "Included". If the **-x** option is used, excluded deltas are listed following the notation "Excluded".

Table 1. 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 |

Notes for Table 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 (i.e., maximum branch number plus one) of level L within release R". Note that 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.
- *** This is used to force creation of the *first* delta in a *new* release.
- # Successor.
- % The -b option is effective only if the b flag (see *admin(1)*) 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, then 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.

Identification Keywords

Identifying information is inserted into the text retrieved from the SCCS file by replacing *identification keywords* with their value wherever they occur. The following keywords can be used in the text stored in an SCCS file:

Keyword Value

- %M% Module name: either the value of the m flag in the file (see *admin(1)*), or if absent, the name of the SCCS file with the leading s. removed.
- %I% SCCS identification (SID) (%R%.%L%.%B%.%S%) of the retrieved text.

| | |
|-----|---|
| %R% | Release. |
| %L% | Level. |
| %B% | Branch. |
| %S% | Sequence. |
| %D% | Current date (YY/MM/DD). |
| %H% | Current date (MM/DD/YY). |
| %T% | Current time (HH:MM:SS). |
| %E% | Date newest applied delta was created (YY/MM/DD). |
| %G% | Date newest applied delta was created (MM/DD/YY). |
| %U% | Time newest applied delta was created (HH:MM:SS). |
| %Y% | Module type: value of the t flag in the SCCS file (see <i>admin</i> (1)). |
| %F% | SCCS file name. |
| %P% | Fully qualified SCCS file name. |
| %Q% | The value of the q flag in the file (see <i>admin</i> (1)). |
| %C% | Current line number. This keyword is intended for identifying messages output by the program such as "this should not have happened" type errors. It is <i>not</i> intended to be used on every line to provide sequence numbers. |
| %Z% | The 4-character string @(#) recognizable by what (see <i>what</i> (1)). |
| %W% | A shorthand notation for constructing <i>what</i> (1) strings for HP-UX system program files. %W% = %Z% %M% horizontal-tab %I% |
| %A% | Another shorthand notation for constructing <i>what</i> (1) strings for non-HP-UX system program files. %A% = %Z% %Y% %M% %I% %Z% |

EXTERNAL INFLUENCES

Environment Variables

LC_CTYPE determines the interpretation of text as single- and/or multi-byte characters.

LC_MESSAGES determines the language in which messages are displayed.

If **LC_CTYPE** or **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang*(5)) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **get** behaves as if all internationalization variables are set to "C". See *environ*(5).

International Code Set Support

Single- and multi-byte character code sets are supported.

DIAGNOSTICS

Use *scchelp*(1) for explanations.

WARNINGS

If the effective user has write permission (either explicitly or implicitly) in the directory containing the SCCS files, but the real user does not, then only one file can be named when the **-e** option is used.

Unexpected results occur when using the **-i** option to merge changes into sections of a file that have been (perhaps inadvertently) deleted and subsequently re-inserted into a file.

An *l-file* cannot be generated when **-g** is used. In other words, **-g -l** does not work.

FILES

Several auxiliary files can be created by **get**. These files are known generically as the *g-file*, *l-file*, *p-file*, and *z-file*. The letter before the hyphen is called the tag. An auxiliary file name is formed from the SCCS file name: the last component of all SCCS file names must be of the form **s.module-name**, the auxiliary

files are named by replacing the leading **s** with the tag. The *g-file* is an exception to this scheme: the *g-file* is named by removing the **s.** prefix. For example, **s.xyz.c**, the auxiliary file names would be **xyz.c**, **l.xyz.c**, **p.xyz.c**, and **z.xyz.c**, respectively.

The *g-file*, which contains the generated text, is created in the current directory (unless the **-p** option is used). A *g-file* is created in all cases, whether or not any lines of text were generated by the **get**. It is owned by the real user. If the **-k** option is used or implied its mode is 644; otherwise its mode is 444. Only the real user need have write permission in the current directory.

The *l-file* contains a table showing which deltas were applied in generating the retrieved text. The *l-file* is created in the current directory if the **-l** option is used; its mode is 444 and it is owned by the real user. Only the real user need have write permission in the current directory.

Lines in the *l-file* have the following format:

1. A blank character if the delta was applied;
* otherwise.
2. A blank character if the delta was applied or was not applied and ignored;
* if the delta was not applied and was not ignored.
3. A code indicating a "special" reason why the delta was or was not applied:
I: Included.
X: Excluded.
C: Cut off (by a **-c** option).
4. Blank.
5. SCCS identification (SID).
6. Tab character.
7. Creation date and time (in the form YY/MM/DD HH:MM:SS).
8. Blank.
9. Login name of person who created *delta*.

The comments and MR data follow on subsequent lines, indented one horizontal tab character. A blank line terminates each entry.

The *p-file* is used to pass information resulting from a **get** with an **-e** option along to *delta*. Its contents are also used to prevent a subsequent execution of **get** with an **-e** option for the same SID until *delta* is executed or the joint edit flag, **j**, (see *admin*(1)) is set in the SCCS file. The *p-file* is created in the directory containing the SCCS file and the effective user must have write permission in that directory. Its mode is 644 and it is owned by the effective user. The format of the *p-file* is: the gotten SID, followed by a blank, followed by the SID that the new delta will have when it is made, followed by a blank, followed by the login name of the real user, followed by a blank, followed by the date-time the **get** was executed, followed by a blank and the **-i** option argument if it was present, followed by a blank and the **-x** option argument if it was present, followed by a new-line. There can be an arbitrary number of lines in the *p-file* at any time; no two lines can have the same new delta SID.

The *z-file* serves as a *lock-out* mechanism against simultaneous updates. Its contents are the binary (2 bytes) process ID of the command (i.e., **get**) that created it. The *z-file* is created in the directory containing the SCCS file for the duration of **get**. The same protection restrictions as those for the *p-file* apply for the *z-file*. The *z-file* is created mode 444.

SEE ALSO

admin(1), **delta**(1), **prs**(1), **sccshelp**(1), **what**(1), **sccsfile**(4).

STANDARDS CONFORMANCE

get: SVID2, SVID3, XPG2, XPG3, XPG4

NAME

getaccess - list access rights to file(s)

SYNOPSIS

getaccess [-u *user*] [-g *user*] *group* [, *group*] ... [-n] *file* ...

getaccess -r [-n] *file* ...

DESCRIPTION

getaccess lists for the specified files the effective access rights of the caller (that is, for their effective user ID, effective group ID, and supplementary groups list). By default, the command prints a symbolic representation of the user's access rights to the named file: **r** or **-** for read/no read, **w** or **-** for write/no write, and **x** or **-** for execute/no execute (for directories, search/no search), followed by the file name.

Options

getaccess recognizes the following options and command-line arguments:

- u *user* List access for the given user instead of the caller. A *user* can be a known user name, a valid ID number, or @, representing the file's owner ID. If information about more than one file is requested, the value of @ can differ for each.
This option sets the user ID only. The access check is made with the caller's effective group ID and supplementary group IDs unless **-g** is also specified.
- g *group* [, *group*] ... List access for the given group(s) instead of the caller's effective group ID and supplementary groups list. A *group* can be a known group name, a valid ID number, or @, representing the file's group ID. If information about more than one file is requested, the value of @ can differ for each.
- r List access using the caller's real user ID, group ID, and supplementary groups list, instead of effective ID values.
- n List access rights numerically (octal digits 0..7 instead of **rw****x**) for each file requested. The bit values **R_OK**, **W_OK**, and **X_OK** are defined in the file `<unistd.h>`.

Checking access using access control lists is described in *acl*(5) and *aclv*(5).

In addition, the write bit is cleared for files on read-only file systems or shared-text programs being executed. The execute bit is not turned off for shared-text programs open for writing because it is not possible to ascertain whether a file open for writing is a shared-text program.

Processes with appropriate privileges have read and write access to all files. However, write access is denied for files on read-only file systems or shared-text programs being executed. Execute access is allowed if and only if the file is not a regular file or the execute bit is set in any of the file's ACL entries.

To use **getaccess** successfully, the caller must have search access in every directory component of the path name of the *file*. **getaccess** verifies search access first by using the caller's effective IDs, regardless of the user and group IDs specified. This is distinct from the case in which the caller can search the path but the user for whom access is being checked does not have access to the file.

Note: a file name argument of **-** has no special meaning (such as standard input) to **getaccess**.

EXTERNAL INFLUENCES**Environment Variables**

LANG determines the language in which messages are displayed.

If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang*(5)) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **getaccess** behaves as if all internationalization variables are set to "C". See *environ*(5).

RETURN VALUE

getaccess returns one of the following values:

- 0** Successful completion.
- 1** **getaccess** was invoked incorrectly or encountered an unknown user or group name. An appropriate message is printed to standard error.

- 2** A file is nonexistent or unreachable (by the caller). **getaccess** prints an appropriate message to standard error, continues, then returns a value of 2 upon completion.

EXAMPLES

The following command prints the caller's access rights to *file1* using the file's group ID instead of the caller's effective group ID and groups list.

```
getaccess -g@ file1
```

Here's how to check access by user **ggd** in groups **red** and **19** to all files in the current directory, with access rights expressed as octal values.

```
getaccess -u ggd -g red,19 -n .* *
```

Here's how to list access rights for all files under **mydir**.

```
find mydir -print | sort | xargs getaccess
```

AUTHOR

getaccess was developed by HP.

FILES

```
/etc/passwd  
/etc/group
```

SEE ALSO

chacl(1), getacl(1), lsacl(1), setacl(1), getaccess(2), glossary(9).

NAME

getacl - list access control lists (ACLs) for files (JFS File Systems only)

SYNOPSIS

`/usr/bin/getacl [-ad] file...`

DESCRIPTION

For each argument that is a regular file, special file, or named pipe, **getacl** displays the owner, group, and the Access Control List (ACL). For each directory argument, **getacl** displays the owner, group, and the ACL and/or the default ACL. Only directories contain default ACLs.

With the **-a** option specified, the filename, owner, group, and the ACL of the file will be displayed. With the **-d** option specified, the filename, owner, group, and the default ACL of the file, if it exists, will be displayed. With options not specified, the filename, owner, group, and both the ACL, and the default ACL, if it exists, will be displayed.

This command may be executed on a file system that does not support ACLs. It will report the ACL consisting of only the owning user, owning group, class and other entries, based on the permission bits.

When multiple files are specified on the command line, a blank line will separate the ACL for each file. The format of an ACL is:

```
# file: filename
# owner: uid
# group: gid
user::perm
user:uid:perm
group::perm
group:gid:perm
class:perm
other:perm
default:user::perm
default:user:uid:perm
default:group::perm
default:group:gid:perm
default:class:perm
default:other:perm
```

The first three lines show the filename, the file owner, and the file owning group. Note that when only the **-d** option is specified, and the file has no default ACL, only these three lines will be displayed.

The **user** entry without a user ID indicates the permissions that will be granted to the owner of the file. One or more additional **user** entries indicate the permissions that will be granted to the specified users. The **group** entry without a group identifier indicates the permissions that will be granted to the owning group of the file. One or more additional **group** entries indicate the permissions that will be granted to the specified groups. The **other** entry indicates the permissions that will be granted to others.

The **default** entries (**default:user**, **default:group**, and **default:other**) may only exist for directories, and indicate the default user, group, and other entries that will be added to a file created within the directory.

The *uid* is a login name, or a user ID if there is no entry for the *uid* in the system's password file; *gid* is a group name, or a group ID if there is no entry for the *gid* in the system's group file; and *perm* is a three character string composed of the letters representing the separate discretionary access rights: **r** (read), **w** (write), **x** (execute/search), or the placeholder character **-**. The *perm* will be displayed in the following order: **rwX**. If a permission is not ACL entry, the placeholder character will appear.

The ACL entries will be displayed in the order in which they will be evaluated when an access check is performed. The default ACL entries which may exist on a directory have no effect on access checks.

The file owner permission bits represent the access that the owning user ACL entry has. The file group class permission bits represent the most access that any additional user entry, additional group entry, or the owning group entry may grant. The file other permission bits represent the access that the other ACL entry has. If a user invokes the **chmod** command and changes the file group class permission bits, the access granted by the additional ACL entries may be restricted.

In order to indicate that the file group class permission bits restrict an ACL entry, **getacl** will display, after each affected entry, text in the form **#effective:perm**, where *perm* will show only the

permissions actually granted.

EXAMPLES

Given file `filea`, with an ACL six entries long, the command

```
$ getacl filea
```

would print:

```
# file: filea
# owner: fletcher
# group: us
user::rwx
user:spy:---
user:archer:rw-
group::r--
class:rw-
other:---
```

Given file `filea`, with an ACL six entries long, after the command `chmod 700 filea` was issued, the command

```
$ getacl filea
```

would print:

```
# file: filea
# owner: fletcher
# group: us
user::rwx
user:spy:---
user:archer:rw- #effective:---
group::r-- #effective:---
class:---
other:---
```

Given directory `fileb`, with an ACL containing default entries, the command

```
$ getacl -d fileb
```

would print:

```
# file: fileb
# owner: fletcher
# group: us
default:user::rwx
default:user:spy:---
default:group::r--
default:other:---
```

Given directory `fileb`, the command

```
$ getacl fileb
```

would print:

```
# file: fileb
# owner: fletcher
# group: us
user::rwx
user:spy:---
user:archer:rw-
group::r--
other:---
default:user::rwx
default:user:spy:---
default:group::r--
default:other:---
```

NOTICES

The output from **getacl** will be in the correct format for input to the **setacl** command. If the output from **getacl** is redirected to a file, the file may be used as input to **setacl**. In this way, a user may easily assign one file's ACL to another file.

FILES

| | |
|--------------------|---------------|
| /etc/passwd | for user IDs |
| /etc/group | for group IDs |

SEE ALSO

acl(2), **aclsort(3C)**, **chmod(1)**, **ls(1)**, **setacl(1)**.

NAME

getconf - get system configuration values

SYNOPSIS

getconf [-v *specification*] *system_var*

getconf [-v *specification*] *system_var* *pathname*

DESCRIPTION

The **getconf** command provides an interface to the *confstr*(3C), *pathconf*(2), and *sysconf*(2) library routines and system calls.

The *system_var* argument specifies the configuration value desired in **confstr()**, **pathconf()**, or **sysconf()**. Use the first synopsis form, for inquiries involving **confstr()**, or **sysconf()** (in the first table below). Use the second synopsis form, for inquiries involving **pathconf()** (in the second table below). For inquiries involving **pathconf()** the *pathname* operand should be specified.

Options

getconf recognizes the following option:

-v *specification* Return configuration variables corresponding to a particular compilation environment supported by HP-UX. If the **-v** option is not specified, *specification* defaults to XBS5_ILP32_OFF32. See table below for possible specifications and meanings.

| Specification | int | long | pointer | off_t |
|----------------------|------------|-------------|----------------|--------------|
| XBS5_ILP32_OFF32 | 32 | 32 | 32 | 32 |
| XBS5_ILP32_OFFBIG | 32 | 32 | 32 | >=64 |
| XBS5_LP64_OFF64 | 32 | 64 | 64 | 64 |
| XBS5_LPBIG_OFFBIG | >=32 | >=64 | >=64 | >=64 |

EXTERNAL INFLUENCES**Environment Variables**

LC_MESSAGES determines the language in which messages are displayed.

If **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang*(5)) is used instead of **LANG**.

If any internationalization variable contains an invalid setting, **getconf** behaves as if all internationalization variables are set to "C". See *environ*(5).

International Code Set Support

Single-byte and multi-byte character code sets are supported.

RETURN VALUE

The error codes returned by **getconf** are:

- 0 Success. A value corresponding to the operand was returned.
- 1 One or more missing operands.
- 2 Operand was not recognized.
- 3 *pathname* either invalid or inaccessible.

EXAMPLES

Request the number of intervals per second:

```
getconf CLK_TCK
```

Request the maximum value of a file's link count:

```
getconf LINK_MAX /etc/passwd
```

Other supported inquiries include the following:

ARG_MAX

_BC_BASE_MAX

BC_DIM_MAX

| | | |
|--------------------------|-----------------------------|----------------------|
| BS_SCALE_MAX | BC_STRING_MAX | CHARCLASS_NAME_MAX |
| CHAR_BIT | CHAR_MAX | CHAR_MIN |
| CHILD_MAX | CLK_TCK | COLL_WEIGHTS_MAX |
| CPU_CHIP_TYPE | CS_MACHINE_IDENT | CS_PARTITION_IDENT |
| CS_PATH | CS_MACHINE_SERIAL | EXPR_NEST_MAX |
| HW_CPU_SUPP_BITS | HW_32_64_CAPABLE | INT_MAX |
| INT_MIN | KERNEL_BITS | LINE_MAX |
| LONG_BIT | LONG_MAX | LONG_MIN |
| MACHINE_IDENT | MACHINE_MODEL | MACHINE_SERIAL |
| MB_LEN_MAX | NGROUPS_MAX | NL_ARGMAX |
| NL_LANGMAX | NL_MSGMAX | NL_NMAX |
| NL_SETMAX | NL_TEXTMAX | NZERO |
| OPEN_MAX | PARTITION_IDENT | PATH |
| _POSIX_ARG_MAX | _POSIX_JOB_CONTROL | _POSIX_NGROUPS_MAX |
| _POSIX_OPEN_MAX | _POSIX_SAVED_IDS | _POSIX_SSIZE_MAX |
| _POSIX_STREAM_MAX | _POSIX_TZNAME_MAX | _POSIX_VERSION |
| POSIX_ARG_MAX | POSIX_CHILD_MAX | POSIX_JOB_CONTROL |
| POSIX_LINK_MAX | POSIX_MAX_CANON | POSIX_MAX_INPUT |
| POSIX_NAME_MAX | POSIX_NGROUPS_MAX | POSIX_OPEN_MAX |
| POSIX_PATH_MAX | POSIX_PIPE_BUF | POSIX_SAVED_IDS |
| POSIX_SSIZE_MAX | POSIX_STREAM_MAX | POSIX_TZNAME_MAX |
| POSIX_VERSION | POSIX2_BC_BASE_MAX | POSIX2_BC_DIM_MAX |
| POSIX2_BC_SCALE_MAX | POSIX2_BC_STRING_MAX | POSIX2_C_BIND |
| POSIX2_C_DEV | POSIX2_C_VERSION | POSIX2_CHAR_TERM |
| POSIX_CHILD_MAX | POSIX2_COLL_WEIGHTS_MAX | POSIX2_EXPR_NEST_MAX |
| POSIX2_FORT_DEV | POSIX2_FORT_RUN | POSIX2_LINE_MAX |
| POSIX2_LOCALEDEF | POSIX2_RE_DUP_MAX | POSIX2_SW_DEV |
| POSIX2_UPE | POSIX2_VERSION | SC_PASS_MAX |
| SC_XOPEN_VERSION | SCHAR_MAX | SCHAR_MIN |
| SHRT_MAX | SHRT_MIN | SSIZE_MAX |
| STREAM_MAX | RE_DUP_MAX | TMP_MAX |
| TZNAME_MAX | UCHAR_MAX | UINT_MAX |
| ULONG_MAX | USHRT_MAX | WORD_BIT |
| XOPEN_VERSION | XOPEN_XCU_VERSION | XOPEN_XPG2 |
| XOPEN_XPG3 | XOPEN_XPG4 | |
| XBS5_ILP32_OFF32_CFLAGS | XBS5_ILP32_OFF32_LDFLAGS | |
| XBS5_ILP32_OFF32_LIBS | XBS5_ILP32_OFF32_LINTFLAGS | |
| XBS5_ILP32_OFFBIG_CFLAGS | XBS5_ILP32_OFFBIG_LDFLAGS | |
| XBS5_ILP32_OFFBIG_LIBS | XBS5_ILP32_OFFBIG_LINTFLAGS | |
| XBS5_LP64_OFF64_CFLAGS | XBS5_LP64_OFF64_LDFLAGS | |
| XBS5_LP64_OFF64_LIBS | XBS5_LP64_OFF64_LINTFLAGS | |

Supported inquiries requiring the second parameter include:

| | | |
|-------------------------|-----------------|------------------|
| LINK_MAX | MAX_CANON | MAX_INPUT |
| NAME_MAX | PATH_MAX | PIPE_BUF |
| _POSIX_CHOWN_RESTRICTED | _POSIX_LINK_MAX | _POSIX_MAX_CANON |
| _POSIX_MAX_INPUT | _POSIX_NO_TRUNC | _POSIX_NAME_MAX |
| _POSIX_PATH_MAX | _POSIX_PIPE_BUF | _POSIX_VDISABLE |
| POSIX_CHOWN_RESTRICTED | POSIX_NO_TRUNC | POSIX_VDISABLE |

AUTHOR

getconf was developed by HP and POSIX.

SEE ALSO

pathconf(2), sysconf(2), confstr(3C).

STANDARDS CONFORMANCE

getconf: POSIX.2, XPG4

NAME

getopt - parse command options

SYNOPSIS

getopt *optstring args*

DESCRIPTION

getopt is used to break up options in command lines for easy parsing by shell procedures and to check for legal options. *optstring* 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 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.

getopt recognizes two hyphens (- -) to delimit the end of the options. If absent, **getopt** places -- at the end of the options.

The most common use of **getopt** is in the shell's **set** command (see the example below) where **getopt** converts the command line to a more easily parsed form. **getopt** writes the modified command line to the standard output.

EXTERNAL INFLUENCES**Environment Variables**

LC_MESSAGES determines the language in which messages are displayed.

If **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable.

If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **getopt** behaves as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single-byte and multi-byte character code sets are supported.

DIAGNOSTICS

getopt prints an error message on the standard error when it encounters an option letter not included in *optstring*.

EXAMPLES

The following code fragment processes the arguments for a command that can take the options **a** or **b**, and the option **o** which requires an argument:

```
set -- `getopt abo: $*`
if [ $? -ne 0 ]; then
    echo $USAGE
    exit 2
fi
while [ $# -gt 0 ]; do
    case $1 in
        -a | -b)
            FLAG=$1
            shift
            ;;
        -o)
            OARG=$2
            shift 2
            ;;
        --)
            shift
            break
            ;;
    esac
done
```

done

This code accepts any of the following as equivalent:

```
cmd -aoarg file file
cmd -a -o arg file file
cmd -oarg -a file file
cmd -a -oarg -- file file
```

WARNINGS

getopt option arguments must not be null strings nor contain embedded blanks.

SEE ALSO

sh(1), getopt(3C).

NAME

getopts - parse utility (command) options

SYNOPSIS

getopts *optstring name* [*arg* ...]

DESCRIPTION

getopts is used to retrieve options and option-arguments from a list of parameters.

Each time it is invoked, **getopts** places the value of the next option in the shell variable specified by the **name** operand and the index of the next argument to be processed in the shell variable **OPTIND**. Whenever the shell is invoked, **OPTIND** is initialized to 1.

When the option requires an option-argument, **getopts** places it in the shell variable **OPTARG**. If no option was found, or if the option that was found does not have an option-argument, **OPTARG** is unset.

If an option character not contained in the *optstring* operand is found where an option character is expected, the shell variable specified by *name* is set to the question-mark (?) character. In this case, if the first character in *optstring* is a colon (:), the shell variable **OPTARG** is set to the option character found, but no output is written to standard error; otherwise, the shell variable **OPTARG** is unset and a diagnostic message is written to standard error. This condition is considered to be an error detected in the way arguments were presented to the invoking application, but is not an error in **getopts** processing.

If an option-argument is missing:

- If the first character of *optstring* is a colon, the shell variable specified by *name* is set to the colon character and the shell variable **OPTARG** is set to the option character found.
- Otherwise, the shell variable specified by *name* is set to the question-mark character, the shell variable **OPTARG** is unset, and a diagnostic message is written to the standard error. This condition is considered to be an error detected in the way arguments are presented to the invoking application, but is not an error in **getopts** processing; a diagnostic message is written as stated, but the exit status is zero.

When the end of options is encountered, **getopts** exits with a return value greater than zero. The shell variable **OPTIND** is set to the index of the first nonoption-argument, where the first -- argument is considered to be an option argument if there are no other non-option arguments appearing before it, or the value \$# + 1 if there are no nonoption-arguments; the *name* variable is set to the question-mark character. Any of the following identifies the end of options: the special option --, finding an argument that does not begin with a -, or encountering an error.

The shell variables **OPTIND** and **OPTARG** are local to the caller of **getopts** and are not exported by default.

The shell variable specified by the *name* operand, **OPTIND**, and **OPTARG** affect the current shell execution environment.

Operands

The following operands are supported:

| | |
|------------------|--|
| <i>optstring</i> | A string containing the option characters recognized by the utility invoking getopts . If a character is followed by a colon (:), the option will be expected to have an argument, which should be supplied as a separate argument. Applications should specify an option character and its option-argument as separate arguments, but getopts will interpret the characters following an option character requiring arguments as an argument whether or not this is done. An explicit null option-argument need not be recognised if it is not supplied as a separate argument when getopts is invoked. The characters question-mark (?) and colon (:) must not be used as option characters by an application. The use of other option characters that are not alphanumeric produces unspecified results. If the option-argument is not supplied as a separate argument from the option character, the value in OPTARG will be stripped of the option character and the -. The first character in <i>optstring</i> will determine how getopts will behave if an option character is not known or an option-argument is missing. |
| <i>name</i> | The name of a shell variable that is set by getopts to the option character that was found. |

getopts by default parses positional parameters passed to the invoking shell procedures. If *args* are given, they are parsed instead of the positional parameters.

EXTERNAL INFLUENCES

Environment Variable

The following environment variable affects the execution of the **getopts** utility:

OPTIND Used by **getopts** as the index of the next argument to be processed.

ERRORS

Whenever an error is detected and the first character in the *optstring* operand is not a colon (:), a diagnostic message will be written to standard error with the following information in an unspecified format:

- The invoking program name will be identified in the message. The invoking program name will be the value of the shell special parameter 0 at the time the **getopts** utility is invoked. A name equivalent to:

```
basename "$0"
```

may be used.

- If an option is found that was not specified in *optstring*, this error will be identified and the invalid option character will be identified in the message.
- If an option requiring an option-argument is found, but an option-argument is not found, this error will be identified and the invalid option character will be identified in the message.

EXAMPLES

Since **getopts** affects the current shell execution environment, it is generally provided as a shell regular built-in. If it is called in a subshell or separate utility execution environment such as one of the following:

```
(getopts abc value "$@")
nohup getopts ...
find -exec getopts ...\;
```

it does not affect the shell variables in the caller's environment.

Note that shell functions share **OPTIND** with the calling shell even though the positional parameters are changed. Functions that use **getopts** to parse their arguments should save the value of **OPTIND** on entry and restore it before returning. However, there will be cases when a function must change **OPTIND** for the calling shell.

The following example script parses and displays its arguments:

```
aflag=
bflag=
while getopts ab: name
do
    case $name in
        a)
            aflag=1;;
        b)
            bflag=1
            bval="$OPTARG";;
        ?)
            printf "Usage: %s: [-a] [-b value] args\n" $0
            exit 2;;
    esac
done
if [ ! -z "$aflag" ] ; then
    printf "Option -a specified\n"
fi
if [ ! -z "$bflag" ] ; then
    printf "Option -b \"%s\" specified\n" "$bval"
fi
shift $((OPTIND -1))
printf "Remaining arguments are: %s\n" "$@"
```

SEE ALSO

getopt(1), ksh(1), sh-posix(1), sh(1), getopt(3C).

STANDARDS CONFORMANCE

getopts: XPG4, POSIX.2

NAME

getprivgrp - get special attributes for group

SYNOPSIS

getprivgrp [-g | *group_name*]

DESCRIPTION

getprivgrp lists the access privileges of privileged groups set by **setprivgrp** (see *setprivgrp(1M)*). If *group_name* is supplied, access privileges are listed for that group only. If the caller is not a member of *group_name*, no information is displayed. If **-g** is used, **getprivgrp** lists access privileges that have been granted to all groups. Otherwise, access privileges are listed for all privileged groups to which the caller belongs.

The super-user is a member of all groups. Access privileges include **RTPRIO**, **RTSCHED**, **MLOCK**, **CHOWN**, **LOCKRDONLY**, **SETRUGID**, and **SERIALIZE**.

AUTHOR

getprivgrp was developed by HP.

SEE ALSO

setprivgrp(1M), *getprivgrp(2)*, *privgrp(4)*.

NAME

gprof - display call graph profile data

SYNOPSIS

gprof [*options*] [*a.out* [*gmon.out* ...]] [*shared_library* *shared_library_profile*]

DESCRIPTION

The **gprof** command produces an execution profile of C++, C, Pascal, and FORTRAN programs. The effect of called routines is incorporated into the profile of each caller. Profile data is taken from the call graph profile file (**gmon.out** default) that is created by programs compiled with the **-G** option of **aCC**, **CC**, **cc**, **pc**, and **f77**. That option also links in versions of the library routines that are compiled for profiling. The symbol table in the named object file (**a.out** default) is read and correlated with the call graph profile file. If more than one profile file is specified, **gprof** output shows the sum of the profile information in the given profile files.

First, a flat profile is given, similar to that provided by **prof** (see *prof(1)*). This listing gives the total execution times and call counts for each function in the program, sorted by decreasing time.

Next, these times are propagated along the edges of the call graph. **gprof** discovers all cycles in the call graph. All calls made into the cycle share the time of that cycle. A second 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 the time of this function and the time of its descendants are propagated to its (direct) call graph parents.

Cycles are also shown, with an entry for the cycle as a whole and a listing of the members of the cycle, each with their contributions to the time and call counts of the cycle.

Shared Library Profiling (32-bit only)

Support for **gprof** profiling of shared libraries is available on 32-bit systems only.

To profile shared libraries, set the environment variable **LD_PROFILE** to the path of the shared library to be profiled. (See *HP-UX Linker and Libraries Online User's Guide* for details.) Do not use the **-G** option to compile programs for shared libraries profiling. Do not link the executable **gcrt0.o** or **mcrt0.o**. This turns on profiling of **a.out**, which is not compatible with profiling of shared libraries. You can either profile your executable or a shared library, but not both.

At the termination of the program, a profile file with the name of the shared library prepended to it is generated by a run-time library. To get the complete listing, provide the **gprof** command with names of the shared library and the profile file for the shared library as arguments.

Options

The **gprof** command recognizes the following options:

- a** Suppress printing statically declared functions. If this option is given, all relevant information about the static function (such as time samples, calls to other functions, and calls from other functions) belongs to the function loaded just before the static function in the **a.out** file.
- b** Suppress printing a description of each field in the profile.
- e name** Suppress printing the graph profile entry for routine *name* and all its descendants (unless they have other ancestors that are not suppressed). More than one **-e** option can be given. Only one *name* can be given with each **-e** option.
- E name** Suppress printing the graph profile entry for routine *name* (and its descendants) as **-e** above, and also exclude the time spent in *name* (and its descendants) from the total and percentage time computations. **-E mcount** **-E mcleanup** is the default.
- f name** Print only the graph profile entry of the specified routine *name* and its descendants. More than one **-f** option can be given. Only one *name* can be given with each **-f** option.
- F name** Print only the graph profile entry of the routine *name* and its descendants (as **-f** above) and also use only the times of the printed routines in total time and percentage computations. More than one **-F** option can be given. Only one *name* can be given with each **-F** option. The **-F** option overrides the **-E** option.

- s** Produce a profile file **gmon.sum** that represents the sum of the profile information in all specified profile files. This summary profile file can be given to subsequent executions of **gprof** (probably also with a **-s** option) to accumulate profile data across several runs of an **a.out** file.
- z** Display routines that have zero usage (as indicated by call counts and accumulated time).

The name of the file created by a profiled program is controlled by the environment variable **GPROFDIR**. If **GPROFDIR** is not set, **gmon.out** is produced in the current directory when the program terminates. If **GPROFDIR=string**, **string/pid.progname** is produced, where *progname* consists of *argv[0]* with any path prefix removed, and *pid* is the program's process ID. If **GPROFDIR** is set to a null string, no profiling output is produced.

EXAMPLES

To profile **libc.sl**:

```
$ cat > test.c
main()
{
printf("hello world\n");
}

$ cc test.c -lc

$ ldd a.out
/usr/lib/libc.2 => /usr/lib/libc.2
/usr/lib/libdld.2 => /usr/lib/libdld.2
/usr/lib/libc.2 => /usr/lib/libc.2

$ export LD_PROFILE=/usr/lib/libc.2

$ ./a.out
hello world

$ unset LD_PROFILE

$ ls libc.2.profile
libc.2.profile

$ ll libc.2.profile
-rw-rw-r-- 1 user1 lang 606464 May 19 10:24 libc.2.profile

$ gprof /usr/lib/libc.2 libc.2.profile
```

WARNINGS

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.

Parents that are not profiled have the time of their profiled children propagated to them, but they appear to be spontaneously invoked in the call graph listing, and do not have their time propagated further. Similarly, signal catchers, even though profiled, 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.

The profiled program must call **exit()** (see *exit(2)*) or return normally, for the profiling information to be saved in the **gmon.out** file.

The following limitations exist for **gprof** shared library profiling:

- Local, static, and hidden functions are not profiled.
- Shared libraries built with **-B symbolic** are not profiled.
- Any function calls made from library initializers are not collected.

Set **LD_PROFILE** to the exact string with which you call **shl_load**. If the library is implicitly loaded, **LD_PROFILE** must match the path encoded in the **a.out**. You can find this value by running the **ldd** command on the executable.

DEPENDENCIES

gprof cannot be used with dynamically linked executables (built with **ld -A** in pre HP-UX 10.20 releases).

AUTHOR

gprof was developed by the University of California, Berkeley.

FILES

| | |
|-------------------------------|--|
| a.out* | Default object file. |
| gmon.out* | Default dynamic call graph and profile. |
| gmon.sum* | Summarized dynamic call graph and profile. |
| /usr/lib/gprof.callg* | Call graph description. |
| /usr/lib/gprof.flat* | Flat profile description. |
| /usr/lib/libgprof32.sl | gprof 32-bit shared library profiler |

SEE ALSO

cc_bundled(1), **prof(1)**, **exit(2)**, **profil(2)**, **crt0(3)**, **monitor(3C)**.

gprof: A Call Graph Execution Profiler; Graham, S.L., Kessler, P.B., McKusick, M.K.;

Proceedings of the SIGPLAN '82 Symposium on Compiler Construction; SIGPLAN Notices; Vol. 17, No. 6, pp. 120-126, June 1982.

HP-UX Linker and Libraries Online User's Guide (See the **ld +help** option).

NAME

grep, egrep, fgrep - search a file for a pattern

SYNOPSIS

Plain call with pattern

```
grep [-E|-F] [-c|-l|-q] [-bhinsvwx] pattern [file ...]
```

Call with (multiple) -e pattern

```
grep [-E|-F] [-c|-l|-q] [-bhinsvwx] -e pattern... [-e pattern] ... [file ...]
```

Call with -f file

```
grep [-E|-F] [-c|-l|-q] [-bhinsvwx] [-f pattern_file] [file ...]
```

Obsolescent:

```
egrep [-cefilnsv] [expression] [file ...]
```

```
fgrep [-cefilnsvx] [strings] [file ...]
```

DESCRIPTION

The **grep** command searches the input text *files* (standard input default) for lines matching a pattern. Normally, each line found is copied to the standard output. **grep** supports the Basic Regular Expression syntax (see *regex*(5)). The **-E** option (**egrep**) supports Extended Regular Expression (ERE) syntax (see *regex*(5)). The **-F** option (**fgrep**) searches for fixed *strings* using the fast Boyer-Moore string searching algorithm. The **-E** and **-F** options treat newlines embedded in the *pattern* as alternation characters. A null expression or string matches every line.

The forms **egrep** and **fgrep** are maintained for backward compatibility. The use of the **-E** and **-F** options is recommended for portability.

Options

- E** Extended regular expressions. Each pattern specified is a sequence of one or more EREs. The EREs can be separated by newline characters or given in separate **-e expression** options. A pattern matches an input line if any ERE in the sequence matches the contents of the input line without its trailing newline character. The same functionality is obtained by using **egrep**.
- F** Fixed strings. Each pattern specified is a sequence of one or more strings. Strings can be separated by newline characters or given in separate **-e expression** options. A pattern matches an input line if the line contains any of the strings in the sequence. The same functionality is obtained by using **fgrep**.
- b** Each line is preceded by the block number on which it was found. This is useful in locating disk block numbers by context. Block numbers are calculated by dividing by 512 the number of bytes that have been read from the file and rounding down the result.
- c** Only a count of matching lines is printed.
- e expression** Same as a simple *expression* argument, but useful when the *expression* begins with a hyphen (-). Multiple **-e** options can be used to specify multiple patterns; an input line is selected if it matches any of the specified patterns.
- f pattern_file** The regular *expression* (**grep** and **grep -E**) or *strings* list (**grep -F**) is taken from the *pattern_file*.
- h** Suppress printing of filenames when searching multiple files.
- i** Ignore uppercase/lowercase distinctions during comparisons.
- l** Only the names of files with matching lines are listed (once), separated by newlines. If standard input is searched, a path name of (**standard input**) will be written, in the POSIX locale. In other locales, (**standard input**) may be replaced by something more appropriate in those locales.
- n** Each line is preceded by its relative line number in the file starting at 1. The line number is reset for each file searched. This option is ignored if **-c**, **-b**, **-l**, or **-q** is specified.

| | |
|-----------|---|
| -q | (Quiet) Do not write anything to the standard output, regardless of matching lines. Exit with zero status upon finding the first matching line. Overrides any options that would produce output. |
| -s | Error messages produced for nonexistent or unreadable files are suppressed. |
| -v | All lines but those matching are printed. |
| -w | Select only those lines containing matches that form whole words. The test is that the matching substring must either be at the beginning of the line, or preceded by a non-word constituent character. Similarly, it must be either at the end of the line or followed by a non-word constituent character. Word-constituent characters are letters, digits, and the underscore. |
| -x | (eXact) Matches are recognized only when the entire input line matches the fixed string or regular expression. |

The file name is output in all the cases in which output is generated if there are more than one input file, unless the **-h** option is specified. Care should be taken when using the characters **\$**, *****, **[**, **^**, **|**, **(**, **)**, and **** in *expression*, because they are also meaningful to the shell. It is safest to enclose the entire *expression* argument in single quotes ('...').

EXTERNAL INFLUENCES

Environment Variables

LANG determines the locale to use for the locale categories when both **LC_ALL** and the corresponding environment variable (beginning with **LC_**) do not specify a locale. If **LANG** is not specified or is set to the empty string, a default of **C** (see *lang(5)*) is used.

LC_ALL determines the locale to use to override any values for locale categories specified by the settings of **LANG** or any environment variables beginning with **LC_**.

LC_COLLATE determines the collating sequence used in evaluating regular expressions.

LC_CTYPE determines the interpretation of text as single byte and/or multi-byte characters, the classification of characters as letters, the case information for the **-i** option, and the characters matched by character class expressions in regular expressions.

LC_MESSAGES determines the language in which messages are displayed.

If any internationalization variable contains an invalid setting, the commands behave as if all internationalization variables are set to **C**. See *environ(5)*.

International Code Set Support

Single-byte and multi-byte character code sets are supported.

RETURN VALUE

Upon completion, **grep** returns one of the following values:

- 0 One or more matches found.
- 1 No match found.
- 2 Syntax error or inaccessible file (even if matches were found).

EXAMPLES

In the Bourne shell (*sh(1)*) the following example searches two files, finding all lines containing occurrences of any of four strings:

```
grep -F 'if
then
else
fi' file1 file2
```

Note that the single quotes are necessary to tell **grep -F** when the strings have ended and the file names have begun.

For the C shell (see *csh(1)*) the following command can be used:

```
grep -F 'if\ then\ else\ fi' file1 file2
```

To search a file named **address** containing the following entries:


```
Ken    112 Warring St.  Apt. A
Judy   387 Bowditch  Apt. 12
Ann    429 Sixth St.
```

the command:

```
grep Judy address
```

prints:

```
Judy   387 Bowditch  Apt. 12
```

To search a file for lines that contain either a **Dec** or **Nov**, use either of the following commands:

```
grep -E '[Dd]ec|[Nn]ov' file
egrep -i 'dec|nov' file
```

Search all files in the current directory for the string **xyz**:

```
grep xyz *
```

Search all files in the current directory subtree for the string **xyz**, and ensure that no error occurs due to file name expansion exceeding system argument list limits:

```
find . -type f -print |xargs grep xyz
```

The previous example does not print the name of files where string **xyz** appears. To force **grep** to print file names, add a second argument to the **grep** command portion of the command line:

```
find . -type f -print |xargs grep xyz /dev/null
```

In this form, the first file name is that produced by **find**, and the second file name is the null file.

WARNINGS

(XPG4 only.) If the **-q** option is specified, the exit status will be zero if an input line is selected, even if an error was detected. Otherwise, default actions will be performed.

If the **-w** option is specified with non-word constituent characters, then the output is unexpected.

SEE ALSO

sed(1), sh(1), regcomp(3C), environ(5), lang(5), regexp(5).

STANDARDS CONFORMANCE

grep: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

egrep: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

fgrep: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

groups - show group memberships

SYNOPSIS

groups [-p] [-g] [-l] [*user*]

DESCRIPTION

groups shows the groups to which the caller or the optionally specified *user* belong. If invoked with no arguments, **groups** prints the current access list returned by **getgroups()** (see *getgroups(2)*).

Each user belongs to a group specified in the password file **/etc/passwd** and possibly to other groups as specified in the files **/etc/group** and **/etc/logingroup**. A user is granted the permissions of those groups specified in **/etc/passwd** and **/etc/logingroup** at login time. The permissions of the groups specified in **/etc/group** are normally available only with the use of **newgrp** (see *newgrp(1)*). If a user name is specified with no options, **groups** prints the union of all these groups.

The **-p**, **-g**, and **-l** options limit the printed list to those groups specified in **/etc/passwd**, **/etc/group**, and **/etc/logingroup**, respectively. If a user name is not specified with any of these options, **cuserid()** is called to determine the default user name (see *cuserid(3S)*).

The printed list of groups is sorted in ascending collation order (see Environment Variables below).

EXTERNAL INFLUENCES**Environment Variables**

LC_COLLATE determines the order in which the output is sorted.

If **LC_COLLATE** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **groups** behaves as if all internationalization variables are set to "C" (see *environ(5)*).

EXAMPLES

Check file **/etc/logingroup** and display all groups to which user **tim** belongs:

```
groups -l tim
```

AUTHOR

groups was developed by the University of California, Berkeley.

FILES

```
/etc/group
/etc/logingroup
/etc/passwd
```

SEE ALSO

id(1), **newgrp(1)**, **getgroups(2)**, **initgroups(3C)**, **cuserid(3S)**, **group(4)**.

NAME

head - give first few lines

SYNOPSIS

head [-c | -l] [-n *count*] [*file* ...]

Obsolescent:

head [-*count*] [*file* ...]

DESCRIPTION

head prints on standard output the first *count* lines of each of the specified files, or of the standard input. If *count* is omitted it defaults to 10.

If multiple *files* are specified, **head** outputs before each file a line of this form:

==> *file* <==

Options

- c The quantity of output is measured in bytes.
- count The number of units of output. This option is provided for backward compatibility (see -n below) and is mutually exclusive of all other options.
- l The quantity of output is measured in lines; this is the default.
- n *count* The number of lines (default) or bytes output. *count* is an unsigned decimal integer. If -n (or -count) is not given, the default quantity is 10. This option provides the same functionality as the -count option, but in a more standard way. Use of the -n option is recommended where portability between systems is important.

EXTERNAL INFLUENCES**Environment Variables**

LC_CTYPE determines the interpretation of text within file as single and/or multi-byte characters.

LC_MESSAGES determines the language in which messages are displayed.

If **LC_CTYPE** or **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**.

If any internationalization variable contains an invalid setting, **head** behaves as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single- and multi-byte character code sets are supported.

WARNINGS

The length of the input lines is limited to {**LINE_MAX**} bytes.

SEE ALSO

tail(1), cat(1), more(1), pg(1).

STANDARDS CONFORMANCE

head: SVID3, XPG4, POSIX.2

NAME

hostname - set or display name of current host system

SYNOPSIS

hostname [*name_of_host*]

DESCRIPTION

The **hostname** command displays the name of the current host, as given in the **gethostname()** system call (see *gethostname(2)*). Users who have appropriate privileges can set the **hostname** by giving the argument *name_of_host*; this is usually done in the startup script **/sbin/init.d/hostname**. The *name_of_host* argument is restricted to **MAXHOSTNAMELEN** characters as defined in **<sys/param.h>**.

The system might be known by other names if networking products are supported. See the node manager documentation supplied with your system.

WARNINGS

If the *name_of_host* argument is specified, the resulting host name change lasts only until the system is rebooted. To change the host name permanently, run the special initialization script **/sbin/set_parms** (see *Using Your HP Workstation*).

Many types of networking services are supported on HP-UX, each of which uses a separately assigned system name and naming convention. To ensure predictable system behavior, it is essential that system names (also called host names or node names) be assigned in such a manner that they do not create conflicts when the various networking facilities interact with each other.

The system does not rely on a single system name in a specific location, partly because different services use dissimilar name formats as explained below. The **hostname** and **uname** commands assign system names as follows:

| Node Name | Command | name Format | Used By |
|---------------|-----------------------------|----------------------|----------------------------------|
| Internet name | hostname <i>name</i> | <i>sys[.x.y.z..]</i> | ARPA and NFS Services |
| UUCP name | uname -S <i>name</i> | <i>sys</i> | uucp and related programs |

where *sys* represents the assigned system name. It is *strongly* recommended that *sys* be identical for all commands and locations and that the optional *.x.y.z..* follow the specified notation for the particular ARPA/NFS environment.

Internet names are also frequently called host names or domain names (which are different from NFS domain names). Refer to *hostname(5)* for more information about Internet naming conventions.

Whenever the system name is changed in any file or by the use of any of the above commands, it should also be changed in all other locations as well. Other files or commands in addition to those above (such as **/etc/uucp/Permissions** if used to circumvent **uname**, for example) may contain or alter system names. To ensure correct operation, they should also use the same system name.

System names are normally assigned by the **/sbin/init.d/hostname** script at start-up, and should not be altered elsewhere.

AUTHOR

hostname was developed by the University of California, Berkeley.

SEE ALSO

uname(1), **gethostname(2)**, **sethostname(2)**, **uname(2)**, **hostname(5)**.

Using Your HP Workstation

NAME

hp - handle special functions of HP 2640 and HP 2621-series terminals

SYNOPSIS

hp [-e] [-m]

DESCRIPTION

hp supports special functions of the Hewlett-Packard HP 2640 and HP 2621 series of terminals, with the primary purpose of producing accurate representations of most **nroff** output. A typical use is:

```
nroff -h files ... | hp
```

Regardless of the hardware options on a given terminal, **hp** tries to do sensible things with underlining and reverse line-feeds. If the terminal has the “display enhancements” feature, subscripts and superscripts can be indicated in distinct ways. If it has the “mathematical-symbol” feature, Greek and other special characters can be displayed.

Options

hp recognizes the following options:

- e Specify that your terminal has the “display enhancements” feature, to make maximal use of the added display modes. Overstruck characters are presented in the Underline mode. Superscripts are shown in Half-bright mode, and subscripts in Half-bright, Underlined mode. If this flag is omitted, **hp** assumes that your terminal lacks the “display enhancements” feature. In this case, all overstruck characters, subscripts, and superscripts are displayed in Inverse Video mode; that is, dark-on-light, rather than light-on-dark.
- m Request minimization of output by removing new-lines. Any contiguous sequence of 3 or more new-lines is converted into a sequence of only 2 new-lines; that is, any number of successive blank lines produces only a single blank output line. This allows you to retain more actual text on the screen.

DIAGNOSTICS

line too long

The representation of a line exceeds 1,024 characters.

RETURN VALUE

hp returns zero for normal termination, and 2 for all errors.

WARNINGS

An “overstriking sequence” is defined as a printing character followed by a backspace followed by another printing character. In such sequences, if either printing character is an underscore, the other printing character is shown underlined or in Inverse Video; otherwise, only the first printing character is shown (again, underlined or in Inverse Video). Nothing special is done if a backspace is adjacent to an ASCII control character. Sequences of control characters (e.g., reverse line-feeds, backspaces) can make text “disappear”; in particular, tables generated by **tbl** that contain vertical lines will often be missing the lines of text that contain the “foot” of a vertical line, unless the input to **hp** is piped through **col** (see **col(1)**).

Although some terminals do support numerical superscript characters, no attempt is made to display them.

SEE ALSO

col(1), neqn(1), nroff(1), tbl(1).

NAME

hyphen - find hyphenated words

SYNOPSIS

hyphen [*files*]

DESCRIPTION

hyphen finds all the hyphenated words ending lines in *files* and prints them on the standard output. If no arguments are given, the standard input is used; thus, **hyphen** can be used as a filter.

EXAMPLES

Prepare an **nroff** hyphenation proofreading file for *textfile*.

```
mm textfile | hyphen
```

WARNINGS

hyphen cannot cope with hyphenated *italics* (i.e., underlined) words; it often misses them completely or mangles them.

hyphen occasionally gets confused, but with no ill effects other than spurious extra output.

SEE ALSO

mm(1), nroff(1).


h

NAME

iconv - code set conversion

SYNOPSIS

```
iconv -f fromcode -t tocode [file ...]
```

DESCRIPTION

iconv converts the encoding of characters in the input files from the *fromcode* code set to the *tocode* code set, and writes the results on standard output. If no input files are given, **iconv** reads from standard input. If **-** appears as an input file name, **iconv** reads standard input at that point. **--** can be used to delimit the end of options (see *getopt(3C)*).

Options

iconv recognizes the following options:

- f *fromcode*** Identify the code set corresponding to option argument *fromcode* as the code set that the input will be converted "from".
- t *tocode*** Identify the code set corresponding to option argument *tocode* as the code set that the input will be converted "to".

The *fromcode* and *tocode* names can be any of the base and alias names listed in the iconv configuration file, `/usr/lib/nls/iconv/config.iconv`. See *iconv(3C)* for details and the configuration file for a list of supported code set names.

EXTERNAL INFLUENCES**Environment Variables**

LANG provides a default value for the internationalization variables that are unset or null. If **LANG** is unset or null, the default value of "C" (see *lang(5)*) is used. If any of the internationalization variables contains an invalid setting, **iconv** will behave as if all internationalization variables are set to "C". See *environ(5)*.

LC_ALL If set to a non-empty string value, overrides the values of all the other internationalization variables.

LC_CTYPE determines the interpretation of text as single and/or multi-byte characters, the classification of characters as printable, and the characters matched by character class expressions in regular expressions. During translation of the file, this variable is superseded by the use of the *fromcode* option-argument.

LC_MESSAGES determines the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.

NLSPATH determines the location of message catalogues for the processing of **LC_MESSAGES**.

International Code Set Support

Single and multi-byte character code sets are supported.

WARNINGS

If an input character does not have a valid equivalent in the code set selected by the **-t** option (the "to" code set), it is mapped to the "galley character", if it has been defined for that conversion. (see *genxlt(1)* and *iconv(3C)*).

If an input character does not belong to the code set selected by the **-f** option (the "from" code set), the command terminates.

EXAMPLES

Convert the contents of file **foo** from code set Roman8 to ISO 8859/1 and store the results in file **bar**.

```
iconv -f roman8 -t iso8859_1 foo > bar
```

FILES

`/usr/lib/nls/iconv/config.iconv` iconv configuration file

AUTHOR

iconv was developed by HP.

SEE ALSO

getopt(3C), iconv(3C).

STANDARDS CONFORMANCE

iconv: XPG2, XPG3, XPG4

NAME

id - print user and group IDs and names

SYNOPSIS

```
id [-u|-g|-G] [-nr] [-P] [user]
```

DESCRIPTION

The **id** command writes a message to standard output, giving the user and group IDs and names for the process. If the effective and real IDs are different, both are printed.

If the process has supplementary group affiliations (see *groups(1)*), the supplementary group affiliations are also written.

If the *user* operand is specified, and the effective user ID of the process is superuser, the user and group IDs of the selected user are written. In this case, effective IDs are assumed to be identical to real IDs.

Options

The following options modify the behavior described above.

- g Display only the group ID. The default is the effective group ID; to modify, use the **-r** option. If the process has supplementary group affiliations that are different from the effective group ID (or the real ID if the **-r** option is used), display each such affiliation on the same line. The default is decimal format; to modify, use the **-n** option.
- G Output all different group IDs (effective, real, and supplementary) only, using the format "%u\n". If there is more than one distinct group affiliation, output each such affiliation, using the format " %u", before the <newline> is output.
- n With A **-u**, **-g**, or **-G**, display the ID name instead of the ID number.
- r With **-u**, **-g**, or **-G**, display the real ID instead of the effective ID.
- u Display only the user ID. The default is the effective user ID; to modify, use the **-r** option. The default is decimal format; to modify, use the **-n** option.
- P Display the process resource group ID. See HP Process Resource Manager in DEPENDENCIES.

EXAMPLES

To display the current user and group data:

```
id
```

produces:

```
uid=1834(allanp) gid=20(users)
```

To display the group ID number for the current process:

```
id -g
```

produces:

```
20
```

To display the group name for the current process:

```
id -gn
```

produces:

```
users
```

To display the user and group data for another user:

```
id ralford
```

produces:

```
uid=329(ralford) gid=20(users)
```

if the effective user ID of the process is superuser. Otherwise, it produces the data for the invoking process.

AUTHOR

id was developed by HP and AT&T.

DEPENDENCIES**HP Process Resource Manager**

The **-P** option requires that the optional HP Process Resource Manager (PRM) software be installed and configured. See *prmconfig(1)* for a description of how to configure HP PRM, and *prmconf(4)* for the definition of the process resource group.

SEE ALSO

groups(1), *logname(1)*, *getuid(2)*.

HP Process Resource Manager: *prmconfig(1)*, *prmconf(4)* in *HP Process Resource Manager User's Guide*.

STANDARDS CONFORMANCE

id: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

ident - identify files in RCS

SYNOPSIS

ident *file* ...

DESCRIPTION

ident searches the named files for all occurrences of the pattern *\$keyword:...\$*, where *keyword* is one of the following:

| | |
|---------------|-----------------|
| Author | Log |
| Date | Revision |
| Header | Source |
| Locker | State |

These patterns are normally inserted automatically by the RCS **co** command, but can also be inserted manually (see *co(1)*).

ident works on text files as well as object files. For example, if the C program in file **f.c** contains:

```
char rcsid[] = "$Header: Header information $";
```

and **f.c** is compiled into **f.o**, the command:

```
ident f.c f.o
```

prints:

```
f.c:
    $Header: Header information $
f.o:
    $Header: Header information $
```

AUTHOR

ident was developed by Walter F. Tichy.

SEE ALSO

ci(1), *co(1)*, *rccs(1)*, *rccsdiff(1)*, *rccsintro(5)*, *rccsmerge(1)*, *rlog(1)*, *rccsfile(4)*.

NAME

idlookup - identify the user of a particular TCP connection

SYNOPSIS

idlookup *host-or-ip-number local-port foreign-port*

DESCRIPTION

idlookup can be used to identify the user at the remote end of a TCP connection, assuming the host at the other end is running an Identification Server.

host-or-ip-number is the name of the host at the other end of the connection, or its IP address.

local-port and *foreign-port* are the port numbers, or service names of the ports at the two ends of the connection.

WARNING

Note that the references to *local-port* and *foreign-port* follow the terminology in RFC931, and are from the point of view of the *server* rather than the user.

AUTHOR

idlookup was originally written by Peter Eriksson. The manual page was originally written by Dave Sheild.

SEE ALSO

sendmail(1M), identd(1M), *RFC931*.

i

NAME

ied - input editor and command history for interactive programs

SYNOPSIS

ied [**-dirt**] [**-h** *file*] [**-s** *size*] [**-p** *prompt*] [**-k** *charmap*] *utility* [*arguments ...*]

DESCRIPTION

ied is a utility command that is intended to act as an interface between the user and an interactive program such as **bc**, **bs**, or Bourne shell, providing most of the line editing and history functionality found in the Korn shell. **ied** interprets the *utility* name as the command to be executed, and passes *arguments* as the arguments to the utility. Subsequent input to *utility* then has access to editing and history functions very similar to those provided by **ksh**.

ied monitors the state of the **pty** it uses to run the command, and, whenever the application it is running, changes the state from the state of the **tty** when **ied** started, **ied** becomes "transparent". This allows programs to do shell escapes to screen-smart programs. In general, **ied** should not in any way interfere with any action taken by any program for which it provides a front end. This includes Korn shell itself: in this case **ied** would provide history for any application that was run by **ksh**, and **ksh** would provide its own independent history. In a useful extreme case, **ied** can be used as a front end to the login shell (which might be **ksh** or **csh**). In this case, all applications that use normal line editing gain line editing and history, sharing a single history. The shell would continue to have its own independent history if it provides such a mechanism.

When **ied** is in its transparent mode, no history is saved. In particular the **ex** mode of **vi** does not use normal line editing (rather, it simulates it) and **ied** cannot provide history in this case. The **Subject:** and address line editing of mailx also cannot be edited with **ied**.

Options

Several options and command-line arguments control **ied**'s operation:

- d** Debug mode. Print information about the operation of the program. It is best used to determine if a program puts **ied** into transparent mode unexpectedly.
- h** *filename* Keep the history in a file named *filename*. If a file of that name already exists and is a history file, the latter part of it (the last *size* lines as specified by the **-s** option) is used as the initial value of the history. If the **-h** option is not used, the environment variable **IEDHISTFILE** is used to supply the name. If neither are present an unnamed temporary file is used, and no initial value is provided.
- i** Force interactive mode. Normally **ied** simply **execs** the command to which it is asked to be a front end when the standard input is not a **tty** (this allows aliases to be used for commands used in shells without interfering with their operation). This option forces **ied** to remain as a front end, and all editing functions are in place. This permits a utility that behaves differently in interactive and batch modes to be driven from a pipe or file in interactive mode. This is particularly useful in testing commands that make this distinction.
- k** *charmap* *charmap* is a file of 256 or fewer lines. The line number in the file is the ordinal of a character as seen as input by **ied**, and the character on the line is the character generated as output (and also used as editing characters). This allows remapping of (ordinary) keys such as for a Dvorak keyboard. Characters must start in column one of each line, and be represented as 1-4 characters followed by a space or the new-line character for the next line. Characters after the space are ignored as comments. Single-character entries represent themselves. Two-character entries where the first character is a circumflex (^) converts the second character to the corresponding control character. Two-character sequences where the first character is backslash (\) use the C language conventions:

| | | | |
|-----------|---------|-----------|-----------|
| \n | newline | \s | space |
| \ | escape | \0 | null |
| \r | return | \f | form feed |

\t tab \v vertical tab
 \b backspace

Three- and four-character sequences must be \nn or \nnn, giving the octal value for the character. If *charmap* is less than 256 lines long, the remaining characters are mapped to themselves.

- p *prompt* Many commands do not prompt when ready for input. **ied** approximates a prompting mechanism for such commands. This is not always perfectly successful, but for many commands it helps. In the worst case, the prompt is interspersed with output in the wrong location. *prompt* is a string as used in the format argument to *printf(3S)*. The only % conversions that can be included are up to one instance of %d which is converted to the sequential number of the command, and any number of occurrences of %% which is treated as a literal % character. Prompting is suppressed when **ied** is operating in transparent mode.
- r This sets "non-raw" mode. Normally **ied** uses its own editing capabilities when reading simple text. This causes **ied** to use tty line discipline most of the time. The disadvantage of the default mode is that more context switches and general processing are required. The advantage is that **ied** is more transparent. For example, to specifically send an end-of-file in the non-raw mode requires that the end-of-file character (usually Ctrl-D) be followed by a carriage return. Similarly the "literal next" function (Ctrl-V) cannot escape the line-erase and line-kill functions in non-raw mode.
- s *size* This option specifies the size of the history buffer. When **ied** is started with an existing history file, approximately the last *size* lines are available to the history mechanism (the number is not guaranteed to be exactly *size*). Other lines in the file are retained until such time as **ied** is started on that history file and it exceeds approximately 4K bytes in size, at which time **ied** discards older entries at the beginning of the file until it is near 4 Kbytes in size. Since this occurs only at startup, history files can grow to be quite large between restarts. Larger values of *size* make the process image larger.
 If -s is not specified, the value of the environment variable **IEDHISTSIZE** is used. If neither is specified, a default is used.
- t Set transparent mode. This forces **ied** to permanently be in transparent mode (as discussed above). It is primarily useful with -i for some classes of automated processing. In particular, it is useful for driving a command if the command takes as input what **ied** would interpret as editing characters. Thus with the appropriate combinations of -i and -t, it is possible to drive an editor such as **vi** or a screen-smart application from a batch file.

Should something go wrong with **ied**, the **SIGQUIT** signal, repeated 3 times, usually aborts **ied**. The exception is the case of a fully transparent application, where **ied** must be killed from another window or terminal. This is really relevant only when there is no way to direct the serviced process to terminate itself.

The editing capabilities of **ied** are essentially those found in **ksh**. Only those that differ from **ksh** are described below. As in **ksh**, the style of editing is determined from the environment variable **VISUAL**, or from **EDITOR** if **VISUAL** is not specified. The value examined should end in **vi**, **emacs**, or **gmacs** to specify an editor type. If it does not, **ied** does no editing, and history is not accessible.

In vi mode:

- J Join lines. Considering the most recently edited line (which is empty immediately after a line is sent to the application) to be the "last line" of the history, the current line being displayed from the history is appended to the end of the last line, and the position in the history is reset to be at the last line which is then displayed. A space is inserted between the old and new text on the last line. The cursor is left on that space. Because **ied**'s understanding of line continuation is minimal, this is useful for editing long statements.
- v Not supported.
- V Not supported.
- # Sends nothing to the application, but inserts the line in the history (useful for adding comments to history file).

`<esc>,*,=` (Filename expansion). Not supported.

`@` Macro expansion. Not supported.

Note however that **ksh** has a rarely-used function `_` that substitutes words from the previous line (this is not the macro `$_`, but rather an editor command). If a preceding *count* is given, it uses the *count*th word of the last line. This is much more useful with **ied**.

In emacs/gmacs mode:

`M-*, M-=, M-<esc>` (filename expansion) Not supported.

Note that the command `M-.` (and its synonym `M-_`) provide the same functionality as the vi mode `_` command.

Macro expansion. Not supported.

`^O` Although supported, it may not always appear correctly on the screen. The `^L` command can be used to redraw the line. See below for the discussion on prompting.

EXAMPLES

Add interactive editing to the **bc** command:

```
ied bc
```

Execute **vi** on **testfile** using commands taken from **script**:

```
cat script | ied -i -t vi testfile
```

Note that without the use of **ied**, **vi** would misbehave because its standard input would not be a terminal device. In this case the `-t` is not required because **vi** puts itself in raw mode, but for an application that does not, `-t` might be required.

The command line

```
ied -i -t grep '^x:' data_file | tee x_lines
```

searches the file **data_file** for lines beginning with **x:**, sending one copy to the terminal and a second to file **x_lines**, just like the command line

```
grep '^x:' data_file | tee x_lines
```

The difference is that in the command line without **ied**, **grep** writes directly to a pipe, and thus buffers its output. If **data_file** is very large and not many lines match the pattern, output to the terminal is delayed. By using **ied**, the output of **grep** goes to a pty instead, which causes **grep** to output each line as it is ready.

WARNINGS

Since **ied** cannot know everything about every application, it is possible that it can become confused, with either the timing or the prompt being out of phase with the application. Since the use of **ied** is never required, it is the user's choice to determine whether the application is more usable with or without **ied**. In general, however, programs that do not confuse **ied** are usually also the most likely to benefit from its use.

ied tries to intuit the currently active prompt when it is not providing one itself. However, this is not always successful. Even when it is successful, the timing of **ied** and the serviced command may occasionally confuse the output. The `^L` commands in both emacs and vi modes redraw the edit line in a consistent fashion that can be used to create the next command.

AUTHOR

ied was developed by HP.

SEE ALSO

ksh(1).

NAME

insertmsg - use findstr(1) output to insert calls to catgets(3C)

SYNOPSIS

```
insertmsg [-h] [-nnumber] [-iamount] [-snumber] stringlist
```

DESCRIPTION

insertmsg examines the file *stringlist*, which is assumed to be the output of **findstr** after subsequent editing to remove any strings that do not need to be localized (see *findstr(1)*). If the **-h** option is specified, **insertmsg** places the following lines at the beginning of each file named in *stringlist*:

```
#ifndef NLS
#define catgets(i,sn,mn,s) (s)
#else NLS
#define NL_SETN number
#include <nl_types.h>
#endif NLS
```

where *number* is a set number defined by the **-s** option; the default is 1. For each string in *stringlist*, **insertmsg** surrounds the string in the corresponding file with an expression of the form:

```
(catgets(catd,NL_SETN,msg_num,"default string" ) )
```

The *default string* is the original string referenced by the line in *stringlist*, and *msg_num* is replaced by the message number assigned to that string. The assigned message numbers begin with the number defined by the **-n** option and are incremented by the amount defined by the **-i** option. The default is 1 for both the starting message number and the increment. If *name.c* is the file to be modified, as specified within the *stringlist* file, **insertmsg** places the modified source in *nl_name.c*. The user must then manually edit the file *nl_name.c* to insert the following statements:

```
nl_catd catd;
catd = catopen("appropriate message catalog",0);
```

The data type *nl_catd* is defined in *<nl_types.h>* and *catd* is a parameter to the calls to *catgets*, which are inserted for each string from *stringlist*.

insertmsg also sends to the standard output a file that can be used as input to **genocat** (see *genocat(1)*).

EXTERNAL INFLUENCES**Environment Variables**

LC_CTYPE determines the interpretation of text as single- and/or multi-byte characters.

LC_MESSAGES determines the language in which messages are displayed.

If **LC_CTYPE** or **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **insertmsg** behaves as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single- and multi-byte character code sets are supported.

DIAGNOSTICS

If **insertmsg** does not find opening or closing double quotes where required in the strings file, it prints **insertmsg exiting : lost in strings file** and aborts. If this happens, check the strings file to ensure that the lines that have been kept there have not been altered.

WARNINGS

If the **-h** option is not used, it may be necessary to manually add the following statement to the file created by **insertmsg**:

```
#include <nl_types.h>
```

insertmsg inserts a pointer to a static area that is overwritten on each call.

The **insertmsg** command is HP proprietary, not portable to other vendors' systems, and will not be provided in future HP-UX releases.

AUTHOR

insertmsg was developed by HP.

SEE ALSO

findstr(1), **gencat(1)**, **catgets(3C)**, **catopen(3C)**.

NAME

iostat - report I/O statistics

SYNOPSIS

iostat [-t] [*interval* [*count*]]

DESCRIPTION

iostat iteratively reports I/O statistics for each active disk on the system. Disk data is arranged in a four-column format:

| Column Heading | Interpretation |
|----------------|----------------------------------|
| device | Device name |
| bps | Kilobytes transferred per second |
| sps | Number of seeks per second |
| mmps | Milliseconds per average seek |

If two or more disks are present, data is presented on successive lines for each disk.

To compute this information, seeks, data transfer completions, and the number of words transferred are counted for each disk. Also, the state of each disk is examined **HZ** times per second (as defined in `<sys/param.h>`) and a tally is made if the disk is active. These numbers can be combined with the transfer rates of each device to determine average seek times for each device.

With the advent of new disk technologies, such as data striping, where a single data transfer is spread across several disks, the number of milliseconds per average seek becomes impossible to compute accurately. At best it is only an approximation, varying greatly, based on several dynamic system conditions. For this reason and to maintain backward compatibility, the milliseconds per average seek (**mmps**) field is set to the value 1.0.

Options

iostat recognizes the following options and command-line arguments:

| | |
|-----------------|--|
| -t | Report terminal statistics as well as disk statistics. Terminal statistics include: |
| tin | Number of characters read from terminals. |
| tout | Number of characters written to terminals. |
| us | Percentage of time system (active processors) has spent in user mode. |
| ni | Percentage of time system (active processors) has spent in user mode running low-priority (<i>nice</i>) processes. |
| sy | Percentage of time system (active processors) has spent in system mode. |
| id | Percentage of time system (active processors) has spent idling. |
| <i>interval</i> | Display successive lines which are summaries of the last <i>interval</i> seconds. The first line reported is for the time since a reboot and each subsequent line is for the last interval only. |
| <i>count</i> | Repeat the statistics <i>count</i> times. |

EXAMPLES

Show current I/O statistics for all disks:

```
iostat
```

Display I/O statistics for all disks every 10 seconds until INTERRUPT or QUIT is pressed:

```
iostat 10
```

Display I/O statistics for all disks every 10 seconds and terminate after 5 successive readings:

```
iostat 10 5
```

Display I/O statistics for all disks every 10 seconds, also show terminal and processor statistics, and terminate after 5 successive readings:

```
iostat -t 10 5
```

WARNINGS

Users of **iostat** must not rely on the exact field widths and spacing of its output, as these will vary depending on the system, the release of HP-UX, and the data to be displayed.

AUTHOR

iostat was developed by the University of California, Berkeley, and HP.

FILES

/usr/include/sys/param.h

SEE ALSO

vmstat(1).

NAME

ipcrm - remove a message queue, semaphore set, or shared memory identifier

SYNOPSIS

ipcrm [*option*]...

DESCRIPTION

The **ipcrm** command removes one or more specified message queue, semaphore set, or shared memory identifiers.

Options

The identifiers are specified by the following *options*:

- m** *shmid* Remove the shared memory identifier *shmid* from the system. The shared memory segment and data structure associated with it are destroyed after the last detach.
- q** *msqid* Remove the message queue identifier *msqid* from the system and destroy the message queue and data structure associated with it.
- s** *semid* Remove the semaphore identifier *semid* from the system and destroy the set of semaphores and data structure associated with it.
- M** *shmkey* Remove the shared memory identifier, created with key *shmkey*, from the system. The shared memory segment and data structure associated with it are destroyed after the last detach.
- Q** *msgkey* Remove the message queue identifier, created with key *msgkey*, from the system and destroy the message queue and data structure associated with it.
- S** *semkey* Remove the semaphore identifier, created with key *semkey*, from the system and destroy the set of semaphores and data structure associated with it.

The details of the removals are described in *msgctl(2)*, *shmctl(2)*, and *semctl(2)*. The identifiers and keys can be found by using **ipcs** (see *ipcs(1)*).

SEE ALSO

ipcs(1), msgctl(2), msgget(2), msgop(2), semctl(2), semget(2), semop(2), shmctl(2), shmget(2), shmop(2).

STANDARDS CONFORMANCE

ipcrm: SVID2, SVID3

NAME

ipcs - report status of interprocess communication facilities

SYNOPSIS

ipcs [-**mq**s] [-**abco**pt] [-**C** *core*] [-**N** *namelist*]

DESCRIPTION

ipcs displays certain information about active interprocess communication facilities. With no options, **ipcs** displays information in short format for the message queues, shared memory segments, and semaphores that are currently active in the system.

Options

The following options restrict the display to the corresponding facilities.

- (none) This is equivalent to **-mq**s.
- m** Display information about active shared memory segments.
- q** Display information about active message queues.
- s** Display information about active semaphores.

The following options add columns of data to the display. See "Column Description" below.

- (none) Display default columns: for all facilities: **T**, **ID**, **KEY**, **MODE**, **OWNER**, **GROUP**.
- a** Display all columns, as appropriate. This is equivalent to **-bco**pt.
- b** Display largest-allowable-size information: for message queues: **QBYTES**; for shared memory segments: **SEGSZ**; for semaphores: **NSEMS**.
- c** Display creator's login name and group name: for all facilities: **CREATOR**, **CGROUP**.
- o** Display information on outstanding usage: for message queues: **CBYTES**, **QNUM**; for shared memory segments: **NATTCH**.
- p** Display process number information: for message queues: **LSPID**, **LRPID**; for shared memory segments: **CPID**, **LPID**.
- t** Display time information: for all facilities: **CTIME**; for message queues: **STIME**, **RTIME**; for shared memory segments: **ATIME**, **DTIME**; for semaphores: **OTIME**.

The following options redefine the sources of information.

- C** *core* Use *core* in place of **/dev/kmem**. *core* can be a core file or a directory created by **savecrash** or **savecore**.
- N** *namelist* Use file *namelist* or the *namelist* within *core* in place of **/stand/vmunix**. It opens a crash dump for reading. Please refer to **cr_open(3)** for more details.

Column Descriptions

The column headings and the meaning of the columns in an **ipcs** listing are given below. The columns are printed from left to right in the order shown below.

- T** Facility type:
 - m** Shared memory segment
 - q** Message queue
 - s** Semaphore
- ID** The identifier for the facility entry.
- KEY** 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** The facility access modes and flags: The mode consists of 11 characters that are interpreted as follows:

The first two characters can be:

- R** A process is waiting on a **msgrcv()**.
- S** A process is waiting on a **msgsnd()**.
- 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.
- The corresponding special flag is not set.

The next 9 characters are interpreted as three sets of three characters each. The first set refers to the owner's permissions, the next to permissions of others in the 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.

- r** Read permission is granted.
- w** Write permission is granted.
- a** Alter permission is granted.
- The indicated permission is not granted.

| | |
|----------------|--|
| OWNER | The login name of the owner of the facility entry. |
| GROUP | The group name of the group of the owner of the facility entry. |
| CREATOR | The login name of the creator of the facility entry. |
| CGROUP | The group name of the group of the creator of the facility entry. |
| CBYTES | The number of bytes in messages currently outstanding on the associated message queue. |
| QNUM | The number of messages currently outstanding on the associated message queue. |
| QBYTES | The maximum number of bytes allowed in messages outstanding on the associated message queue. |
| LSPID | The process ID of the last process to send a message to the associated message queue. |
| LRPID | The process ID of the last process to receive a message from the associated message queue. |
| STIME | The time the last msgsnd() message was sent to the associated message queue. |
| RTIME | The time the last msgrcv() message was received from the associated message queue. |
| CTIME | The time when the associated facility entry was created or changed. |
| NATTCH | The number of processes attached to the associated shared memory segment. |
| SEGSZ | The size of the associated shared memory segment. |
| CPID | The process ID of the creating process of the shared memory segment. |
| LPID | The process ID of the last process to attach or detach the shared memory segment. |
| ATIME | The time the last shmat() attach was completed to the associated shared memory segment. |
| DTIME | The time the last shmdt() detach was completed on the associated shared memory segment. |
| NSEMS | The number of semaphores in the set associated with the semaphore entry. |
| OTIME | The time the last semop() semaphore operation was completed on the set associated with the semaphore entry. |

WARNINGS

ipcs produces only an approximate indication of actual system status because system processes are continually changing while **ipcs** is acquiring the requested information.

Do not rely on the exact field widths and spacing of the output, as these will vary depending on the system, the release of HP-UX, and the data to be displayed.

FILES

| | |
|----------------------------|-----------------------|
| <code>/dev/kmem</code> | Kernel virtual memory |
| <code>/etc/group</code> | Group names |
| <code>/etc/passwd</code> | User names |
| <code>/stand/vmunix</code> | System namelist |

SEE ALSO

`msgop(2)`, `semop(2)`, `shmop(2)`.

STANDARDS CONFORMANCE

`ipcs`: SVID2, SVID3

NAME

join - relational database operator

SYNOPSIS

join [*options*] *file1 file2*

DESCRIPTION

join forms, on the standard output, a join of the two relations specified by the lines of *file1* and *file2*. If *file1* or *file2* is -, the standard input is used.

file1 and *file2* must be sorted in increasing collating sequence (see Environment Variables below) on the fields on which they are to be joined; normally the first in each line.

The output contains one line for each pair of lines in *file1* and *file2* that have identical join fields. The output line normally consists of the common field followed by the rest of the line from *file1*, then the rest of the line from *file2*.

The default input field separators are space, 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 space.

Some of the below options use the argument *n*. This argument should be a 1 or a 2 referring to either *file1* or *file2*, respectively.

Options

- a *n* In addition to the normal output, produce a line for each unpairable line in file *n*, where *n* is 1 or 2.
- e *s* Replace empty output fields by string *s*.
- j *m* Join on field *m* of both files. The argument *m* must be delimited by space characters. This option and the following two are provided for backward compatibility. Use of the -1 and -2 options (see below) is recommended for portability.
- j1 *m* Join on field *m* of *file1*.
- j2 *m* Join on field *m* of *file2*.
- o *list* Each output line comprises the fields specified in *list*, each element of which has the form *n.m*, where *n* is a file number and *m* is a field number. The common field is not printed unless specifically requested.
- t *c* Use character *c* as a separator (tab character). Every appearance of *c* in a line is significant. The character *c* is used as the field separator for both input and output.
- v *file_number* Instead of the default output, produce a line only for each unpairable line in *file_number*, where *file_number* is 1 or 2.
- 1 *f* Join on field *f* of file 1. Fields are numbered starting with 1.
- 2 *f* Join on field *f* of file 2. Fields are numbered starting with 1.

EXTERNAL INFLUENCES**Environment Variables**

LC_COLLATE determines the collating sequence **join** expects from input files.

LC_CTYPE determines the alternative blank character as an input field separator, and the interpretation of data within files as single and/or multi-byte characters. **LC_CTYPE** also determines whether the separator defined through the -t option is a single- or multi-byte character.

If **LC_COLLATE** or **LC_CTYPE** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **join** behaves as if all internationalization variables are set to "C" (see *environ(5)*).

International Code Set Support

Single- and multi-byte character code sets are supported with the exception that multi-byte-character file names are not supported.

EXAMPLES

The following command line joins 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 the collating sequence defined by the `LC_COLLATE` or `LANG` environment variable on the group ID fields.

```
join -1 4 -2 3 -o 1.1 2.1 1.6 -t: /etc/passwd /etc/group
```

The following command produces an output consisting all possible combinations of lines that have identical first fields in the two sorted files *sf1* and *sf2*, with each line consisting of the first and third fields from `sorted_file1` and the second and fourth fields from `sorted_file2`:

```
join -j1 1 -j2 1 -o 1.1,2.2,1.3,2.4 sorted_file1 sorted_file2
```

WARNINGS

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 `join`, `sort`, `comm`, `uniq`, and `awk` are incongruous.

Numeric filenames may cause conflict when the `-o` option is used immediately before listing filenames.

AUTHOR

`join` was developed by OSF and HP.

SEE ALSO

`awk(1)`, `comm(1)`, `sort(1)`, `uniq(1)`.

STANDARDS CONFORMANCE

`join`: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

kdestroy - destroy Kerberos tickets

SYNOPSIS

kdestroy [-q] [-c *cache_filename*]

DESCRIPTION

The **kdestroy** utility destroys the user's active Kerberos authorization tickets by writing zeros to the specified credentials cache that contains them. If the credentials cache is not specified, the default credentials cache is destroyed.

Options

- q Run quietly. Normally **kdestroy** beeps if it fails to destroy the user's tickets. The -q flag suppresses this behavior.
- c *cache_filename* Use *cache_filename* as the credentials ticket cache name and location. If this option is not used, the default cache name and location is used.

The default credentials cache may vary between systems. If the **KRB5CCNAME** environment variable is set, its value is used to name the default ticket cache.

Most installations recommend that you place the **kdestroy** command in your **.logout** file, so that your tickets are destroyed automatically when you log out.

Note

For DCE operations use **/opt/dce/bin/kdestroy**.

Environment

kdestroy uses the following environment variable:

KRB5CCNAME Location of the credentials ticket cache.

WARNINGS

Only the tickets in the specified credentials cache are destroyed. Separate ticket caches are used to hold root instance and password changing tickets. Even these should be destroyed. It is recommended that to keep all user tickets in a single credentials cache.

FILES

/tmp/krb5cc_{uid} Default credentials cache. {uid} is the decimal UID of the user.

AUTHOR

kdestroy was developed by the Massachusetts Institute of Technology.

SEE ALSO

kerberos(9), kinit(1), klist(1).

(HP-UX C-Kermit)

NAME

kermit - C-Kermit 7.0 communications software for serial and network connections: modem dialing, file transfer and management, terminal connection, character-set translation, numeric and alpha paging, and script programming

SYNOPSIS

kermit [*command-file*] [*options ...*]

DESCRIPTION

Kermit is a family of file transfer, management, and communication software programs from the Kermit Project at Columbia University available for most computers and operating systems. The version of Kermit for Hewlett-Packard HP-UX, called **C-Kermit**, supports both serial connections (direct or dialed) and TCP/IP connections.

C-Kermit can be thought of as a user-friendly and powerful alternative to **cu**, **tip**, **uucp**, **ftp**, **telnet**, **rlogin**, **expect**, and even your shell; a single package for both network and serial communications, offering automation, convenience, and language features not found in the other packages, and having a great deal in common with its cousins, C-Kermit on other UNIX platforms, Kermit 95 for Windows 95, Windows 98, Windows NT and 2000, and OS/2; MS-DOS Kermit for PCs with DOS and Windows 3.x, and IBM Mainframe Kermit-370 for VM/CMS, MVS/TSO, and CICS. C-Kermit itself also runs on Digital VMS, Data General AOS/VS, Stratus VOS, OS-9, QNX, Plan 9, the Commodore Amiga, and elsewhere. Together, C-Kermit, Kermit 95, MS-DOS Kermit, and IBM Mainframe Kermit offer a consistent and nearly universal approach to inter-computer communications.

C-Kermit 7.0 is Copyright (C) 1985, 2000 by the Trustees of Columbia University in the City of New York. For use and redistribution rights, see the C-Kermit COPYING.TXT file or give the C-Kermit COPYRIGHT command (summary: no license is required for own use; no license is required for distribution with Open Source operating systems; a license is required for certain other forms of redistribution).

C-Kermit 7.0 is included with HP-UX by Hewlett-Packard in partnership with the Kermit Project at Columbia University.

C-Kermit 6.0 is thoroughly documented in the book *Using C-Kermit* by Frank da Cruz and Christine M. Gianone, Digital Press, Second Edition, 1997; see REFERENCES at the end of this manual page. This manual page is not a substitute for the book. If you are a serious user of C-Kermit, particularly if you plan to write C-Kermit script programs, you should purchase the manual. Book sales are the primary source of funding for the nonprofit Kermit Project.

Any new features added since the most recent edition of the book was published are documented in the online file *ckermi2.upd* until such time as the Third Edition of the book is ready. Hints, tips, limitations, restrictions are listed in *ckcker.txt* (general C-Kermit) and *ckuker.bwr* (UNIX-specific); see FILES below. Please consult all of these references before reporting problems or asking for technical support.

Kermit software is available for hundreds of different computers and operating systems from Columbia University. For best file-transfer results, please use C-Kermit in conjunction with real Columbia University Kermit software on other computers, such as Kermit 95 for Windows 95 and NT or MS-DOS Kermit for DOS 3.x or Windows. See CONTACTS below.

MODES OF OPERATION

C-Kermit can be used in two "modes": remote and local. In **remote mode**, you connect to the HP-UX system from a desktop computer and transfer files between your desktop computer and HP-UX C-Kermit. In that case, connection establishment (dialing, TELNET connection, etc.) is handled by the Kermit program on your desktop computer.

In **local mode**, C-Kermit establishes a connection to another computer by direct serial connection, by dialing a modem, or by making a network connection. When used in local mode, C-Kermit gives you a terminal connection to the remote computer, using your actual terminal, emulator, or UNIX workstation terminal window or console driver for specific terminal emulation.

C-Kermit also has two types of commands: the familiar UNIX-style command-line options, and an interactive dialog with a prompt. **Command-line options** give you access to a small but useful subset of C-Kermit's features for terminal connection and file transfer, plus the ability to pipe files into or out of Kermit for transfer.

Interactive commands give you access to dialing, script programming, character-set translation, and, in general, detailed control and display, as well as automation, of all C-Kermit's features. Interactive commands can also be collected into command files or macros. C-Kermit's command and script language is

portable to many and diverse platforms.

STARTING C-KERMIT

You can start C-Kermit by typing `/usr/bin/kermit`, or just `kermit` if your PATH includes `/usr/bin`, possibly followed by command-line options. If there are no "action options" on the command line (explained below), C-Kermit starts in interactive command mode; you will see a greeting message and then the "C-Kermit>" prompt. If you do include action options on the command line, C-Kermit takes the indicated actions and then exits directly back to UNIX. Either way, C-Kermit executes the commands in its initialization file, `/usr/share/lib/kermit/ckermmit.ini`, before it executes any other commands, unless you have included the `'-Y'` (uppercase) command-line option, which means to skip the initialization file, or you have included the `'-y filename'` option to specify an alternative initialization file.

FILE TRANSFER

Here is the most common scenario for Kermit file transfer. Many other methods are possible, most of them more convenient, but this basic method should work in all cases.

- Start Kermit on your local computer and establish a connection to the remote computer. If C-Kermit is on your local computer, use the sequence `SET MODEM TYPE modem-name`, `SET LINE device-name`, `SET SPEED bits-per-second`, and `DIAL phone-number` if you are dialing; `SET LINE` and `SPEED` for direct connections; `SET NETWORK network-type` and `SET HOST host-name-or-address` for network connections.
- `SET` any other necessary communication parameters, such as `PARITY`, `DUPLEX`, and `FLOW-CONTROL`.
- Give the `CONNECT` command.
- Log in to the remote computer.
- Start Kermit on the remote computer, give it any desired `SET` commands for file-, communication-, or protocol-related parameters. If you will be transferring binary files, give the command `SET FILE TYPE BINARY` to the Kermit program that will be sending them.
- To **download** a file or file group, give the remote Kermit a `SEND` command, following by a filename or "wildcard" file specification, for example:

```
send oofa.txt           # (send one file)
send oofa.*             # (send a group of files)
```

To **upload** a file or files, give the remote Kermit a `RECEIVE` command. The sending Kermit will tell the receiving Kermit the name (and other attributes) of each file.

- Escape back to the Kermit program on your local (desktop) computer. If your local computer is running C-Kermit, type `Ctrl-\c` (Control-backslash followed by the letter 'c') (on NeXT workstations, use `Ctrl-]c`). If MS-DOS or Kermit 95, use `Alt-x` (hold down the Alt key, press 'x'). Now you should see your local Kermit program's prompt.
- If you will be transferring binary files, give the command `SET FILE TYPE BINARY` to the Kermit program that is sending the files.
- If you are *downloading* files, tell the local Kermit program to `RECEIVE`. If you are *uploading*, give your local Kermit program a `SEND` command, specifying a filename or wildcard file specification. In other words, tell the *remote* Kermit program what to do first, `SEND` or `RECEIVE`, then escape back to the *local* Kermit and give it the opposite command, `RECEIVE` or `SEND`.
- When the transfer is complete, give a `CONNECT` command. Now you are talking to Kermit on the remote computer again. Type `EXIT` to get back to the command prompt on the remote computer. When you are finished using the remote computer, log out and then (if necessary) escape back to Kermit on your local computer. Then you can make another connection or `EXIT` from the local Kermit program.

Note that other methods can be used to simplify the file-transfer process: **client/server operation**, in which all commands are given to the client and passed on automatically to the server, and **autodownload** (and upload), in which the remote Kermit initiates file transfers automatically through your terminal emulator.

The file transfer protocol defaults in C-Kermit 7.0, unlike those for earlier releases, favor speed over robustness, on the assumption that connections in these times are usually reliable (over TCP/IP and/or

(HP-UX C-Kermit)

error-correcting modems with hardware flow control). If you experience file transfer failures, use the CAUTIOUS or ROBUST commands to choose more conservative (and therefore slower) protocol settings. For fine tuning of performance, you can choose specific packet lengths, window sizes, and control-character prefixing strategies as explained in Chapter 12 of the manual, *Using C-Kermit*.

If you are accessing a remote host where C-Kermit resides via Telnet or other connection that is guaranteed reliable from end to end, and both Kermits support it (C-Kermit 7.0 does), a new "streaming" form of the Kermit protocol is used automatically to give ftp-like speeds (the limiting factor being the overhead from the remote Telnet or Rlogin server and/or PTY driver).

OTHER FEATURES

C-Kermit includes features too numerous to be explained in a man page. For further information about connection establishment, modem dialing, networks, terminal connection, key mapping, logging, file transfer options and features, troubleshooting, client/server operation, character-set translation during terminal connection and file transfer, "raw" up- and downloading of files, macro construction, script programming, convenience features, and shortcuts, plus numerous tables, examples, and illustrations, please consult *Using C-Kermit*.

GETTING HELP

C-Kermit has extensive built-in help. You can find out what commands exist by typing ? at the C-Kermit> prompt. You can type HELP at the C-Kermit> prompt for "getting-started" message, or HELP followed by the name of a particular command for information about that command, for example:

```
help send
help set file
```

You can type ? anywhere within a command to get brief help about the current command field. You can also type the INTRO command to get a brief introduction to C-Kermit, and the MANUAL command to access this (or another) manual page. Finally, you can use the SUPPORT command for instructions on obtaining technical support.

ENTERING COMMANDS

You can use upper or lower case for interactive-mode commands, but remember that UNIX filenames are case-sensitive. You can abbreviate commands as long as the abbreviation matches only one possibility. While typing a command, you can use the following editing characters:

- Delete, Backspace, or Rubout erases the rightmost character.
- Ctrl-W erases the rightmost "word".
- Ctrl-U erases the current command line.
- Ctrl-R redisplay the current command.
- Ctrl-P recalls a previous command (scrolls back in command buffer).
- Ctrl-N scrolls forward in a scrolled-back command buffer.
- Ctrl-C cancels the current command.
- Tab, Esc, or Ctrl-I tries to complete the current keyword or filename.
- ? gives help about the current field.

To enter the command and make it execute, press the Return or Enter key.

BACKSLASH NOTATION

Within an interactive command, the \ character (backslash) is a prefix used to enter special quantities, including ordinary characters that would otherwise be illegal. At the end of a line, \ or - (dash) makes the next line a continuation of the current line. Other than that, the character following the \ identifies what the special quantity is:

| | |
|----------|---|
| % | A user-defined simple (scalar) variable such as %a or %1 |
| & | an array reference such as &a[3] |
| \$ | an environment variable such as \$(TERM) |
| v (or V) | a built-in variable such as \v(time) |
| f (or F) | a function such as \Fsubstring(\%a,3,2) |
| s (or S) | compact substring notation, macronames, like \s(foo[3:12]) |
| : | compact substring notation, all variables, like \:(a[3:12]) |
| d (or D) | a decimal (base 10) number (1 to 3 digits, 0..255) such as \d27 |
| o (or O) | an octal (base 8) number (1 to 3 digits, 0..377) such as \o33 |
| x (or X) | a hexadecimal (base 16) number (2 digits, 00..ff) like \x1b |

| | |
|-----------------|--|
| \ | the backslash character itself |
| b (or B) | the BREAK signal (OUTPUT command only) |
| l (or L) | a Long BREAK signal (OUTPUT only) |
| n (or N) | a NUL (0) character (OUTPUT only) |
| a decimal digit | a 1-, 2-, or 3-digit decimal number, such as \27 |
| { } | used for grouping, e.g. \{27}123 |
| anything else: | following character taken literally. |

Note that numbers turn into the character with that binary code (0-255), so you can use \7 for a bell, \13 for carriage return, \10 for linefeed. For example, to have C-Kermit send a BELL to your screen, type:

```
echo \7
```

COMMAND LIST

The commands most commonly used, and important for beginners to know, are marked with "*":

Program Management:

| | |
|---------------|--|
| BACK | Return to previous directory. |
| BROWSE | Invoke Web browser. |
| * CD | Change Directory |
| PWD | Print Working Directory. |
| CHECK | See if the given feature is configured. |
| CLOSE | Close a connection or a log or other local file. |
| COMMENT | Introduce a full-line comment. |
| COPYRIGHT | Display copyright notice. |
| DATE | Display date and time. |
| * EXIT | Leave the program, return to UNIX. |
| * HELP | Display a help message for a given command. |
| * INTRO | Print a brief introduction to C-Kermit. |
| KERMIT | Give command-line options at the prompt. |
| LOG | Open a log file -- debugging, packet, session, transaction. |
| PUSH | Invoke local system's interactive command interpreter. |
| QUIT | Synonym for EXIT. |
| REDO | Re-execute a previous command. |
| RUN | Run a program or system command. |
| SET COMMAND | Command-related parameters: bytesize, recall buffer size. |
| SET PROMPT | The C-Kermit programs' interactive command prompt. |
| SET EXIT | Items related to C-Kermit's action upon exit or SET LINE/HOST. |
| SHOW EXIT | Display SET EXIT parameters. |
| SHOW FEATURES | Show features that C-Kermit was built with. |
| SHOW VERSIONS | Show version numbers of each source module. |
| SUPPORT | Find out how to get technical support. |
| SUSPEND | Suspend Kermit (use only if shell supports job control!). |
| * SHOW | Display values of SET parameters. |
| * TAKE | Execute commands from a file. |
| VERSION | Display the C-Kermit program version number. |
| Z | Synonym for SUSPEND. |
| * Ctrl-C | Interrupt a C-Kermit command in progress. |
| Ctrl-Z | Synonym for SUSPEND. |
| ; or # | Introduce a full-line or trailing comment. |
| ! or @ | Synonym for RUN. |
| < | Synonym for REDIRECT. |

Connection Establishment and Release:

| | |
|----------|--|
| * DIAL | Dial a telephone number. |
| PDIAL | Partially dial a telephone number. |
| * LOOKUP | Lookup a phone number, test dialing rules. |
| ANSWER | Wait for a phone call and answer it when it comes. |
| * HANGUP | Hang up the phone or network connection. |
| EIGHTBIT | Shortcut to set all i/o to 8 bits. |
| PAD | Command for X.25 PAD (SunOS / Solaris / VOS only). |
| PING | Check status of remote TCP/IP host. |

| | |
|------------------|--|
| REDIAL | Dial the most recently DIALEd number again. |
| LOG CONNECTIONS | Keep a record of each connection. |
| REDIRECT | Redirect standard i/o of command to communication connection. |
| PIPE | Make a connection through an external command or program. |
| SET CARRIER | Treatment of carrier on terminal connections. |
| * SET DIAL | Parameters related to modem dialing. |
| * SET FLOW | Communication line flow control: AUTO, RTS/CTS, XON/XOFF, etc. |
| * SET HOST | Specify remote network host name or address. |
| * SET LINE | Specify serial communication device name, like /dev/cul0p0. |
| SET PORT | Synonym for SET LINE. |
| * SET MODEM TYPE | Specify type of modem on SET LINE device, like USR. |
| * SET NETWORK | Network type, X.25 (SunOS / Solaris / VOS only) or TCP/IP. |
| SET TCP | Specify TCP protocol options (advanced). |
| SET TELNET | Specify TELNET protocol options. |
| SET X.25 | Specify X.25 connection parameters (SunOS / Solaris / VOS only). |
| SET PAD | X.25 X.3 PAD parameters (SunOS / Solaris / VOS only). |
| * SET PARITY | Character parity (none, even, etc.) for communications. |
| * SET SPEED | Serial communication device speed, e.g. 2400, 9600, 57600. |
| SET SERIAL | Set serial communications data size, parity, stop bits. |
| SET STOP-BITS | Set serial communications stop bits. |
| SHOW COMM | Display all communications settings. |
| SHOW CONN | Display info about current connection. |
| SHOW DIAL | Display SET DIAL values. |
| SHOW MODEM | Display modem type, signals, etc. |
| SHOW NETWORK | Display network-related items. |
| * TELNET | = SET NETWORK TCP/IP, SET HOST ..., CONNECT. |
| RLOGIN | Makes an RLOGIN connection (requires privilege). |
| TELOPT | Send a TELNET option negotiation (advanced). |
| CLOSE | Close the current connection. |

Terminal Connection:

| | |
|---------------|---|
| * C | Special abbreviation for CONNECT. |
| * CONNECT | Establish a terminal connection to a remote computer. |
| LOG SESSION | Record terminal session. |
| SET COMMAND | Bytesize between C-Kermit and your keyboard and screen. |
| * SET DUPLEX | Specify which side echoes during CONNECT. |
| SET ESCAPE | Prefix for "escape commands" during CONNECT. |
| SET KEY | Key redefinitions in CONNECT mode. |
| SET TERMINAL | Terminal connection items: bytesize, character-set, echo, etc. |
| SHOW ESCAPE | Display current CONNECT-mode escape character. |
| SHOW KEY | Display keycode and assigned value or macro. |
| SHOW TERMINAL | Display SET TERMINAL items. |
| * Ctrl-\ | CONNECT-mode escape character, followed by another character: C to return to C-Kermit> prompt. B to send BREAK signal. ? to see other options. |

File Transfer:

| | |
|---------------------|---|
| ADD SEND-LIST | Add a file specification to the SEND-LIST. |
| ADD BINARY-PATTERNS | Add a pattern to the binary file pattern list. |
| ADD TEXT-PATTERNS | Add a pattern to the text file pattern list. |
| ASSOCIATE | A file character-set with a transfer character-set. |
| LOG SESSION | Download a file with no error checking. |
| * SEND | Send a file or files. |
| MSEND | Multiple SEND - accepts a list of files, separated by spaces. |
| MOVE | SEND and then delete source file(s) if successful. |
| MMOVE | Multiple MOVE - accepts a list of files, separated by spaces. |
| MAIL | SEND a file to other Kermit, to be delivered as e-mail. |
| RESEND | Continue a incomplete SEND. |
| PSEND | Send part of a file. |
| * RECEIVE | Passively wait for files to arrive from other Kermit. |

| | |
|-----------------|--|
| * R | Special abbreviation for RECEIVE. |
| * S | Special abbreviation for SEND. |
| GET | Ask server to send the specified file(s). |
| MGET | Like GET but accepts a list of files. |
| REGET | Continue a incomplete download from a server. |
| G | Special abbreviation for GET. |
| FAST | Shortcut for fast file-transfer settings. |
| CAUTIOUS | Shortcut for medium file-transfer settings. |
| ROBUST | Shortcut for conservative file-transfer settings. |
| SET ATTRIB | Control transmission of file attributes. |
| * SET BLOCK | Choose error-checking level, 1, 2, or 3. |
| SET BUFFERS | Size of send and receive packet buffers. |
| SET PREFIX | Which control characters to "unprefix" during file transfer. |
| SET DELAY | How long to wait before sending first packet. |
| SET DESTINATION | DISK, PRINTER, or SCREEN for incoming files. |
| * SET FILE | Transfer mode (type), character-set, collision action, etc. |
| * SET RECEIVE | Parameters for inbound packets: packet-length, etc. |
| SET REPEAT | Repeat-count compression parameters. |
| SET RETRY | Packet retransmission limit. |
| SET SEND | Parameters for outbound packets: length, etc. |
| SET HANDSHAKE | Communication line half-duplex packet turnaround character. |
| SET LANGUAGE | Enable language-specific character-set translations. |
| PATTERNS | Turn off filename-pattern-based text/binary mode switching. |
| SET SESSION-LOG | File type for session log, text or binary. |
| SET TRANSFER | File transfer parameters: character-set, display, etc. |
| SET TRANSMIT | Control aspects of TRANSMIT command execution. |
| SET UNKNOWN | Specify handling of unknown character sets. |
| * SET WINDOW | File transfer packet window size, 1-31. |
| SHOW ATTRIB | Display SET ATTRIBUTE values. |
| SHOW CONTROL | Display control-character prefixing map. |
| * SHOW FILE | Display file-related settings. |
| SHOW PROTOCOL | Display protocol-related settings. |
| SHOW LANGUAGE | Display language-related settings. |
| SHOW TRANSMIT | Display SET TRANSMIT values. |
| * STATISTICS | Display statistics about most recent file transfer. |
| TRANSMIT | Send a file with no error checking. |
| XMIT | Synonym for TRANSMIT. |

SEND Command switches:

| | |
|---------------|---|
| /AS-NAME: | Name to send file under. |
| /AFTER: | Send files modified after date-time. |
| /BEFORE: | Send files modified before date-time. |
| /BINARY | Send in binary mode. |
| /COMMAND | Send from standard output of a command. |
| /DELETE | Delete file after successfully sending. |
| /EXCEPT: | Don't send files whose names match given pattern(s). |
| /FILTER: | Pass file contents through given filter program. |
| /FILENAMES: | Specify how to send filenames. |
| /LARGER-THAN: | Send files larger than given size. |
| /LIST: | Send files whose names are listed in given file. |
| /MAIL: | Send file(s) as e-mail to given address. |
| /MOVE-TO: | Move source file to given directory after successfully sending. |
| /NOT-AFTER: | Send files modified not after given date-time. |
| /NOT-BEFORE: | Send files modified not before given date-time. |
| /PATHNAMES: | Specify how to send pathnames. |
| /PRINT: | Send files to be printed. |
| /PROTOCOL: | Send files using given protocol. |
| /QUIET | Don't display file-transfer progress. |
| /RECOVER | Recover interrupted transfer from point of failure. |
| /RECURSIVE | Send a directory tree. |
| /RENAME-TO: | Rename files as specified after successfully sending. |

| | |
|----------------|--|
| /SMALLER-THAN: | Send files smaller than given size. |
| /STARTING-AT: | Send file starting at given byte number. |
| /SUBJECT: | Subject for SEND /MAIL. |
| /TEXT | Send in text mode. |

GET and RECEIVE Command switches:

| | |
|-------------|--|
| /AS-NAME: | Store incoming file under given name. |
| /BINARY | Receive in binary mode if transfer mode not specified. |
| /COMMAND: | Send incoming file data to given command. |
| /EXCEPT: | Don't accept incoming files whose names match. |
| /FILENAMES: | How to treat incoming file names. |
| /FILTER: | Filter program for incoming file data. |
| /MOVE-TO: | Where to move a file after successful receipt. |
| /PATHNAMES: | How to treat incoming path names. |
| /PROTOCOL: | Protocol to use for receiving (RECEIVE only). |
| /RENAME-TO: | New name for file after successful receipt. |
| /QUIET: | Suppress file-transfer display. |
| /TEXT | Receive in text mode if transfer mode not specified. |

Switches only for GET:

| | |
|------------|---|
| /DELETE | Tells server to delete each file after successful transmission. |
| /RECOVER | Resume interrupted file transfer from point of failure. |
| /RECURSIVE | Tells server to send a directory tree. |

File Management:

| | |
|---------------------|---|
| * CD | Change Directory. |
| * PWD | Display current working directory. |
| COPY | Copy a file. |
| * DELETE | Delete a file or files. |
| * DIRECTORY | Display a directory listing. |
| EDIT | Edit a file. |
| MKDIR | Create a directory. |
| PRINT | Print a local file on a local printer. |
| PURGE | Remove backup files. |
| RENAME | Change the name of a local file. |
| RMDIR | Remove a directory. |
| SET PRINTER | Choose printer device. |
| SPACE | Display current disk space usage. |
| SHOW CHARACTER-SETS | Display character-set translation info. |
| TRANSLATE | Translate a local file's character set. |
| TYPE | Display a file on the screen. |
| TYPE /PAGE | Display a file on the screen, pausing after each screenful. |
| XLATE | Synonym for TRANSLATE. |

Client/Server Operation:

| | |
|---------------|---|
| BYE | Terminate a remote Kermit server and log out its job. |
| DISABLE | Disallow access to selected features during server operation. |
| E-PACKET | Send an Error packet. |
| ENABLE | Allow access to selected features during server operation. |
| FINISH | Instruct a remote Kermit server to exit, but not log out. |
| G | Special abbreviation for GET. |
| GET | Get files from a remote Kermit server. |
| QUERY | (Same as REMOTE QUERY) |
| RETRIEVE | Like GET but server deletes files after. |
| REMOTE xxx | Command for server, can be redirected with > or . |
| REMOTE ASSIGN | (RASG) Assign a variable. |
| REMOTE CD | (RCD) Tell remote Kermit server to change its directory. |
| REMOTE COPY | (RCOPY) Tell server to copy a file. |
| REMOTE DELETE | (RDEL) Tell server to delete a file. |
| REMOTE DIR | (RDIR) Ask server for a directory listing. |
| REMOTE EXIT | (REXIT) Ask the server program to exit. |

| | |
|---------------|--|
| REMOTE HELP | (RHELP) Ask server to send a help message. |
| REMOTE HOST | (RHOST) Ask server to ask its host to execute a command. |
| REMOTE KERMIT | (RKER) Send an interactive Kermit command to the server. |
| REMOTE LOGIN | Authenticate yourself to a remote Kermit server. |
| REMOTE LOGOUT | Log out from a Kermit server previously LOGIN'd to. |
| REMOTE MKDIR | (RMKDIR) Tell the server to create a directory. |
| REMOTE PRINT | (RPRINT) Print a local file on the server's printer. |
| REMOTE PWD | (RPWD) Ask server to reveal its current (working) directory. |
| REMOTE QUERY | (RQUERY) Get value of a variable. |
| REMOTE RENAME | (RENAME) Tell server to rename a file. |
| REMOTE RMDIR | (RRMDIR) Tell server to remove a directory. |
| REMOTE SET | Send a SET command to a remote server. |
| REMOTE SPACE | Ask server how much disk space it has left. |
| REMOTE TYPE | Ask server to display a file on your screen. |
| REMOTE WHO | Ask server for a "who" or "finger" listing. |
| SERVER | Be a Kermit server. |
| SET SERVER | Parameters for server operation. |
| SHOW SERVER | Show SET SERVER, ENABLE/DISABLE items. |

Script programming:

| | |
|-----------|--|
| ASK | Prompt the user, store user's reply in a variable. |
| ASKQ | Like ASK, but does not echo (useful for passwords). |
| ASSERT | Evaluate condition and set SUCCESS/FAILURE accordingly. |
| ASSIGN | Assign an evaluated string to a variable or macro. |
| CLEAR | Clear communication device input buffer or other item. |
| CLOSE | Close the connection, or a log or other file. |
| DECLARE | Declare an array. |
| DECREMENT | Subtract one (or other number) from a variable. |
| DEFINE | Define a variable or macro. |
| DO | Execute a macro ("DO" can be omitted). |
| ECHO | Display text on the screen. |
| ELSE | Used with IF. |
| END | A command file or macro. |
| EVALUATE | An arithmetic expression. |
| FAIL | Set FAILURE. |
| FOPEN | Open a local file. |
| FREAD | Read from a file opened with FOPEN. |
| FWRITE | Write to an FOPEN'd file. |
| FSEEK | Seeks to given position in FOPEN'd file. |
| FCLOSE | Close an FOPEN'd file. |
| FOR | Execute commands repeatedly in a counted loop. |
| FORWARD | GOTO in the forward direction only. |
| GETC | Issue a prompt, get one character from keyboard. |
| GETOK | Ask question, get Yes or No answer, set SUCCESS or FAILURE. |
| GOTO | Go to a labeled command in a command file or macro. |
| IF | Conditionally execute the following command. |
| INCREMENT | Add one (or other number) to a variable. |
| INPUT | Match characters from another computer against a given text. |
| LOCAL | Declares local variables in a macro. |
| MINPUT | Like INPUT, but allows several match strings. |
| MSLEEP | Sleep for given number of milliseconds. |
| OPEN | Open a local file for reading or writing. |
| OUTPUT | Send text to another computer. |
| O | Special abbreviation for OUTPUT. |
| PAUSE | Do nothing for a given number of seconds. |
| READ | Read a line from a local file into a variable. |
| REINPUT | Reexamine text previously received from another computer. |
| RETURN | Return from a user-defined function. |
| SCREEN | Screen operations - clear, position cursor, etc. |
| SCRIPT | Execute a UUCP-style login script. |
| SET ALARM | Set a timer to be used with IF ALARM; SHOW ALARM shows it. |

| | |
|----------------|--|
| SET CASE | Treatment of alphabetic case in string comparisons. |
| SET COMMAND | QUOTING turns on/off interpretation of backslash notation. |
| SET COUNT | For counted loops. |
| SET INPUT | Control behavior of INPUT command. |
| SET MACRO | Control aspects of macro execution. |
| SET TAKE | Control aspects of TAKE file execution. |
| SHIFT | Shift macro arguments left the given number of places. |
| SHOW ARGUMENTS | Display arguments to current macro. |
| SHOW ARRAYS | Display information about active arrays. |
| SHOW COUNT | Display current COUNT value. |
| SHOW FUNCTIONS | List names of available \f() functions. |
| SHOW GLOBALS | List defined global variables \%a.\%z. |
| SHOW MACROS | List one or more macro definitions. |
| SHOW SCRIPTS | Show script-related settings. |
| SHOW VARIABLES | Display values all \v() variables. |
| SLEEP | Sleep for given number of seconds. |
| SORT | Sort an array (many options). |
| STATUS | Show SUCCESS or FAILURE of previous command. |
| STOP | Stop executing macro or command file, return to prompt. |
| SUCCEED | Set SUCCESS. |
| SWITCH | Execute selected command(s) based on value of variable. |
| TAKE | Execute commands from a file. |
| UNDEFINE | Undefine a variable. |
| WAIT | Wait for the specified modem signals. |
| WHILE | Execute commands repeatedly while a condition is true. |
| WRITE | Write material to a local file. |
| WRITE-LINE | Write a line (record) to a local file. |
| Writeln | Synonym for WRITE-LINE. |
| XECHO | Like ECHO but no CRLF at end. |
| XIF | Extended IF command. |

BUILT-IN VARIABLES

Built-in variables are referred to by \v(name), can be used in any command, usually used in script programming. They cannot be changed. Type SHOW VARIABLES for a current list.

| | |
|----------------|--|
| \v(argc) | Number of arguments in current macro |
| \v(args) | Number of program command-line arguments |
| \v(blockcheck) | Current SET BLOCK-CHECK type |
| \v(browser) | Current Web browser |
| \v(browsopts) | Current Web browser options |
| \v(browsurl) | Most recent Web browser site (URL) |
| \v(byteorder) | Hardware byte order |
| \v(charset) | Current file character-set |
| \v(cmdbufsize) | Size of command buffer |
| \v(cmdfile) | Name of current command file, if any |
| \v(cmdlevel) | Current command level |
| \v(cmdsource) | Where command are currently coming from, macro, file, etc. |
| \v(cols) | Number of screen columns |
| \v(connection) | Connection type: serial, tcp/ip, etc. |
| \v(count) | Current COUNT value |
| \v(cps) | Speed of most recent file transfer in chars per second |
| \v(cpu) | CPU type C-Kermit was built for |
| \v(crc16) | 16-bit CRC of most recent file transfer |
| \v(ctty) | Device name of controlling terminal |
| \v(d\$ac) | SET DIAL AREA-CODE value |
| \v(d\$cc) | SET DIAL COUNTRY-CODE value |
| \v(d\$ip) | SET DIAL INTL-PREFIX value |
| \v(d\$lc) | SET DIAL LD-PREFIX value |
| \v(d\$px) | SET DIAL PBX-EXCHANGE value |
| \v(date) | Date as 8 Feb 1993 |
| \v(day) | Day of week |

(HP-UX C-Kermit)

| | |
|---------------------------------|---|
| <code>\v(dialcount)</code> | Current value of DIAL retry counter |
| <code>\v(dialnumber)</code> | Phone number most recently dialed |
| <code>\v(dialresult)</code> | Most recent dial result message or code from modem |
| <code>\v(dialstatus)</code> | Return code from DIAL command (0 = OK, 22 = BUSY, etc) |
| <code>\v(dialsuffix)</code> | Current SET DIAL SUFFIX value |
| <code>\v(dialtype)</code> | Code for type of call most recently placed |
| <code>\v(directory)</code> | Current/default directory |
| <code>\v(download)</code> | Current download directory if any |
| <code>\v(editor)</code> | Your preferred editor |
| <code>\v(editfile)</code> | File most recently edited |
| <code>\v(editopts)</code> | Options for editor |
| <code>\v(errno)</code> | Current "errno" (system error number) value |
| <code>\v(errmsg)</code> | Error message string associated with errno |
| <code>\v(escape)</code> | Decimal ASCII value of CONNECT-mode escape character |
| <code>\v(evaluate)</code> | Result of most recent EVALUATE command |
| <code>\v(exitstatus)</code> | Current EXIT status (0 = good, nonzero = something failed) |
| <code>\v(filename)</code> | Name of file currently being transferred |
| <code>\v(filename)</code> | Number of file currently being transferred (1 = first, etc) |
| <code>\v(filespec)</code> | Filespec given in most recent SEND/RECEIVE/GET command |
| <code>\v(fsize)</code> | Size of file most recently transferred |
| <code>\v(ftype)</code> | SET FILE TYPE value (text, binary) |
| <code>\v(herald)</code> | C-Kermit's program herald |
| <code>\v(home)</code> | Home directory |
| <code>\v(host)</code> | Computer host name (computer where C-Kermit is running) |
| <code>\v(hwparity)</code> | SET PARITY HARDWARE setting (if any) |
| <code>\v(input)</code> | Current INPUT buffer contents |
| <code>\v(inchar)</code> | Character most recently INPUT |
| <code>\v(incount)</code> | How many characters arrived during last INPUT |
| <code>\v(inidir)</code> | Directory where initialization file was found |
| <code>\v(inmatch)</code> | [M]INPUT material that matched given \fpattern(). |
| <code>\v(instatus)</code> | Status of most recent INPUT command |
| <code>\v(intime)</code> | How long it took most recent INPUT to succeed (msec) |
| <code>\v(inwait)</code> | Most recent [M]INPUT time limit |
| <code>\v(ipaddress)</code> | IP address of C-Kermit's computer if known |
| <code>\v(kbchar)</code> | Keyboard character that interrupted PAUSE, INPUT, etc. |
| <code>\v(line)</code> | Current communications device, set by LINE or HOST |
| <code>\v(local)</code> | 0 if in remote mode, 1 if in local mode |
| <code>\v(lockdir)</code> | UUCP lockfile directory on this platform |
| <code>\v(lockpid)</code> | Process ID found in lockfile when port is in use |
| <code>\v(macllevel)</code> | Current macro stack level |
| <code>\v(macro)</code> | Name of currently executing macro, if any |
| <code>\v(math_e)</code> | Floating-point constant e |
| <code>\v(math_pi)</code> | Floating-point constant pi |
| <code>\v(math_precision)</code> | Floating point number precision (digits) |
| <code>\v(minput)</code> | Result of most recent MINPUT command |
| <code>\v(model)</code> | Computer hardware model if known |
| <code>\v(modem)</code> | Current modem type |
| <code>\v(m_aa_off)</code> | Modem command to turn autoanswer off |
| <code>\v(m_aa_on)</code> | Modem command to turn autoanswer on |
| <code>\v(m_xxxxx)</code> | (many other modem commands) |
| <code>\v(m_sig_xx)</code> | Value of modem signal xx |
| <code>\v(name)</code> | Name by which C-Kermit was called (kermit, wermit, etc) |
| <code>\v(ndate)</code> | Current date as 19930208 (yyyymmdd) |
| <code>\v(nday)</code> | Numeric day of week (0 = Sunday) |
| <code>\v(newline)</code> | System-independent newline character or sequence |
| <code>\v(ntime)</code> | Current local time in seconds since midnight (noon = 43200) |
| <code>\v(osname)</code> | Operating System name |
| <code>\v(osrelease)</code> | Operating System release |
| <code>\v(osversion)</code> | Operating System version |
| <code>\v(packetlen)</code> | Current SET RECEIVE PACKET-LENGTH value |
| <code>\v(parity)</code> | Current parity setting |

| | |
|----------------|---|
| \v(pexitstat) | Exit status of most recently forked process |
| \v(pid) | C-Kermit's process ID |
| \v(platform) | Specific machine and/or operating system |
| \v(program) | Name of this program ("C-Kermit") |
| \v(protocol) | Currently selected file transfer protocol |
| \v(p_8bit) | Current 8th-bit prefix (Kermit protocol) |
| \v(p_ctl) | Current control-character prefix (Kermit protocol) |
| \v(p_rpt) | Current repeat-count prefix (Kermit protocol) |
| \v(query) | Result of most recent REMOTE QUERY command |
| \v(return) | Most recent RETURN value |
| \v(rows) | Number of rows on the terminal screen |
| \v(sendlist) | Number of entries in SEND-LIST |
| \v(serial) | Serial port settings in 8N1 format |
| \v(speed) | Current speed, if known, or "unknown" |
| \v(startup) | Current directory when C-Kermit was started |
| \v(status) | 0 or 1 (SUCCESS or FAILURE of previous command) |
| \v(sysid) | Code for platform ID of C-Kermit's computer (U1=UNIX) |
| \v(system) | UNIX (name of operating system family) |
| \v(terminal) | Terminal type |
| \v(test) | C-Kermit test version, if any (e.g. Beta.10) |
| \v(textdir) | Where C-Kermit thinks its text files are |
| \v(tfsiz) | Total size of file group most recently transferred |
| \v(time) | Time as 13:45:23 (hh:mm:ss) |
| \v(tmpdir) | Temporary directory |
| \v(trigger) | Most recent string to trigger return from CONNECT |
| \v(ttyfd) | File descriptor of current communication device |
| \v(ty_xx) | Used internally by TYPE |
| \v(userid) | User ID of person running C-Kermit |
| \v(version) | Numeric version of Kermit, e.g. 501190. |
| \v(window) | Current window size (SET WINDOW value) |
| \v(xferstatus) | Status of most recent file transfer |
| \v(xfermsg) | Error message, if any, terminating most recent transfer |
| \v(xfer_XXX) | Various statistics from last file transfer. |
| \v(xprogram) | C-Kermit |
| \v(xversion) | Same as \v(version) |

BUILT-IN FUNCTIONS

Builtin functions are invoked as \Fname(args), can be used in any command, and are usually used in script programs. Type SHOW FUNCTIONS for a current list. Type "help function <name>" for a description of the arguments and return value, for example, **help function basename**.

COMMAND LINE OPTIONS

C-Kermit accepts commands (or "options") on the command line, in the time-honored UNIX style. Alphabetic case is significant. All options are optional. If one or more action options are included, Kermit exits immediately after executing the command-line options, otherwise it enters interactive command mode.

```
kermit [filename] [-x arg [-x arg]...[-yyy]...]
```

where:

filename is the name of a command file to execute,

-x is an option requiring an argument,

-y an option with no argument.

Actions:

| | |
|----------|---|
| -s files | send files |
| -s - | send files from stdin |
| -r | receive files |
| -k | receive files to stdout |
| -x | enter server mode |
| -O | like -x but exits after one transaction |
| -f | finish remote server |

| | |
|----------|--|
| -g files | get remote files from server (quote wildcards) |
| -G files | like -g but sends file to standard output |
| -a name | alternate file name, used with -s, -r, -g |
| -c | connect (before file transfer), used with -l or -j |
| -n | connect (after file transfer), used with -l or -j |

Settings:

| | |
|---------|--|
| -l line | communication line device (to make a serial connection) |
| -l n | open file descriptor of communication device |
| -j host | TCP/IP network host name (to make a network connection) |
| -J host | connect like TELNET, exit when connection closes |
| -l n | open file descriptor of TCP/IP connection (n = number) |
| -X | X.25 network address |
| -Z | open file descriptor of X.25 connection |
| -o n | X.25 closed user group call info |
| -u | X.25 reverse-charge call |
| -q | quiet during file transfer |
| -I | connection is reliable (e.g. TCP or X.25) |
| -8 | 8-bit clean |
| -0 | 100% transparency in CONNECT mode (and no escaping back) |
| -i | transfer files in binary mode |
| -T | transfer files in text mode |
| -P | send/accept literal path (file) names |
| -b bps | serial line speed, e.g. 1200 |
| -m name | modem type, e.g. hayes |
| -p x | parity, x = e,o,m,s, or n |
| -t | half duplex, xon handshake |
| -e n | receive packet-length |
| -v n | window size |
| -L | used with -s to select recursive directory transfer |
| -Q | Quick file-transfer settings |
| -w | write over files of same name, do not backup old file |
| -D n | delay n seconds before sending a file |
| -V | "manual mode" = SET FILE PATTERNS OFF, SET TRANSFER MODE MANUAL. |

Other:

| | |
|-----------|---|
| -y name | alternate init file name |
| -Y | Skip init file |
| -R | Advise C-Kermit it will be used only in remote mode |
| -d | log debug info to file debug.log |
| -S | Stay, do not exit, after action command |
| -C "cmds" | Interactive-mode commands, comma-separated |
| -z | Force foreground operation |
| -B | Force background (batch) operation |
| -h | print command-line option help screen |
| = | Ignore all text that follows |
| -- | Same as = |

COMMAND LINE EXAMPLES

Remote-mode example (C-Kermit is on the far end):

```
kermit -v 4 -i -s oofa.bin
```

sends file oofa.bin in binary mode (-i) using a window size of 4 (-v 4).

Local-mode example (C-Kermit makes the connection):

```
kermit -l /dev/tty0p0 -b 19200 -c -r -n
```

makes a 19200-bps direct connection out through /dev/tty0p0, CONNECTs (-c) so you can log in and, presumably start a remote Kermit program and tell it to send a file, then it RECEIVES the file (-r), then it CONNECTs back (-n) so you can finish up and log out.

(HP-UX C-Kermit)

For dialing out, you must specify a modem type, and you might have to use a different device name:

```
kermit -m hayes -l /dev/cul0p0 -b 2400 -c -r -n
```

FILES

| | |
|---|--|
| <code>\$HOME/.mykerlerc</code> | Your personal C-Kermit customization file. |
| <code>\$HOME/.kdd</code> | Your personal dialing directory. |
| <code>\$HOME/.ksd</code> | Your personal services directory. |
| <code>/usr/share/lib/kermit/READ.ME</code> | Overview of HP-UX C-Kermit, please read |
| <code>/usr/share/lib/kermit/COPYING.TXT</code> | Copyright, permissions, disclaimer |
| <code>/usr/share/lib/kermit/ckermid.ini</code> | System-wide initialization file |
| <code>/usr/share/lib/kermit/ckermid.ini</code> | Sample customization file |
| <code>/usr/share/lib/kermit/ckermid.kdd</code> | Sample dialing directory |
| <code>/usr/share/lib/kermit/ckermid.ksd</code> | Sample services directory |
| <code>/usr/share/lib/kermit/ckermid2.txt</code> | Updates to "Using C-Kermit" 2nd Ed |
| <code>/usr/share/lib/kermit/ckcbwr.txt</code> | C-Kermit "beware" file - hints & tips |
| <code>/usr/share/lib/kermit/ckubwr.txt</code> | UNIX-specific beware file |
| <code>/usr/share/lib/kermit/ck*.txt</code> | Other plain-text documentation |
| <code>/usr/share/lib/kermit/ckdemo.ksc</code> | Macros from "Using C-Kermit" |
| <code>/usr/share/lib/kermit/ckevt.ksc</code> | Ditto |
| <code>/usr/share/lib/kermit/ckepager.ksc</code> | Alpha pager script |
| <code>/var/spool/locks/LCK..*</code> | UUCP lockfiles |

To make **personalized customizations**, copy the file `/usr/share/lib/kermit/ckermid.ini` file to your home directory, make any desired changes, and rename it to `.mykerlerc`.

You may also create a personalized **dialing directory** like the sample one in `/usr/share/lib/kermit/ckermid.kdd`. Your personalized dialing directory should be stored in your home directory as `.kdd` and your personal network directory as `.ksd`. See Chapters 5 and 6 of *Using C-Kermit* for details.

And you may also create a personalized **services directory** like the sample one in `/usr/share/lib/kermit/ckermid.ksd`. Your personalized services directory should be stored in your home directory as `.ksd`. See Chapter 7 of *Using C-Kermit* for instructions.

The demonstration files illustrate C-Kermit's script programming constructs; they are discussed in chapters 17-19 of the book. You can run them by typing the appropriate TAKE command at the C-Kermit> prompt, for example: `take /usr/share/lib/kermit/ckdemo.ini`.

AUTHORS

Frank da Cruz, Columbia University, with contributions from hundreds of other volunteer programmers all over the world. See Acknowledgements in *Using C-Kermit*.

REFERENCES

Frank da Cruz and Christine M. Gianone,
Using C-Kermit, Second Edition, 1997, 622 pages, Digital Press / Butterworth-Heinemann, 225 Wildwood Street, Woburn, MA 01801, USA. ISBN 1-55558-164-1. (In the USA, call +1 800 366-2665 to order Digital Press books.) Also available in a German edition from Verlag Heinze Heise, Hannover.

Frank da Cruz,
Kermit, A File Transfer Protocol, Digital Press / Butterworth-Heinemann, Woburn, MA, USA (1987). ISBN 0-932376-88-6. The Kermit file transfer protocol specification.

Christine M. Gianone,
Using MS-DOS Kermit, Digital Press / Butterworth-Heinemann, Woburn, MA, USA (1992). ISBN 1-5558-082-3. Also available in a German edition from Heise, and a French edition from Heinz Schiefer & Cie, Versailles.

Kermit News,
Issues 4 (1990) and 5 (1993), Columbia University, for detailed discussions of Kermit file transfer performance.

DIAGNOSTICS

The diagnostics produced by *C-Kermit* itself are intended to be self-explanatory. In addition, every command returns a SUCCESS or FAILURE status that can be tested by IF FAILURE or IF SUCCESS. In addition, the program itself returns an exit status code of 0 upon successful operation or nonzero if any of

various operations failed.

BUGS

See the comp.protocols.kermit.* newsgroups on Usenet for discussion, or the files, ckcker.bwr and ckuker.bwr, for a list of bugs, hints, tips. etc. Report bugs via e-mail to **kermit-support@columbia.edu**. Visit <http://www.columbia.edu/kermit/support.html> for details about tech support.

CONTACTS

For more information about Kermit software and documentation, visit the Kermit Web site:

<http://www.columbia.edu/kermit/>

Or write to:

The Kermit Project
Columbia University
612 West 115th Street
New York, NY 10025-7221
USA

Or send e-mail to **kermit@columbia.edu**. Or call +1 212 854-3703. Or fax +1 212 663-8202.

k

NAME

keylogin - decrypt and store secret key with keyserv

SYNOPSIS

```
/usr/bin/keylogin [ -r ]
```

DESCRIPTION

The **keylogin** command prompts for a password, and uses it to decrypt the user's secret key. The key may be found in the **/etc/publickey** 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 **/etc/nsswitch.conf** file (see *nsswitch.conf(4)*). Once decrypted, the user's secret key is stored by the local key server process, *keyserv(1M)*. This stored key is used when issuing requests to any secure RPC services, such as NIS+. The program *keylogout(1)* can be used to delete the key stored by **keyserv**.

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)*.

Options

-r Update the **/etc/.rootkey** 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 *keyserv(1M)*). The **-r** option should be used by the administrator when the host's entry in the publickey database has changed, and the **/etc/.rootkey** file has become out-of-date with respect to the actual key pair stored in the publickey database. The permissions on the **/etc/.rootkey** file are such that it may be read and written by the super-user but by no other user on the system.

AUTHOR

keylogin was developed by Sun Microsystems, Inc.

FILES

/etc/.rootkey Super-user's secret key

SEE ALSO

chkey(1), *keylogout(1)*, *login(1)*, *keyserv(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]`

DESCRIPTION

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.

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.

Options

-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 that are started by the super-user.

AUTHOR

keylogout was developed by Sun Microsystems, Inc.

SEE ALSO

at(1), *chkey*(1), *login*(1), *keylogin*(1), *keyserv*(1M), *newkey*(1M), *publickey*(4).

k

NAME

keysh - context-sensitive softkey shell

SYNOPSIS

keysh

DESCRIPTION

keysh is an extension of the standard Korn-shell (for a description of the basic Korn-shell functionality, see *ksh*(1)).

keysh uses hierarchical softkey menus and context-sensitive help to aid users in building command-lines, combining the power of the Korn-shell with the ease-of-use of a menu system.

And **keysh** is entirely data-driven, allowing its menus and help to be easily extended as needed.

Note that during **keysh** invocation, the environment variable **\$TERM** must specify the terminal type, as defined in the *terminfo*(4) database (see ENVIRONMENT VARIABLES below).

COMMAND ENTRY

keysh continually parses the command-line and always presents the user with an appropriate set of *current choices* on the softkey labels.

The user can select these softkeys to create readable *softkey commands* on the command-line. **keysh** automatically translates these softkey commands into equivalent *HP-UX commands* prior to executing them.

Alternatively, the user can ignore the softkeys altogether in favor of entering the traditional HP-UX commands directly, as when using the Korn-shell.

During command entry, **keysh** ordinarily displays a *status-line* near the bottom of the screen. This status-line contains information such as the host name, current directory, and time and date.

Whenever the user *must* perform an action to complete the current softkey command, **keysh** temporarily displays a *prompt message* in place of the status-line. This message briefly describes the required action.

Softkey Types

keysh presents four basic softkey types:

- Help--** Selecting the **--Help--** softkey causes **keysh** to display help information associated with the next selected softkey, rather than actually performing its action.
- More--** If there are more current choices than there are softkeys, **keysh** breaks the choices into banks and displays a special **--More--** softkey along with the first bank. Selecting the **--More--** softkey causes **keysh** to display the next bank of softkeys in sequence, eventually cycling back to the first.
- <param>** *parameter* softkeys are displayed as a name enclosed between a pair of less-than and greater-than symbols. They indicate that the user-supplied text (such as a file name) should be entered into the command-line at that point, rather than actually selecting the softkey. (Actually selecting the softkey only causes **keysh** to display a hint message on the status line; the command-line remains unchanged.)
- option** All other softkeys are *option* softkeys that can be used to insert the corresponding command or option name into the command-line.

Softkeys can be selected from left to right.

Editing The Command-Line

keysh supports the normal Korn-shell command-line editing modes. In addition, **keysh** also recognizes the cursor movement and editing keys found on most terminals, as defined in the *terminfo*(4) database. These include:

- <Clear display>** Clear the screen and command-line. If the screen is scrolled, clear only from the cursor position to the end of scrolling memory.
- <Clear line>** Clear from the cursor position to the end of the command-line.
- <Delete line>** Clear the entire command-line.

| | |
|--------------------|---|
| <Insert line> | Translate any softkey commands in the current command-line and then edit the result. |
| <Delete char> | Delete the character under the cursor. |
| <Insert char> | Toggle between insert and overwrite modes. |
| <Up/Down arrow> | Recall the previous/next command from the history buffer. |
| <Left/Right arrow> | Move the cursor left/right. |
| <Home up/down> | Move the cursor to the beginning/end of the command-line. |
| <Tab> | If no <Insert line> key is present, perform the <Insert line> function (see above). Otherwise, if no --Help-- softkey is present, perform the --Help-- function (also see above). Otherwise, perform the normal tab function. |
| <Backtab> | Move the cursor to the beginning of the previous word. |
| <Ctrl-L> | Redraw the lower lines of the screen and restore any necessary terminal modes. |

Visible Softkey Commands

If the **visibles** configuration option is enabled (see CONFIGURATION below), **keysh** displays a list of configured softkey commands on the softkey labels whenever it is expecting a new command. This is the the top-level softkey menu.

If the user selects one of these softkey commands, **keysh** inserts its command name into the command-line then displays a sub-menu listing the command's major parameters and/or options.

The user can then (from left to right) select option softkeys and/or enter text in place of parameter softkeys. **keysh** automatically navigates the hierarchical softkey menu, always presenting the user with an appropriate set of current choices on the softkey labels.

Note that **keysh** automatically redisplay the top-level softkey menu when it detects that a command separator (such as a pipe or semi-colon) has been entered, thus allowing the user to use softkeys for subsequent commands on the command-line as well as the first.

Invisible Softkey Commands

If the **invisibles** configuration option is enabled (see CONFIGURATION below) and **keysh** recognizes a traditional HP-UX command being entered, it gives the user one last chance to use the softkeys by again presenting an appropriate set of current choices on the softkey labels. As with the top-level softkey menu options, the user can choose to ignore the softkeys in favor of entering the traditional HP-UX options directly.

Backup Softkeys

If the **backups** configuration option is enabled (see CONFIGURATION below), **keysh** displays the *backup softkeys* and programs the terminal function keys appropriately whenever it has no other softkeys to display (such as when a command is running). These provide the traditional static softkey control which many users may be used to.

Traditional HP-UX Commands

If the user enters a traditional HP-UX command when **keysh** is displaying its top-level softkey menu, **keysh** simply displays the backup softkeys and allows the user to proceed.

If **keysh** subsequently detects a command separator, it again redisplay the top-level softkey menu.

Softkey Command Syntax Errors

Many softkey commands present the user with a set of softkey options from which exactly one (or at least one) *must* be selected. If the user fails to do this, **keysh** treats it as a syntax error, displaying an error message and not accepting the command until the error has been corrected.

Similarly, many softkey commands require that the user enter one or more softkey parameters before the command is semantically complete. If the user fails to do this, **keysh** again treats it as a syntax error.

Softkey Command Redirections

The user can append redirection symbols (such as a less-than or greater-than symbol followed by a file name) following a softkey command. These are appended *verbatim* to the translated HP-UX command.

USING KEYSH WITH TERMINAL SESSION MANAGER

When operating under the Terminal Session Manager (see *tsm(1)*), **keysh** displays the **tsm** softkeys instead of the backup softkeys. If desired, this interaction can be overridden by setting the **\$KEYTSM** environment variable (see ENVIRONMENT VARIABLES below).

When operating under **tsm**, **keysh** also automatically displays the **tsm** window number in the status-line.

CONFIGURATION

All **keysh** configuration functions are accessed through the top-level **Keysh_config** softkey command or **kc** built-in command. These functions include:

- adding, placing, and deleting softkeys,
- specifying backup softkeys,
- selecting global options,
- selecting status-line items,
- restarting keysh,
- writing configuration changes, and
- undoing other configuration changes.

Each time the user changes **keysh**'s configuration, **keysh** automatically updates the user's **\$HOME/.keyshrc** file. Upon subsequent invocations, **keysh** automatically reconfigures itself as configured previously.

Adding, Placing, And Deleting Softkeys

Any of the standard softkeys (see STANDARD SOFTKEY DEFINITIONS below) can be added to the top-level softkey menu using the **kc softkey add** command. If desired, an alternate softkey label may be specified (usually in place of a cryptic HP-UX command name) using the **with_label** option.

By default, added softkeys are placed at the end of the last **--More--** bank of the top-level softkey menu. This placement can be overridden using the **and_place** option of the **kc softkey add** command or using the **kc softkey move** command.

In addition to the standard softkeys, custom softkeys can also be added from custom softkey files using the **from_user** or **from_file** options. For a description of the softkey file format, see *softkeys(4)*.

Note that any time a softkey is added from a particular softkey file, all of the remaining softkeys from that file are automatically loaded for use as invisible softkey commands. All softkeys from a file can also be loaded for use as invisible softkey commands using the **kc softkey add invisibles** command.

Any of the softkeys in the top-level softkey menu can be deleted using the **kc softkey delete** command.

Specifying Backup Softkeys

Backup softkeys are typically specified in the user's **\$HOME/.softkeys** file. The basic backup softkey definition line resembles:

```
backup softkey "<softkey>" literal "<string>";
```

Where **<softkey>** is the softkey label to display and **<string>** is the text string to program the terminal function key with. A maximum of eight backup softkeys can be specified.

Note that backup softkeys must be explicitly added using the **kc softkey add backups** command before **keysh** can program them.

Selecting Global Options

Various global options can be configured using the **kc option** command, including:

- | | |
|-------------------|--|
| backups | Enable or disable the programming of the backup softkeys. |
| help | Enable or disable the --Help-- softkey. |
| invisibles | Enable or disable the recognition of invisible softkey commands. |
| prompts | Enable or disable the automatic generation of prompt messages. When enabled, keysh displays a prompt message whenever the user <i>must</i> perform an action to complete the current softkey command. This message briefly describes the required action. |

- selectors** Enable or disable the use of keyboard selectors. When enabled, **keysh** displays an upper-case selector character in each softkey label. Typing the unquoted (upper-case) character selects the softkey just as if its corresponding function key had been pressed. Quoting the selector character in any way restores its traditional meaning. Selector keys are intended to be used on terminals that do not support a sufficient number of softkeys.
- translations** Enable or disable the display of HP-UX command translations.
- visibles** Enable or disable the presentation and recognition of visible softkey commands.

Selecting Status-Line Items

Various information items can be configured into the status-line displayed at the bottom of the screen using the **kc status_line** command, including:

- host_name** The host name.
- user_name** The user name.
- current_dir** The current directory.
- mail_status** The mail status based on the **\$MAIL** environment variable (i.e., **No mail**, **You have mail**, or **You have new mail**).
- date** The date.
- time** The time of day.

In addition, the **\$KEYSH** environment variable, if set, is always displayed first in the status-line.

Restarting Keysh

keysh can be forced to reread the **\$HOME/.keyshrc** file with the **kc restart** command. This command is typically used to update a **keysh** to a new configuration specified in another window.

keysh can also be forced to remove the **\$HOME/.keyshrc** file and restart from the default user configuration with the **kc restart default** command.

Writing Configuration Changes

keysh can be forced to rewrite the **\$HOME/.keyshrc** file with the **kc write** command.

Undoing Other Configuration Changes

keysh can also be forced to rewrite the **\$HOME/.keyshrc** file with its original contents, thus undoing all configuration changes made since **keysh** was invoked, using the **kc undo** command.

Scaling Keysh Functionalities

keysh provides a scalable set of functionalities which can be tailored to suit personal preferences.

For users who are familiar with the HP-UX command names (though not necessarily with the command options) or for users who prefer to usually have the **tsm** softkeys visible, the command **kc options visibles off** prevents **keysh** from displaying its top-level softkey menu while waiting for a command; instead, it displays the backup softkeys or **tsm** softkeys, as appropriate. (**keysh** start-up time can then be decreased significantly by editing the **\$HOME/.keyshrc** file and removing the lines which add visible softkeys.)

For users who are also familiar with the HP-UX command options, the command **kc options invisibles off** prevents **keysh** from displaying softkey menus for invisible softkey commands, also.

And for users who have no need for the backup softkeys, the command **kc options backups off** prevents **keysh** from ever programming the backup softkeys.

Note that if **visibles**, **invisibles**, and **backups** are all turned off, **keysh** performs *no* softkey processing at all. **keysh** effectively transforms into a Korn-shell which displays a status-line and recognizes the cursor movement and editing keys.

EXAMPLES

To add the `od` (see `od(1)`) softkey to the end of the top-level softkey menu and label it `Octal_dump`,

```
kc softkey add od with_label Octal_dump
```

To add the `paste(1)` softkey to the beginning of the top-level softkey menu and label it `Paste`,

```
kc softkey add paste and_place as_first_softkey
```

To add the custom emacs softkey from the file `~rpt/.softkeys` to the top-level softkey menu immediately before the `ls` (see `ls(1)`) softkey,

```
kc softkey add emacs from_user rpt and_place before_softkey ls
```

To add all invisible softkeys from the file `~rpt/.softkeys`,

```
kc softkey add invisibles from_user rpt
```

To add the backup softkeys from the file `$HOME/.softkeys`,

```
kc softkey add backups
```

To delete the `Edit_file` softkey from the top-level softkey menu,

```
kc softkey delete Edit_file
```

To disable the `--Help--` softkey,

```
kc options help off
```

To configure the user name into the status-line,

```
kc status_line user_name on
```

To configure the exit-value of the last command executed into the status-line,

```
KEYSH="\${?#0}"
```

To list the ten largest files in the current directory,

```
ls long_format | sort_lines numerically reverse_order \
  starting_at_field 5 | head
```

STANDARD SOFTKEY DEFINITIONS

`Copy_files`, `Move_files`, `Print_files`, `Set_file_attribs`, `Switch`.

`adjust`, `ar`, `bdf`, `cal`, `cancel`, `cat`, `cd`, `cdb`, `chatr`, `chgrp`, `chmod`, `chown`, `cmp`, `col`, `comm`, `cpio`, `cut`, `dd`, `df`, `diff`, `dircmp`, `disable`, `du`, `elm`, `enable`, `exit`, `find`, `fold`, `grep`, `head`, `jobs`, `kill`, `lp`, `lpstat`, `ls`, `mailx`, `make`, `man`, `mkdir`, `more`, `nm`, `nroff`, `od`, `paste`, `pg`, `pr`, `ps`, `remsh`, `rlogin`, `rm`, `rmdir`, `sdiff`, `set`, `shar`, `sort`, `tail`, `tar`, `tee`, `touch`, `tr`, `umask`, `uname`, `vi`, `wc`, `who`, `write`, `xd`, `xdb`.

ENVIRONMENT VARIABLES

| | |
|----------------|---|
| TERM | Specifies the terminal type, as defined in the <code>terminfo(4)</code> database. This variable must be either part of keysh 's invocation environment or it must be set within one of the standard Korn-shell start-up files. |
| COLUMNS | Specifies the number of columns in the terminal screen if different than the <code>terminfo(4)</code> default. |
| LINES | Specifies the number of lines in the terminal screen if not the same as the <code>terminfo(4)</code> default. |
| PAGER | Specifies the preferred pager to be used to display help. The default is <code>more</code> (see <code>more(1)</code>). |
| TZ | Specifies the time-zone to be used for time and date representations on the status-line. The default is <code>en_US.roman8</code> . |
| KEYBEL | Specifies the character sequence sent to the terminal by keysh to ring the bell. The default is <code>^G</code> . |
| KEYENV | Specifies an alternate keysh configuration file. The default is <code>\$HOME/.keyshrc</code> . |
| KEYESC | Specifies the maximum allowable delay between characters (in milliseconds) if they are to be treated as part of a terminal escape sequence. The default is 350 ms. |

| | |
|---------------|--|
| KEYKSH | If set, specifies that keysh should mimic the behavior of the Korn-shell as closely as possible. No softkeys or status-line are displayed. This mode is particularly useful over slow modem lines. |
| KEYLOC | If set, specifies that keysh should leave the terminal keypad in local mode while commands are being entered. This mimics the behavior of the Korn-shell. |
| KEYPS1 | If set, specifies that keysh should not reset the initial values of \$PS1 , \$PS2 , and \$PS3 . Note that \$PS1 must be a constant character string in order for keysh to recognize it and provide subsequent softkey assistance. |
| KEYSH | Specifies arbitrary text to be included in the keysh status-line. |
| KEYSIM | If set, specifies that keysh should always simulate softkey labels and not use the built-in labels on HP terminals. |
| KEYTSM | If set, specifies that keysh should <i>not</i> use the tsm softkeys when tsm is running. In this case, the user can either use the tsm hotkey , the backup softkeys, or the Switch softkey command (see STANDARD SOFTKEY DEFINITIONS above) to switch tsm windows. |

KSH DIFFERENCES

keysh is an extension of *ksh*(1) with the following exceptions:

Screen Updates

keysh optimizes its display output to take advantage of available terminal capabilities. Unlike the Korn-shell which often has to redraw large portions of the command-line, **keysh** can simply insert or delete characters at the appropriate screen position.

This makes **keysh** significantly faster over slow modem lines, especially if the **\$KEYKSH** environment variable is set (see ENVIRONMENT VARIABLES above).

Emacs-Mode Editing

The new **<ESC>v** command performs the function of the vi-mode **v** command.

An initial **^N** command recalls the history line *following* the history line executed as the previous command. This provides an easy mechanism to repeat a sequence of history commands.

gmacs editing mode is not supported; **emacs** editing mode follows the GNU emacs (18.54) definition of **^T**.

The **^@** and **<ESC>n ^K** commands are not supported.

The **M- <letter>** and **M-] <letter>** alias functions are not supported (in lieu of true softkey support).

Vi-Mode Editing

The new **o** command performs the function of the emacs-mode **^O** command.

An initial **j** command recalls the history line *following* the history line executed as the previous command. This provides an easy mechanism to repeat a sequence of history commands.

The **|** command is not supported.

The **@<letter>** alias function is not supported (in lieu of true softkey support).

The **u** command performs an emacs-style nested undo; **u<space>** performs a traditional vi-style undo.

WARNINGS

keysh requires that the **\$TERM** environment variable be set appropriately in your **\$HOME/.profile** file. It also requires that **\$LINES** and **\$COLUMNS** be set appropriately if running on a non-standard size terminal. Otherwise, an error message or a garbled screen display results.

keysh requires that option softkeys be selected from left to right. When editing a command-line, it is possible to back up and insert a softkey out-of-order -- resulting in a command error.

keysh initializes **\$PS1**, **\$PS2**, and **\$PS3** and types them *read-only* — do not change them. Instead, use **\$KEYSH** to display additional status information.

keysh normally maintains the **\$HOME/.keyshrc** file without user intervention; however, start-up errors may occasionally occur and persist. In this case, either execute the command **kc restart default** (to remove the file and revert to the default user configuration) or execute the command **kc write** (to rewrite the file with the current configuration).

keysh assumes that HP-UX commands are not heavily aliased; otherwise unexpected command translations may occur.

keysh neglects the effects of the Korn-shell expansion mechanisms when counting command-line parameters, causing it to occasionally underestimate the true number of parameters specified. The <ESC>* emacs-mode or vi-mode editing command can often be used to pre-expand these parameters.

The <ESC>**v** emacs-mode editing command and **v** vi-mode editing command cannot be used to edit (pre-translated) softkey commands, since no subsequent command translation can occur.

Adding a large number of softkeys can cause **keysh** to overflow a 1-Mbyte Korn-shell data size limitation, causing disconcerting behavior.

keysh can only program the function keys on terminals whose *terminfo*(4) entry defines the **pfkey** capability; similarly, it can only use hardware softkey labels on terminals whose *terminfo*(4) entry defines the **pln** capability (along with specifying **lh** equal to 2).

The default value for **\$KEYESC** was chosen to provide reasonable response in both local and networked environments. If **keysh** misinterprets quickly typed emacs-mode or vi-mode editing commands as terminal escape sequences, it may be necessary to decrease this value.

Specifying a **\n** (new-line) in the literal key sequence for a backup softkey causes undesired results on HP terminals; use a **\r** (carriage-return) instead.

keysh does not display **tsm** softkeys when simulating softkey labels.

A limited number of environment variables and arguments are exported to the pager when displaying help.

EXTERNAL INFLUENCES

Environment Variables

LANG determines the language in which softkeys and messages are displayed.

LC_TIME determines the format and contents of date and time strings in the status-line.

International Code Set Support

Single-byte character code sets are supported.

AUTHOR

keysh was developed by HP and AT&T.

FILES

| | |
|---------------------------------------|---|
| /usr/bin/keysh | main executable |
| /usr/lib/keysh/builtins | Keysh_config softkey definition file |
| /usr/lib/keysh/\$LANG/softkeys | standard softkey definitions file |
| /usr/lib/keysh/\$LANG/keyshrc | default user configuration file |
| /usr/lib/nls/\$LANG/keysh.cat | message catalog |
| \$HOME/.keyshrc | user configuration file |
| \$HOME/.softkeys | user softkey definitions file |

SEE ALSO

ksh(1), tsm(1), softkeys(4), terminfo(4).

NAME

kill - send a signal to a process; terminate a process

SYNOPSIS

kill [-s *signame*] *pid* ...
kill [-s *signum*] *pid* ...
kill -1

Obsolescent Versions:

kill -*signame* *pid* ...
kill -*signum* *pid* ...

DESCRIPTION

The **kill** command sends a signal to each process specified by a *pid* process identifier. The default signal is **SIGTERM**, which normally terminates processes that do not trap or ignore the signal.

pid is a process identifier, an unsigned or negative integer that can be one of the following:

- > 0 The number of a process.
- = 0 All processes, except special system processes, whose process group ID is equal to the process group ID of the sender.
- ==1 All processes, except special system processes, if the user has appropriate privileges. Otherwise, all processes, except special system processes, whose real or effective user ID is the same as the user ID of the sending process.
- <-1 All processes, except special system processes, whose process group ID is equal to the absolute value of *pid* and whose real or effective user ID is the same as the user of the sending process.

Process numbers can be found with the **ps** command (see *ps(1)*) and with the built-in **jobs** command available in some shells.

Options

kill recognizes the following options:

- 1 (ell) List all values of *signame* supported by the implementation. No signals are sent with this option. The symbolic names of the signals (without the **SIG** prefix) are written to standard output, separated by spaces and newlines.
- s *signame* Send the specified signal name. The default is **SIGTERM**, number 15. *signame* can be specified in upper- and/or lowercase, with or without the **SIG** prefix. These values can be obtained by using the -1 option. The symbolic name **SIGNULL** represents signal value zero. See "Signal Names and Numbers" below.
- s *signum* Send the specified decimal signal number. The default is 15, **SIGTERM**. See "Signal Names and Numbers" below.
- signame* (Obsolescent.) Equivalent to -s *signame*.
- signum* (Obsolescent.) Equivalent to -s *signum*.

Signal Names and Numbers

The following table describes a few of the more common signals that can be useful from a terminal. For a complete list and a full description, see the header file **<signal.h>** and the manual entry *signal(5)*.

| <i>signum</i> | <i>signame</i> | Name | Description |
|---------------|----------------|-----------|--|
| 0 | SIGNULL | Null | Check access to <i>pid</i> |
| 1 | SIGHUP | Hangup | Terminate; can be trapped |
| 2 | SIGINT | Interrupt | Terminate; can be trapped |
| 3 | SIGQUIT | Quit | Terminate with core dump; can be trapped |
| 9 | SIGKILL | Kill | Forced termination; cannot be trapped |
| 15 | SIGTERM | Terminate | Terminate; can be trapped |

| | | | |
|----|----------------|---------------|--------------------------------------|
| 24 | SIGSTOP | Stop | Pause the process; cannot be trapped |
| 25 | SIGTSTP | Terminal stop | Pause the process; can be trapped |
| 26 | SIGCONT | Continue | Run a stopped process |

SIGNULL (0), the null signal, invokes error checking but no signal is actually sent. This can be used to test the validity or existence of *pid*.

SIGTERM (15), the (default) terminate signal, can be trapped by the receiving process, allowing the receiver to execute an orderly shutdown or to ignore the signal entirely. For orderly operations, this is the preferred choice.

SIGKILL (9), the kill signal, forces a process to terminate immediately. Since **SIGKILL** cannot be trapped or ignored, it is useful for terminating a process that does not respond to **SIGTERM**.

The receiving process must belong to the user of the sending process, unless the user has appropriate privileges.

As a single special case, the continue signal **SIGCONT** can be sent to any process that is a member of the same session as the sending process.

RETURN VALUE

Upon completion, **kill** returns with one of the following values:

- 0 At least one matching process was found for each *pid* operand, and the specified signal was successfully processed for at least one matching process.
- >0 An error occurred.

EXAMPLES

The command:

```
kill 6135
```

signals process number 6135 to terminate. This gives the process an opportunity to exit gracefully (removing temporary files, etc.).

The following equivalent commands:

```
kill -s SIGKILL 6135
kill -s KILL 6135
kill -s 9 6135
kill -SIGKILL 6135
kill -KILL 6135
kill -9 6135
```

terminate process number 6135 abruptly by sending a **SIGKILL** signal to the process. This tells the kernel to remove the process immediately.

WARNINGS

If a process hangs during some operation (such as I/O) so that it is never scheduled, it cannot die until it is allowed to run. Thus, such a process may never go away after the kill. Similarly, defunct processes (see *ps(1)*) may have already finished executing, but remain on the system until their parent reaps them (see *wait(2)*). Using **kill** to send signals to them has no effect.

Some non-HP-UX implementations provide **kill** only as a shell built-in command.

DEPENDENCIES

This manual entry describes the external command `/usr/bin/kill` and the built-in **kill** command of the POSIX shell (see *sh-posix(1)*). Other shells, such as C and Korn (see *csh(1)* and *ksh(1)* respectively), also provide **kill** as a built-in command. The syntax for and output from these built-ins may be different.

SEE ALSO

csh(1), *ksh(1)*, *ps(1)*, *sh(1)*, *sh-bourne(1)*, *sh-posix(1)*, *kill(2)*, *wait(2)*, *signal(5)*.

STANDARDS CONFORMANCE

kill: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

kinit - obtain and cache the Kerberos ticket-granting ticket

SYNOPSIS

```
kinit [-l life_time] [-s start_time] [-v] [-p] [-f] [-k [-t keytab_filename]] [-r renewable_life]
[-R] [-c cache_filename] [-S service-name] [principal]
```

DESCRIPTION

kinit obtains and caches an initial ticket-granting ticket for the *principal*.

Options

- l *life_time*** Requests a ticket with the lifetime value defined in *life_time*. The value for *life_time* must be followed immediately by one of the following delimiters:
- | | |
|----------|---------|
| s | seconds |
| m | minutes |
| h | hours |
| d | days |
- For example, as in **kinit -l 90m** for 90 minutes. You cannot mix units; a value of **3h30m** will result in an error.
- If the **-l** option is not specified, the default ticket lifetime (configured by each site) is used. Specifying a ticket lifetime longer than the maximum ticket lifetime (configured by each site) results in a ticket with the maximum lifetime.
- s *start_time*** Requests a postdated ticket, valid starting at *start_time*. The format for *start_time* is the same as the **-l** option, one of the following: seconds, minutes, hours, or days. Postdated tickets are issued with the *invalid* flag set, and need to be fed back to the Kerberos KDC (Key Distribution Center) before use.
- v** Requests that the ticket granting ticket in the cache (with the *invalid* flag set) be passed to the KDC for validation. If the ticket is within its requested time range, the cache is replaced with the validated ticket.
- p** Requests proxiabable tickets.
- f** Requests forwardable tickets.
- r *renewable_life*** Requests renewable tickets, with a total lifetime of *renewable_life*. The format for *renewable_life* is the same as the **-l** option, one of the following: seconds, minutes, hours, or days.
- R** Requests renewal of the ticket-granting ticket. Note that an expired ticket cannot be renewed, even if the ticket is still within its renewable life.
- k [-t *keytab_filename*]** Requests a host ticket, obtained from a key in the local host's keytab file. The name and location of the keytab file may be specified with the **-t *keytab_filename*** option; otherwise the default name and location will be used.
- c *cache_filename*** Uses *cache_filename* as the credentials ticket cache name and location. If this option is not used, the default cache name and location are used.
- The default credentials cache may vary between systems. If the **KRB5CCNAME** environment variable is set, its value is used to name the default ticket cache. Any existing contents of the cache are destroyed by **kinit**.
- S *service_name*** Specifies an alternate service name to use when getting initial tickets.
- principal* Uses the principal name from an existing cache if there is one.

Note

For DCE operations use **/opt/dce/bin/kinit**.

Environment

kinit uses the following environment variable:

KRB5CCNAME Location of the credentials ticket cache.

FILES

`/tmp/krb5cc_{uid}` Default credentials cache. `{uid}` is the decimal UID of the user.
`/etc/krb5.keytab` Default location for the local host's keytab file.

AUTHOR

`kinit` was developed by the Massachusetts Institute of Technology.

SEE ALSO

`kdestroy(1)`, `kerberos(9)`, `klist(1)`, `libkrb5(3)`.

NAME

klist - list cached Kerberos tickets

SYNOPSIS

klist [-e] [[-c] [-f] [-s] [*cache_filename*]] [-k [-t] [-K] [*keytab_filename*]]

DESCRIPTION

klist lists the Kerberos principal and Kerberos tickets held in a credentials cache, or the keys held in a keytab file.

Options

- e Displays the encryption types of the session key and the ticket for each credential in the credential cache, or each key in the keytab file.
- c List tickets held in a credentials cache. This is the default if neither -c nor -k is specified.
- f Shows the flags present in the credentials, using the following abbreviations:
 - F** Forwardable
 - f** forwarded
 - P** Proxiable
 - p** proxy
 - D** postdateable
 - d** postdated
 - R** Renewable
 - I** Initial
 - i** invalid
- s Causes **klist** to run silently (produce no output), but still sets the exit status depending on whether it finds the credentials cache. The exit status is '0' if **klist** finds a credentials cache, and the exit status is '1' if it does not.
- k List keys held in a keytab file.
- t Display the time entry timestamps for each keytab entry in the keytab file.
- K Display the value of the encryption key of the keytab entry in the keytab file.

If *cache_filename* or *keytab_filename* is not specified, **klist** will display the credentials in the default credentials cache or keytab file as appropriate. If the **KRB5CCNAME** environment variable is set, its value is used to name the default ticket cache.

Note

For DCE operations use `/opt/dce/bin/klist`.

Environment

klist uses the following environment variable:

KRB5CCNAME Location of the credentials (ticket) cache.

FILES

- `/tmp/krb5cc_{uid}` Default credentials cache. {uid} is the decimal UID of the user.
- `/etc/krb5.keytab` Default location of the keytab file.

AUTHOR

klist was developed by the Massachusetts Institute of Technology.

SEE ALSO

kdestroy(1), kerberos(9), kinit(1).

NAME

kpasswd - change a user's Kerberos password

SYNOPSIS

kpasswd [*principal*]

DESCRIPTION

The **kpasswd** command is used to change a Kerberos principal's password. **kpasswd** prompts for the current Kerberos password, which is used to obtain a **changepw** ticket from the KDC (Key Distribution Center) for the user's Kerberos realm. If **kpasswd** successfully obtains the **changepw** ticket, the user is prompted twice for the new password, and the password is changed.

If the principal is governed by a policy that specifies the length and/or number of character classes required in the new password, the new password must conform to the policy. The five character classes are lower case, upper case, numbers, punctuation, and all other characters.

Options

principal Changes the password for the Kerberos principal, *principal*. **kpasswd** uses the principal name from an existing cache if there is one. If not, the principal is derived from the identity of the user invoking the **kpasswd** command.

Note

kpasswd looks first for **kpasswd_server = host:port** in the [realms] section of the **krb5.conf** file under the current realm. If that is missing, **kpasswd** looks for the **admin_server** entry, and substitutes 464 for the port.

FILES

/etc/krb5.conf Kerberos configuration file.

AUTHOR

kpasswd was developed by the Massachusetts Institute of Technology.

SEE ALSO

kerberos(9), krb5.conf(4).

NAME

ksh, rksh - shell, the standard/restricted command programming language

SYNOPSIS

ksh [-aefhikmnoprstuvx] [+aefhikmnoprstuvx] [-o *option*] ... [+o *option*] ... [-c *string*] [*arg* ...]

rksh [-aefhikmnoprstuvx] [+aefhikmnoprstuvx] [-o *option*] ... [+o *option*] ... [-c *string*] [*arg* ...]

DESCRIPTION

ksh is a command programming language that executes commands read from a terminal or a file. **rksh** is a restricted version of the command interpreter **ksh**, used to set up login names and execution environments whose capabilities are more controlled than those of the standard shell. See *Invoking ksh* and *Special Commands* sections later in this entry for details about command line options and arguments, particularly the **set** command.

Definitions**metacharacter**

One of the following characters:

; & () | < > *new-line space tab*

blank

A tab or space character.

identifier

A sequence of letters, digits, or underscores starting with a letter or underscore. Identifiers are used as names for **functions** and **named parameters**.

word

A sequence of characters separated by one or more non-quoted metacharacters .

command

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.

special command

A command that is carried out by the shell without creating a separate process. Often called “built-in commands”. Except for documented side effects, most special commands can be implemented as separate utilities.

#

The # character is interpreted as the beginning of a comment. See *Quoting* below.

Commands

A **simple-command** is a sequence of blank-separated words that can be preceded by a parameter 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 **value** of a simple-command 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).

A **pipeline** is a sequence of one or more **commands** separated by |. The standard output of each command except the last is connected by a pipe (see *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 in the pipeline.

A **list** is a sequence of one or more pipelines separated by ;, &, &&, or ||, and optionally terminated by ;, &, or |&. Of these five symbols, ;, &, and |& have equal precedence. && and || have a higher but also equal precedence. A semicolon (;) causes sequential execution of the preceding pipeline; an ampersand (&) causes asynchronous execution of the preceding pipeline (that is, the shell does not 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 (known as a **co-process**). 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 later. The symbol && (|) causes the *list* following it to be executed only if the preceding pipeline returns a zero (non-zero) value. An arbitrary number of new-lines can appear in a *list*, instead of semicolons, to delimit commands.

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 **for** is executed, *identifier* is set to the next *word* taken from the **in** *word* list. If **in** *word* ... is omitted, **for** executes the **do** *list* once for each positional parameter 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 on standard error (file descriptor 2), the set of *words*, each preceded by a number. If **in** *word* ... is omitted, 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 starts with the number of one of the listed *words*, the value of the parameter *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 parameter *identifier* is set to null. The contents of the line read from standard input is saved in the parameter **REPLY**. The *list* is executed for each selection until a **break** or end-of-file (*eof*) is encountered.

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 identical to 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 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, **if** 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, **while** returns a zero exit status; **until** can be used in place of **while** to negate the loop termination test.

(*list*) Execute *list* in a separate environment. If two adjacent open parentheses are needed for nesting, a space must be inserted to avoid arithmetic evaluation as described below.

{ *list* ; } Execute *list*, but not in a separate environment. Note that { is a keyword and requires a trailing blank 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*. Note that [[and]] are keywords and require blanks between them and *expression*.

function *identifier* { *list* ; }*identifier* () { *list* ; }

Define a function referred to by *identifier*. The body of the function is the *list* of commands between { and } (see *Functions* below).

time *pipeline* *pipeline* is executed and the elapsed time, user time, and system time are printed on standard error. Note that the **time** keyword can appear anywhere in the *pipeline* to time the entire *pipeline*. To time a particular command in a *pipeline*, see *time(1)*.

The following keywords are recognized only 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 subsequent 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, is tested for additional aliases. If the last character of the alias value is a **blank**, the word following the alias is also checked for alias substitution. Aliases can be used to redefine special built-in commands, but cannot be used to redefine the keywords 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 subshells but must be reinitialized for separate invocations of the shell (see *Invoking ksh* below).

Aliasing is performed when scripts are read, not while they are executed. Therefore, for it to take effect, **alias** must be executed before the command referring to the alias is read.

Aliases are frequently used as a shorthand for full path names. An option to the aliasing facility allows the value of the alias to be automatically set to the full path name of the corresponding command. These aliases are called **tracked aliases**. The value of a tracked alias is defined the first time the identifier is read and becomes undefined each time the **PATH** variable is reset. These aliases remain tracked so that the next reference redefines the value. Several tracked aliases are compiled into the shell. The **-h** option of the **set** command converts each command name that is an *identifier* into a tracked alias.

The following **exported aliases** are compiled into 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 -'
stop='kill -STOP'
suspend='kill -STOP $$'
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, the word up to a **/** is checked to see if it matches a user name in the **/etc/passwd** file. 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 **~**, alone or before a **/**, is replaced by the value of the **HOME** parameter. A **~** followed by a **+** or **-** is replaced by the value of the parameter **PWD** and **OLDPWD**, respectively. In addition, tilde substitution is attempted when the value of a parameter 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 back single quotes (accent grave) (**`command`**) can 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 (built-ins) that do not perform I/O redirection are carried out without creating a separate process. However, command substitution of a function creates a separate process to execute the function and all commands (built-in or otherwise) in that function.

An arithmetic expression enclosed in double parenthesis preceded by a dollar sign (**\$((expression))**) is replaced by the value of the arithmetic expression within the double parenthesis (see *Arithmetic Evaluation* below for a description of arithmetic expressions).

Parameter Substitution

A **parameter** is an **identifier**, one or more digits, or any of the characters *****, **@**, **#**, **?**, **-**, **\$**, and **!**. A **named parameter** (a parameter denoted by an identifier) has a value and zero or more attributes. Named parameters can be assigned values and attributes by using the **typeset** special command. Attributes supported by **ksh** are described later with the **typeset** special command. Exported parameters pass values and attributes to the environment.

The shell supports a limited one-dimensional array facility. An element of an array parameter 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 named parameter with a valid subscript is legal and an array is created if necessary. Referencing an array without a subscript is equivalent to referencing the first element.

The value of a named parameter can also be assigned by writing:

```
name=value [name=value] ...
```

If the `-i` integer attribute is set for *name*, the *value* is subject to arithmetic evaluation as described below.

Positional parameters, parameters denoted by a number, can 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}` Substitute the value of the parameter, if any. Braces are required when *parameter* is followed by a letter, digit, or underscore that should not be interpreted as part of its name or when a named parameter is subscripted. If *parameter* is one or more digits, it is a positional parameter. A positional parameter of more than one digit must be enclosed in braces. If *parameter* is `*` or `@` all the positional parameters, starting with `$1`, are substituted (separated by a field separator character). If an array *identifier* with subscript `*` or `@` is used, the value for each element is substituted (separated by a field separator character). The shell reads all the characters from `${` to the matching `}` as part of the same word even if it contains braces or metacharacters.

`${#parameter}` If *parameter* is `*` or `@`, the number of positional parameters is substituted. Otherwise, the length of the value of the *parameter* is substituted.

`${#identifier[*]}` Substitute the number of elements in the array *identifier*.

`${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*; then substitute the value of the parameter. Positional parameters cannot 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, a standard message is printed.

`${parameter:+word}` If *parameter* is set and is non-null, substitute *word*; otherwise substitute nothing.

`${parameter#pattern}`
`${parameter##pattern}`

If the shell *pattern* matches the beginning of the value of *parameter*, the value of this substitution is the value of the *parameter* with the matched portion deleted; otherwise the value of this *parameter* substituted. In the former case, the smallest matching pattern is deleted; in the latter case, the largest matching pattern is deleted.

`${parameter%pattern}`
`${parameter%%pattern}`

If the shell *pattern* matches the end of the value of *parameter*, the value of *parameter* with the matched part is deleted; otherwise substitute the value of *parameter*. In the former, the smallest matching pattern is deleted; in the latter, the largest matching pattern is deleted.

In the above, *word* is not evaluated unless it is used as the substituted string. Thus, 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 to determine whether or not *parameter* is set.

The following parameters are set automatically by the shell:

| | |
|---------|--|
| # | The number of positional parameters in decimal. |
| - | Options 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 <i>environment</i> . 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. |
| COLUMNS | If this variable is set, its value is used to define the width of the edit window for the shell edit modes and for printing select lists. In a windowed environment, if the shell detects that the window size has changed, the shell updates the value of COLUMNS . |
| 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. |
| LINES | If this variable is set, the value is used to determine the column length for printing select lists. select lists print vertically until about two-thirds of LINES lines are filled. In a windowed environment, if the shell detects that the window size has changed, the shell updates the value of LINES . |
| 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 parameter is evaluated, 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 parameter is set by the select statement and by the read special command when no arguments are supplied. |
| SECONDS | Each time this parameter is referenced, the number of seconds since shell invocation is returned. If this parameter is assigned a value, the value returned upon reference is the value that was assigned plus the number of seconds since the assignment. |

The following parameters are used by the shell:

| | |
|----------|---|
| CDPATH | The search path for the cd command. |
| EDITOR | If the value of this variable ends in emacs , gmacs , or vi and the VISUAL variable is not set, the corresponding option is turned on (see set in <i>Special Commands</i> below). |
| ENV | If this parameter is set, parameter substitution is performed on the value to generate the path name of the script to be executed when the shell is invoked (see <i>Invoking ksh</i> below). This file is typically used for <i>alias</i> and <i>function</i> definitions. |
| FCEDIT | The default editor name for the fc command. |
| FPATH | The search path for function definitions. This path is searched when a function with the -u attribute is referenced and when a command is not found. If an executable file is found, then it is read and executed in the current environment. |
| IFS | Internal field separators, normally <i>space</i> , <i>tab</i> , and <i>new-line</i> that are used to separate command words resulting from command or parameter substitution, and for separating words with the special command read . The first character of the IFS parameter is used to separate arguments for the "\$*" substitution (see <i>Quoting</i> below). |
| HISTFILE | If this parameter is set when the shell is invoked, its value is the path name of the file that is used to store the command history. The default value is \$HOME/.sh_history . If the user has appropriate privileges and no HISTFILE is given, then no history file is used (see <i>Command Re-entry</i> below). |
| HISTSIZE | If this parameter is set when the shell is invoked, the number of previously entered commands accessible to 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. |

| | |
|------------------|--|
| 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 arrival of mail in the specified file. |
| MAILCHECK | This variable specifies how often (in seconds) the shell checks for changes in the modification time of any of the files specified by the MAILPATH or MAIL parameters. The default value is 600 seconds. When the time has elapsed the shell checks before issuing the next prompt. |
| MAILPATH | A list of file names separated by colons (:). If this parameter is set, 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 to be printed, in which case the message undergoes parameter and command substitution with the parameter \$_ defined as the name of the changed file. The default message is you have mail in \$_ . |
| PATH | The search path for commands (see <i>Execution</i> below). The user cannot change PATH if executing rksh (except in the .profile file). |
| PS1 | The value of this parameter is expanded for parameter substitution, to define the primary prompt string which, by default, is \$ followed by a space character. The character ! in the primary prompt string is replaced by the command number (see <i>Command Re-entry</i> below). To include a ! in the prompt, use !! . |
| PS2 | Secondary prompt string, by default > followed by a space character. |
| PS3 | Selection prompt string used within a select loop, by default #? followed by a space character. |
| PS4 | The value of this variable is expanded for parameter substitution and precedes each line of an execution trace. If PS4 is unset, the execution trace prompt is + followed by a space character. |
| SHELL | The path name of the shell is kept in the environment. When invoked, the shell is restricted if the value of this variable contains an r in the basename. |
| TMOUT | If set to a value greater than zero, the shell terminates if a command is not entered within the prescribed number of seconds after issuing the PS1 prompt. |
| VISUAL | Invokes the corresponding option when the value of this variable ends in <i>emacs</i> , <i>gmacs</i> , or <i>vi</i> (see set in <i>Special Commands</i> below). |

The shell gives default values to **PATH**, **PS1**, **PS2**, **MAILCHECK**, **TMOUT**, and **IFS**. **HOME**, **SHELL**, **ENV**, and **MAIL** are never set automatically by the shell (although **HOME**, **SHELL**, and **MAIL** are set by *login*(1)).

Blank Interpretation

After parameter and command substitution, the results of substitution are scanned for field separator characters (found in **IFS**), and split into distinct arguments where such characters are found. **ksh** retains explicit null arguments (or **' '**) but removes implicit null arguments (those resulting from *parameters* that have no values).

File Name Generation

Following substitution, each command *word* is processed as a pattern for file name expansion unless the **-f** option has been **set**. The form of the patterns is the Pattern Matching Notation defined by *regex*(5). The word is replaced with sorted file names matching the pattern. If no file name is found that matches the pattern, the word is left unchanged.

In addition to the notation described in *regex*(5), **ksh** recognizes composite patterns made up of one or more pattern lists 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 can be **quoted** (i.e., made to stand for itself) by preceding it with a ****. The pair **\new-line** is ignored. 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 occurs and \ quotes the characters \, `, ", and \$. \$* and \$@ have identical meanings 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 "\$1\$d\$2d...", where *d* is the first character of the IFS parameter, whereas "\$@" is equivalent to "\$1" "\$2" Inside back single quote (accent grave) marks (` `) \ quotes the characters \, `, and \$. If the back single quotes occur within double quotes, \ also quotes the character " .

The special meaning of keywords or aliases can be removed by quoting any character of the keyword. The recognition of function names or special command names listed below cannot be altered by quoting them.

Arithmetic Evaluation

The ability to perform integer arithmetic is provided with the special command **let**. Evaluations are performed using long arithmetic. Constants take 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, base 10 is used.

An arithmetic expression uses the same syntax, precedence, and associativity of expression of 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 beginning with ((, all characters until the matching)) are treated as a quoted expression. More precisely, ((...)) is equivalent to **let** "...".

Prompting

When used interactively, the shell prompts with the 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, the secondary prompt (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:

| | |
|-------------------------------|---|
| -a <i>file</i> | True if <i>file</i> exists. |
| -b <i>file</i> | True if <i>file</i> exists and is a block special file. |
| -c <i>file</i> | True if <i>file</i> exists and is a character special file. |
| -d <i>file</i> | True if <i>file</i> exists and is a directory. |
| -f <i>file</i> | True if <i>file</i> exists and is an ordinary file. |
| -g <i>file</i> | True if <i>file</i> exists and is has its setgid bit set. |
| -h <i>file</i> | True if <i>file</i> exists and is a symbolic link. |
| -k <i>file</i> | True if <i>file</i> exists and is has its sticky bit set. |
| -n <i>string</i> | True if length of <i>string</i> is non-zero. |
| -o <i>option</i> | True if option named <i>option</i> is on. |
| -p <i>file</i> | True if <i>file</i> exists and is a fifo special file or a pipe. |
| -r <i>file</i> | True if <i>file</i> exists and is readable by current process. |
| -s <i>file</i> | True if <i>file</i> exists and has size greater than zero. |
| -t <i>fildev</i> | True if file descriptor number <i>fildev</i> is open and associated with a terminal device. |
| -u <i>file</i> | True if <i>file</i> exists and is has its setuid bit set. |
| -w <i>file</i> | True if <i>file</i> exists and is writable by current process. |
| -x <i>file</i> | True if <i>file</i> exists and is executable by current process. If <i>file</i> exists and is a directory, the current process has permission to search in the directory. |
| -z <i>string</i> | True if length of <i>string</i> is zero. -L <i>file</i> True if <i>file</i> exists and is a symbolic link. |
| -O <i>file</i> | True if <i>file</i> exists and is owned by the effective user ID of this process. |
| -G <i>file</i> | True if <i>file</i> exists and its group matches the effective group ID of this process. |
| -S <i>file</i> | True if <i>file</i> exists and is a socket. |
| <i>file1</i> -nt <i>file2</i> | True if <i>file1</i> exists and is newer than <i>file2</i> . |

| | |
|---|--|
| <i>file1</i> -ot <i>file2</i> | True if <i>file1</i> exists and is older than <i>file2</i> . |
| <i>file1</i> -ef <i>file2</i> | True if <i>file1</i> and <i>file2</i> exist and refer to the same file. |
| <i>string</i> = <i>pattern</i> | True if <i>string</i> matches <i>pattern</i> . |
| <i>string</i> != <i>pattern</i> | True if <i>string</i> does not match <i>pattern</i> . |
| <i>string1</i> < <i>string2</i> | True if <i>string1</i> comes before <i>string2</i> based on ASCII value of their characters. |
| <i>string1</i> > <i>string2</i> | True if <i>string1</i> comes after <i>string2</i> based on ASCII value of their characters. |
| <i>exp1</i> -eq <i>exp2</i> | True if <i>exp1</i> is equal to <i>exp2</i> . |
| <i>exp1</i> -ne <i>exp2</i> | True if <i>exp1</i> is not equal to <i>exp2</i> . |
| <i>exp1</i> -lt <i>exp2</i> | True if <i>exp1</i> is less than <i>exp2</i> . |
| <i>exp1</i> -gt <i>exp2</i> | True if <i>exp1</i> is greater than <i>exp2</i> . |
| <i>exp1</i> -le <i>exp2</i> | True if <i>exp1</i> is less than or equal to <i>exp2</i> . |
| <i>exp1</i> -ge <i>exp2</i> | True if <i>exp1</i> is greater than or equal to <i>exp2</i> . |

A compound expression can be constructed from these primitives by using any of the following, listed in decreasing order of precedence.

| | |
|---|---|
| (expression) | True, if <i>expression</i> is true. Used to group expressions. |
| ! expression | True if <i>expression</i> is false. |
| <i>expression1</i> && <i>expression2</i> | True, if <i>expression1</i> and <i>expression2</i> are both true. |
| <i>expression1</i> <i>expression2</i> | True, if either <i>expression1</i> or <i>expression2</i> is true. |

Input/Output

Before a command is executed, its input and output can be redirected using a special notation interpreted by the shell. The following can appear anywhere in a simple-command or can precede or follow a command and are not passed on to the invoked command. Command and parameter substitution occurs 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 <i>word</i> as standard input (file descriptor 0). |
| >word | Use file <i>word</i> as standard output (file descriptor 1). If the file does not exist, it is created. If the file exists, and the noclobber option is on, an error occurs; otherwise, the file is truncated to zero length. |
| > word | Sames as > , except that it overrides the noclobber option. |
| >>word | Use file <i>word</i> as standard output. If the file exists, output is appended to it (by first searching for the end-of-file); otherwise, the file is created. |
| <>word | Open file <i>word</i> for reading and writing as standard input. If the file does not exist it is created. |
| <<[-]word | The shell input is read up to a line that matches <i>word</i> , or to an end-of-file. No parameter substitution, command substitution, or file name generation is performed on <i>word</i> . The resulting document, called a here-document , becomes the standard input. If any character of <i>word</i> is quoted, no interpretation is placed upon the characters of the document. Otherwise, parameter and command substitution occurs, <i>\new-line</i> is ignored, and <i>\</i> must be used to quote the characters <i>\</i> , <i>\$</i> , <i>`</i> , and the first character of <i>word</i> . If <i>-</i> is appended to << , all leading tabs are stripped from <i>word</i> and from the document. |
| <&digit | The standard input is duplicated from file descriptor <i>digit</i> (see <i>dup(2)</i>). |
| >&digit | The standard output is duplicated to file descriptor <i>digit</i> (see <i>dup(2)</i>). |
| <&- | The standard input is closed. |
| >&- | The standard output is closed. |
| <&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, the file descriptor number cited 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.

Redirection order is significant because the shell evaluates redirections referencing file descriptors in terms of the currently open file associated with the specified file descriptor at the time of evaluation. For example:

```
... 1>fname 2>&1
```

first assigns file descriptor 1 (standard output) to file *fname*, then assigns file descriptor 2 (standard error) to the file assigned to file descriptor 1; i.e., *fname*. On the other hand, if the order of redirection is reversed as follows:

```
... 2>&1 1>fname
```

file descriptor 2 is assigned to the current standard output (user terminal unless a different assignment is inherited). File descriptor 1 is then reassigned to file *fname* without changing the assignment of file descriptor 2.

The input and output of a *co-process* can be moved to a numbered file descriptor allowing other commands to write to them and read from them using the above redirection operators. If the input of the current *co-process* is moved to a numbered file descriptor, another *co-process* can be started.

If a command is followed by **&** and job control is inactive, 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 passed to an executed program much like a normal argument list. The names must be **identifiers** and the values are character strings. The shell interacts with the environment in several ways. When invoked, the shell scans the environment and creates a parameter for each name found, gives it the corresponding value, and marks it *export*. Executed commands inherit the environment. If the user modifies the values of these parameters or creates new ones by using the **export** or **typeset -x** commands, the values 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 can 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 can be augmented by prefixing it with one or more parameter assignments. A parameter assignment argument takes the form *identifier=value*. For example,

```
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 by a percent sign).

If the **-k** option is set, all parameter assignment arguments are placed in the environment, even if they occur after the command name. The following echo statement prints **a=b c**. After the **-k** option is set, the second echo statement prints only **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** keyword (described in the *Commands* section above) is used to define shell functions. Shell functions are read 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 except that command substitution of a function creates a new process. Functions share all files and present working directory with the caller. Traps caught by the caller are reset to their default action inside the function. If a function does not catch or specifically ignore a trap condition, the function terminates and the condition is 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.

Function identifiers can be listed with the **+f** option of the **typeset** special command. Function identifiers and the associated text of the functions can be listed with the **-f** option. 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 reinvoking the shell. Functions that must be defined across separate invocations of the shell should be placed in the **ENV** file.

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 resembling:

```
[1] 1234
```

indicating that job number 1 was started asynchronously and had one (top-level) process whose process ID was 1234.

If you are running a job and want to do something else, type the suspend character (usually **^Z** (Ctrl-Z)) to send a STOP signal to the current job. The shell then indicates that the job has been 'Stopped', and prints another prompt. The state of this job can be manipulated by using the **bg** command to put it in the background, running other commands (while it is stopped or running in the background), and eventually restarting or returning the job to the foreground by using the **fg** command. A **^Z** takes effect immediately and resembles an interrupt, since pending output and unread input are discarded when **^Z** is typed.

A job run in the background stops if it tries to read from the terminal. Background jobs normally are allowed to produce output, but can be disabled by giving the **stty tostop** command. If the user sets this **tty** option, background jobs stop when trying to produce output.

There are several ways to refer to jobs in the shell. A job can be referred to by the process ID of any process in 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 when a process changes state. It informs the user when a job is blocked and prevented from further progress, but only just before it prints a prompt.

When the monitor mode is on, each background job that completes triggers any trap set for **CHLD**.

Attempting to leave the shell while jobs are running or stopped produces the warning, **You have stopped (running) jobs**. Use the **jobs** command to identify them. An immediate attempt to exit again terminates the stopped jobs; the shell does not produce a warning the second time.

Signals

The **INT** and **QUIT** signals for an invoked command are ignored if the command is followed by **&** and the **monitor** option is off. 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

Substitutions are made each time a command is executed. If the command name matches one of the *Special Commands* listed below, it is executed within the current shell process. Next, **ksh** checks the command name to determine whether it matches one of the user-defined functions. If it does, **ksh** saves the positional parameters and then sets them to the arguments of the *function* call. The positional parameter 0 is set to the function name. When the *function* completes or issues a **return**, **ksh** restores the positional parameter list and executes any trap set on **EXIT** within the function. The value of a *function* is the value of the last command executed. A function is executed in the current shell process. If a command name is not a **special command** or a user-defined **function**, **ksh** creates a process and attempts to

execute the command using **exec** (see *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:** (specifying **/usr/bin** and the current directory in that order). Note that the current directory is specified by a null path name which can appear immediately after the equals sign, between colon delimiters, or at the end of the path list. The search path is not used if the command name contains a /. Otherwise, each directory in the path is searched for an executable file. If the file has execute permissions but is not a directory or an executable object code file, it is assumed to be a script file, which is a file of data for an interpreter. If the first two characters of the script file are **#!**, **exec** (see *exec(2)*) expects an interpreter path name to follow. **exec** then attempts to execute the specified interpreter as a separate process to read the entire script file. If a call to **exec** fails, **/usr/bin/ksh** is spawned to interpret the script file. All non-exported aliases, functions, and named parameters are removed in this case. If the shell command file does not have read permission, or if the *setuid* and/or *setgid* bits are set on the file, the shell executes an agent to set up the permissions and execute the shell with the shell command file passed down as an open file. A parenthesized command is also 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 writable. A shell can access the commands of all **interactive** shells that 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 no editor program is specified as an argument to **fc**, the value of the **FCEDIT** parameter is used. If **FCEDIT** is not defined, **/usr/bin/ed** is used. The edited command 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 -**, typing **r bad=good c** re-executes the most recent command that starts with the letter **c** and replaces the first occurrence of the string **bad** with the string **good**.

The history file will be trimmed when all of the following conditions occurs:

Its size is greater than four kilobytes.

The number of commands in it is more than **HISTSIZE**.

The file has not been modified in the last ten minutes.

The user has write permission for the directory in which the history file resides.

If any one of the above conditions does not occur, the history file will not be trimmed. When the history file is trimmed, the latest **HISTSIZE** commands will be available in the history file.

Special Commands

The following simple-commands are executed in the shell process. They permit input/output redirection. Unless otherwise indicated, file descriptor 1 is the default output location and the exit status, when there are no syntax errors, is zero. Commands that are preceded by **%** or **%%** 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. A zero exit code is returned.

% . file [arg ...]

Read and execute commands from *file* and return. 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. It is not necessary that the execute permission bit be set for *file*.

%% 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** option is used to set and list tracked aliases. The value of a tracked alias is the full path name corresponding to the given *name*. The value of a tracked alias becomes undefined when the value of **PATH** is reset, but the alias remains tracked. Without the **-t** option, for each *name* in the argument list for which no *value* is given, the name and value of the alias is printed. The **-x** option is used to set or print exported aliases. An exported alias is defined across sub-shell environments. Alias returns true unless a *name* is given for which no alias has been defined.

bg [job ...] Puts the specified *jobs* into the background. The current job is put in the background if *job* is unspecified. See *Jobs* for a description of the format of *job*.

% break [n] Exit from the enclosing **for**, **while**, **until**, or **select** loop, if any. If *n* is specified, break *n* levels.

% continue [n]
Resume the next iteration of the enclosing **for**, **while**, **until**, or **select** loop. If *n* is specified, resume at the *n*-th enclosing loop.

cd [-L|-P] [arg]
cd old new
This command can take 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 **-L** option (default) preserves logical naming when treating symbolic links. **cd -L ..** moves the current directory one path component closer to the root directory. The **-P** option preserves the physical path when treating symbolic links. **cd -P ..** changes the working directory to the parent directory of the current directory. The shell parameter **HOME** is the default *arg*. The parameter **PWD** is set to the current directory. The shell parameter **CDPATH** defines the search path for the directory containing *arg*. Alternative directory names are separated by a colon (:). If **CDPATH** is null or undefined, the default value is 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 /, the search path is not used. Otherwise, each directory in the path is searched for *arg*. See also **cd(1)**.

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 cannot be executed by **rksh**.

echo [arg ...]
See *echo(1)* for usage and description.

% eval [arg ...]
Reads the arguments as input to the shell and executes the resulting command(s).

% exec [arg ...]
Parameter assignments remain in effect after the command completes. 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 can 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 opened with this mechanism are closed when invoking another program.

% exit [n]
Causes the 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 end-of-file also causes the shell to exit, except when a shell has the *ignoreeof* option set (see *set* below).

%% export [name [=value] ...]
The given *names* are marked for automatic export to the *environment* of subsequently executed commands.

fc [-eename] [-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 typed at the terminal. The arguments *first* and *last* can be specified as a

number or string. A given string is used to locate the most recent command. A negative number is used to offset the current command number. The **-l** option causes the commands to be listed on standard output. Otherwise, the editor program *ename* is invoked on a file containing these keyboard commands. If *ename* is not supplied, the value of the parameter **FCEDIT** (default **/usr/bin/ed**) is used as the editor. Once editing has ended, the commands (if any) are executed. If *last* is omitted, only the command specified by *first* is used. If *first* is not specified, the default is the previous command for editing and **-16** for listing. The **-r** option reverses the order of the commands and the **-n** option suppresses command numbers when listing. In the latter, the *command* is re-executed after the substitution *old=new* is performed.

fg [*job* ...] Brings each *job* into the foreground in the order specified. If no *job* is specified, 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 **+** preceding it 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 also *getopts*(1).

jobs [**-lnp**] [*job* ...]

Lists information about each given job; or all active jobs if *job* is omitted. The **-l** option lists process ids in addition to the normal information. The **-n** option only displays jobs that have stopped or exited since last notified. The **-p** option causes only the process group to be listed. See *Jobs* for a description of the format of *job*.

kill [**-sig**] *process* ...

Sends either the **TERM** (terminate) signal or the specified signal to the specified jobs or processes. Signals are given either by number or name (as given in *signal*(5), stripped of the prefix **SIG**). The signal names are listed by **kill -l**. No default exists; merely typing **kill** does not affect the current job. If the signal being sent is **TERM** (terminate) or **HUP** (hangup), the job or process is sent a **CONT** (continue) signal when stopped. The *process* argument can be either a process ID or job. If the first argument to **kill** is a negative integer, it is interpreted as a *sig* argument and not as a process group. See also *kill*(1).

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 newgrp arg**...

print [**-Rnrpsu** [*n*]] [*arg* ...]

The shell output mechanism. With no options or with option **-** or **--** the arguments are printed on standard output as described by *echo*(1). Raw mode, **-R** or **-r**, ignores the escape conventions of *echo*. The **-R** option prints all subsequent arguments and options other than **-n**. The **-p** option causes the arguments to be written onto the pipe of the process spawned with **|&** instead of standard output. The **-s** option causes the arguments to be written onto the history file instead of standard output. The **-u** option can be used to specify a one-digit file descriptor unit number *n* on which the output is to be placed. The default is 1. If the option **-n** is used, no new-line character is added to the output.

pwd [**-L** | **-P**]

With no arguments prints the current working directory (equivalent to **print -r - \$PWD**). The **-L** option (default) preserves the logical meaning of the current directory and **-P** preserves the physical meaning of the current directory if it is a symbolic link. See the special **cd** command, *cd*(1), *ln*(1), and *pwd*(1).

read [-prsu[*n*]] [*name*] [?*prompt*] [*name* ...]

The shell input mechanism. One line is read and broken up into words using the characters in **IFS** as separators. In **-r** raw mode, \ at the end of a line does not signify line continuation. The first word is assigned to the first *name*, the second word to the second *name*, etc., with remaining words 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** option is present, the input is saved as a command in the history file. The option **-u** can be used to specify a one-digit file descriptor unit 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, **REPLY** is used as the default *name*. The return code is 0, unless an end-of-file is encountered. An end-of-file with the **-p** option causes cleanup for this process so that another process can be spawned. If the first argument contains a ?, the remainder of this word is used as a **prompt** when the shell is interactive. If the given file descriptor is open for writing and is a terminal device, the prompt is placed on this unit. Otherwise the prompt is issued on file descriptor 2. The return code is 0, unless an end-of-file is encountered. See also *read(1)*.

%% **readonly** [*name*[=*value*] ...]

The given *names* are marked read-only and these names cannot be changed by subsequent assignment.

% **return** [*n*]

Causes a shell **function** to return to the invoking script with the return status specified by *n*. If *n* is omitted, the return status is that of the last command executed. Only the low 8 bits of *n* are passed back to the caller. If **return** is invoked while not in a **function** or executing a script by the . (dot) built-in command, it has the same effect as an **exit** command.

set [±*aefhkmnopstuvx* | ±*o option*] ... [±*A name*] [*arg* ...]

The following options are used for this command:

- 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 defined parameters are automatically exported.
- e If the shell is non-interactive and if a command fails, execute the **ERR** trap, if set, and exit immediately. This mode is disabled while reading profiles.
- f Disables file name generation.
- h Each command whose name is an **identifier** becomes a tracked alias when first encountered.
- k All parameter assignment arguments (not just those that precede the command name) are placed in the environment for a command.
- m Background jobs are run in a separate process group and a line is printed upon completion. The exit status of background jobs is reported in a completion message. This option is turned on automatically for interactive shells.
- n Read commands and check them for syntax errors, but do not execute them. The **-n** option is ignored for interactive shells.
- o The **-o** argument takes any of several *option* names, but only one *option* can be specified with each **-o** option. If none is supplied, the current option settings are printed. The **-o** argument *option* names follow:

| | |
|------------------|---|
| allexport | Same as -a. |
| bgnice | All background jobs are run at a lower priority. |
| errexit | Same as -e. |
| emacs | Activates an emacs -style in-line editor for command entry. |
| gmacs | Activates a gmacs -style in-line editor for command entry. |
| ignoreeof | The shell does 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. Requires > 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 | Activates the insert mode of a vi -style in-line editor until you press the ESC key which puts you in move mode. A return sends the line. |
| viraw | Each character is processed as it is typed in vi mode. |
| xtrace | Same as -x . |
| -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 (gid) is not equal to the real uid (gid). Turning this off causes the effective uid and gid to be set to the real uid and gid. |
| -s | Sort the positional parameters. |
| -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 options and stops examining arguments for options. |
| -- | Do not change any of the options; useful in setting \$1 to a value beginning with - . If no arguments follow this option, the positional parameters are unset. |

Using **+** instead of **-** before a option causes the option to be turned off. These options can also be used when invoking the shell. The current set of options can be examined by using **\$-**.

Unless **-A** is specified, the remaining *arg* arguments are positional parameters and are assigned consecutively to **\$1**, **\$2**, If neither arguments nor options are given, the values of all names are printed on the standard output.

% shift [*n*] The positional parameters from **\$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 **\$#**.

test [*expr*] Evaluate conditional expression *expr*. See *test(1)* for usage and description. The arithmetic comparison operators are not restricted to integers. They allow any arithmetic expression. Four additional primitive expressions are allowed:

| | |
|--------------------------------------|--|
| -L <i>file</i> | True if <i>file</i> is a symbolic link. |
| <i>file1</i> -nt <i>file2</i> | True if <i>file1</i> is newer than <i>file2</i> . |
| <i>file1</i> -ot <i>file2</i> | True if <i>file1</i> is older than <i>file2</i> . |
| <i>file1</i> -ef <i>file2</i> | True if <i>file1</i> has the same device and i-node number as <i>file2</i> . |

% 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 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 name of the signal. Trap commands are executed in signal number order. Any attempt to set a trap on a signal that was ignored upon entering the current shell has no effect. If *arg* is omitted or is **-**, all traps for *sig* are reset to their original values. If *arg* is the null string, this signal is ignored by the shell and by the commands it invokes. If *sig* is **DEBUG**, *arg* is executed after each command. If *sig* is **ERR**, *arg* is executed whenever a command has a non-zero exit code. If *sig* is **0** or **EXIT** and the **trap** statement is executed inside the body of a function, the command *arg* is executed after the function completes. If *sig* is **0** or **EXIT** for a **trap** set outside any function, 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 [**±LRZfilrtux** [*n*]] [*name*[= *value*]] ...

Parameter assignments remain in effect after the command completes. When invoked

inside a function, a new instance of the parameter *name* is created. The parameter value and type are restored when the function completes. The following list of attributes can be specified:

- 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 *name* is assigned, the value is filled on the right with blanks or truncated, if necessary, to fit into the field. Leading zeros are removed if the -Z option is also set. The -R option 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 parameter is reassigned. The -L option is turned off.
- Z Right justify and fill with leading zeros if the first non-blank character is a digit and the -L option 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 Cause *name* to refer to function names rather than parameter names. No assignments can be made to the *name* declared with the **typeset** statement. The only other valid options are -t (which turns on execution tracing for this function) and -x (which allows the function to remain in effect across shell procedures executed in the same process environment).
- 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 Convert all uppercase characters to lowercase. The uppercase -u option is turned off.
- r Any given *name* is marked "read only" and cannot be changed by subsequent assignment.
- t Tag the named parameters. Tags are user definable and have no special meaning to the shell.
- u Convert all lowercase characters to uppercase characters. The lowercase -l option is turned off.
- x Mark any given *name* for automatic export to the environment of subsequently executed commands.

Using + instead of - causes these options to be turned off. If no *name* arguments are given but options are specified, a list of names (and optionally the values) of the parameters that have these options set is printed. Using + instead of - retains the values to be printed. If neither names nor options are given, the names and attributes of all parameters are printed.

ulimit [*n*] If *n* is given, impose a size limit of *n* 512 byte blocks on files written by child processes (files of any size can be read). If *n* is not given, the current limit is printed.

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. See also **umask**(1).

unalias *name* ...

The parameters given by the list of *names* are removed from the *alias* list.

unset [-f] *name* ...

The parameters given by the list of *names* are unassigned; that is, their values and attributes are erased. Read-only variables cannot be unset. If the -f option is set, *names* refer to function names. Unsetting **ERRNO**, **LINENO**, **MAILCHECK**, **OPTARG**, **OPTIND**, **RANDOM**, **SECONDS**, **TMOU**T, and **_** removes their special meaning even if they are subsequently assigned to.

% wait [*job*] Wait for the specified *job* to terminate or stop, and report its status. This status becomes the return code for the **wait** command. If *job* is not given, **wait** waits for all currently active child processes to terminate or stop. The termination status returned is that of the last process. See *Jobs* for a description of the format of a *job*.

whence [-pv] *name* ...

For each *name*, indicate how it would be interpreted if used as a command name. The **-v** option produces a more verbose report. The **-p** option does a path search for *name* even if *name* is an alias, a function, or a reserved word.

Invoking ksh

If the shell is invoked by **exec** (see *exec(2)*), and the first character of argument zero (\$0) is -, the shell is assumed to be a login shell and commands are read first from **/etc/profile**. The expression **\${HOME:-.}/.profile** is then evaluated and an attempt to open the resulting filename is made. If the file is opened successfully, the file is read. Next, commands are read from the file named by performing parameter substitution on the value of the environment parameter **ENV**, if the file exists. If the **-s** option is not present and *arg* is, a path search is performed on the first *arg* to determine the name of the script to execute. When running **ksh** with *arg*, the script *arg* must have read permission and any *setuid* and *getgid* settings are ignored. Commands are then read as described below. The following options are interpreted by the shell when it is invoked:

- c** *string* If the **-c** option is present, commands are read from *string*.
- s** If the **-s** option is present or if no arguments remain, commands are read from the standard input. Shell output, except for the output of some of the *Special Commands* listed above, is written to file descriptor 2.
- i** If the **-i** option is present or if the shell input and output are attached to a terminal, the shell is interactive. In this case SIGTERM is ignored (so that **kill 0** does not kill an interactive shell) and SIGINT +1 is caught and ignored (so that **wait** is interruptible). In all cases, SIGQUIT is ignored by the shell. (See *signal(5)*.)
- r** If the **-r** option is present, the shell is a restricted shell.

The remaining options and arguments are described under the **set** command above.

rksh Only

rksh is used to set up login names and execution environments where 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 forbidden:

- Changing directory (see *cd(1)*)
- Setting the value of **SHELL**, **ENV**, or **PATH**
- Specifying path or command names containing /
- Redirecting output (>, >|, <>, and >>)

The restrictions above are enforced after the **.profile** and **ENV** files are interpreted.

When a command to be executed is found to be a shell procedure, **rksh** invokes **ksh** to execute it. Thus, the end-user is provided with shell procedures accessible to the full power of the standard shell, while being restricted to a limited menu of commands. This scheme assumes that the end-user does not have write and execute permissions in the same directory.

When a shell procedure is invoked from **rksh**, the shell interpreter specified with the **#!** magic inherits all the restricted features of **rksh**. So, the shell procedures written for execution under **rksh** with the intent of utilizing the full power of the standard shell should not specify an interpreter with **#!**.

These rules effectively give the writer of the **.profile** file complete control over user actions, by performing guaranteed set-up actions and leaving the user in an appropriate directory (probably not the login directory).

The system administrator often sets up a directory of commands (usually **/usr/rbin**) that can be safely invoked by **rksh**. HP-UX systems provide a restricted editor **red** (see *ed(1)*), suitable for restricted users.

COMMAND-LINE EDITING

In-line Editing Options

Normally, each command line typed at a terminal device is followed by a new-line (carriage-return or line-feed). If either the **emacs**, **gmacs**, or **vi** option is set, the user can edit the command line. 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 character must overwrite the current character on the screen. ADM terminal users should set the "space/advance" switch to "space". Hewlett-Packard terminal users should set the straps to "bcGHxZ etX".

The editing modes enable the user to look through a window at the current line. The default window width is 80, unless the value of `COLUMNS` is defined. If the line is longer than the window width minus two, a mark displayed at the end of the window notifies the user. The mark is a `>`, `<`, or `*` if the line extends respectively on the right, left, or both side(s) of the window. As the cursor moves and reaches the window boundaries, the window is centered about the cursor.

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 invoked by either the `emacs` or `gmacs` option. Their sole difference is their handling of `^T`. To edit, the user moves the cursor to the point needing correction and inserts or deletes characters or words. All editing commands are control characters or escape sequences. The notation for control characters is circumflex (`^`) followed by the character. For example, `^F` is the notation for Ctrl-F. This is entered by pressing the `f` key while holding down the Ctrl (control) key. The Shift key is *not* pressed. (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 only 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 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. |
| <code>erase</code> | (User defined erase character as defined by the <code>stty(1)</code> command, usually <code>^H</code> or <code>#</code> .) Delete previous character. |
| <code>^D</code> | Delete current character. |
| <code>eof</code> | End-of-file character, normally <code>^D</code> , terminates the shell if the current line is null. |
| <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 interrupt character is <code>^?</code> (DEL, the default) this command does not work). |
| <code>^T</code> | Transpose current character with next character in <code>emacs</code> mode. Transpose two previous characters in <code>gmacs</code> mode. |
| <code>^C</code> | Capitalize current character. |
| <code>M-c</code> | Capitalize current word. |
| <code>M-l</code> | Change the current word to lowercase. |
| <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, delete from the given position up to the cursor. If preceded by a numerical parameter whose value is greater than the current cursor position, from the cursor up to the given position. |
| <code>^W</code> | Kill from the cursor to the mark. |
| <code>M-p</code> | Push the region from the cursor to the mark on the stack. |
| <code>kill</code> | (User-defined kill character, as defined by the <code>stty(1)</code> command, usually <code>^G</code> or <code>@</code> .) Kill the entire current line. If two <code>kill</code> characters are entered in succession, all subsequent consecutive kill characters cause a line feed (useful when using paper terminals). |
| <code>^Y</code> | Restore last item removed from line (yank item back to the line). |
| <code>^L</code> | Line feed and print current line. |
| <code>^@</code> | (Null character) Set mark. |
| <code>M-space</code> | (Meta space) Set mark. |
| <code>^J</code> | (New line) Execute the current line. |
| <code>^M</code> | (Return) Execute the current line. |

| | |
|-----------------|--|
| ^P | Fetch previous command. Each time ^P is entered, the next previous command in the history list is accessed. |
| ^N | Fetch next command. Each time ^N is entered the next command in the history list is accessed. |
| M-< | Fetch the least recent (oldest) history line. |
| M-> | Fetch the most recent (youngest) history line. |
| ^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 <i>string</i> is preceded by a ^ , the matched line must begin with <i>string</i> . If <i>string</i> is omitted, 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 from the history file the next line relative to current line. |
| 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-b , M-c , M-d , M-f , M-h , M-l and M-^H . |
| M-letter | Softkey. User's alias list is searched for an alias by the name <i>_letter</i> and if an alias of this name is defined, its value is inserted on the input queue. This <i>letter</i> must not be one of the above meta-functions. |
| 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-__ | Same as M-. |
| M-* | Attempt file-name generation on the current word. |
| 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 as 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 can 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. Entering a command puts you into **input** mode. To edit, the user enters **control** mode by pressing ESC and moves the cursor to the point needing correction, then inserts or deletes characters or words. Most control commands accept an optional repeat *count* prior to the command.

In vi mode on most systems, canonical processing is initially enabled and the command is echoed again if the speed is 1200 baud or greater and contains any control characters, or if 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.

Setting the **viraw** option always disables canonical processing on the terminal. This mode is implicit for systems that do not support two alternate end-of-line delimiters, and can be helpful for certain terminals.

Input Edit Commands

By default the editor is in input mode.

| | |
|--------------|---|
| <i>erase</i> | Delete previous character. (<i>erase</i> is a user-defined erase character, as defined by the <i>stty</i> (1) command, usually ^H or # .) |
| ^W | Delete the previous blank separated word. |
| ^D | Terminate the shell. |
| ^V | Escape next character. Editing characters, erase or kill characters can be entered in a command line or in a search string if preceded by a ^V . ^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 move the cursor. The designation *[count]* causes a repetition of the command the cited number of times.

| | |
|----------------------------------|--|
| <i>[count]</i> l | Cursor forward (right) one character. |
| <i>[count]</i> w | Cursor forward one alphanumeric word. |
| <i>[count]</i> W | Cursor to the beginning of the next word that follows a blank. |
| <i>[count]</i> e | Cursor to end of word. |
| <i>[count]</i> E | Cursor to end of the current blank-delimited word. |
| <i>[count]</i> h | Cursor backward (left) one character. |
| <i>[count]</i> b | Cursor backward one word. |
| <i>[count]</i> B | Cursor to preceding blank separated word. |
| <i>[count]</i> | Cursor to column <i>count</i> . Default is 1. |
| <i>[count]</i> f <i>c</i> | Find the next character <i>c</i> in the current line. |
| <i>[count]</i> F <i>c</i> | Find the previous character <i>c</i> in the current line. |
| <i>[count]</i> t <i>c</i> | Equivalent to f followed by <i>h</i> . |
| <i>[count]</i> T <i>c</i> | Equivalent to F followed by <i>l</i> . |
| <i>[count]</i> ; | Repeats the last single character find command, f , F , t , or T . |
| <i>[count]</i> , | Reverses the last single character find command. |
| 0 | Cursor to start of line. |
| ^ | Cursor to first nonblank character in line. |
| \$ | Cursor to end of line. |

Search Edit Commands

These commands access your command history.

| | |
|-------------------------|---|
| <i>[count]</i> k | Fetch previous command. Each time k is pressed, the next earlier command in the history list is accessed. |
| <i>[count]</i> - | Equivalent to k . |
| <i>[count]</i> j | Fetch next command. Each time j is entered, the next later command in the history list is accessed. |
| <i>[count]</i> + | Equivalent to j . |
| <i>[count]</i> G | The command number <i>count</i> is fetched. The default is the first command in the history list. |
| / <i>string</i> | 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 is used. |
| ? <i>string</i> | Same as / but search in the forward direction. |
| n | Search for next match of the last pattern to / or ? commands. |
| N | Search for next match of the last pattern to / or ? , but in reverse direction. |
| | Search history for the <i>string</i> entered by the previous / command. |

Text Modification Edit Commands

These commands 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 \$a . |
| <i>[count]</i> c <i>motion</i> | Move cursor to the character position specified by <i>motion</i> , deleting all characters between the original cursor position and new position, and enter input mode. If <i>motion</i> is c , the entire line is deleted and input mode entered. |
| c <i>[count]</i> <i>motion</i> | |
| C | Delete the current character through the end of line and enter input mode. Equivalent to c\$. |
| S | Equivalent to cc . |
| D | Delete the current character through end of line. Equivalent to d\$. |
| <i>[count]</i> d <i>motion</i> | Move cursor to the character position specified by <i>motion</i> , deleting all characters between the original cursor position and new position. If <i>motion</i> is d , the entire line is deleted. |
| d <i>[count]</i> <i>motion</i> | |
| i | Enter input mode and insert text before the current character. |
| I | Insert text before the beginning of the line. Equivalent to the two-character sequence 0i . |
| <i>[count]</i> P | Place the previous text modification before the cursor. |
| <i>[count]</i> p | Place the previous text modification after the cursor. |
| R | Enter input mode and replace characters on the screen with characters you type in overlay fashion. |

| | |
|-----------------------------|--|
| <code>[count]rc</code> | Replace the current character with <i>c</i> . |
| <code>[count]x</code> | Delete current character. |
| <code>[count]X</code> | Delete preceding character. |
| <code>[count].</code> | Repeat the previous text modification command. |
| <code>[count]~</code> | Invert the case of the current character and advance the cursor. |
| <code>[count]_</code> | Causes the <i>count</i> word of the previous command to be appended at the current cursor location and places the editor in input mode at the end of the appended text. The last word is used if <i>count</i> is omitted. |
| <code>*</code> | Appends an <code>*</code> to the current word and attempts file name generation. If no match is found, the bell rings. If a match is found, the word is replaced by the matching string and the command places the editor in input mode. |
| <code>ESC</code> | |
| <code>\</code> | Attempt file name completion on the current word. 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 <code>/</code> is appended if the file is a directory and a space is appended if the file is not a directory. |
| Other Edit Commands | |
| <code>[count]ymotion</code> | |
| <code>y[count]motion</code> | Yank current character through character that <i>motion</i> would move the cursor to and puts them into the delete buffer. The text and cursor are unchanged. |
| <code>Y</code> | Yanks from current position to end of line. Equivalent to <code>y\$</code> . |
| <code>u</code> | Undo the last text modifying command. |
| <code>U</code> | Undo all the text modifying commands performed on the line. |
| <code>[count]v</code> | Returns the command <code>fc -e \${VISUAL:-\${EDITOR:-vi}}</code> <i>count</i> in the input buffer. If <i>count</i> is omitted, the current line is used. |
| <code>^L</code> | Line feed and print current line. Has effect only in control mode. |
| <code>^J</code> | (New line) Execute the current line, regardless of mode. |
| <code>^M</code> | (Return) Execute the current line, regardless of mode. |
| <code>#</code> | Equivalent to <code>I#</code> followed by Return . Sends the line after inserting a <code>#</code> in front of the line and after each new-line. Useful for inserting the current command line in the history list without executing it. |
| <code>=</code> | List the filenames that match the current word if an asterisk were appended to it. |
| <code>@letter</code> | The user's alias list is searched for an alias by the name <code>_letter</code> and if an alias of this name is defined, its value is inserted on the input queue for processing. |

EXTERNAL INFLUENCES

Environment Variables

`LC_COLLATE` determines the collating sequence used in evaluating pattern matching notation for file name generation.

`LC_CTYPE` determines the classification of characters as letters, and the characters matched by character class expressions in pattern matching notation.

If `LC_COLLATE` or `LC_CTYPE` is not specified in the environment or is set to the empty string, the value of `LANG` is used as a default for each unspecified or empty variable. If `LANG` is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of `LANG`. If any internationalization variable contains an invalid setting, **ksh** behaves as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single-byte character code sets are supported.

RETURN VALUE

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 (also see the **exit** command above). If the shell is being used non-interactively, execution of the shell file is abandoned. Runtime errors detected by the shell are reported by printing the command or function name and the error condition. If the line number on which the error occurred is greater than one, the line number is also printed in brackets ([*n*]) after the command or function name.

WARNINGS

File descriptors 10 and 54 through 60 are used internally by the Korn Shell. Applications using these and forking a subshell should not depend upon them surviving in the subshell or its descendants.

If a command which is a **tracked alias** 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 continues to load and execute the original command. Use the **-t** option of the **alias** command to correct this situation.

If you move the current directory or one above it, **pwd** may not give the correct response. Use the **cd** command with a full path name to correct this situation.

Some very old shell scripts contain a caret (^) as a synonym for the pipe character (|). Note however, **ksh** does not recognize the caret as a pipe character.

If a command is piped into a shell command, all variables set in the shell command are lost when the command completes.

Using the **fc** built-in command within a compound command causes the entire command to disappear from the history file.

The built-in command **. file** reads the entire file before any commands are executed. Therefore, **alias** and **unalias** commands in the file do not apply to any functions defined in the file.

Traps are not processed while the shell is waiting for a foreground job. Thus, a trap on **CHLD** is not executed until the foreground job terminates.

The **export** built-in command does not handle arrays properly. Only the first element of an array is exported to the *environment*.

Background processes started from a non-interactive shell cannot be accessed by using job control commands.

In an international environment, character ordering is determined by the setting of **LC_COLLATE**, rather than by the binary ordering of character values in the machine collating sequence. This brings with it certain attendant dangers, particularly when using range expressions in file name generation patterns. For example, the command,

```
rm [a-z]*
```

might be expected to match all file names beginning with a lowercase alphabetic character. However, if dictionary ordering is specified by **LC_COLLATE**, it would also match file names beginning with an uppercase character (as well as those beginning with accented letters). Conversely, it would fail to match letters collated after **z** in languages such as Danish or Norwegian.

The correct (and safe) way to match specific character classes in an international environment is to use a pattern of the form:

```
rm [[:lower:]]*
```

This uses **LC_CTYPE** to determine character classes and works predictably for all supported languages and codesets. For shell scripts produced on non-internationalized systems (or without consideration for the above dangers), it is recommended that they be executed in a non-NLS environment. This requires that **LANG**, **LC_COLLATE**, etc., be set to "C" or not set at all.

Be aware that the value of the **IFS** variable in the user's environment affects the behavior of scripts.

ksh implements command substitution by creating a pipe between itself and the command. If the root file system is full, the substituted command cannot write to the pipe. As a result, the shell receives no input from the command, and the result of the substitution is null. In particular, using command substitution for variable assignment under such circumstances results in the variable being silently assigned a **NULL** value.

AUTHOR

ksh was developed by AT&T.

FILES

| | |
|-------------------|--|
| /etc/passwd | to find home directories |
| /etc/profile | read to set up system environment |
| /etc/suid_profile | security profile |
| \$HOME/.profile | read to set up user's custom environment |
| /tmp/sh* | for here-documents |

SEE ALSO

cat(1), cd(1), echo(1), env(1), getopts(1), kill(1), pwd(1), read(1), test(1), time(1), umask(1), vi(1), dup(2), exec(2), fork(2), gtty(2), pipe(2), stty(2), signal(5), umask(2), ulimit(2), wait(2), rand(3C), a.out(4), profile(4), environ(5), lang(5), regexp(5).

NAME

ktutil - Kerberos keytab file maintenance utility

SYNOPSIS

ktutil

DESCRIPTION

The **ktutil** command invokes a subshell from which an administrator can read, write, or edit entries in a Kerberos V5 keytab or V4 srvtab file.

ktutil Commands

| | |
|--|---|
| list | Display the current key list. Alias: l |
| read_kt <i>keytab_filename</i> | Read the Kerberos V5 keytab file, <i>keytab_filename</i> , into the current key list. Alias: rkt |
| read_st <i>srvtab_filename</i> | Read the Kerberos V4 srvtab file, <i>srvtab_filename</i> , into the current key list. Alias: rst |
| write_kt <i>keytab_filename</i> | Write the current key list into the Kerberos V5 keytab file, <i>keytab_filename</i> . Alias: wkt |
| write_st <i>srvtab_filename</i> | Write the current key list into the Kerberos V4 srvtab file, <i>srvtab_filename</i> . Alias: wst |
| clear_list | Clear the current key list. Alias: clear |
| delete_entry <i>slot</i> | Delete the entry in slot number <i>slot</i> from the current key list. Alias: delete |
| list_requests | Display a listing of available commands. Aliases: lr, ? |
| quit | Quit or exit from ktutil . Aliases: exit, q |

FILES

| | |
|-------------------------|--------------------------------------|
| /etc/krb5.keytab | Default location of the keytab file. |
| /etc/srvtab | Default location of the srvtab file |

AUTHOR

ktutil was developed by the Massachusetts Institute of Technology.

SEE ALSO

kerberos(5).

k

NAME

kvno - print key version numbers of Kerberos principals

SYNOPSIS

kvno [-e *etype*] *service1*, [*service2*, ...]

DESCRIPTION

kvno acquires a service ticket for the specified Kerberos principals and prints out the key version numbers of each principal.

Options

-e *etype* Specifies the encryption type which will be requested for the session key of all the services named on the command line. This is useful in certain backward compatibility situations. The value of *etype* can be one of DES-CBC-CRC, DES-CBC-RAW or DES-CBC-MD4.

service1,service2 Service name(s) or principal name(s).

Environment Variables

kvno uses the following environment variable:

KRB5CCNAME Location of the credentials ticket cache.

FILES

/tmp/krb5cc_{uid} Default location of the credentials cache. {*uid*} is the decimal UID of the user.

AUTHOR

kvno was developed by FundsXpress, INC.

SEE ALSO

kdestroy(1), kerberos(5), kinit(1), krb5.conf(4), libkrb5(3).

k

NAME

last, **lastb** - indicate last logins of users and ttys

SYNOPSIS

```
/usr/bin/last [-R] [-number] [-f file] [name ...] [tty ...]
/usr/bin/lastb [-R] [-number] [-f file] [name ...] [tty ...]
```

DESCRIPTION

The **last** command searches backwards through the file **/var/adm/wtmp** (which contains a record of all logins and logouts) for information about a user, a tty, or any group of users and ttys. Arguments specify names of users or ttys of interest. The names of ttys can be given fully or abbreviated. For example, **last 0** is the same as **last tty0**. If multiple arguments are given, the information that 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. The **last** command prints the sessions of the specified users and ttys, most recent first, indicating when the session began, the duration of the session, and the tty on which the session took place. **last** indicates if the session is still in progress or if it was cut short by a reboot.

The pseudo-user **reboot** logs each time the system reboots. Thus, **last reboot** is a useful command for evaluating the relative time between system reboots.

If **last** is interrupted, it indicates how far the search has progressed in **wtmp**. If interrupted by a quit signal (generated by a Ctrl-\\), **last** indicates how far the search has progressed, then continues the search.

The **lastb** command searches backwards through the database file **/var/adm/btmp** to display bad login information. Access to **/var/adm/btmp** should be restricted to users with appropriate privileges (owned by and readable only by **root**) because it may contain password information.

Options

The **last** and **lastb** commands recognize the following options and arguments:

- (none) If no arguments are specified, **last** prints a record of all logins and logouts in reverse order, most recent first.
- R** When used with **last** and **lastb**, **-R** displays the user's host name as it is stored in the files **/var/adm/wtmp** and **/var/adm/btmp**, respectively. The host name is displayed between the tty name and the user's login time.
- number** Limits the report to *number* of lines.
- f file** Use *file* as the name of the accounting file instead of **/var/adm/wtmp** or **/var/adm/btmp**.

AUTHOR

last was developed by the University of California, Berkeley and HP.

FILES

```
/var/adm/btmp  Bad login database
/var/adm/wtmp  Login database
```

SEE ALSO

login(1), utmp(4).

NAME

lastcomm - show last commands executed in reverse order

SYNOPSIS

lastcomm [*commandname*] ... [*username*] ... [*terminalname*] ...

DESCRIPTION

lastcomm gives information on previously executed commands. If no arguments are specified, **lastcomm** prints information about all the commands recorded in the accounting file, **/var/adm/pacct** during the current accounting file's lifetime. If called with arguments, only accounting entries with a matching command name, user name, or terminal name are printed. For example, to produce a listing of all executions of commands named **a.out** by user **root** on terminal **ttyd0** use:

```
lastcomm a.out root ttyd0
```

For each process entry, the following are printed.

- Name of the user who ran the process.
- Flags, as accumulated by the accounting facilities in the system.
- Command name under which the process was called.
- Amount of cpu time used by the process (in seconds).
- What time the process started.

Flags are encoded as follows:

- S** Command was executed by a user who has appropriate privileges.
- F** Command ran after a fork, but without a following *exec*.
- D** Command terminated with the generation of a **core** file.
- X** Command was terminated with the signal SIGTERM.

FILES

/var/adm/pacct current file for per-process accounting

AUTHOR

lastcomm was developed by the University of California, Berkeley.

SEE ALSO

last(1), acct(4), acctsh(1M), core(4).

NAME

ld - link editor

SYNOPSIS**The link editor. Common options:**

```
ld [-b bdmnqrstvxzEGIOPQVZ] [-a search] [-c filename] [-dynamic]
[-e epsym] [-h symbol] ... [-l x | file] ... [-l: library]
[-m] [-noshared] [-o outfile] [-u symbol] ... [-y symbol] ...
[-B bind] [-D offset] [-L dir] ... [-R offset] [-Pd] [-PD file] [-PF file] [+b path_list]
[+compat] [+copyobjdebug] [+df file]
[+e symbol] [+ee symbol] [+fb] [+fbu] [+gst]
[+gstsize size] [+h internal_name] [+help] [+k] [+n]
[+nocopyobjdebug] [+noobjdebug] [+objdebugonly] [+pd size]
[+pgm name] [+pi size] [+s] [+std] [+tools] [+v[no]shlibunsats]
[+vallcompatwarnings] [+v[no]compatwarnings] [+FP flag]
[+I symbol] [+O[no]fastaccess] [+O[no]procelim] [ +Oreusedir= dir ]
[+Oselectivepercent n] [+Oselectivesize size] [+OselectiveO3]
[+Ostaticprediction]
```

32 Bit (SOM) Link Editor options:

```
ld [-NST] [-A name] [-C n] [-F1] [-Fw] [-Fz] [+cdp oldpath:newpath]
[+cg pathname] [+dpv] [+filter shared_library_path]
[+gstbuckets size] [+nosmartbind] [+plabel_cache flag]
```

64 Bit (ELF) Link Editor options:

```
ld [-k filename] [+fini function] [+ild] [+ildnowarn] [+ildpad percentage] [+ildrelink]
[+init function] [+interp file] [+no]allowunsats] [+no]forceload]
[+hideallsymbols] [+nodefaultmap] [+noenvvar] [+nodynhash]
[+nosectionmerge] [+paddata pagesize] [+padtext pagesize] [+pdzero]
[+stripunwind] [+vtype type]
```

DESCRIPTION

ld takes one or more object files or libraries as input and combines them to produce a single (usually executable) file. In doing so it resolves references to external symbols, assigns final addresses to procedures and variables, revises code and data to reflect new addresses (a process called "relocation"), and updates symbolic debug information when present in the file. By default, ld produces an executable file that can be run by the HP-UX loader `exec()` (see `exec(2)`). Alternatively, the linker can generate a relocatable file that is suitable for further processing by ld (see `-r` below). It can also generate a shared library (see `-b` below). The linker marks the output file non-executable if there are any duplicate symbols or any unresolved external references remain. ld may or may not generate an output file (see `+k` option) if any other errors occur during its operation.

ld recognizes three kinds of input files: object files created by the compilers, assembler, or linker (also known as `.o` files), shared libraries created by the linker, and archives of object files (called archive libraries). An archive library contains a table of all the externally-visible symbols from its component object files. (The archiver command `ar(1)` creates and maintains this index.) ld uses this table to resolve references to external symbols.

ld processes files in the same order as they appear on the command line. It includes code and data from an archive library element if and only if that object module provides a definition for a currently unresolved reference within the user's program. It is common practice to list libraries following the names of all simple object files on the command line.

Code and data from shared libraries is never copied into an executable program. The dynamic loader `/usr/lib/dld.sl` on 32-bit links is invoked at startup time by the startup file `crt0.o` if a program uses shared libraries. Identical copies of `crt0.o` can be found in either `/usr/ccs/lib/crt0.o` or `/opt/langtools/lib/crt0.o`. For 64-bit mode, the dynamic loader is `/usr/lib/pa20_64/dld.sl` and is invoked by `exec` for those programs that use shared libraries. `crt0.o` is not required in shared bound links. The dynamic loader attaches each required library to the process and resolves all symbolic references between the program and its libraries.

The text segment of a shared library is shared among all processes that use the library; each process using the library receives its own copy of the data segment. If `pxldb -s on` has been run on the executable that loads the library, the text segment of a shared library is mapped privately for each process running the

executable. **ld** recursively examines the dependencies of shared libraries used by a program that was created by **ld**. If **ld** does not find a supporting shared library at the path recorded in the dependency list of a shared library, and if the dependency is the result of an **-l** argument used when the shared library was created, **ld** will search all the directories that it would search for a library that was specified with **-l** (see **-L** and **LPATH**).

Environment Variables

Arguments can be passed to the linker through the **LDOPTS** environment variable as well as on the command line. The linker gets the value of **LDOPTS** and places its contents before any arguments on the command line.

The **LD_PXDB** environment variable defines the full execution path for the debug preprocessor **pxdb**. The default value is **/opt/langtools/bin/pxdb**. **ld** invokes **pxdb** on its output file if that file is executable and contains debug information. To defer invocation of **pxdb** until the first debug session, set **LD_PXDB** to **/bin/true**.

The **LPATH** environment variable can be used to specify default directories to search for library files. See the **-l** option.

Options

The common **ld** options are listed first, followed by the options supported only in a 32-bit linker, and then the options only supported in a 64-bit linker.

- a search** Specify whether shared or archive libraries are searched with the **-l** option. The value of *search* should be one of **archive**, **shared**, **archive_shared**, **shared_archive**, or **default**. This option can appear more than once, interspersed among **-l** options, to control the searching for each library. The default is to use the shared version of a library if one is available, or the archive version if not.

If either **archive** or **shared** is active, only the specified library type is accepted.
If **archive_shared** is active, the archive form is preferred, but the shared form is allowed.
If **shared_archive** is active, the shared form is preferred but the archive form is allowed.
- b** Create a shared library rather than a normal executable file. Object files processed with this option must contain **position-independent code** (PIC). See the discussion of PIC in *cc(1)*, *CC(1)* (part of the optional C++ compiler documentation), *f77(1)*, *pc(1)*, *as(1)*, and *Linker and Libraries Online User Guide*.
- c filename** Read **ld** options from a file. Each line contains zero or more arguments separated by white space. Each line in the file, including the last line, must end with a newline character. A **#** character implies that the rest of the line is a comment. To escape a **#** character, use the sequence **##**.
- d** Force definition of “common” storage; that is, assign addresses and sizes, for **-r** output.
- e epsym** Set the default entry point address for the output file to be that of the symbol *epsym*. (This option only applies to executable files.)
- h symbol** Prior to writing the symbol table to the output file, mark this name as “local” so that it is no longer externally visible. This ensures that this particular entry will not clash with a definition in another file during future processing by **ld**.

More than one *symbol* can be specified, but **-h** must precede each one. If used when building a shared library or program, this option prevents the named symbol from being visible to the dynamic loader.
- l_x** Search a library **lib_x.a** or **lib_x.sl**, where *x* is one or more characters. The current state of the **-a** option determines whether the archive (**.a**) or shared (**.sl**) version of a library is searched. Because a library is searched when its name is encountered, the placement of a **-l** is significant. By default, 32-bit libraries are located in **/usr/lib** and **/usr/ccs/lib**. 64-bit libraries are located in **/usr/lib/pa20_64**. If the environment variable **LPATH** is present in the user's environment, it should contain a colon-separated list of directories to search. These directories are searched instead of the default directories, but **-L** options can still be

used. If a program uses shared libraries, the dynamic loader `/usr/lib/dld.sl` for 32-bit or `/usr/lib/ia20_64/dld.sl` for 64-bit will attempt to load each library from the same directory in which it was found at link time (see the `+s` and `+b` options).

- l *library*** Search the library specified. Similar to the `-l` option except the current state of the `-a` option is not important. The library name can be any valid filename. (Note that previous releases required that the library name contain the prefix `lib` and end with a suffix of `.a` or `.sl`.)
 - m** This option produces a load map on the standard output.
 - n** This option is accepted but ignored by the 64-bit `ld`. Generate an executable output file with file type `SHARE_MAGIC`. This is the default. This option is incompatible with `-N` and `-q`.
 - o *outfile*** Produce an output object file named *outfile* (`a.out` if `-o outfile` is not specified).
 - q** This option is ignored for 64-bit links. Generate an executable output file with file type `DEMAND_MAGIC`. This option is incompatible with `-n`, `-N`, and `-Q`.
 - r** Retain relocation information in the output file for subsequent re-linking. The `ld` command does not report undefined symbols. This option cannot be used when building a shared library (`-b`) or in conjunction with `-A` or `+ild` incremental linking options.
 - s** Strip the output file of all symbol table, relocation, and debug support information. This might impair or prevent the use of a symbolic debugger on the resulting program. This option is incompatible with `-r`. (The `strip(1)` command also removes this information.) This option is incompatible with `+ild`. (The incremental linking requires the parts of the output load module which are stripped out with `-s` option.)
 - t** Print a trace (to standard output) of each input file as `ld` processes it.
 - u *symbol*** Enter *symbol* as an undefined symbol in the symbol table. The resulting unresolved reference is useful for linking a program solely from object files in a library. More than one *symbol* can be specified, but each must be preceded by `-u`.
 - v** Display verbose messages during linking. For each library module that is loaded, the linker indicates which symbol caused that module to be loaded.
 - x** Strip local symbols from the output file. This reduces the size of the output file without impairing the effectiveness of object file utilities. This option is incompatible with the `-r` option. This option is incompatible with the `+ild` option. (The incremental linking requires the parts of the output load module which are stripped out with the `-x` option.)
- Note: use of `-x` might affect the use of a debugger.
- y *symbol*** Indicate each file in which *symbol* appears. More than one *symbol* can be specified, but each must be preceded by `-y`.
 - z** Arrange for run-time dereferencing of null pointers to produce a `SIGSEGV` signal. (This is the complement of the `-Z` option.)
 - B *bind*** Select run-time binding behavior of a program using shared libraries or the binding preference in building a shared library. The most common values for *bind* are:

deferred

Bind addresses on first reference rather than at program start-up time. This is the default.

immediate

Bind addresses of all symbols immediately upon loading the library. Commonly followed by `-B nonfatal` to allow procedure calls that cannot be resolved at program start-up to be resolved on first reference.

Since `-B nonfatal` suppresses messages about unresolved symbols, also specify `-B verbose` to display those messages.

See the example below.

nonfatal

If also using **-B immediate**, for code symbols that could not be bound at program startup, defer binding them until they are referenced. See description of **-B immediate** above.

Since **-B nonfatal** suppresses messages about unresolved symbols, also specify **-B verbose** to display those messages.

restricted

Causes the search for a symbol definition to be restricted to those symbols that were visible when the library was loaded.

symbolic

Used only when building a shared library (with the **-b** option), this option causes all unresolved symbols inside a library to be resolved internally if possible. By default, unresolved symbols are resolved to the most visible definition in the library or outside of the library.

verbose

Display verbose messages when binding symbols. This is the default except when **-B nonfatal** is specified. In that case, **-B verbose** must be explicitly specified to get verbose messages.

See the **+help** option or the *HP-UX Linker and Libraries User's Guide* manual for more information on the uses of binding modes.

-D offset

Set the origin (in hexadecimal) for the data segment.

When used with the **+ild** option, if you change the offset after the initial incremental link, the linker performs an initial incremental link automatically.

-E

This option is accepted but ignored by the 64-bit **ld**. Mark all symbols defined by a program for export to shared libraries. By default, **ld** marks only those symbols that are actually referenced by a shared library seen at link time.

-G

Strip all unloadable data from the output file. This option is typically used to strip debug information.

-I

Instrument the code to collect profile information upon execution. The profile data gathered during program execution can be used in conjunction with the **-P** option. 32-bit programs linked with this option should use the startup file `/opt/langtools/lib/icrt0.o`. This option should not be used with the **-P**, **-A**, **-O**, **+ild**, or **+O** options.

-L dir

Search for `libx.a` or `libx.sl` in `dir` before looking in default locations. More than one directory can be specified, but each must be preceded by **-L**. The **-L** option is effective only if it precedes the **-l** option on the command line.

-O

Turn on linker optimizations. Currently the optimizations include the elimination of unnecessary **ADDIL** instructions from the code in the executable file (32-bit only), and the removal of dead procedures.

-O is passed to the linker by the compilers when the **+O4** compiler option is selected.

This option is incompatible with the **+ild** option.

For more details on linker optimizations refer to the **+help** option or the *HP-UX Linker and Libraries User's Guide* manual.

-P

Examine the data file produced by an instrumented program (see the **-I** option) to perform profile based optimizations on the code. This option should not be used with the **-A** or **+ild** options.

-Q

Ignored for 64-bit links. Generate an executable output file with file type **EXEC_MAGIC** or **SHARE_MAGIC**, depending on whether **-N** or **-n** is specified. This is the default. This option is incompatible with **-q**.

-R offset

Set the origin (in hexadecimal) for the text (i.e., code) segment.

When used with the **+ild** option, if you change the offset after the initial incremental link, the linker performs an initial incremental link automatically.

- V** Output a message giving information about the version of **ld** being used.
- Z** Allow run-time dereferencing of null pointers. See the discussions of **-Z** and *pointers* in *cc(1)*. (This is the complement of the **-z** option.)
- Pd** Reorder debuggable functions. Ordinarily **-P** does not reorder functions from **.o** files with debugging information, because reordering renders them non-debuggable. This option overrides this and reorders the functions.
- PD filename** Save link order file generated by **fdp** during linking with **-P** option into user-specified file. This option is incompatible with the **+ild** option.
- PF filename** Indicate to the linker to use the specified file for the linker file instead of generating it using **/usr/ccs/bin/fdp**. This option is incompatible with the **+ild** option.
- dynamic** This allows the linker to create a program which can use shared libraries. This is the default for 64-bit links unless **+compat** is used. In 32-bit mode, the linker creates a static executable if there are no shared libraries on the link line.
- noshared** This option forces the linker to create a fully archive bound program.
- +b path_list** Specify a colon-separated list of directories (embedded path) to be searched at program run-time to locate shared libraries needed by the executable output file that were specified with either the **-l** or **-l:** option. An argument consisting of a single colon (**:**) indicates that **ld** should build the list using all the directories specified by the **-L** option and the **LPATH** environment variable (see the **+s** option).
- +compat** This option is ignored for 32-bit links. This option turns on compatibility mode in the linker — 64-bit links mimic behavior of 32-bit links).
- +copyobjdebug** Copy objdebug space.
- +df file** Used together with the **-P** option, this option specifies that *file* should be used as the profile database file. The default value is **flow.data**. See the discussion of the **FLOW_DATA** environment variable for more information. This option is incompatible with the **+ild** option.
- +e symbol** When building a shared library or program, mark the symbol for export to the dynamic loader. Only symbols explicitly marked are exported. When building a shared library, calls to symbols that are not exported are resolved internally.
- +ee symbol** This option is similar to the **+e** option in that it exports a symbol. However, unlike the **+e** option, **+ee** does not alter the visibility of any other symbol in the file. In a 32-bit link or **+compat** mode 64-bit link, it has the effect of exporting the specified symbol without hiding any of the symbols exported by default. In a 64 **+std** link, all symbols are exported by default, so **+ee** is not necessary to make a symbol visible. However, it has the additional side effect of identifying the symbol as necessary, so that it will not be removed when using dead code elimination (**+Oprocelim**). Of course, **+ee** still retains its export behavior if an option such as **+hideallsymbols** is also given.
- +fb** Instructs the linker to run the fastbind tool on the executable it has produced. The executable should be linked with shared libraries. For more details refer to *fastbind(1)*, the **+help** option, or the *HP-UX Linker and Libraries User's Guide* manual. This option is incompatible with the **+ild** option.
- +fbu** Pass the **-u** option to the fastbind tool. For more details refer to *fastbind(1)*, the **+help** option, or the *HP-UX Linker and Libraries User's Guide* manual. This option is incompatible with the **+ild** option.
- +gst** Enable the global symbol table hash mechanism, used to look up values of symbol import/export entries. The **+gst** and related options provide performance enhancements through use of global symbol table which improves searching for exported symbols. See *dld.sl(5)* and the *HP-UX Linker and Libraries Online User Guide* for more information.
- +gstsize size** Request a particular hash array size using the global symbol table hash mechanism. The default value is 1103. The value can be overridden at runtime by setting the

`_HP_DLDOPTS` environment variable to the value `-syntab_size prime number`. You can set the value using `chatr +gstsize size file`.

+h *internal_name*

When building a shared library, record *internal_name* as the name of the library. When the library is used to link another executable file (program or shared library), this *internal_name* is recorded in the library list of the resulting output file instead of the file system pathname of the input shared library.

That is, if **+h** is not used, the shared library does not have an internal name and when an executable is built with the shared library, the linker records the library name that it looks at.

If *internal_name* is a fully-qualified pathname, it is recorded **as is** in the library list of any executable file it is subsequently linked against. *internal_name* is a relative pathname or no directory component was specified, *internal_name* is **appended** to the file system directory component of the input shared library in the library list of any executable file it is subsequently linked against.

If more than one **+h** option is seen on the link line, the first one is used and a warning message is emitted.

+help

Starts the help window utility *HP-UX Linker and Libraries Online User Guide* which comes with some HP compilers. (You must be running the X window system and your `DISPLAY` environment variable must be set to the name of your workstation or X terminal.) For more information, refer to the *HP-UX Linker and Libraries User's Guide* manual. See *manuals(5)* for ordering information.

+k

Direct the linker to only create an executable if there were no errors encountered during the link. If there were errors found (system errors or unresolved references), the output file will be removed.

+n

Causes the linker to load all object modules before searching any archive or shared libraries. Then it searches the archive and shared libraries specified on the command line in left to right order. Repeats the left to right search of the libraries on the command line until there are no more unsatisfied symbols, or the last search added no new definitions. This option is useful if two libraries are specified that have symbol dependencies on each other.

+nocopyobjdebug

Do not copy objdebug space. Use this option (with **-r** object files on the link line) to suppress the default behavior of copying LINKMAP space to the executable file .

+noobjdebug

Override the **+objdebug** compiler option, and copy all debug information to the executable file.

When used with **-g**, **+objdebug** leaves debug information in the object files instead of copying it to the executable file at link time, resulting in shorter link times and smaller executables. The compile-time default, **+noobjdebug**, copies the debug information to the executable file.

When you specify **-g** when compiling, the compiler places symbolic debugging information into the object files. By default, the linker calls **pxdb** which compacts this debug information and copies it to the executable file. When **+objdebug** was used at compile time, the linker leaves the debug information in the object files. To debug the executable file, the HP WDB debugger must have access to the object files. If you move the object files, use HP WDB's **objdir** command to tell it where the object files are. (The HP DDE debugger does not support this option.) The **+objdebug** option reduces link time and the size of the executable file by avoiding this copying of debug information.

The compile-time default is **+noobjdebug**. If the linker detects any object files that were compiled with **+objdebug**, it leaves the debug information in those files. Any object files not compiled with **+objdebug** have their debug information copied into the executable file. You can leave debug information in some object files and not in others.

When archive members are not compiled with **+objdebug** (and they contain debug information, that is, they were compiled with **-g**), the debug information from the archive members is copied into the **a.out** by default. The linker does not copy debug information from shared libraries into the **a.out**, regardless of whether they are compiled with **+objdebug** or **+noobjdebug**.

Use the **+noobjdebug** option when linking to explicitly tell the linker to copy all debug information to the executable file, even from files compiled with **+objdebug**.

+objdebugonly

Ignore debug information from non-objdebug objects or archives and proceed in **+objdebug** mode. This option can be passed from the C or C++ compiler command line as **-Wl,+objdebugonly**. If you are debugging only files compiled with the **+objdebug** option, **+objdebugonly** can improve link time by instructing the linker to bypass the processing of debug information from files compiled with **+noobjdebug**. With the **+objdebugonly** option, the linker suppresses the call to **pxdb**.

+pd size

Request a particular virtual memory page size that should be used for data. Sizes of **4K**, **16K**, **64K**, **256K**, **1M**, **4M**, **16M**, **64M**, **256M**, **D**, and **L** are supported. A size of **D** allows the kernel to choose what page size should be used. A size of **L** will result in using the largest page size available. The actual page size may vary if the requested size cannot be fulfilled.

+pgm name

Used together with the **-P** option, this option specifies that *name* should be used as the look-up name in the profile database file. The default is the basename of the output file (specified by the **-o** option.) This option is incompatible with the **+ild** option.

+pi size

Request a particular virtual memory page size that should be used for instructions. See the **+pd** option for additional information.

+s

Indicates that at run-time, the shared library loader can use the environment variable **SHLIB_PATH** and **LD_LIBRARY_PATH** (64-bit only) to locate shared libraries needed by the executable output file that were specified with either the **-l** or **-L** option. The environment variables should be set to a colon-separated list of directories. If both **+s** and **+b** are used, their relative order on the command line indicates which path list will be searched first (see the **+b** option).

+std

This option is ignored for 32-bit links. Turns on standard mode, which is the default in 64-bit mode. Options turned on with this option are: **-dynamic**. Options turned off or ignored when this option is specified are: **+compat**, **+noenvvar**, **-noshared**.

+vallcompatwarnings

This option is accepted but ignored by the 64-bit **ld**. Show more detail for any warnings about compatibility issues. By default, only a terse message is printed. See the **WARNINGS** section below for further details.

+v[no]compatwarnings

This option is accepted but ignored by the 64-bit **ld**. Enable [disable] printing warnings about compatibility issues between systems. This includes any functionality which may change in future releases. The default is **+vcompatwarnings**. See the **WARNINGS** section below for further details.

+v[no]shlibunsats

Enable [disable] printing a list of unsatisfied symbols used by shared libraries. The default is **+vnoshlibunsats**. Some unsatisfied symbols reported by the linker are not required at run time because the modules which reference the symbols are not used.

+FP flag

Specify how the environment for floating-point operations should be initialized at program start-up. By default, all behaviors are disabled. The following flags are supported (upper case flag enables; lower case flag disables):

V (v) Trap on invalid floating-point operations

Z (z) Trap on divide by zero

- O (o) Trap on floating-point overflow
- U (u) Trap on floating-point underflow
- I (i) Trap on floating-point operations that produce inexact results.
- D (d) Enable sudden underflow (flush to zero) of denormalized values.

Note: Enabling sudden underflow is an undefined operation on PA-RISC 1.0-based systems, but it is defined on all subsequent versions of PA-RISC. Selecting this flag enables sudden underflow only if it is available on the processor being used at run-time.

To dynamically change these settings at run-time, see *fesettrapenable(3M)*.

+I *symbol* Specify the name of the initializer function when building a shared library. A shared library may have multiple initializers specified. Initializers are executed in the order that they are specified. You can specify more than one initializer, but each must be preceded by **+I**. For more details on the initializer function, refer to the **+help** option or the *HP-UX Linker and Libraries User's Guide* manual.

+O[no]fastaccess

Enable [disable] fast access to global data. The linker rearranges the data to further reduce the number of **ADDIL** instructions in the executable.

This optimization may reveal some subtle programming errors related to assumptions about data layout. This optimization can occur at optimization levels 2, 3 and 4. The default is **+Onofastaccess** at optimization levels 2 and 3, and **+Ofastaccess** at optimization level 4.

Using the **+Ofastaccess** option in the presence of debug information generates a warning message and can cause corruption of the debug information. Avoid using **+Ofastaccess** with the **-g** option.

This option is accepted but ignored by the 64-bit **ld**.

+O[no]procelim

Enable [disable] the elimination of procedures that are not referenced by the application. The default is **+Onoprocelim**. Procedure elimination can occur at any optimization level, including level 0. For more details refer to the **+help** option or the *HP-UX Linker and Libraries User's Guide* manual. This option is incompatible with the **+ild** option.

+Oreusedir= *dir*

Specify the name of the directory to use as the reuse repository for the object code reuse feature. This directory holds the compiled object files for reuse. *dir* can be an absolute or relative path to directory where you invoked the linker. The **+Oreusedir** option shortens link time by not recompiling intermediate object code to object code when using **+O4** or profile-based optimization.

+Oselectivepercent *n*

Instructs the interprocedural optimizer driver to pass the first *n* percent of the object files to the high level optimizer for interprocedural optimizations such as inlining.

+Oselectivesize *size*

Instructs the interprocedural optimizer driver to pass the first *k* routines to the high level optimizer for interprocedural optimization where the size of *k* routines are approaching but less than *size*.

+OselectiveO3

Instructs the interprocedural optimizer driver to compile the routines not included in the **+O4** list to be compiled at **+O3**.

+Ostaticprediction

Meaningful only on PA 2.0 architecture, this option sets the branch prediction bit in the output executable file's auxiliary header.

32 Bit Link Editor Options

-A *name*

This option specifies incremental loading. Linking is arranged so that the resulting object can be read into an already executing program. The argument *name* specifies a

file whose symbol table provides the basis for defining additional symbols. Only newly linked material is entered into the text and data portions of **a.out**, but the new symbol table reflects all symbols defined before and after the incremental load. Also, the **-R** option can be used in conjunction with **-A**, which allows the newly linked segment to commence at the corresponding address. The default starting address is the old value of **_end**. The **-A** option is incompatible with **-r** and **-b**. Also note that a program that dynamically loads code with **ld -A** cannot use shared libraries. See the **+help** option or the *HP-UX Linker and Libraries User's Guide* manual for a description of this option.

- Cn** Set the maximum parameter-checking level to *n*. The default maximum is 3. See the language manuals for the meanings of the parameter-checking level.
- F1** Force load all member objects of an archive library. If you do not use **-F1**, the linker only loads the needed archive members. This option is useful for creating a shared library from an archive library, or when you need to load archive member that define symbols needed by shared libraries.
- Fw** Don't emit unwind tables. Do not use this option if your compiler or tools require unwind tables.
- Fz** Disable the linker feature that translates some calls to `$$dyncall_external` to calls to `$$dyncall`.
- N** Generate an executable output file with file type **EXEC_MAGIC**. This option is incompatible with **-n** and **-q**. This option causes the data to be placed immediately following the text, and makes the text writable. Files of this type cannot be shared.
- S** Generate an Initial Program Loader (IPL) auxiliary header for the output file, instead of the default HP-UX auxiliary header.
- T** Save the load data and relocation information in temporary files instead of in memory during linking. This option reduces the virtual memory requirements of the linker. If the **TMPDIR** environment variable is set, the temporary files are created in the specified directory, rather than in **/var/tmp**.

+cdp *oldpath:newpath*

Replace the recorded path for a shared library in the **a.out**. In 32-bit mode, **ld** records the absolute path names of any shared libraries searched at link time in the **a.out** file. When the program begins execution, the dynamic loader attaches any shared libraries that were searched at link time. Although you can use the **+b** and/or **+s** linker options to direct the dynamic loader to directories to search for the shared libraries, the dynamic loader, as a last resort, searches for the shared libraries in its absolute, recorded path in the **a.out**. You can specify more than one shared library *oldpath:newpath*, but each must be preceded by the **+cdp** option.

- +cg** *pathname* Specify the use of *pathname* as the code generator for compiling ISOMs to SOMs. See the discussion of profile based optimization in the *HP-UX Linker and Libraries Online User's Guide* for more information.

- +dpv** Display verbose messages regarding procedures which have been removed due to dead procedure elimination. The symbol name, input object file, and the size (in bytes) of the deleted procedure are displayed. The total size (in bytes) of the deleted procedures is also displayed.

+filter *shared_library_path*

Enables the shared library filter mechanism, which allows you divide a large library into a "filter" and several "implementation" libraries for more efficient organization of shared libraries. *shared_library_path* specifies the location of the filter library. See the *HP-UX Linker and Libraries User's Guide* for more information.

+gstbuckets *size*

Request a particular number of buckets per entry using the global symbol table hash mechanism. The default value is 3. The value can be overridden at runtime by setting the **_HP_DLDOPTS** environment variable to the value **-symtab_buckets number**. You can set the value using **chatr +gstbuckets size file**.

+nosmartbind

Disable SmartBind when binding a shared library. With this option enabled, the

linker places all symbols in the link into a single SmartBind module instead of placing each `.o` file in its own module.

+plabel_cache

Enable the label caching mechanism. Use this option with the **+gst** option.

This option is only effective with C++. In C++ applications, the dynamic loader needs to repetitively access PLABEL information (import stub). In order to make this access faster, the dynamic loader uses the global symbol table structure to also contain PLABEL entries. This behavior is enabled when the PLABEL_CACHE flag is set in the `dl_header` structure (enabled **ld +plabel_cache enable a.out** or **chatr +plabel_cache enable a.out**).

64-bit Link Editor Options

-k filename **filename** specifies a mapfile that describes the output file memory map.

The user specified mapfile specifications are permitted with the **+ild** option. But you should not modify the mapfile after the initial incremental link. If the mapfile is modified after the initial link, an initial incremental link is performed automatically.

Please refer to *HP-UX Linker and Libraries User's Guide* for more information. Also see **+nodefaultmap**.

+[no]allowunsats

+allowunsats Does not flag errors if the resulting output file has unsatisfied symbols. This is the default for relocatable links and shared library builds. **+noallowunsats** Flags an error if the resulting output file has unsatisfied symbols. This is the default for program files.

+fini function_name

Specify the terminator function.

+ild

Specify incremental linking.

If the output file does not exist, or if it was created without the **+ild** option, the linker performs an initial incremental link. The output file produced is suitable for subsequent incremental links. The incremental link option is valid for both executable and shared library links.

The following options are compatible with the **+ild** option with limitations:

-D offset , -R offset

Set the origin for the data and text segments. If you change the offset after the initial incremental link, the linker performs an initial incremental link automatically.

-k mapfile

provide a non-default mapfile. The user specified mapfile specifications are permitted with the **+ild** option. But you should not modify the mapfile after the initial incremental link. If the mapfile is modified after the initial link, an initial incremental link is performed automatically.

If you specify one of the following incompatible **ld** option with **+ild**, the linker emits a warning message and ignores the **+ild** option.

-r create a relocatable object file.

Strip options:

-s and **-x** strip the output file. (The incremental linking requires the parts of the output load module which are stripped out with these options.)

Optimization options:

-I, -O, -P, -PD, -PF, +df file, +fb, +fbu, +fbs, +pgm name, +Oprocelim

+ildnowarn Suppress incremental-linking related warnings. By default, the linker issues all incremental-linking related warnings. This option is ignored if used without **+ild** or **+ildrelink**.

+ildpad percentage

Control the amount of padding *percentage* the incremental linker allocates, relative to

sizes of object file structures being padded. By default the linker allocates 25% of padding space. This option is ignored if used without **+ild** or **+ildrelink**.

- +ildrelink** Perform an initial incremental link, regardless of the output load module.
In certain situations (for example, internal padding space is exhausted) the incremental linker is forced to perform an initial incremental link. The **+ildrelink** option allows you to avoid such unexpected initial incremental links by periodically rebuilding the output file.
- +init *function_name***
Specify the initializer function.
- +interp *file*** Change the **dld** path to use the argument as the "interpreter" program instead of the **dld.sl**.
- +*[no]*forceload**
+forceload loads all object files from archive libraries. **+noforceload** is the default — loads only the required object files from archive libraries. The mode that is selected, either explicitly or by default, remains on until it is changed.
- +hideallsymbols**
This option is used to prevent all the symbols from being exported unless explicitly exported with the **+e**.
- +nodefaultmap**
This option tells the linker not to use the default memory map. The user needs to supply a mapfile through the **-k** linker option.
- +nodynhash** Disables the default linker behavior of the **+gst** option to create the .dynhash section for executables or shared libraries. Use this option to eliminate generation of pre-computed hash table information for a library or an executable that is rarely used with the global symbol table lookup scheme or for which the overhead of storing pre-computed hash values is too high. This option has no effect when used with the **-r** option.
- +noenvvar** Instructs the dynamic loader to not look at the LD_LIBRARY_PATH and SHLIB_PATH environment variables at runtime. This is turned on if **ld +compat** is specified. This is turned off by default or if **ld +std** is specified. See **+compat** or **+std**. Generally, this option is used for secure programs (e.g. setuid).
- +nosectionmerge**
Used with the **-r** option to allow procedures to be positioned independently. The default is to merge all procedures into a single section.
- +paddata *pagesize***
Pads the data segment to a multiple of *pagesize* with zeros. This can improve page allocation, thus reduce TLB misses.
- +padtext *pagesize***
Pads the text segment to a multiple of *pagesize* with zeros. This can improve page allocation, thus reduce TLB misses.
- +pdzero** Generates a zero page of 4K bytes at the beginning of the data segment. It also sets the default starting address of the data segment to 0x8000000000000000. This option improves performance by reducing TLB misses by allowing the kernel to allocate bigger data pages.
- +stripunwind**
Do not output the unwind table. This creates smaller executable sizes. Use this option if you do not need the unwind table for debugging or C++ exception handling.
- +tools** Request that the application be linked for profiling with CXperf.
- +vtype *type*** Produces verbose output about the link operation. *type* can have the following values:
 - files**
Dump info about each object file loaded.
 - heap**
Dump info about the size of the heap used by a link.

libraries

Dump info about libraries searched.

procelim

Dump info about sections that have been rejected by the **+Oprocelim** option

sections

Dump info about each input section added to the output file.

symbols

Dump info about global symbols referenced/defined from/in the input files.

all Dumps all of the above info. Same as **-v**.

Defaults

Unless otherwise directed, **ld** names its output **a.out**. The **-o** option overrides this. Executable output files are of type **SHARE_MAGIC**. The default state of **-a** is to search shared libraries if available, archive libraries otherwise. The default bind behavior is **deferred**.

The default value of the **-Z/-z** option is **-Z**.

For 64-bit mode, **+std** is on by default.

Incremental linking with ld (64-bit Mode ONLY)

In the edit-compile-link-debug development cycle, link time is a significant component. The incremental linker (available through the **+ild** and **+ildrelink** options) can reduce the link time by taking advantage of the fact that you can reuse most of the previous version of the program and that the unchanged object files do not need to be processed. The incremental linker allows you to insert object code into an output file (executable or shared library) that you created earlier, without relinking the unmodified object files. The time required to relink after the initial incremental link depends on the number of modules you modify.

The linker performs the following different modes of linking:

- *normal link*: the default operation mode in which the linker links all modules.
- *initial incremental link*: the mode entered when you request an incremental link, but the output module created by the incremental linker does not exist, or it exists but the incremental linker is unable to perform an incremental update.
- *incremental link*: the mode entered when you request an incremental link, an output module created by the incremental linker exists, and the incremental linker does not require an initial incremental link.

Incremental links are usually much faster than regular links. On the initial link, the incremental linker requires about the same amount of time that a normal link process requires, but subsequent incremental links can be much faster than a normal link. A change in one object file in a moderate size link (tens of files, several megabytes total) normally is about 10 times faster than a regular **ld** link. The incremental linker perform as many incremental links as allocated padding space and other constraints permit. The cost of the reduced link time is an increase in the size of the executable or shared library.

The incremental linker allocates padding space for all components of the output file. Padding makes modules larger than those modules linked by **ld**. As object files increase in size during successive incremental links, the incremental linker can exhaust the available padding. If this occurs, it displays a warning message and does a complete initial incremental link of the module. When an object file changes, the incremental linker not only replaces the content of that file in the executable or shared library being linked, but also adjusts references to all symbols defined in the object file and referenced by other objects. This is done by looking at relocation records saved in the incrementally linked executable or shared library.

On the initial incremental link, the linker processes the input object files and libraries in the same way as the normal link. In addition to the normal linking process, the incremental linker saves information about object files, global symbols, and relocations, and pads sections in the output file for expansion. On subsequent incremental links, the linker uses timestamps and file sizes to determine which object files have changed, and updates those modules.

Under certain conditions, the incremental linker cannot perform incremental links. When this occurs, the incremental linker automatically performs an initial incremental link to restore the process. In the following situations, the linker automatically performs an initial incremental link of the output file:

- Changed linker command line, where the linker command line does not match the command line stored in the output file. (With the exceptions of the verbose and tracing options)
- Any of the padding spaces have been exhausted.
- Modules have been modified by the **ld -s** or **ld -x** options or tools (for example, **strip(1)**). The incremental linking requires the parts of the output load module which are stripped out with these options.
- Incompatible incremental linker version, when you run a new version of the incremental linker on an executable created by an older version.
- New working directory, where the incremental linker performs an initial incremental link if current directory changes.
- Archive or shared libraries are added/removed to/from the linker command line.
- Objects are added/removed to/from the linker command line.

See the *Online Linker and Libraries User's Guide* (**ld +help**) for more information.

Archive Library Processing

The incremental linker searches an archive library if there are unsatisfied symbols. It extracts all archive members satisfying unsats and processes them as new object files. If an archive library is modified, the linker replaces the modified archive library.

An object file extracted from an archive library in the previous link remains in the output load module even if all references to symbols defined in the object file have been removed. The linker removes these object files when it performs the next initial incremental link.

Shared Library Processing

In an initial incremental link, the linker scans shared library symbol tables and resolves unsats the same way it would in a regular link. In incremental links, the linker does not process shared libraries and their symbol tables at all and does not report shared library unsats. The dynamic loader detects them at run time. If any of the shared libraries on the command line was modified, the linker reverts to an initial incremental link.

Performance

Performance of the incremental linker may suffer greatly if you change a high percentage of object files.

The incremental linker may not link small programs much faster, and the relative increase in size of the executable is greater than that for larger programs.

Generally, the linker needs to scan through all shared libraries on a link line in order to determine all the unsats, even in incremental links. This process may slow down incremental links. The incremental linker does not scan shared libraries and leaves detection of shared library unsats to the dynamic loader.

Do not use the incremental linker to create final production modules. Because it reserves additional padding space, modules created by the incremental linker are considerably larger than those created in regular links.

Notes

Any program that modifies an executable (for example, **strip(1)**), may affect the ability of **ld** to perform an incremental link. When this happens, the incremental linker issues a message and performs an initial incremental link.

Third-party tools that work on object files may have unexpected results on modules produced by the incremental linker.

EXTERNAL INFLUENCES

Environment Variables

The following internationalization variables affect the execution of **ld**:

FLOW_DATA

An instrumented executable (see the **-I** option) writes out the profile data to a database file named **flow.data** in the current directory. The name and location of this file can be specified by setting **FLOW_DATA** to the desired path name. The profile data is stored in the database file under a look-up name that is the same as the basename of the executable file specified at run-time. A single **flow.data** file can hold profile data for multiple program files.

LANG

Determines the locale category for native language, local customs and coded character set in the absence of **LC_ALL** and other **LC_*** environment variables. If **LANG** is not specified or is set to the empty string, a default of **C** (see *lang(5)*) is used instead of **LANG**.

LC_ALL

Determines the values for all locale categories and has precedence over **LANG** and other **LC_*** environment variables.

LC_MESSAGES

Determines the locale that should be used to affect the format and contents of diagnostic messages written to standard error.

LC_NUMERIC

Determines the locale category for numeric formatting.

LC_CTYPE

Determines the locale category for character handling functions.

NLSPATH

Determines the location of message catalogs for the processing of **LC_MESSAGES**.

If any internationalization variable contains an invalid setting, **ld** behaves as if all internationalization variables are set to **C**. See *environ(5)*.

In addition, the following environment variable affects **ld**:

TMPDIR

Specifies a directory for temporary files (see *tmpnam(3S)*).

DIAGNOSTICS

ld returns zero when the link is successful. A non-zero return code indicates that an error occurred.

EXAMPLES

Link part of a C program for later processing by **ld**. (Note the **.o** suffix for the output object file; this is an HP-UX convention for indicating a linkable object file):

```
ld -r file1.o file2.o -o prog.o
```

On 32-bit, link a simple FORTRAN program for use with the **dde** symbolic debugger (see *dde(1)*). Output file name is **a.out** since there is no **-o** option in the command line.

```
ld /usr/ccs/lib/crt0.o ftn.o -lc1 -lisamstub \
    -lc /opt/langtools/lib/end.o
```

To do this on 64-bit:

```
ld ftn.o -lc1 -lisamstub \
    -lc /opt/langtools/lib/pa20_64/end.o
```

On 64-bit, link a shared bound program in standard mode. Note that **crt0.o** is not specified because for shared links, it is no longer necessary.

```
ld himom.o +std -lc
```

On 64-bit, link a compatibility mode program. **crt0.o** is included because this is an archive link.

```
ld /opt/langtools/lib/pa20_64/crt0.o himom.o +compat -a archive -lc
```

Create a shared library:

```
ld -b -o libfunc.sl func1.o func2.o func3.o
```

Create a shared library with an internal name:

```
ld -b -o libfoo1.1 foo1.o foo2.o +h libfoo1.1
ln -s libfoo1.1 libfoo1.sl
cc -g mytest.c -L. -lfoo1
chatr a.out
...
```

```
shared library list:
dynamic ./libfoo1.1
```


If you do not use `+h`, the shared library does not have an internal name. The linker does not check whether `.sl` is a symbolic link. It records the library name that it looks at, if it does not have the internal name.

```
chatr a.out
...
shared library list:
dynamic /libfoo1.sl
```

On 32-bit, link a program with `libfunc.sl` but use the archive version of the C library. Specify the immediate binding mode together with the nonfatal modifier and allow verbose diagnostics to be displayed:

```
ld /usr/ccs/lib/crt0.o -B immediate -B nonfatal -B verbose \
  program.o -L . -lfunc -a archive -lc
```

To do this on 64-bit:

```
ld -B immediate -B nonfatal -B verbose \
  program.o -L . -lfunc -a archive -lc
```

On 32-bit, link a Pascal program:

```
ld /usr/ccs/lib/crt0.o main.o -lc1 -lm -lc
```

Note that in the above examples, `/usr/ccs/lib/crt0.o` can be replaced by `/opt/langtools/lib/crt0.o`.

WARNINGS

`ld` recognizes several names as having special meanings. The symbol `__end` is reserved by the linker to refer to the first address beyond the end of the program's address space. Similarly, the symbol `__edata` refers to the first address beyond the initialized data, and the symbol `__etext` refers to the first address beyond the program text. The symbols `end`, `edata`, and `etext` are also defined by the linker, but only if the program contains a reference to these symbols and does not define them (see *end(3C)* for details). On 32-bit, the symbol `__tdsize` is the total thread local storage size required by the program or shared library.

On 64-bit, the linker defines a few more symbols. The symbol `__TLS_SIZE` is the total thread local storage size. The symbol `__FPU_STATUS` is the initial status of the FPU status register. The symbol `__SYSTEM_ID` is the largest architecture revision level used by any compilation unit.

The linker treats a user definition of any of the symbols listed here as an error.

Through its options, the link editor gives users great flexibility. However, those who invoke the linker directly must assume some added responsibilities. Input options should ensure the following properties for programs:

- When the link editor is called through `cc(1)`, a start-up routine is linked with the user's program. This routine calls `exit(2)` after execution of the main program. If users call `ld` directly, they must ensure that the program always calls `exit()` rather than falling through the end of the entry routine.
- When linking for use with the symbolic debugger `dde`, the user must ensure that the program contains a routine called `main`. Also, the user must link in the file `/opt/langtools/lib/end.o` on 32-bit and `/opt/langtools/lib/pa20_64/end.o` on 64-bit as the last file named on the command line.

There is no guarantee that the linker will pick up files from archive libraries and include them in the final program in the same relative order that they occur within the library.

The linker emits warnings where ever it detects any compatibility issues. Among other things, these issues include architectural ones, as well as functionality that may change over time. Some of these include:

- Linking a PA 2.0 object file, which will not run on a PA 1.x system.
- Incremental loading with the `-A` option.
- Procedure call parameter and return type checking, including the `-C` option.
- Symbols with the same name but different types, such as CODE and DATA.
- Checking of unsatisfied symbols by the linker, which sometimes skips certain object files from an archived library. This warning is only given if the `-v` option is also provided.

- Versioning of objects within a shared library.

These messages can be turned off with the **+vnocompatwarnings** option.

As noted in the *Options* section, certain options no longer exist in a 64-bit linker. They are:

- **-q**
- **-A**
- **-C**
- **-E**
- **-Q**
- **-S**
- **-T**
- **-X**
- **+dpv**

Some restrictions exist for 64-bit mode with executables built with shared libraries and with S-bit set (through **+s** enabled). The behavior of 64-bit executables does not match the behavior of equivalent 32-bit executables or executables built with **+compat**. For any **setuid** or **setgid** programs, **dld** disables any dynamic library searching through environment variables, **SHLIB_PATH** for 32-bit and 64-bit mode, and **LD_LIBRARY_PATH** for 64-bit mode. If a library only exists in the directory specified in **SHLIB_PATH** (or **LD_LIBRARY_PATH**), you get the runtime error "library not found" if the program is a **setuid** program (that is, **uid** not equal to **euid** or **gid** not equal to **egid**) and it depends on that library. **dld** uses the dynamic path lookup (with **SHLIB_PATH**) only if the following conditions are satisfied: **getuid () == geteuid () && getgid () == getegid ()**. That is, if the **uid** or **gid** does not match its effective counterpart, **dld** searches only the recorded library path. As a result, **dld** does not check the **SHLIB_PATH** which causes the runtime error "library not found".

AUTHOR

ld was developed by AT&T, the University of California, Berkeley, and HP.

FILES

| | |
|---|--|
| /usr/lib/lib* | 32-bit system archive and shared libraries |
| /usr/lib/pa20_64/lib* | 64-bit system archive and shared libraries |
| /usr/ccs/lib* | 32-bit development archive and shared libraries |
| /opt/langtools/lib/pa20_64 | 64-bit development object files, archive and shared libraries |
| a.out | output file |
| /usr/lib/dld.sl | 32-bit dynamic loader |
| /usr/lib/pa20_64/dld.sl | 64-bit dynamic loader |
| /opt/langtools/lib/end.o | for use with the 32-bit dde debugger |
| /opt/langtools/lib/pa20_64/end.o | for use with the 64-bit dde debugger |
| /usr/ccs/lib/crt0.o | 32-bit run-time startup file |
| /opt/langtools/lib/crt0.o | Identical to /usr/ccs/lib/crt0.o |
| /opt/langtools/lib/pa20_64/crt0.o | 64-bit run-time startup file |
| /usr/ccs/lib/dyncall.o | 32-bit only. Used with -A option links |
| /opt/langtools/lib/mcrt0.o | 32-bit run-time startup with profiling (see <i>prof(1)</i>) |
| /opt/langtools/lib/pa20_64/mcrt0.o | 64-bit run-time startup with profiling |
| /usr/lib/milli.a | 32-bit millicode library automatically searched by ld |
| /usr/lib/pa20_64/milli.a | 64-bit millicode library automatically searched by ld |
| /opt/langtools/lib/gcrt0.o | run-time start-up with profiling (see <i>gprof(1)</i>) |
| /opt/langtools/lib/pa20_64/gcrt0.o | 64-bit run-time start-up with profiling (see <i>gprof(1)</i>) |

| | |
|---|--|
| <code>/opt/langtools/lib/icrt0.o</code> | 32-bit run-time start-up with profiling (see discussion of profile based optimization above.) |
| <code>/opt/langtools/lib/pa20_64/icrt0.o</code> | 64-bit run-time start-up with profiling (see discussion of profile based optimization above.) |
| <code>/usr/lib/nls/\$LANG/ld.cat</code> | message catalog |
| <code>/var/tmp/ld*</code> | temporary files |
| <code>flow.data</code> | file containing profile data generated by running an instrumented executable |
| <code>/usr/ccs/bin/fdp</code> | program for creating the procedure link order from a profile database file created by an instrumented executable; forked by the <code>-P</code> option |
| <code>/opt/langtools/lbin/ucomp</code> | PA-RISC code generator |

SEE ALSO**Profiling and Debugging Tools:**

| | |
|-----------------------|---|
| <code>adb(1)</code> | absolute debugger |
| <code>gprof(1)</code> | display call graph profile data |
| <code>prof(1)</code> | display profile data |
| <code>dde(1)</code> | C, C++, FORTRAN, and Pascal symbolic debugger |

System Tools:

| | |
|--------------------------|--|
| <code>aCC(1)</code> | invoke the HP-UX aC++ compiler |
| <code>ar(1)</code> | create archived libraries |
| <code>CC(1)</code> | invoke the HP-UX C++ compiler |
| <code>cc(1)</code> | invoke the HP-UX C compiler |
| <code>chatr(1)</code> | change program's internal attributes |
| <code>exec(2)</code> | execute a file |
| <code>f77(1)</code> | invoke the HP-UX FORTRAN 77 compiler |
| <code>f90(1)</code> | invoke the HP-UX Fortran 90 compiler |
| <code>fastbind(1)</code> | invoke the fastbind tool |
| <code>nm(1)</code> | print name list of object file |
| <code>pc(1)</code> | invoke the HP-UX Pascal compiler |
| <code>strip(1)</code> | strip symbol and line number information from an object file |

Miscellaneous:

| | |
|------------------------|---|
| <code>a.out(4)</code> | assembler, compiler, and linker output |
| <code>ar(4)</code> | archive format |
| <code>crt0(3)</code> | execution startup routine |
| <code>dld.sl(5)</code> | dynamic loader |
| <code>end(3C)</code> | symbol of the last locations in program |

Texts and Tutorials:

| | |
|---|--|
| <i>HP-UX Linker and Libraries Online User Guide</i> | (See the <code>+help</code> option) |
| <i>HP-UX Linker and Libraries User's Guide</i> | (See <i>manuals(5)</i> for ordering information) |

STANDARDS CONFORMANCE

ld: SVID2, SVID3, XPG2, XPG4

NAME

ldd - list dynamic dependencies of executable files or shared libraries

SYNOPSIS

ldd [-b] [-d] [-r] [-s] [-v] *filename...*

DESCRIPTION

ldd is a command that can list the dynamic dependencies of incomplete executable files or shared libraries.

ldd lists verbose information about dynamic dependencies and symbol references. If the object file is an executable file, **ldd** lists all shared libraries that would be loaded as a result of executing the file. If it is a shared library, **ldd** lists all shared libraries that would be loaded as a result of loading the library.

ldd uses the same algorithm as the dynamic loader (`/usr/lib/dld.sl` and `/usr/lib/pa20_64/dld.sl`) to locate the shared libraries at runtime. See "Dynamic Path List" in *dld.sl*(5) for more information.

Options

ldd recognizes the following options:

- b *PA32-only*. Used in conjunction with **-d** and/or **-r** to force `dld.sl` to bind all dependent libraries and report unsats. By default, the smartbind mechanism in `dld.sl` only binds libraries whose symbols are explicitly referenced.
- d Check reference to data symbols.
- r Check reference to data and code symbols.
- s Display the search path used to locate the shared libraries.
- v Display all dependency relationships.

EXTERNAL INFLUENCES**Environment Variables**

ldd uses the following environment variables to locate shared libraries.

LD_LIBRARY_PATH

64-bit mode: A colon-separated list of path names which defines the search path for shared libraries at runtime. See "Dynamic Path List" in *dld.sl*(5) for more information.

SHLIB_PATH

A colon-separated list of path names which defines the search path for shared libraries at runtime. See "Dynamic Path List" in *dld.sl*(5) for more information.

The following internationalization variables affect the execution of **ldd**:

LANG

Determines the locale category for native language, local customs and coded character set in the absence of **LC_ALL** and other **LC_*** environment variables. If **LANG** is not specified or is set to the empty string, a default of **C** (see *lang*(5)) is used instead of **LANG**.

LC_ALL

Determines the values for all locale categories and has precedence over **LANG** and other **LC_*** environment variables.

LC_MESSAGES

Determines the locale that should be used to affect the format and contents of diagnostic messages written to standard error.

LC_NUMERIC

Determines the locale category for numeric formatting.

LC_CTYPE

Determines the locale category for character handling functions.

NLSPATH

Determines the location of message catalogs for the processing of **LC_MESSAGES**.

If any internationalization variable contains an invalid setting, **ldd** behaves as if all internationalization variables are set to **C**. See *environ*(5).

DIAGNOSTICS

ldd prints the record of shared library path names to stdout. The optional list of symbol resolution problems are printed to stderr.

ldd returns zero when the operation is successful. A non-zero return code indicates that an error occurred.

EXAMPLES

By default **ldd** prints a simple dynamic path information. This consists of the dependencies recorded in the executable (or the shared library) followed by the physical location where these libraries are found.

```
ldd a.out
./libx.sl =>      ./libx.sl
libc.2 =>         /lib/pa20_64/libc.2
libdl.1 =>        /lib/pa20_64/libdl.1
```

The **-v** option causes **ldd** to print dependency relationship along with the dynamic path information.

```
ldd -v a.out
find library=./libx.sl; required by a.out
./libx.sl =>      ./libx.sl
find library=libc.2; required by a.out
libc.2 =>         /lib/pa20_64/libc.2
find library=libdl.1; required by /lib/pa20_64/libc.2
libdl.1 =>        /lib/pa20_64/libdl.1
```

The **-r** option to **ldd** causes it to analyze all symbol references and print information about unsatisfied code and data symbols.

```
ldd -r a.out
./libx.sl =>      ./libx.sl
libc.2 =>         /lib/pa20_64/libc.2
libdl.1 =>        /lib/pa20_64/libdl.1
symbol not found: vall (./libx.sl)
symbol not found: count (./libx.sl)
symbol not found: func1 (./libx.sl)
```

WARNINGS

ldd does not list shared libraries explicitly loaded using *dlopen(3C)* or *shl_load(3X)*.

FILES

| | |
|-------------------------------------|--|
| a.out | output file |
| /usr/lib/dld.sl | 32-bit PARISC dynamic loader |
| /usr/lib/pa20_64/dld.sl | 64-bit PARISC dynamic loader |
| /usr/ccs/lib/lddstub | 32-bit dummy executable loaded to check the dependencies of shared libraries |
| /usr/ccs/lib/pa20_64/lddstub | 64-bit dummy executable loaded to check the dependencies of shared libraries |
| /usr/lib/nls/\$LANG/ldd.cat | message catalog |

SEE ALSO**System Tools:**

ld(1) invoke the link editor

Miscellaneous:

a.out(4) assembler, compiler, and linker output
dld.sl(5) dynamic loader

Texts and Tutorials:

HP-UX Linker and Libraries User's Guide

NAME

leave - remind you when you have to leave

SYNOPSIS

leave [*hhmm*]

DESCRIPTION

The **leave** command waits until the specified time, then reminds you to leave. You are reminded 5 minutes and 1 minute before the actual time, at the time, and every minute thereafter. When you log off, **leave** exits.

The time of day is in the form *hhmm*, where *hh* is a time in hours (which can range from 0 through 11 or 0 through 24 hours), and *mm* is the number of minutes after the specified hour. If the value of *hh* is greater than 11 (24-hour clock time), the specified value is reduced by 12 to a new value in the range of 0 through 11, thus ensuring that the alarm time is always set to activate within the next 12 hours. For example, if *hhmm* is 1350 and the current time is 4:00 PM (1600), the 1350 value is changed to 150 and the alarm is set for 1:50 AM, nine hours and 50 minutes later. On the other hand, if it is 9:00 AM and *hhmm* is specified as 2200 (10:00 PM), the value used is converted to 1000 and the alarm is set for one hour later instead of 13 hours as specified.

If no argument is provided, **leave** prompts with

When do you have to leave?

A reply of newline causes **leave** to exit; otherwise the reply is assumed to be a time. This form is suitable for inclusion in a **.login** or **.profile** file.

The **leave** command ignores interrupts, quits, and terminate signals. To get rid of it you should either log off or use **kill -9** giving its process ID.

EXAMPLES

The command

leave 1204

sends an alarm (a beep) to your terminal to remind you that you have to leave at 12:04 and reminds you that you are late at one minute intervals after 12:04.

WARNINGS

The **leave** command checks to see if a user has logged out by checking the **/etc/utmp** file every 100 seconds. If a user logs out and logs back in to the same tty before **leave** makes its periodic check, **leave** may not know that the user has logged out.

AUTHOR

leave was developed by the University of California, Berkeley.

FILES

/etc/utmp

SEE ALSO

calendar(1).

NAME

lifcp - copy to or from LIF files

SYNOPSIS

lifcp [-Txxx] [-Lxxx] [-vxxx] [-a] [-b] [-i xxx] [-r] [-t] *file1 file2*

lifcp [-Txxx] [-Lxxx] [-vxxx] [-a] [-b] [-i xxx] [-r] [-t] [*file1 file2 ...*] *directory*

DESCRIPTION

lifcp copies a LIF file to an HP-UX file, an HP-UX file to a LIF file, or a LIF file to another LIF file. It also copies a list of (HP-UX/LIF) files to a (LIF/HP-UX) directory. The last name on the argument list is the destination file or directory.

Options can be used singly or combined in any order before the file names. The space between option and argument is optional.

- Txxx Used only when copying files to a LIF volume. This option forces the file type of the LIF directory entry to be set to the argument given. The argument can be decimal, octal or hex, using standard "C" notation.
- Lxxx Used only when copying files to a LIF volume. This option will set the "last volume flag" to xxx (0 or 1). The default "last volume flag" is 1.
- vxxx Used only when copying files to a LIF volume. This option sets the "volume number" to xxx. The default "volume number" is one.
- a This option forces a ASCII mode of copying regardless of the file type. When copying in ASCII mode from HP-UX to LIF the default file type is BINARY (1). (For details on available modes of copying refer to *lif*(4)). This option has no effect when copying from LIF to LIF.
- b This option forces a BINARY mode of copying regardless of the file type. When copying in BINARY mode from HP-UX to LIF the default file type is BINARY (-2). (For details on available modes of copying refer to *lif*(4)). This option has no effect when copying from LIF to LIF.
- i xxx Used only when copying files to a LIF volume. This option sets the "implementation" field of the LIF directory entry to the argument given. The argument value can be decimal, octal or hex, using standard "C" notation. The "implementation" field can only be set for file types -2001 to -100000 (octal). The "implementation" field is set to zero for all interchange file types and for file types -2 to -200 (octal). Note that the "implementation" value controls the attributes of the LIF file with regard to protection and record sizes. **lifs -l** or **lifs -i** can be used to determine the "implementation" value of a file.
- r Forces RAW mode copying regardless of file type. When copying in RAW mode from HP-UX to LIF the default file type is BIN (-23951). -T option overrides the default file type. (various modes of copying are explained in *lif*(4).) -r option has no effect in LIF to LIF copy operations.
- t causes HP-UX file names to be translated to a name acceptable by a LIF utility; that is, all lowercase letters are converted to uppercase and all other characters except numerics are changed to an underscore (_). If the HP-UX file name starts with a nonletter, the file name is preceded by the capital letter X. Thus, for example, if two files named colon (:) and semicolon (;), were copied, both of them would be translated to X_. File names are truncated to a maximum of 10 characters. When copying a LIF file to an HP-UX or LIF file, -t has no effect. Omitting -t causes an error to be generated if an improper name is used.

The default copying modes when copying from LIF to HP-UX are summarized in the following table:

| File Type | Default Copying Mode |
|-----------|----------------------|
| ASCII | ASCII |
| BINARY | BINARY |
| BIN | RAW |
| other | RAW |

When copying from HP-UX to LIF, the default copying mode is ASCII and an ASCII file is created.

When copying from LIF to LIF, if no options are specified, then all the LIF directory fields and file contents are duplicated from source to destination.

A LIF file name is recognized by the embedded colon (:) delimiter (see *lif(4)* for LIF file naming conventions). A LIF directory is recognized by a trailing colon. If an HP-UX file name containing a colon is used, the colon must be escaped with two backslash characters (\\) (the shell removes one of them).

The file name - (dash) is interpreted to mean standard input or standard output, depending on its position in the argument list. This is particularly useful if the data requires nonstandard translation. When copying from standard input, if no other name can be found, the name "STDIN" is used.

LIF file naming conventions are known only to the LIF utilities. Since file name expansion is done by the shell, this mechanism cannot be used for expanding LIF file names.

Do not mount the special file while using *lifcp*.

DIAGNOSTICS

lifcp returns exit code 0 if the file is copied successfully. Otherwise it prints a diagnostic and returns nonzero.

EXAMPLES

Copy HP-UX file **abc** to LIF file **CDE** on LIF volume **lifvol** which is actually an HP-UX file initialized to be a LIF volume:

lifcp abc lifvol:CDE

Copy all the HP-UX files in the current directory to the LIF volume **lifvol** which is present in the parent directory. File names are translated to appropriate LIF file names.

lifcp -t * ../lifvol:

Copy all the HP-UX object files in the current directory to the LIF volume *lifvol*. Copying mode is RAW and LIF file types are set to -5555.

lifcp -r -T -5555 -t *.o lifvol:

Copy all the object files in the current directory to the LIF volume **lifvol**. Copying mode is BINARY and LIF BINARY files are created.

lifcp -r -T 0xffff961 -i 0x20200080 bdat lifvol:BDAT

Copy a BDAT file, without a password, from a BASIC WorkStation to an HP-UX LIF volume **lifvol**. Note that **-i** controls protection and record size attributes. The file type for a BDAT file is -5791 (or 0xffff961) and its record size is 256 bytes per record.

lifcp -b *.o lifvol:

Copy all files in the current directory to the LIF volume **lifvol** in the **root** directory. Copying mode is RAW and LIF file types are set to BIN.

lifcp -r -t * /lifvol:

Copy file **abc**: to LIF file **CDE** in **lifvol**.

lifcp abc\\: lifvol:CDE

Copy files **abc** and **def** to LIF files **ABC** and **DEF** within **lifvol**.

lifcp -t abc def lifvol:

Copy LIF file **ABC** within **lifvol** to file **ABC** within current directory.

lifcp lifvol:ABC .

Copy standard input to LIF file **A_FILE** on LIF volume **/dev/dsk/c0t6d0**.

lifcp - /dev/dsk/c0t6d0:A_FILE

Copy LIF file **ABC** in **lifvol** to LIF file **CDE** on **/dev/dsk/c0t6d0** .

lifcp lifvol:ABC /dev/dsk/c0t6d0:CDE

Copy the output of *pr* to the LIF file **ABC**.

pr abc | lifcp - lifvol:ABC

Copy the output of *pr* to the LIF volume **lifvol**. LIF file **STDIN** is created since no files names are specified.

pr abc | lifcp - lifvol:

Copy LIF file **ABC** in **lifvol** to standard output.

lifcp lifvol:ABC -

Copy all files in current directory to LIF files of the same name on LIF volume **lifvol** (may cause errors if file names in the current directory do not obey LIF naming conventions!).

lifcp * ../lifvol:

DEPENDENCIES

Series 700/800

The following option is also supported:

-Knnn forces each file copied in to begin on a *nnn* × 1024-byte boundary from the beginning of the volume. This is useful when files are used for Series 700/800 boot media. This option has no effect when copying from a LIF volume.

AUTHOR

lifcp was developed by the Hewlett-Packard Company.

SEE ALSO

lifinit(1), lifls(1), lifrename(1), lifrm(1), lif(4).

NAME
lifinit - write LIF volume header on file

SYNOPSIS
lifinit [-vnnn] [-dnnn] [-n string] [-snnn] [-lnnn] [-ennn] file

DESCRIPTION
lifinit writes a LIF volume header on a volume or file.

- Options**
lifinit recognizes the following options and command-line arguments which can appear in any order:
- vnnn Sets volume size to nnn bytes. If nnn is not a multiple of 256, it is rounded down to the next such multiple.
 - dnnn Sets directory size to nnn file entries. If nnn is not an integer multiple of 8, it is rounded up to next such multiple.
 - n string Sets the volume name to be string. If the -n option is not specified, the volume name is set to the last component of the path name specified by file. A legal LIF volume name is 6 characters long and is limited to uppercase letters (A-Z), digits (0-9) and the underscore character (_). The first character (if any) must be a letter. The utility automatically performs translation to create legal LIF volume names. Therefore, all lowercase letters are converted to uppercase, and all other characters except numeric and underscore are replaced with a capital letter X. If the volume name does not start with a letter, the volume name is preceded by a capital letter X. The volume name is also right-padded with spaces or truncated as needed to be six characters long. If -n is used with no string, the default volume name is set to 6 spaces.
 - snnn set the initial system load (ISL) start address to nnn in the volume label. This is useful when building boot media for Series 700/800 systems.
 - lnnn specifies the length in bytes of the ISL code in the LIF volume.
 - ennn set the ISL entry point to nnn bytes from the beginning of the ISL. For example, specifying -e3272 means that the ISL entry point is 3272 (decimal) bytes from the beginning of the ISL object module.
 - Knnn forces the directory start location to be the nearest multiple of nnn × 1024 bytes from the beginning of the volume. This is necessary for booting Series 700/800 systems from LIF media.

If file does not exist, a regular HP-UX disk file is created and initialized.

The default values for volume size are 256 kilobytes for regular files, and the actual capacity of the device for device files.

The default directory size is a function of the volume size. A percentage of the volume size is allocated to the volume directory as follows:

| Volume Size | Directory Size |
|-------------|----------------|
| < 2MB | ~1.3% |
| > 2MB | ~0.5% |

Each directory entry occupies 32 bytes of storage. The actual directory space is subject to the rounding rules stated above.

Do not mount the special file while using lifinit.

RETURN VALUE
lifinit returns exit code 0 if the volume is initialized successfully. Otherwise it prints a diagnostic message and returns nonzero.

EXAMPLES
Initialize file x to be a LIF volume containing 500 000 bytes with 10 directory file entries:

```
lifinit -v500000 -d10 x
```

Initialize device /dev/rdisk/c0t6d0 as a LIF volume using default initialization conditions (device must not be a mounted file system device):

```
lifinit /dev/rdisk/c0t6d0
```

WARNINGS

To prevent media corruption, do not terminate *lifinit* once it has started executing.

AUTHOR

lifinit was developed by HP.

SEE ALSO

lifcp(1), lifls(1), lifrename(1), lifrm(1), lif(4).

NAME

lifls - list contents of a LIF directory

SYNOPSIS

lifls [**option**] *name*

DESCRIPTION

lifls lists the contents of a LIF directory on standard output. The default output format lists file names in multiple columns (similar to *ls*(1), except unsorted) if standard output is a character special file. If standard output is not a tty device, the output format is one file name per line. *name* is a path name to an HP-UX file containing a LIF volume and optional file name. If *name* is a volume name, the entire volume is listed. If *name* is of the form *volume:file*, only the file is listed. The following options are available, and only one option should be specified with a given command:

- l List in long format, giving volume name, volume size, directory start, directory size, file type, file size, file start, "implementation" field (in hex), date created, last volume, and volume number.
- C Force multiple column output format regardless of standard output type.
- L Return the content of the "last volume flag" in decimal.
- i Return the content of the "implementation" field in hex.
- v Return the content of the "volume number" in decimal.
- b *blist*
Report only on files using block numbers specified on the command line in *blist*, a comma separated list of block numbers in DEV_BSIZE units.

Do not mount the special file while using lifls.

DIAGNOSTICS

lifls returns zero if the directory was listed successfully. Otherwise it prints a diagnostic and returns nonzero.

EXAMPLES

lifls -C /dev/rdisk/c0t6d0

AUTHOR

lifls was developed by HP.

SEE ALSO

lifcp(1), lifinit(1), lifrename(1), lifrm(1), lif(4).

NAME

lifrename - rename LIF files

SYNOPSIS

lifrename *oldfile newfile*

DESCRIPTION

oldfile is a full LIF file specifier (see *lif(4)* for details) for the file to be renamed (e.g. **liffile:A_FILE**). *newfile* is new name to be given to the file (only the file name portion). This operation does not include copy or delete. Old file names must match the name of the file to be renamed, even if that file name is not a legal LIF name.

*Do not mount the special file while using **lifrename**.*

DIAGNOSTICS

lifrename returns zero if the file name is changed successfully. Otherwise it prints a diagnostic and returns nonzero.

EXAMPLES

```
lifrename liffile:A_FILE B_FILE
lifrename /dev/dsk/c0t6d0:ABC CDE
```

AUTHOR

lifrename was developed by HP.

SEE ALSO

lifcp(1), lifinit(1), lifls(1), lifrm(1), lif(4).

NAME

lifrm - remove a LIF file

SYNOPSIS

lifrm *file1* ... *filen*

DESCRIPTION

lifrm removes one or more entries from a LIF volume. File name specifiers are as described in *lif*(4).

Do not mount the special file while using lifrm.

DIAGNOSTICS

lifrm returns zero if the file is removed successfully. Otherwise it prints a diagnostic and returns nonzero.

EXAMPLES

lifrm **liffile:MAN**
lifrm **/dev/rdisk/c0t6d0:F**

AUTHOR

lifrm was developed by HP.

SEE ALSO

lifcp(1), lifinit(1), lifls(1), lifrename(1), lif(4).

NAME

`line` - read one line from user input

SYNOPSIS

`line [-t timeout]`

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.

Options

`line` recognizes the following command-line option:

- `-t timeout` Timeout after *timeout* seconds where *timeout* is an integer value (if a non-integer value is specified, it is converted to an integer; i.e., rounded down). A blank is required between `-t` and the *timeout* argument. This option is not documented in POSIX and other industry standards, and should not be used in portable applications.

EXTERNAL INFLUENCES**International Code Set Support**

Single- and multi-byte character code sets are supported.

EXAMPLES

The following lines in a shell script prompt for a file name and display information about the file:

```
echo 'Enter file name: \c'
reply='line'
ls -l $reply
```

To limit the response time to 10 seconds, use the form:

```
reply='line -t 10'
```

then test for no response. If no response occurs before timeout expires, a default behavior should be provided.

WARNINGS

This command is likely to be withdrawn from X/Open standards. Applications using this command might not be portable to other vendors' systems. As an alternative `read` is recommended.

SEE ALSO

`sh(1)`, `read(2)`.

STANDARDS CONFORMANCE

`line`: SVID2, SVID3, XPG2, XPG3

NAME

listusers - display user login data

SYNOPSIS

listusers [-g *groups*] [-l *logins*]

DESCRIPTION

The **listusers** command displays data concerning user logins. The output shows the user login and the **/etc/passwd** comment field value (e.g., user name, etc.). The default displays data about all user logins.

Options

The **listusers** command supports the following options:

- g *groups* Display all users belonging to *groups*, sorted by login. A comma separated list specifies multiple groups.
- l *logins* Display the requested *logins*. A comma separated list specifies multiple logins.

A user login has a UID of 100 or greater.

When the -l and -g options are combined, a user login is only displayed once, even though the login may belong to multiple specified groups.

EXAMPLES

List all user logins.

```
listusers
```

List all user logins in the group **cmds** and the users **bob**, **john** and **otto**, removing all duplicates.

```
listusers -g cmds -l bob,john,otto
```

FILES

/etc/passwd
/etc/group

SEE ALSO

passwd(1), logins(1M), group(4), passwd(4).

STANDARDS COMPLIANCE

listusers: SVID3

NAME

ln - link files and directories

SYNOPSIS

```
ln [-f] [-i] [-s] file1 new_file
ln [-f] [-i] [-s] file1 [file2 ...] dest_directory
ln [-f] [-i] [-s] directory1 [directory2 ...] dest_directory
```

DESCRIPTION

The **ln** command links:

- *file1* to a new or existing *new_file*,
- *file1* to a new or existing file named *file1* in existing *dest_directory*,
- *file1*, *file2*, ... to new or existing files of the same name in existing *dest_directory*,
- *directory1*, *directory2*, ... to new directories of the same name in existing *dest_directory*,
- or it creates symbolic links between files or between directories.

If links are to *dest_directory*, corresponding file or directory names in that directory are linked to *file1*, *file2*, ..., or *directory1*, *directory2*, ..., etc., as appropriate. If two or more existing files or directories (excluding destination file name *new_file*) are specified, the destination must be a directory. If *new_file* already exists as a regular file (or link to another file), its contents (or the existing link) and its ACL are destroyed only if the **-f** option is specified. The ACL on the *new_file* after the link is the same as that of the *source_file* file.

If the **-f** and **-i** options are specified and the link being created is the name of an existing link or ordinary file and the access permissions of the file forbid writing, **ln** asks permission to overwrite the file. If the access permissions of the directory forbid writing, **ln** aborts and returns with the error message:

```
cannot unlink new_file
```

(even if the file is an ordinary file and not a link to another file). When asking for permission to overwrite an existing file or link, **ln** prints the mode (see *chmod(2)* and *Access Control Lists* below), followed by the first letters of the words **yes** and **no** in the current native language, prompting for a response, and reading one line from the standard input. If the response is affirmative and is permissible, the operation occurs; if not, the command proceeds to the next source file, if any.

Hard links are created with the same ownerships and permissions as the file or directory to which they are linked. If ownership or permissions are changed on a link or file, the same changes appear on corresponding hard links. The **ln** command does not permit hard links to a directory.

Symbolic links are created with the ownership of the creator and the permissions are of the creator's current umask. Once created, the symbolic link ownership and permissions will not change, since the mode and ownership of the symbolic link is ignored by the system.

If *file1* is a file and *new_file* is a link to an existing file or an existing file with other links, *new_file* is disassociated from the existing file and links and linked to *file1*. When **ln** creates a link to a new or existing file name, ownerships and permissions are always identical to those for the file to which it is linked. If **chown**, **chgrp**, or **chmod** is used to change ownership or permissions of a file or link, the change applies to the file and all associated links. The last modification time and last access time of the file and all associated links are identical (see *chown(1)* and *chmod(1)*).

For a discussion of symbolic links, see *symlink(4)*.

Options

The **ln** command recognizes the following options:

- f** Force existing destination path names to be removed to allow the link.
- i** Write a prompt to the standard error output requesting confirmation for each link that would overwrite an existing file. This option takes effect only if used in conjunction with the **-f** option.
- s** Cause **ln** to create symbolic links instead of the usual hard links. A symbolic link contains the name of the file to which it is linked. The referenced file is used when an **open()** operation is performed on the link (see *open(2)*). A **stat()** on a symbolic link returns the linked-

to file; an `lstat()` must be performed to obtain information about the link (see `stat(2)`). A `readlink()` call can be used to read the contents of the symbolic link (see `readlink(2)`). Symbolic links may span file systems and refer to directories.

Access Control Lists (ACLs)

If optional ACL entries are associated with *new_file*, `ln` displays a plus sign (+) after the access mode when asking permission to overwrite the file.

If *new_file* is a new file, it inherits the access control list of *file1*, altered to reflect any difference in ownership between the two files (see `acl(5)` and `aclv(5)`). In JFS file systems, new files created by `ln` do not inherit their parent directory's default ACL entries (if any), but instead retain their original ACLs.

EXTERNAL INFLUENCES

Environment Variables

`LC_CTYPE` determines the interpretation of text as single byte and/or multibyte characters.

`LANG` and `LC_CTYPE` determine the local language equivalent of `y` (for yes/no queries).

`LANG` determines the language in which messages are displayed.

If `LC_CTYPE` is not specified in the environment or is set to the empty string, the value of `LANG` is used as a default for each unspecified or empty variable. If `LANG` is not specified or is set to the empty string, a default of `C` (see `lang(5)`) is used instead of `LANG`. If any internationalization variable contains an invalid setting, `ln` behaves as if all internationalization variables are set to `C`. See `environ(5)`.

International Code Set Support

Single byte and multibyte character code sets are supported.

EXAMPLES

The following command creates `file1` and `file2` in `dest_dir`, which are linked back to the original files `file1` and `file2`:

```
ln -f file1 file2 dest_dir
```

If `file1` and/or `file2` exists in the destination directory, it is removed and replaced by a link to `file1` or `file2`, respectively. If existing file `file1` or `file2` is a link to another file or a file with links, the existing file remains. Only the link is broken and replaced by a new link to `file1` or `file2`.

WARNINGS

`ln` does not create hard links across file systems.

DEPENDENCIES

NFS

Access control lists of networked files are summarized (as returned in `st_mode` by `stat()`), but not copied to the new file. When using `ln` on such files, a + is not printed after the mode value when asking for permission to overwrite a file.

AUTHOR

`ln` was developed by AT&T, the University of California, Berkeley and HP.

SEE ALSO

`cp(1)`, `cpio(1)`, `mv(1)`, `rm(1)`, `link(1M)`, `readlink(2)`, `stat(2)`, `symlink(2)`, `symlink(4)`, `acl(5)`, `aclv(5)`.

STANDARDS CONFORMANCE

`ln`: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

locale - get locale-specific (NLS) information

SYNOPSIS

```
locale [ -a [ 32 | 64 ] | -A | -m ]
locale [ -ck ] name ...
```

DESCRIPTION

The **locale** command displays information about the current locale or about available locales.

When invoked without arguments, **locale** displays the name and actual or implied value of each of the locale-related environment variables in the order shown below, one per line:

```
LANG
LC_CTYPE
LC_COLLATE
LC_MONETARY
LC_NUMERIC
LC_TIME
LC_MESSAGES
LC_ALL
```

An actual value is the value the variable actually has in the user's environment. An implied value is derived from the value of another variable. Implied values are displayed enclosed in double quotes, while actual values are unquoted.

The determination of implied values is that if the variable **LC_ALL** is present and has a non-null value, that is the actual value for **LC_ALL**, and all of the other variables take its value as an implied value. If **LC_ALL** is not set, all of the **LC_*** variables that are set are shown with their value as an actual value. Any that have no value are shown with the value of the **LANG** environment variable as their implied value. **LC_ALL** is displayed as **LC_ALL=\n** if it has no value.

The **locale** command can take multiple arguments, which may be locale category names, locale keywords, or the special word **charmap** (see *localedef(1M)* for a description of locale keywords and charmaps). If an argument is a keyword, the value associated with that keyword in the current environment is displayed and possibly other information, depending on selected options. If an argument is a category name (i.e., **LC_***), the values of all keywords defined in that category are displayed. If an argument is the special word **charmap**, the charmap file (if any) that was used in the definition of the current locale is displayed.

Non-printable characters are printed as hexadecimal values in the form,

```
\xhh
```

except that if a different escape character has been defined for the locale, it is displayed instead of the "\".

Options

The following options are available:

- a** List all available locales. These are the possible meaningful values that can be assigned to **LANG** or any of the **LC_*** variables on the system. They are dependent upon which locales have been installed on the system. By default on a 32-bit system, the locales in **/usr/lib/nls/loc/locales** are listed. By default on a 64-bit system, the locales in **/usr/lib/nls/loc/pa20_64/locales** are listed. This option takes 32 (for ILP32, 32-bit int, long, pointer, 32-bit offset) or 64 (for LP64, 64-bit long, pointer, 64-bit offset) as its argument.
 - a** Display 32-bit locales for 32-bit and 64-bit systems.
 - a 32** Display 32-bit locales for 32-bit and 64-bit systems.
 - a 64** Display only 64-bit locales on a 64-bit system. If executed on a 32-bit system, an error message is returned.
- A** List 32-bit locales on a 32-bit system. List both 32-bit and 64-bit bit locales on a 64-bit system.
- m** Display a list of available charmap files on the system. See *localedef(1M)* for a definition of charmap files and their usage.

- c** Display names of locale categories that have been selected either explicitly or by giving a keyword contained therein. This option may be used with the **-k** option.
- k** Display names of keywords that have been selected either explicitly or by providing their containing category as an argument. Keyword names and values are displayed as:
 <keyword>=<value>
 Without the **-k** option, only the values are displayed. This option can be used with the **-c** option.
- name* Specify the locale category name, locale keyword, or the special word charmap.

EXTERNAL INFLUENCES

Environment Variables

LANG provides a default value for the internationalization variables that are unset or null. If **LANG** is unset or null, the default value of "C" (see *lang(5)*) is used. If any of the internationalization variables contains an invalid setting, **locale** will behave as if all internationalization variables are set to "C". See *environ(5)*.

LC_ALL, when set to a non-empty string value, overrides the values of all other internationalization variables.

LC_CTYPE determines the interpretation of text as single and/or multi-byte characters, the classification of characters as printable, and the characters matched by character class expressions in regular expressions.

LC_MESSAGES determines the locale that should be used to affect the format and content of diagnostic messages written to standard error, and informative messages written to standard output.

NLSPATH determines the location of message catalog for the processing of **LC_MESSAGES**.

International Code Set Support

Single- and multi-byte character code sets are supported.

RETURN VALUE

The **locale** command exits with one of the following values:

- 0 All requested information was found and displayed successfully.
- >0 An error occurred in either finding or displaying the information.

EXAMPLES

If the locale environment variables are set as:

```
LANG=fr_FR.iso88591
LC_COLLATE=C
```

the command:

```
locale
```

gives the following output:

```
LANG=fr_FR.iso88591
LC_CTYPE="fr_FR.iso88591"
LC_COLLATE=C
LC_MONETARY="fr_FR.iso88591"
LC_NUMERIC="fr_FR.iso88591"
LC_TIME="fr_FR.iso88591"
LC_MESSAGES="fr_FR.iso88591"
LC_ALL=
```

The command:

```
LC_ALL=POSIX locale -ck decimal_point
```

produces:

```
LC_NUMERIC decimal_point="."
```

If **LANG** is set to **POSIX** and no other locale variables are set, the command:

```
    locale LC_NUMERIC
produces:
    "."
    ""
    ""
```

which correspond to the keywords *decimal_point*, *thousands_sep*, *grouping*, and *alt_digit*.

SEE ALSO

localedef(1M), localeconv(3C), nl_langinfo(3C), setlocale(3C), charmap(4), localedef(4), environ(5), lang(5).

STANDARDS CONFORMANCE

locale: XPG4, POSIX.2

NAME

lock - reserve a terminal

SYNOPSIS

lock

DESCRIPTION

lock requests a password from the user, then prints **LOCKED** on the terminal and refuses to relinquish the terminal until the password is repeated. If the user forgets the password, the only recourse is to log in elsewhere and kill the lock process.

NAME

logger - make entries in the system log

SYNOPSIS

logger [-t *tag*] [-p *pri*] [-i] [-f *file*] [*message*...]

DESCRIPTION

The **logger** command provides a program interface to the **syslog()** system log module (see *syslog(3C)*).

A message can be given on the command line, which is logged immediately, or a file is read and each line is logged. If no *file* or *message* is specified, the contents of the standard input are logged.

Options

The **logger** command recognizes the following command-line options and arguments:

- t *tag* Mark every line in the log with the specified *tag*. The default is the value returned by **getlogin()** (see *getlogin(3C)*). If **getlogin()** returns NULL, **syslog** is the default.
- p *pri* Enter the message with the specified priority. The priority can be specified numerically or as a *facility.level* pair. For example, -p **local3.info** logs the message or messages as informational level in the **local3** facility. The default is **user.notice**.
- i Log the process ID of the logger process with each line.
- f *file* Log the contents of the specified file.
- message* The message to log; if not specified, the file specified by the -f option or standard input is logged.

EXTERNAL INFLUENCES**Environment Variables**

LC_MESSAGES determines the language in which messages are displayed.

If **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**.

If any internationalization variable contains an invalid setting, **logger** behaves as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single- and multi-byte character code sets are supported.

EXAMPLES

Send the message **System rebooted** to the **syslogd** daemon:

```
logger System rebooted
```

Send output from the **users** command (see *users(1)*) to the **syslogd** daemon, marked as level **info** and facility **local0**. The message is tagged with the string **USERS**:

```
users | logger -p local0.info -t USERS
```

Send the message **System going down immediately!!!** to the **syslog** daemon, at the **emerg** level and **user** facility:

```
logger -p user.emerg "System going down immediately!!!"
```

WARNINGS

The **logger** command has no effect if the **syslogd** daemon (see *syslogd(1M)*) is not running on the system.

Messages written in locales other than the POSIX/C locale are not supported.

AUTHOR

logger was developed by the University of California, Berkeley.

SEE ALSO

syslogd(1M), getlogin(3C), syslog(3C).

STANDARDS CONFORMANCE

logger: XPG4, POSIX.2

NAME

login - sign on; start terminal session

SYNOPSIS

login [*name* [*env-var*] ...]

DESCRIPTION

The **login** command is used at the beginning of each terminal session to properly identify a prospective user. **login** can be invoked as a user command or by the system as an incoming connection is established. **login** can also be invoked by the system when a previous user shell terminates but the terminal does not disconnect.

If **login** is invoked as a command, it must replace the initial command interpreter (the user's login shell). This is accomplished with the shell command

```
exec login
```

The user's login name is requested, if it is not specified on the command line, and the corresponding password is obtained, if required, with the following prompts:

```
login:
Password:
```

Terminal echo is turned off (where possible) during password entry to prevent written records of the password. If the account does not have a password, and the authentication profile for the account requires one, **login** invokes **pam_chauthtok()** to establish one for the account. On a trusted system, **login** displays the last successful and unsuccessful login times and terminal devices.

As a security precaution, some installations use an option that requires a second "dialup" password. This occurs only for dialup connections, and is requested with the prompt:

```
dialup password:
```

Both passwords must be correct for a successful login (see *dialups(4)* for details on dialup security).

If password aging is activated, the user's password may have expired. **pam_chauthtok()** is invoked to change the password. In an untrusted environment, the user is required to re-login after a successful password change (see *passwd(1)*).

After three unsuccessful login attempts, a **HANGUP** signal is issued. If a login is not successfully completed within a certain period of time (for example, one minute), the terminal is silently disconnected.

After a successful login, the accounting files are updated, user and group IDs, group access list, and working directory are initialized, and the user's command interpreter (shell) is determined from corresponding user entries in the files **/etc/passwd** and **/etc/logingroup** (see *passwd(4)* and *group(4)*). If **/etc/passwd** does not specify a shell for the user name, **/usr/bin/sh** is used by default. **login** then forks the appropriate shell by using the last component of the shell path name preceded by a **-** (for example, **-sh** or **-ksh**). When the command interpreter is invoked with its name preceded by a minus in this manner, the shell performs its own initialization, including execution of profile, login, or other initialization scripts.

For example, if the user login shell is the Bourne, Korn, or POSIX shell (see *sh-bourne(1)*, *ksh(1)*, or *sh-posix(1)*, respectively), the shell executes the profile files **/etc/profile** and **\$HOME/.profile** if they exist (and possibly others as well). Depending on what these profile files contain, messages regarding mail in the user's mail file or any messages the user may have received since the user's last login may be displayed.

If the command name field is *****, a **chroot()** to the directory named in the directory field of the entry is performed. At that point, **login** is re-executed at the new level, which must have its own root structure, including a **/usr/bin/login** command and an **/etc/passwd** file.

For the normal user, the basic environment variables (see *environ(5)*) are initialized to:

```
HOME=login_directory
LOGNAME=login_name
MAIL=/var/mail/login_name
PATH=/usr/bin
SHELL=login_shell
```

login_directory, *login_name*, and *login_shell* are taken from the corresponding fields of the **passwd** file entry (see *passwd(4)*).

For superuser, **PATH** is set to:

```
PATH=: /usr/sbin: /usr/bin: /sbin
```

In the case of a remote login, the environment variable **TERM** is also set to the remote user's terminal type.

The environment can be expanded or modified by supplying additional arguments to **login**, either at execution time or when **login** requests the user's login name. The arguments can take either the form *value* or *varname=value*, where *varname* is a new or existing environment variable name and *value* is a value to be assigned to the variable.

An argument in the first form (without an equals sign) is placed in the environment as if it were entered in the form

```
Ln=value
```

where *n* is a number starting at 0 that is incremented each time a new variable name is required.

An argument in the second form (with an equals sign) is placed into the environment without modification.

If the variable name (**Ln** or *varname*) already appears in the environment, the new value replaces the older one.

There are two exceptions. The variables **PATH** and **SHELL** cannot be changed. This prevents users logged in with restricted shell environments from spawning secondary shells that are not restricted.

Both **login** and **getty** understand simple single-character quoting conventions. Typing a backslash in front of a character quotes it and allows the inclusion of such things as spaces and tabs.

If **/var/adm/btmp** is present, all unsuccessful login attempts are logged to that file. This feature is disabled if the file is not present. The **lastb** command, (see *last(1)*), displays a summary of bad login attempts for users with read access to **btmp**.

If the **/etc/securetty** file is present, login security is in effect, i.e., **root** is allowed to log in successfully only on the ttys listed in this file. Restricted ttys are listed by device name, one per line. Valid tty names are dependent on the installation. An example is

```
console
tty01
ttya1
etc.
```

Note that this feature does not inhibit a normal user from using the **su** command (see *su(1)*).

HP-UX Smart Card Login

If the user account is configured to use a Smart Card, the user password is stored in the card. This password has characteristics identical to a normal password stored on the system.

In order to login using a Smart Card account, the card must be inserted into the Smart Card reader. The user is prompted for a PIN (personal identification number) instead of a password during authentication. The prompts are:

```
login:
Enter PIN:
```

The password is retrieved automatically from the Smart Card when a valid PIN is entered. Therefore, it is not necessary to know the password, only the PIN.

The card is locked if an incorrect PIN is entered three consecutive times. It may be unlocked only by the card issuer.

SECURITY FEATURES

On a trusted system, **login** prohibits a user from logging in if any of the following is true:

- The password for the account has expired and the user cannot successfully change the password.
- The password lifetime for the account has passed.
- The time between the last login and the current time exceeds the time allowed for login intervals.

- The administrative lock on the account has been set.
- The maximum number of unsuccessful login attempts for the account has been exceeded.
- The maximum number of unsuccessful login attempts for the terminal has been exceeded.
- The administrative lock on the terminal has been set.
- The terminal has an authorized user list and the user is not on it.
- The terminal has time of day restrictions and the current time is not within the allowable period.

On a trusted system, **login** allows superuser to log in on the console unless **/etc/securetty** exists and does not contain **console**.

Refer to the **/etc/default/security** file in the *security(4)* man page for detailed information on configurable parameters that affect the behavior of this command. Currently supported parameters are:

```
ABORT_LOGIN_ON_MISSING_HOMEDIR
NOLOGIN
NUMBER_OF_LOGINS_ALLOWED
```

EXTERNAL INFLUENCES

Environment Variables

```
HOME      User's home directory.
MAIL      Where to look for mail.
PATH      Path to be searched for commands.
SHELL     Which command interpreter is being used.
TERM      User's terminal type.
varname   User-specified named variables.
Ln        User-specified unnamed variables.
```

DIAGNOSTICS

The following diagnostics appear if the associated condition occurs:

.rhosts is a soft link

The personal equivalence file is a symbolic link.

Bad .rhosts ownership

The personal equivalence file is not owned by the local user or by a user with appropriate privileges.

Bad group id

setgid() failed (see *setuid(2)*).

Bad user id

setuid() failed (see *setuid(2)*).

Cannot open password file

Consult system administrator.

Locuser too long

The indicated string was too long for **login**'s internal buffer.

Login incorrect

User name and password cannot be matched.

No /usr/bin/login or /etc/login on root

Attempted to log in to a subdirectory root that does not have a subroot login command. That is, the **passwd** file entry had shell path *****, but the system cannot find a **login** command under the given home directory.

No directory

Consult system administrator.

No Root Directory

Attempted to log in to a subdirectory root that does not exist. That is, the **passwd** file entry had shell path *****, but the system cannot **chroot()** to the given home directory.

No shell

The user shell (**/usr/bin/sh** if shell name is null in **/etc/passwd**) could not be started with the **exec** command. Consult system administrator.

No utmp entry. You must exec "login" from the lowest level "sh"

Attempted to execute **login** as a command without using the shell's **exec** internal command or from other than the initial shell. The current shell is terminated.

Remuser too long

The indicated string was too long for **login**'s internal buffer.

Terminal type too long

The indicated string was too long for **login**'s internal buffer.

Unable to change to directory *name*

Cannot **chdir** to the user's home directory.

Your password has expired. Choose a new one

Password aging is enabled and the user's password has expired.

WARNINGS

If **/etc/group** is linked to **/etc/loggingroup**, and group membership for the user trying to log in is managed by the Network Information Service (NIS), and no NIS server is able to respond, **login** waits until a server does respond.

DEPENDENCIES

Pluggable Authentication Modules (PAM)

PAM is an Open Group standard for user authentication, password modification, and validation of accounts. In particular, **pam_authenticate()** is invoked to perform all functions related to **login**. This includes retrieving the password, validating the account, and displaying error messages. **pam_chauthtok()** is invoked during password expiration or establishment.

HP Process Resource Manager

If the optional HP Process Resource Manager (PRM) software is installed and configured, the login shell is launched in the user's initial process resource group. If the user's initial group is not defined, the shell runs in the user default group (**PRMID=1**). See **prmconfig(1)** for a description of how to configure HP PRM, and **prmconf(4)** for a description of how the user's initial process resource group is determined.

AUTHOR

login was developed by AT&T and HP.

FILES

| | |
|------------------------------|---|
| \$HOME/.profile | Personal profile (individual user initialization) |
| \$HOME/.rhosts | Personal equivalence file for the remote login server |
| /etc/d_passwd | Dialup security encrypted passwords |
| /etc/dialups | Lines which require dialup security |
| /etc/hosts.equiv | System list of equivalent hosts allowing logins without passwords |
| /etc/loggingroup | Group file — defines group access lists |
| /etc/motd | Message-of-the-day |
| /etc/passwd | Password file — defines users, passwords, and primary groups |
| /etc/profile | System profile (initialization for all users) |
| /etc/securetty | List of valid ttys for root login |
| /etc/utmp | Users currently logged in |
| /tcb/files/auth/*/* | The trusted system password database |
| /var/adm/btmp | History of bad login attempts |
| /var/adm/wtmp | History of logins, logouts, and date changes |
| /var/mail/login_name | Mailbox for user <i>login_name</i> |
| /etc/default/security | Security defaults configuration file |

SEE ALSO

csh(1), groups(1), ksh(1), last(1), mail(1), newgrp(1), passwd(1), sh(1), sh-bourne(1), sh-posix(1), su(1),
getty(1M), initgroups(3C), dialups(4), group(4), passwd(4), profile(4), security(4), utmp(4), environ(5).

HP Process Resource Manager

prmconfig(1), prmconf(4) in *HP Process Resource Manager Users Guide*.

Pluggable Authentication Modules (PAM)

pam_acct_mgmt(3), pam_authenticate(3), pam_chauthtok(3).

HP-UX Smart Card Login

scpin(1), scsync(1).

NAME

logname - get login name

SYNOPSIS

logname

DESCRIPTION

logname writes the user's login name to standard output. The login name is equivalent to that returned by **getlogin()** (see *getlogin(3C)*).

EXTERNAL INFLUENCES**Environment Variables**

LANG determines the language in which diagnostic messages are displayed.

FILES

/etc/profile

SEE ALSO

env(1), login(1), getlogin(3C), logname(3C), environ(5).

STANDARDS CONFORMANCE

logname: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

lorder - find ordering relation for an object library

SYNOPSIS

lorder [*files*]

DESCRIPTION

The input consists of one or more object or archive library *files* (see *ar(1)*) placed on the command line or read from standard input. The standard output is a list of pairs of object file names, meaning that the first file of the pair refers to external identifiers defined in the second. Output can be processed by **tsort** to find an ordering of a library suitable for one-pass access by **ld** (see *tsort(1)* and *ld(1)*). Note that the link editor **ld** is capable of multiple passes over an archive in the archive format and does not require that **lorder** be used when building an archive. Using the **lorder** command may, however, allow for a slightly more efficient access of the archive during the link edit process.

The symbol table maintained by **ar** allows **ld** to randomly access symbols and files in the archive, making the use of **lorder** unnecessary when building archive libraries (see *ar(1)*).

EXTERNAL INFLUENCES**Environment Variables**

The following internationalization variables affect the execution of **lorder**:

LANG

Determines the locale category for native language, local customs and coded character set in the absence of **LC_ALL** and other **LC_*** environment variables. If **LANG** is not specified or is set to the empty string, a default of **C** (see *lang(5)*) is used instead of **LANG**.

LC_ALL

Determines the values for all locale categories and has precedence over **LANG** and other **LC_*** environment variables.

LC_COLLATE

Determines the locale category for character collation.

LC_CTYPE

Determines the locale category for character handling functions.

LC_MESSAGES

Determines the locale that should be used to affect the format and contents of diagnostic messages written to standard error.

LC_NUMERIC

Determines the locale category for numeric formatting.

NLSPATH

Determines the location of message catalogues for the processing of **LC_MESSAGES**.

If any internationalization variable contains an invalid setting, **lorder** behaves as if all internationalization variables are set to **C**. See *environ(5)*.

International Code Set Support

Single- and multi-byte character code sets are supported.

EXAMPLES

Build a new library from existing **.o** files:

```
ar cr library `lorder *.o | tsort`
```

When creating libraries with so many objects that the shell cannot properly handle the ***.o** expansion, the following technique may prove useful:

```
ls | grep '.o$' | lorder | tsort | xargs ar cq library
```

WARNINGS

Object files whose names do not end with **.o** are overlooked, even when contained in library archives. Their global symbols and references are attributed to some other file.

FILES

*/var/tmp/*symref* temporary files
*/var/tmp/*symdef*

SEE ALSO**System Tools:**

ar(1) create archived libraries
ld(1) invoke the link editor

Miscellaneous:

tsort(1) produce an ordered list of items (topological sort)

STANDARDS CONFORMANCE

lorder: SVID2, SVID3, XPG2, XPG4

NAME

lp, lpalt, cancel - print/alter/cancel requests on an LP printer or plotter

SYNOPSIS

```
lp [-c] [-ddest] [-m] [-nnumber] [-ooption] [-ppriority] [-s] [-ttitle] [-w] [file ...]
lpalt id [-ddest] [-i] [-m] [-nnumber] [-ooption] [-ppriority] [-s] [-ttitle] [-w]
cancel [id ...] [printer ...] [-a] [-e] [-i] [-uuser]
```

DESCRIPTION

The **lp** command queues files for printing. The **lpalt** command changes information in a queued request. The **cancel** command deletes a queued request.

lp Command

The **lp** command arranges for the named files, *file ...*, and associated information (collectively called a *request*) to be queued for output to a printer or plotter in the LP (line printer) subsystem. The process is called printing, regardless of the actual output device.

lp associates a unique identifier with each request and writes it to standard output, using the message:

```
request id is dest-sequence (fileinfo)
```

The request ID is *dest-sequence*, which can be used later to alter, cancel, or find the status of the request (see **lpalt** and **cancel** below, and **lpstat**(1)).

For example, in the following message,

```
request id is pr471f8e-2410 (1 file)
```

the request ID is **pr471f8e-2410**.

lp Options and Arguments

lp recognizes the following options and arguments. The keyletter options can be specified in any order. The *file ...* names must be last. Blanks are not permitted between a keyletter and its argument.

file ... Print each named file. If no file names are specified, standard input is assumed. The hyphen symbol (-) also specifies standard input and can be intermixed on the command line with file names. Files are printed in the same order in which they are specified.

-c Copy the named files to LP subsystem spooling directories.

Normally, the files are linked into a spool directory. The ownership and mode of the linked files remain unchanged. If the **-c** option is given, or linking is not possible (perhaps because the printer is not physically attached to the local system or cluster), the files are copied into the spool directories. The ownership and mode of the copies are set to allow read access to owner **lp** and group **bin** only.

If the files are linked rather than copied, any changes made to the named files after the request is made but before it is printed will be reflected in the printed output. Standard input is always copied instead of linked.

-ddest Select *dest* as the printer or class of printers that is to do the printing. If *dest* is a printer, the request will be printed only on that specific printer. If *dest* is a class, the request will be printed on the first available printer that is a member of the class. Under certain conditions (printer unavailability, file space limitation, etc.), requests for a specific *dest* might not be accepted (see **accept**(1M) and **lpadmin**(1M)).

If the **-d** option is omitted, *dest* is taken from the environment variable **LPDEST**. If that variable is unset or empty, the default queue is used, if one has been defined. If there is no default queue, or **LPDEST** is set but invalid, **lp** issues an error message and the request is not queued. Printer and class names and the default queue are defined by your LP subsystem administrator (see **lpadmin**(1M) and **lpstat**(1)).

-m Send a mail message (see **mail**(1)). to the user after the request has been printed. By default, no mail is sent upon normal completion of the print request.

-nnumber Print *number* copies of the output. The default is 1.

-ooption Specify a printer-dependent *option*. You can specify several printer options by repeating the **-o** option. For information about the options that are available for a printer supported

on your system, see the interface script for the printer name in the `/etc/lp/interface` directory.

- ppriority** Set the priority of the print request. *priority* must be in the range 0 (lowest priority) to 7 (highest priority). The priority is used to select the next spooled file for the targeted printer or class of printers. If the priority is less than the *fence*, the minimum priority set for the printer, the print request is deferred until the fence is lowered or the priority is raised. The default for a printer queue is the default priority set by the **lpadmin** or **lpfence** command (see *lpadmin*(1M) and *lpsched*(1M)). The default for a class queue is the highest default priority among printers in the class.
- s** Suppress standard output messages from **lp** such as "request id is ...". Error messages are still displayed on standard error.
- t title** Print *title* on the banner page of the output.
- w** Write a message to the user's terminal after the request has been printed. If the user is not logged in or (for remote printing) if **rlpdaemon** (see *rlpdaemon*(1M)) is not running on the user's local system, mail will be sent instead.

lpalt Command

The **lpalt** command alters a request made by a previous **lp** command, if it is not currently printing. (To requeue a currently printing request, use the **disable** command (see *enable*(1)) to stop the printer.)

lpalt Options

lpalt recognizes the following options and arguments, which can be specified in any order. Blanks are not permitted between a keyletter and its argument.

- id* Specifies the request to be altered. *id* is a request ID returned by **lp** or **lpalt**.
- d dest** Requeue the request to the named printer or class *dest*. A new unique request ID is written to standard output.
- i** Alter only local requests.
- m** Send mail upon normal completion of the print request.
- n number** Change the number of copies to *number*.
- o option** Specify a printer-dependent *option*. You can specify several printer options by repeating the **-o** option. All **-o** options from previous **lp** and **lpalt** commands for this request ID are deleted.
- ppriority** Change the request's priority to *priority*.
- s** Suppress standard output messages from **lpalt** such as "request id is ...". Error messages are still displayed on standard error.
- t title** Change the *title* on the banner page of the output.
- w** Write a message to the user's terminal after the request has been printed. If the user is not logged in or (for remote printing) if **rlpdaemon** (see *rlpdaemon*(1M)). is not running on the user's local system, mail will be sent instead.

cancel Command

The **cancel** command cancels requests that were made with the **lp** command, even if they are currently printing. At least one *id* or *printer* must be specified.

The cancellation of a request that is currently printing frees the printer to print its next available request.

cancel Options and Arguments

cancel recognizes the following options and arguments, which can be specified in any order. Blanks are not permitted between a keyletter and its argument.

- id ...* Specifies one or more requests. *id* is a request ID returned by **lp** or **lpalt**.
- printer ...* Specifies one or more printers. *printer* is the name of a printer, not a class. Either cancel the request that is currently printing on each *printer*, or, if an **-a**, **-e**, or **-u** option is specified, specify the printer on which to perform the corresponding operation.

- a** Remove all requests the user owns on each *printer*. The owner is determined by the user's login name and the host name of the machine where the **lp** command was invoked.
- e** Empty the spool queue of all requests for each *printer*. Only users with appropriate privileges can use this option.
- i** Cancel only local requests.
- u user** Remove any requests belonging to *user*. You can repeat the **-u** option to specify more users. Only users with appropriate privileges can use this option.

Printing Overview

A printer can print requests from one or two destination queues: its own private queue and an optional class queue, which can serve one or more printers. The destination queues are set up with the **lpadmin** command. The **lp** command places a printing request into a printer or class destination queue as directed by a user. The **lpsched** scheduler directs the requests from the destination queues to the printers. The **accept** and **reject** commands control whether **lp** can place requests in the destination queues. The **enable** and **disable** commands control whether **lpsched** can send a queued request to a printer. If a printer has two queues and one queue is rejecting requests, users can still direct requests to the other destination queue and have the requests printed. **lpstat** reports the current status of the destination queues and the scheduler. See *enable(1)*, *lpstat(1)*, *accept(1M)*, and *lpadmin(1M)*.

EXTERNAL INFLUENCES

Environment Variables

LANG determines the locale to use for the locale categories when both **LC_ALL** and the corresponding environment variable (beginning with **LC_**) do not specify a locale. If **LANG** is not set or is set to the empty string, a default of "C" (see *lang(5)*) is used.

LC_ALL determines the locale to use to override any values for locale categories specified by the setting of **LANG** or any environment variables beginning with **LC_**.

LC_CTYPE determines the locale for interpretation of sequences of bytes of text data as characters (e.g., single- versus multibyte characters in arguments and input files).

LC_MESSAGES determines the language in which messages are displayed.

LPDEST determines the output device or destination. If the **LPDEST** environment variable is not set, the **PRINTER** environment variable is used. The **-d dest** option takes precedence over **LPDEST**. Results are undefined when **-d** is not specified and **LPDEST** contains a value that is not a valid device or destination name.

PRINTER determines the output device or destination. If the **LPDEST** and **PRINTER** environment variables are not set, an unspecified output device is used. The **-d dest** option and the **LPDEST** environment variable takes precedence over **PRINTER**. Results are undefined when **-d** is not specified, **LPDEST** is unset, and **PRINTER** contains a value that is not a valid device or destination name.

If any internationalization variable contains an invalid setting, the commands behave as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single- and multibyte character code sets are supported.

RETURN VALUE

Exit values are:

- 0 Successful completion.
- >0 Error condition occurred.

EXAMPLES

For a HP2934A printer named **lp2**, configured with an interface script that defines the **-c** option to cause the printer to print in a compressed mode, use the following command to print **myfile** with compressed print on **lp2**:

```
lp -dlp2 -oc myfile
```

lp can be used at the end of a pipeline to print the results of a previous command. It is commonly used with the **pr** command (see *pr(1)*) to print formatted output. For a default printer, to format file **.profile** into pages and print three copies of it:

```
pr .profile | lp -n3
```

WARNINGS

A remote print request can be altered or canceled only by the user who requested it, and only from the system from which the the original **lp** command was issued.

If the restrict cancel feature (see *lpadmin(1M)*) is enabled for the specified printer, a user can only alter or cancel requests owned by the user.

For a remote system, **lpalt** cannot change *dest* and *priority*.

FILES

| | |
|--------------------------|---|
| /etc/lp | Directory of spooler configuration data |
| /etc/lp/interface | Directory of active LP device interface scripts |
| /usr/lib/lp | Directory of model and font file directories |
| /var/adm/lp | Directory of spooler log files |
| /var/spool/lp | Directory of LP spooling files and directories |

SEE ALSO

enable(1), lpstat(1), mail(1), slp(1), accept(1M), lpadmin(1M), lpana(1M), lpsched(1M), rcancel(1M), rlp(1M), rlpdaemon(1M), rlpstat(1M).

STANDARDS CONFORMANCE

lp: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

cancel: SVID2, SVID3, XPG4

NAME

lpfilter, divpage, fontdl, lprpp, plotdvr, printstat, reverse - filters invoked by lp interface scripts

SYNOPSIS

```
/usr/sbin/divpage [-p | -l] [-h | -q] [-nFontID] filename
/usr/sbin/fontdl [-nFontID] [-l] [-p] filename
/usr/sbin/lprpp [-i] [-o] [-e] [-l $nn$ ] [-n] [-p]
/usr/sbin/plotdvr -l request_id -u username [-e] [-f] [-i] filename
/usr/sbin/printstat -l request_id -u username filename
/usr/sbin/reverse [-l page_length]
```

Remarks

The structure of these filters is currently under review. They may become obsolete or be restructured in a future release.

DESCRIPTION

Various filters are used by the **lp** subsystem to obtain specialized behavior for specific types of devices or data. This entry describes currently supported filters.

A number of these filters use a specified *username* and *filename* to determine the location of the user who originated the print message.

The *filename* is used to determine the *hostname* of the system where the request originated, and must have the form *[dirname]/d?A nnnhostname* or *[dirname]/dA nnnhostname*, where *dirname* is not a path name, but only the name of the basename's parent directory. *filename* meets this requirement when it is set to **\$6** in the interface script for the printer.

divpage

Provides capabilities for printing multiple pages per sheet and selection of built-in fonts.

Options:

- p** Set mode to portrait (default).
- l** Set mode to landscape.
- h** Print half pages (default).
- q** Print quarter pages.
- nFontID** Use font number *FontID*. Default is 0. Causes the string $\text{\textasciix{FontID}}$ to be sent to the printer.

fontdl

fontdl downloads the font contained in *filename* to a printer connected to standard output.

Options:

- nFontID** Specifies the ID number by which the font will be referenced. Default is 0.
- l** Specify landscape mode. Default is portrait.
- p** Specifies proportional spacing. Default is fixed.

lprpp

This is a filter that converts backspace overstrike to line overprint with horizontal print positioning to enhance bold print. This functionality is required on printers such as the LaserJet, which cannot produce bold print by overstriking.

Options:

- i** Converts *<ANYCHAR>* to *<ANYCHAR><BACKSPACE>* to italicize *ANYCHAR*. Also properly italicizes overstruck (bold) characters. Does *not* work correctly for "hashed-overstrike" such as:
 $\text{\textasciix{ANYCHAR}\text{\textasciix{BACKSPACE}\text{\textasciix{DIFFERENTCHAR}\text{\textasciix{BACKSPACE}}}}$

- o** Prints only the odd numbered pages. Used with **-e** for double-sided printing.
- e** Print only the even numbered pages. Used with **-o** for double sided printing.
- l *nn*** Specifies the page length, in lines. Default is 60 unless **-n** or **-p** is selected, in which case it is 66.
- n** Specifies nroff mode for printing output of the *nroff* command. Prints 66 lines per page with the first line appearing on logical line 4 of the printer.
- p** Specifies pr mode for printing output from the *pr* command. Prints 66 lines per page with the first line appearing on logical line 3 of the printer.

plotdvr

HP-GL plotter filter. This filter scans the data for "PG" commands (paper feed). When this data is encountered, the filter strips it from the data stream and informs the requesting user of the need to change paper in the plotter.

Options:

- l *request_id*** Specifies the printer request ID and is used in various messages regarding the plot request.
- u *username*** The requesting user's login name, used to communicate with the user regarding the request.
- e** Specifies the use of escape sequences, rather than HP-GL commands, to determine plotter status.
- f** Plot without stopping for paper changes. The "PG" commands are not stripped from the data stream and the user is not notified of them. This option is used on plotters capable of automatic page feed.
- i** Prevents initialization of the plotter.

printstat

Interrogates an RS232 printer as to its status, and does not return until the printer is ready. If the printer is *off-line*, *out of paper*, or *disconnected*, the submitter of the print request is notified of this condition periodically until it is corrected. When the printer is ready to print, the command exits.

Standard input and standard output must both be connected to the serial printer device.

This program uses the *send-status* command `^c?^D1?` to determine status. Not all serial printers respond to this command. Only the following configurations support this command:

| Printer | Comments |
|--------------|---------------------------|
| LaserJet | Not supported |
| LaserJetII | Supported |
| LaserJetIID | Requires HP 26013A module |
| LaserJetIIP | Not supported |
| LaserJetIII | Requires HP 26013A module |
| LaserJet2000 | Not supported |

Options:

- l *request-id*** Print request ID used in various communications with the user.
- u *username*** The requesting user's login name, used to communicate with the user regarding the request.

reverse

Prints the data appearing on the standard input in reverse page order to the standard output. This command can handle up to 2000 pages.

Options:

- l *page_length*** Specifies the page length, in lines. Default is 66.

DIAGNOSTICS

Error and diagnostic messages appear on the printed output, on the user's terminal, or are mailed to the user, depending on circumstances.

WARNINGS

There is little consistency in the interface to these filters.

SEE ALSO

lp(1), lpadmin(1M).

Managing Systems and Workgroups.

NAME

lpstat - report line printer status information

SYNOPSIS

lpstat [-**drst**] [-**a**[*list*]] [-**c**[*list*]] [-**o**[*list*]] [-**p**[*list*]] [-**u**[*list*]] [-**v**[*list*]] [**ID**...]

DESCRIPTION

The **lpstat** utility writes to standard output information about the current status of the line printer system.

If no arguments are given, **lpstat** writes the status of all requests made to **lp** by the user that are still in the output queue.

OPTIONS

The **lpstat** utility supports the XBD specification, Section 10.2, Utility Syntax Guidelines, except the option-arguments are optional and cannot be presented as separate arguments.

Some of the options below can be followed by an optional list that can be in one of two forms: a list of items separated from one another by a comma, or a quoted list of items separated from one another by a comma or one or more blank characters, or combinations of both. See EXAMPLES.

The omission of a list following such options causes all information relevant to the option to be written to standard output; for example:

```
lpstat -o
```

writes the status of all output requests that are still in the output queue.

- a**[*list*] Write the acceptance status of destinations for output requests. The *list* argument is a list of intermixed printer names and class names.
- c**[*list*] Write the class names and their members. The *list* argument is a list of class names.
- d** Write the system default destination for output requests.
- o**[*list*] Write the status of output requests. The *list* argument is a list of intermixed printer names, class names and request IDs.
- p**[*list*] Write the status of printers. The *list* argument is a list of printer names.
- r** Write the status of the line printer request scheduler.
- s** Write a status summary, including the status of the line printer scheduler, the system default destination, a list of class names and their members and a list of printers and their associated devices.
- t** Write all status information.
- u**[*list*] Write the status of output requests for users. The *list* argument is a list of login names.
- v**[*list*] Write the names of printers and the pathnames of the devices associated with them. The *list* argument is a list of printer names.

OPERANDS

The following operand is supported:

ID A request ID, as returned by **lp**.

STDIN

Not used.

INPUT FILES

None.

ENVIRONMENT VARIABLES

The following environment variables affect the execution of **lpstat**:

LANG Provide a default value for the internationalisation variables that are unset or null. If **LANG** is unset or null, the corresponding value from the implementation-specific default locale will be used. If any of the internationalisation variables contains an invalid setting, the utility will behave as

| | |
|--------------------|---|
| | if none of the variables had been defined. |
| LC_ALL | If set to a non-empty string value, override the values of all the other internationalisation variables. |
| LC_CTYPE | Determine the locale for the interpretation of sequences of bytes of text data as characters (for example, single- as opposed to multi-byte characters in arguments). |
| LC_MESSAGES | Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error, and informative messages written to standard output. |
| LC_TIME | Determine the format of date and time strings output when displaying line printer status information with the -a , -o , -p , -t , or -u options. |
| NLSPATH | Determine the location of message catalogues for the processing of LC_MESSAGES . |
| TZ | Determine the timezone used with date and time strings. |

ASYNCHRONOUS EVENTS

Default.

STDOUT

The standard output is a text file containing the information described in **OPTIONS**, in an unspecified format.

STDERR

Used only for diagnostic messages.

OUTPUT FILES

None.

EXTENDED DESCRIPTION

None.

EXIT STATUS

The following exit values are returned:

- 0 Successful completion.
- >0 An error occurred.

CONSEQUENCES OF ERRORS

Default.

APPLICATION USAGE

The **lpstat** utility cannot reliably determine the status of print requests in all conceivable circumstances. When the printer is under the control of another operating system or resides on a remote system across a network, it need not be possible to determine the status of the print job after it has left the control of the local operating system. Even on local printers, spooling hardware in the printer may make it appear that the print job has been completed long before the final page is printed.

EXAMPLES

1. Obtain the status of two printers, the pathnames of two printers, a list of all class names and the status of the request named HiPri-33:

```
lpstat -plaser1,laser4 -v"laser2 laser3" -c HiPri-33
```

2. Obtain user print job status using the obsolescent mixed blank and comma form:

```
lpstat -u"ddg,gmv, maw"
```

FUTURE DIRECTIONS

A version of **lpstat** that fully supports the XBD specification, Section 10.2, Utility Syntax Guidelines may be introduced in a future issue.

SEE ALSO

cancel(1), lp(1).

CHANGE HISTORY

First released in Issue 2.

Issue 3

The operation of this utility in an 8-bit transparent manner has been noted.

The operation of this utility in an internationalised environment has been described.

Issue 4

Format reorganised.

Exceptions to Utility Syntax Guidelines conformance noted.

Internationalised environment variable support mandated.

STANDARDS CONFORMANCE

lpstat: SVID2, SVID3, XPG2, XPG3, XPG4

HP-UX EXTENSIONS

DESCRIPTION

Any arguments that are not *options* are assumed to be request *ids* (as returned by **lp**). **lpstat** prints the status of such requests. *options* can appear in any order and can be repeated and intermixed with other arguments.

- i** Inhibit the reporting of remote status.
- o[list]** Also see the **-i** option.
- t** Print all status information. Same as specifying **-r**, **-s**, **-a**, **-p**, **-o**. See the **-i** option.

Security Restriction

Only users who have the **lp** subsystem authorization or the **printqueue** secondary subsystem authorization can view the entire queue. Unauthorized users can view only their own jobs whose sensitivity levels are dominated by the user's current sensitivity level.

The **allowmacaccess** privilege allows viewing jobs at higher sensitivity levels.

EXAMPLES

Check whether your job is queued:

```
lpstat
```

Check the relative position of a queued job:

```
lpstat -t
```

Verify that the job scheduler is running:

```
lpstat -r
```

FILES

```
/var/spool/lp/*
/var/adm/lp/*
/etc/lp/*
/usr/lib/lp/*
```

SEE ALSO

enable(1), lp(1), rlpstat(1M).

STANDARDS CONFORMANCE

lpstat: SVID2, SVID3, XPG2, XPG3, XPG4

NAME

ls, lc, l, ll, lsf, lsr, lsx - list contents of directories

SYNOPSIS

```
ls [-abcdefgilmnopqrstuxACFLRl] [names]
lc [-abcdefgilmnopqrstuxACFLRl] [names]
l [ls_options] [names]
ll [ls_options] [names]
lsf [ls_options] [names]
lsr [ls_options] [names]
lsx [ls_options] [names]
```

DESCRIPTION

For each directory argument, the **ls** command lists the contents of the directory. For each file argument, **ls** repeats its name and any other information requested. The output is sorted in ascending collation order by default (see Environment Variables below). 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.

If you are a user with appropriate privileges, all files except **.** and **..** are listed by default.

There are three major listing formats. The format chosen depends on whether the output is going to a login device (determined by whether output device file is a **tty** device), and can also be controlled by option flags. The default format for a login device is to list the contents of directories in multicolumn format, with entries sorted vertically by column. (When individual file names (as opposed to directory names) appear in the argument list, those file names are always sorted across the page rather than down the page in columns because individual file names can be arbitrarily long.) If the standard output is not a login device, the default format is to list one entry per line. The **-C** and **-x** options enable multicolumn formats, and the **-m** option enables stream output format in which files are listed across the page, separated by commas. 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 each output line. If this variable is not set, the **terminfo** database is used to determine the number of columns, based on the environment variable **TERM**. If this information cannot be obtained, 80 columns is assumed.

The command **lc** functions the same as **ls** except that the **lc** default output is columnar, even if output is redirected.

Options

ls recognizes the following options:

- a List all entries; usually entries whose names begin with a period (**.**) are not listed.
- b List nonprinting characters in the octal **\ddd** notation.
- c Use time of last modification of the inode (file created, mode changed, etc.) for sorting (**-t**) or printing (**-l** (ell)).
- d If an argument is a directory, list only its name (not its contents); often used with **-l** (ell) to get the status of a directory.
- e List the extent attributes of the file. If any of the files has a extent attribute, this option lists the extent size, space reserved and allocation flags. This option must be used with the **-l** (ell) option.
- f Interpret each argument as a directory and list the name found in each slot. This option disables **-l** (ell), **-r**, **-s**, and **-t**, and enables **-a**; the order is the order in which entries appear in the directory.
- g Same as **-l** (ell), except that only the group is printed (owner is omitted). If both **-l** (ell) and **-g** are specified, the owner is not printed.
- i For each file, list the inode number in the first column of the report. When used in multicolumn output, the number precedes the file name in each column.
- l (ell) List in long format, giving mode, number of links, owner, group, size in bytes, and time of last modification for each file (see further DESCRIPTION and Access Control Lists below). If the time of last modification is greater than six months ago, or any time in the future, the year is

substituted for the hour and minute of the modification time. If the file is a special file, the size field 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.

- m** Stream output format.
- n** The same as **-l**, (ell) except that the owner's UID and group's GID numbers are printed, rather than the associated character strings.
- o** The same as **-l**, (ell) except that only the owner is printed (group is omitted). (If both **-l** (ell) and **-o** are specified, the group is not printed).
- p** Put a slash (/) after each file name if that file is a directory.
- q** List nonprinting characters in file names as the character (?).
- r** Reverse the order of sort to get reverse (descending) collation or oldest first, as appropriate.
- s** List size in blocks, including indirect blocks, for each entry. The first entry listed is the total number of blocks in the directory. When used in multicolumn output, the number of blocks precedes the file name in each column. The number of indirect blocks in a file is filesystem dependent.
- t** Sort by time modified (latest first) before sorting alphabetically.
- u** Use time of last access instead of last modification for sorting (**-t** option) or printing (**-l** (ell) option).
- x** List multicolumn output with entries sorted across rather than down the page.
- A** The same as **-a**, except that the current directory `.` and parent directory `..` are not listed. For a user with appropriate privileges, this flag defaults to on, and is turned off by **-A**.
- C** List multicolumn output with entries sorted down the columns.
- F** After each file name, put one of:
 - A slash (/) if the file is a directory or a symbolic link to a directory.
 - An asterisk (*) if the file is executable;
 - An at-sign (@) if the file is a symbolic link to a file;
 - A vertical bar (|) if the file is a fifo.
- L** If the argument is a symbolic link, list the file or directory to which the link refers rather than the link itself.
- R** Recursively list subdirectories encountered.
- 1** (one) List the file names in single column format regardless of the output device. This forces single column format to the user's terminal.

Specifying more than one of the options in the following mutually exclusive pairs is not considered an error: **-C** and **-l** (ell), **-m** and **-l** (ell), **-x** and **-l** (ell), **-C** and **-1** (one), and **-c** and **-u**.

ls is known by several shorthand-version names for the various formats:

- l** is equivalent to **ls -m**
- ll** is equivalent to **ls -l** (ell)
- lsf** is equivalent to **ls -F**
- lsr** is equivalent to **ls -R**
- lsx** is equivalent to **ls -x**

The shorthand notations are implemented as links to **ls**. Option arguments to the shorthand versions behave exactly as if the long form above had been used with the additional arguments.

Mode Bits Interpretation (-l option)

The mode printed in listings produced by the **-l** (ell) option consists of 10 characters, for example, **-rwxr-xr-x**.

The first character indicates the entry type:

- b** Block special file
- c** Character special file
- d** Directory

l Symbolic link
n Network special file
p Fifo (also called a "named pipe") special file
s Socket
- Ordinary file

The next 9 characters are interpreted as three sets of three characters each which identify access and execution permissions for the owner, group, and others categories, as described in *chmod(1)*. The **-** indicates the permission is not granted. The various permissions can be put together in any combination, except that the **x**, **s**, **S**, **t**, and **T** characters are mutually exclusive, as implied below.

-r----- Read by owner
--w----- Write by owner
---x----- Execute (or search directory) by owner; do not set user ID on execution
---s----- Execute/search by owner; set user ID on execution
---S----- No execute/search by owner; set user ID on execution
----r----- Read by group
----w----- Write by group
-----x--- Execute/search by group; do not set group ID on execution
-----s--- Execute/search by group; set group ID on execution
-----S--- No execute/search by group; set group ID on execution
-----r-- Read by others
-----w- Write by others
-----x Execute/search by others; do not set sticky bit on execution
-----t Execute/search by others; set sticky bit on execution
-----T No execute/search by others; set sticky bit on execution

The mode characters are interpreted as follows:

- Deny all permissions in the corresponding position.
r Grant read permission to the corresponding user class.
w Grant write permission to the corresponding user class.
x Grant execute (or search in directory) permission to the corresponding user class.
s Grant execute (search) permission to the corresponding user class. Execute the file as if by the owner (set user ID, SUID) or group (set group ID, SGID), as indicated by position.
S Deny execute (search) permission to the corresponding user class. Execute the file as if by the owner (set user ID, SUID) or group (set group ID, SGID), as indicated by position.
t Grant execute (search) permission to others. The "sticky" (save text image) bit is set (see the description of **S_ISVTX** in *chmod(2)*).
T Deny execute (search directory) permission to others. The "sticky" (save text image) bit is set.

When an option is specified that results in a listing of directory and/or file sizes in bytes or blocks (such as the **-s** or **-l** (ell) option), a total count of blocks, including indirect blocks, is also printed at the beginning of the listing.

Access Control Lists (ACLs)

If a file has optional ACL entries, the **-l** (ell) option displays a plus sign (+) after the file's permissions. The permissions shown are a summary representation of the file's access control list, as returned by **stat()** in the **st_mode** field (see *stat(2)*). To list the contents of an access control list, use the **lsacl** command (see *lsacl(1)* and *acl(5)*) for HFS file systems, or the **getacl** command (see *getacl(1)* and *aclv(5)*) for JFS file systems.

EXTERNAL INFLUENCES

Environment Variables

If the **COLUMNS** variable is set, **ls** uses the width provided in determining positioning of columnar output.

LANG determines the locale to use for the locale categories when both **LC_ALL** and the corresponding environment variable (beginning with **LC_**) do not specify a locale. If **LANG** is not set or is null, it defaults to **C** (see *lang(5)*).

LC_COLLATE determines the order in which the output is sorted.

LC_CTYPE determines which characters are classified as nonprinting for the **-b** and **-q** options, and the interpretation of single- and/or multibyte characters within file names.

LC_TIME determines the date and time strings output by the **-g**, **-l** (ell), **-n**, and **-o** options.

LC_MESSAGES determines the language in which messages (other than the date and time strings) are displayed.

If any internationalization variable contains an invalid setting, they all default to **C** (see *environ*(5)).

International Code Set Support

Single- and multibyte character code sets are supported.

RETURN VALUE

ls exits with one of the following values:

- 0 All input files were listed successfully.
- >0 **ls** was aborted because errors occurred when accessing files. The following conditions cause an error:
 - Specified file not found.
 - User has no permission to read the directory.
 - Process could not get enough memory.
 - Invalid option specified.

EXAMPLES

Print a long listing of all the files in the current working directory (including the file sizes). List the most recently modified (youngest) file first, followed by the next older file, and so forth, to the oldest. Files whose names begin with a **.** are also printed.

```
ls -alst
```

WARNINGS

Setting options based on whether the output is a login (*tty*) device is undesirable because **ls -s** is very different from **ls -s | lp**. On the other hand, not using this setting makes old shell scripts that used **ls** almost inevitably fail.

Nonprinting characters in file names (without the **-b** or **-q** option) may cause columnar output to be misaligned.

DEPENDENCIES

NFS

The **-l** (ell) option does not display a plus sign (+) after the access permission bits of networked files to represent existence of optional access control list entries.

AUTHOR

ls was developed by AT&T, the University of California, Berkeley and HP.

FILES

| | |
|------------------------------------|---|
| /etc/group | For group IDs for -l (ell) and -g . |
| /etc/passwd | For user IDs for -l (ell) and -o . |
| /usr/share/lib/terminfo/?/* | For terminal information. |

SEE ALSO

chmod(1), **find**(1), **getacl**(1), **lsacl**(1), **stat**(2), **acl**(5), **aclv**(5).

STANDARDS CONFORMANCE

ls: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

lsacl - list access control lists (ACLs) of files

SYNOPSIS

`/usr/bin/lsacl [-l] file...`

DESCRIPTION

lsacl lists access control lists (ACLs) of one or more files in symbolic, “short” form, one file’s ACL per line of output, followed by the file name; see *acl(5)* for ACL syntax.

Options

lsacl recognizes the following option:

- l** Print ACLs in long form. Each file’s ACL can be more than one line long, and is always preceded by file name, colon, and newline. Consecutive file names are separated by blank lines.

If a hyphen (-) is given as a file name argument, **lsacl** prints the ACL for the standard input. By default, it prints the file name as -. For the **-l** option it prints a file name of **<stdin>**.

Unlike **ls**, **lsacl** cannot list ACLs of files in subdirectories automatically or list the ACL entries of the files in the current directory by default. A good way to do the latter is:

```
lsacl *
```

or

```
lsacl .* *
```

EXTERNAL INFLUENCES**Environment Variables**

LANG determines the language in which messages are displayed.

If **LANG** is not specified or is set to the empty string, a default of “C” (see *lang(5)*) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **lsacl** behaves as if all internationalization variables are set to “C”. See *environ(5)*.

RETURN VALUE

If **lsacl** succeeds, it returns zero. If invoked incorrectly, it returns a value of 1. If **lsacl** is unable to read the ACL for a specified file, it prints an error message to standard error, continues, and later returns a value of 2.

EXAMPLES

List the ACL for the file **dir/file1**:

```
lsacl dir/file1
```

List ACLs for all files in the current directory, in long form.

```
lsacl -l .* *
```

List ACLs for all files under **mydir**:

```
find mydir -print | sort | xargs lsacl
```

DEPENDENCIES

lsacl will fail when the target file resides on a file system which does not support ACLs.

NFS:

lsacl is not supported on remote files.

AUTHOR

lsacl was developed by HP.

SEE ALSO

chacl(1), *getaccess(1)*, *ls(1)*, *getacl(2)*, *acl(5)*.

NAME

m4 - macro processor

SYNOPSIS

m4 [*options*] [*file* ...]

DESCRIPTION

m4 is a macro processor intended as a front end for Ratfor, C, and other languages. Each of the argument files is processed in order; if there are no files, or if a file name is -, standard input is read. The processed text is written to standard output.

Options

m4 recognizes the following options:

- e Operate interactively. Interrupts are ignored and the output is unbuffered. Using this mode may be very difficult.
- s Enable line sync output for the C preprocessor (#line ...)
- B*int* Change the size of the push-back and argument collection buffers from the default of 4,096.
- H*int* Change the size of the symbol table hash array from the default of 199. The size should be prime.
- S*int* Change the size of the call stack from the default of 100 slots. Macros take three slots, and nonmacro arguments take one.
- T*int* Change the size of the token buffer from the default of 512 bytes.

To be effective, the options listed above must appear before any file names and before any -D or -U options.

- D*name*[=*val*] Define *name* as *val* or as null if *val* is omitted.
- U*name* Undefine *name*.

Macro Calls

Macro calls have the form:

name(*arg1*, *arg2*, . . . , *argn*)

The left parenthesis () 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 alphabetic letters, digits, and underscore (_); the first character cannot be a digit.

Leading unquoted blanks, tabs, and newlines are ignored while collecting arguments. Left and right single quotes (' and ') are used to quote strings. The value of a quoted string is the string stripped of the quotes.

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 which 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.

Built-In Macro Names

m4 makes available the following built-in macros. They can be redefined, but, once this is done, the original meaning is lost. Their values are null unless otherwise stated.

- changeocom** Change left and right comment markers from the default # and newline. With no arguments, the comment mechanism is effectively disabled. With one argument, the left marker becomes the argument and the right marker becomes newline. With two arguments, both markers are affected. Comment markers may be up to five characters long.
- 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 (i.e., ' and ').

| | |
|-----------------|---|
| decr | Returns the value of its argument decremented by 1. |
| define | The second argument is installed as the value of the macro whose name is the first argument. Each occurrence of <code>\$n</code> in the replacement text, where <i>n</i> is a digit, is replaced by the <i>n</i> th argument. Argument 0 is the name of the macro; missing arguments are replaced by the null string; <code>\$#</code> is replaced by the number of arguments; <code>\$*</code> is replaced by a list of all the arguments separated by commas; <code>\$@</code> is equivalent to <code>\$*</code> , but each argument is quoted (with the current quotes). |
| defn | Returns the quoted definition of its arguments. It is useful for renaming macros, especially built-ins. |
| divert | m4 maintains 10 output streams, numbered 0 to 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. |
| divnum | Returns the value of the current output stream. |
| dnl | Reads and discards characters up to and including the next newline. |
| dumpdef | Prints current names and definitions, for the named items, or for all if no arguments are given. |
| errprint | Prints its argument on the diagnostic output file. |
| eval | Evaluates its argument as an arithmetic expression, using 32-bit arithmetic. Operators include <code>+</code> , <code>-</code> , <code>*</code> , <code>/</code> , <code>%</code> , <code>**</code> (exponentiation), bitwise <code>&</code> , <code> </code> , <code>^</code> , and <code>~</code> , relationals, and parentheses. Octal and hexadecimal 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. |
| hpux | Is a predefined object with a null value. |
| 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 on HP-UX system versions of m4 . |
| ifelse | 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. |
| include | Returns the contents of the file named in the argument. |
| 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. |
| index | Returns the position in its first argument where the second argument begins (zero origin), or -1 if the second argument does not occur. |
| len | Returns the number of characters in its argument. |
| m4exit | Causes immediate exit from m4 . Argument 1, if given, is the exit code; the default is 0. |
| m4wrap | Argument 1 is pushed back at final EOF; for example: <code>m4wrap('cleanup')</code> |
| maketemp | Fills in a string of XXXXXX in its argument with the current process ID. |
| popdef | Removes current definition of its arguments, exposing the previous one, if any. |
| pushdef | Similar to define , but saves any previous definition. |
| 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. |
| sinclude | Identical to include , except that it says nothing if the file is inaccessible. |
| 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. |

| | |
|-----------------|--|
| syscmd | Executes the HP-UX system command given in the first argument. No value is returned. |
| sysval | Is the return code from the last call to syscmd . |
| traceoff | Turns off trace globally and for any macros specified. Macros specifically traced by traceon can be untraced only by specific calls to traceoff . |
| traceon | With no arguments, turns on tracing for all macros (including built-ins). Otherwise, turns on tracing for named macros. |
| 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. |
| undefine | Removes the definition of the macro named in its argument. |
| undivert | 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. |

(XPG4 only.) It is an error to specify an argument containing any non-numeric character for the built-in-macros: **decr**, **divert**, **incr**, **m4exit**, **substr**, **undivert**, and **eval**.

SEE ALSO

cpp(1), ratfor(1).

STANDARDS CONFORMANCE

m4: SVID2, SVID3, XPG2, XPG3, XPG4

NAME

hp9000s200, hp9000s300, hp9000s500, hp9000s800, pdp11, u3b, u3b2, u3b5, u3b10, u370, vax - provide truth value about processor type

SYNOPSIS

hp9000s200
hp9000s300
hp9000s400
hp9000s500
hp9000s700
hp9000s800
hp-mc680x0
hp-pa
pdp11
u3b
u3b2
u3b5
u3b10
u370
vax

DESCRIPTION

The following commands return a true value (exit code 0) if the a processor type matches the command name. Otherwise a false value (exit code non-zero) is returned. These commands are commonly used within **make** makefiles and shell procedures to improve portability of applications (see *make(1)*).

| Command | True for: | Command | True for: |
|------------|-------------------------|---------|--------------------------|
| hp9000s200 | Series 200 | pdp11 | PDP-11/45 or PDP-11/70 |
| hp9000s300 | Series 300 | u3b | 3B20 computer |
| hp9000s400 | Series 400 | u3b2 | 3B2 computer |
| hp9000s500 | Series 500 | u3b5 | 3B5 computer |
| hp9000s700 | Series 700 | u3b10 | 3B10 computer |
| hp9000s800 | Series 800/700 | u370 | IBM System/370 computer |
| hp-mc680x0 | Series 200, 300, or 400 | vax | VAX-11/750 or VAX-11/780 |
| hp-pa | Series 700 or 800 | | |

EXAMPLES

Given a shell script that must behave differently when run on an HP 9000 Series 700 or 800 system, select the correct code segment to be executed:

```
if hp9000s800
then
    # system is Series 700 or 800.
    if hp9000s700
    then
        # System is Series 700
        Series 700 code fragment goes here
    else
        # System is Series 800
        Series 800 code fragment goes here
    fi
fi
```

WARNINGS

hp9000s800 always returns true on both Series 800 and Series 700 systems. Therefore, when using this command in scripts to determine hardware type, always use both **hp9000s800** and **hp9000s700** in the appropriate sequence to ensure correct results (see **EXAMPLES**).

machid(1) will no longer provide support for future machines beyond the Series 800 and Series 700 systems. Decisions should be based on the hardware and software configuration information returned by *getconf(1)*.

SEE ALSO

getconf(1), make(1), sh(1), test(1), true(1).

NAME

mail, rmail - send mail to users or read mail

SYNOPSIS

mail **[+]** **[-epqr]** **[-f file]**

mail **[-dt]** *person* ...

rmail **[-dt]** *person* ...

Remarks:

See *mailx*(1) and *elm*(1) for an enhanced user mail interface.

DESCRIPTION

The **mail** command, when used without arguments, prints the user's mail, message-by-message, in last-in, first-out order.

For each message, **mail** prints a **?** prompt and reads a line from the standard input to determine the disposition of the message. Commands that automatically proceed to the next message exit from **mail** if **mail** already on the last message.

Commands

mail supports the following commands:

| | |
|---------------------|---|
| <new-line> | Go on to next message. Exit if already on last message. |
| + | Same as <new-line>. |
| n | Same as <new-line>. |
| d | Delete message and go on to next message. |
| p | Print message again. |
| - | Go back to previous message. |
| s [<i>files</i>] | Save message in the named <i>files</i> (default is mbox), mark the message for deletion from the user's <i>mailfile</i> , and proceed to next message. |
| y [<i>files</i>] | Same as s [<i>files</i>]. |
| w [<i>files</i>] | Save message without its header (the "From ..." line), in the named <i>files</i> (default is mbox), mark the message for deletion, and go on to next message. |
| m <i>person</i> ... | Mail the message to each named <i>person</i> , mark the message for deletion, and go on to next message. |
| q | Put undeleted mail back in the mailfile and stop. |
| EOT (Ctrl-D) | Same as q. |
| x | Abort. Leave original mailfile unchanged and stop. |
| ! <i>command</i> | Escape to the command interpreter and execute <i>command</i> . |
| ? | Print a command summary. |
| * | Same as ?. |

Command-Line Options

The following command-line options alter printing of the mail:

| | |
|----|--|
| + | Cause messages to be printed in first-in, first-out order. |
| -e | Suppresses printing of mail and returns the exit value: <div style="margin-left: 40px;"> 0 = Mail present 1 = No mail 2 = Other error </div> |
| -p | Prints all mail without prompting for disposition. |
| -q | Causes mail to terminate if an interrupt is received. Normally an interrupt only causes the termination of the printing of the current message. |

- r** Same as **+**.
- f file** Causes **mail** to use *file* (for example, **mbox**) instead of the default *mailfile*.
- t** Causes the outbound message to be preceded by each *person* the mail is sent to. A *person* is usually a user name recognized by **login** (see *login(1)*). If a *person* being sent mail is not recognized, or if **mail** is interrupted during input, the file **dead.letter** will be saved to allow editing and resending. Note that **dead.letter** is regarded as a temporary file in that it is recreated every time needed, erasing the previous contents of **dead.letter**.
- d** Causes **mail** to deliver mail directly. This isolates **mail** from making routing decisions, and allows it to be used as a local delivery agent. Typically this option is used by auto-routing facilities when they deliver mail locally.

When *persons* are named, **mail** takes the standard input up to an end-of-file (or up to a line consisting of just a **.**) and adds it to each *person's* *mailfile*. The message is preceded by the sender's name and a postmark.

To denote a recipient on a remote system, prefix *person* by the system name and exclamation mark (see *uucp(1)*). Everything after the first exclamation mark in *person* is interpreted by the remote system. In particular, if *person* contains additional exclamation marks, it can denote a sequence of machines through which the message is to be sent on the way to its ultimate destination. For example, specifying **a!b!cde** as a recipient's name causes the message to be sent to user **b!cde** on system **a**. System **a** then interprets that destination as a request to send the message to user **cde** on system **b**. This might be useful, for instance, if the sending system can access system **a** but not system **b**. **mail** does not use **uucp** if the remote system is the local system name (i.e., *localsystem/user*).

The **mailfile** can be manipulated in two ways to alter the function of **mail**. The *other* permissions of the file can be read-write, read-only, or neither read nor write to allow different levels of privacy. If changed to other than the default, the file is preserved, even when empty, to perpetuate the desired permissions. The file can also contain the first line:

Forward to person

which causes all mail sent to the owner of the **mailfile** to be forwarded to *person*. This is especially useful for forwarding all of a person's mail to a given machine in a multiple-machine environment. In order for forwarding to work properly the **mailfile** should have "mail" as group ID, and the group permission should be read-write.

rmail only permits the sending of mail. **uucp** uses **rmail** as a security precaution.

When a user logs in, the command **mail -e** can be used to detect the presence of mail, if any, and so indicate. When terminating, **mail** produces a notification message if new mail arrived while **mail** was running.

EXTERNAL INFLUENCES

Environment Variables

LC_TIME determines the format and contents of the displayed date and time strings.

If **LC_TIME** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **mail** behaves as if all internationalization variables are set to "C". See *environ(5)*.

When set, the **TMPDIR** environment variable specifies a directory to be used for temporary files, overriding the default directory **/tmp**.

International Code Set Support

Between HP-UX systems, single- and multi-byte character code sets are supported within mail text. Headers are restricted to characters from the 7-bit USASCII code set (see *ascii(5)*).

WARNINGS

Conditions sometimes result in a failure to remove a lock file.

After an interrupt, the next message may not be printed. To force printing, type a **p**.

Lines that look like postmarks in the message (that is, "From ...") are preceded by **>**.

mail treats a line consisting solely of a dot (.) as the end of the message, except when the **rmail -d** command is used.

The maximum allowable line length in mail messages is **BUFSIZ** bytes as defined in **/usr/include/stdio.h**. If line length exceeds this limit, **mail** truncates the line starting at beginning-of-line, and uses only the trailing **BUFSIZ** characters.

Using two separate mail programs to access the same mail file simultaneously (usually inadvertently from two separate windows) can cause unpredictable results.

Some sites that have programs that adhere strictly to RFC-822 will fail to deliver a message if any of the recipient fields below is missing.

```
To:
resent-to:
cc:
resent-cc:
bcc:
resent-bcc:
```

You can add the RFC-822 commands into the mail program buffer/editor. For instance:

```
mail user1@domain.com
From: user2@domain.com
Subject: This is a test
To: test_address@domain.com

This is a test
```

FILES

| | |
|-------------------------|--|
| /var/mail/*.lock | Lock for mail directory |
| dead.letter | Unmailable text |
| /tmp/ma* | Temporary file |
| \$MAIL | Variable containing path name of mailfile |
| \$HOME/mbox | Saved mail |
| /etc/passwd | To identify sender and locate persons |
| /var/mail | Directory for incoming mail (mode 775, group ID mail) |
| /var/mail/ user | Incoming mail for <i>user</i> ; that is, the mailfile (mode 660, group ID mail) |

SEE ALSO

login(1), mailx(1), uucp(1), write(1).

STANDARDS CONFORMANCE

mail: SVID2, SVID3, XPG2, XPG3

rmail: SVID2, SVID3

NAME

mailfrom - summarize mail folders by subject and sender

SYNOPSIS

mailfrom [-hnQStv] [-s *status*] [*folder*|*username*]...

DESCRIPTION

The **mailfrom** command reads one or more mail folders and outputs one line per message in the form:

from [*subject*]

where *from* is the name of the person the message is from, and *subject* is the subject of the message, if present. If **mailfrom** determines that the message is from you, the *from* portion will read **To user**, where *user* is the user the message was sent to. This happens when you receive a copy of a letter you sent.

The default folder is your incoming mailbox, **/var/mail/yourloginname**. See the Operands subsection below.

Options

mailfrom recognizes the following options:

- h** Print a brief help message summarizing the options.
- n** Number the messages using the same numbering scheme used by **readmail**.
- Q** Very quiet mode. Only error messages are produced. This option is useful in shell scripts, where only the success or failure of the program is important, and output is not desired.
- q** Quiet mode. Output only a one-line summary for each mailbox or folder.
- s** Add a summary of the number of messages by message status in each mailbox or folder. To get the summary only, use this with the **-q** option. The summary has the form:

Folder contains:
New messages: *n*
Unread messages: *u*
Read messages: *r*

 If an item count, *n*, *r*, or *u* is zero, the line is omitted.
- s status** Only display headers from messages with the given status. *status* can be one of **new**, **old**, **read**, or **unread**. **old** and **unread** are equivalent. The **-s** option can be repeated to print header information from more than one category, for example, only new and unread messages. The values can be abbreviated to their first letters. The default is all messages.
- t** Tidy mode. If the *from* field is long enough to displace the *subject* field from its normal start column, move the subject down onto the next line.
- v** Verbose mode. Print a descriptive header before listing the contents of each mailbox or folder.

Operands

mailfrom recognizes the following optional operands:

folder|*username*

A file name or the name of a mail user on your system. You can use the **=filename** format to specify a folder in your mail directory, defined by the **maildir** string variable in your **elmr**c configuration file.

mailfrom searches for the value as a file name relative to your current directory. Then, if the file name is not an absolute path, it searches for the value relative to the incoming mailbox directory, **/var/mail**. The first file found is selected. You must have read access to the file.

EXIT STATUS

mailfrom returns the following values:

- 0 Messages matching *status* are present.
- 1 No messages matching *status* are present, but there are some messages.
- 2 There are no messages at all.
- 3 An error occurred.

If multiple mailboxes or folders are specified, the exit status only applies to the last one examined. This can be used in scripts to determine what kind of mail a user has.

EXAMPLES

Display header information from all the messages in your mailbox.

```
mailfrom
```

Display header information from all new messages in your mailbox.

```
mailfrom -s new
```

Assuming you have the proper file permissions to read **guest**'s mail, print out header information from all new and unread messages in **guest**'s incoming mailbox.

```
mailfrom -s new -s unread guest
```

Print only a summary of how many new, unread, and read messages are in your incoming mailbox.

```
mailfrom -q -S
```

FILES

`$HOME/.elm/elmrc` Your **elm** configuration file.

`/var/mail` Directory of incoming mailboxes.

AUTHOR

mailfrom was developed by HP.

SEE ALSO

`elm(1)`, `mail(1)`, `mailx(1)`, `readmail(1)`.

NAME

mailq - prints the mail queue

SYNOPSIS

mailq [-v]

DESCRIPTION

mailq prints a summary of the mail messages queued for future delivery.

The first line printed for each message shows the internal identifier used on this host for the message, the size of the message in bytes, the date and time the message was accepted into the queue, and the envelope sender of the message. The second line shows the error message that caused this message to be retained in the queue; it will not be present if the message is being processed for the first time. The following lines show message recipients, one per line.

mailq is identical to **sendmail -bp**.

The options are as follows:

Options

-v Print verbose information. This adds the priority of the message and a single character indicator (“+” or blank) indicating whether a warning message has been sent on the first line of the message. Additionally, extra lines may be intermixed with the recipients indicating the “controlling user” information; this shows who will own any programs that are executed on behalf of this message and the name of the alias this command expanded from, if any.

RETURN VALUE

The **mailq** utility exits 0 on success, and >0 if an error occurs.

AUTHOR

mailq was developed by the University of California, Berkeley, and originally appeared in 4.0BSD.

FILES

/var/spool/mqueue/* mail queue files

SEE ALSO

sendmail(1M).

NAME

mailstats - print mail traffic statistics

SYNOPSIS

mailstats [-o] [-C *cffile*] [-f *stfile*]

DESCRIPTION

mailstats reads and interprets the **sendmail** statistics file, **/etc/mail/sendmail.st**, then prints out the mail traffic statistics. If **/etc/mail/sendmail.st** exists, **sendmail** collects statistics about your mail traffic and stores them in the statistics file. This file does not grow.

Statistics are gathered on a per-mailer basis for each mailer defined in the **sendmail** configuration file. Statistics are kept on the number of messages and the number of bytes for all inbound and outbound traffic.

The **mailstats** utility displays the time at which statistics collection was started on the first line. Then, the statistics of each mailer is displayed on a single line, each with the following white space separated fields (see the "EXAMPLES" section below):

| | |
|-------------------|-------------------------------------|
| M | The mailer number. |
| msgsfr | Number of messages from the mailer. |
| bytes_from | Kbytes from the mailer. |
| msgsto | Number of messages to the mailer. |
| bytes_to | Kbytes to the mailer. |
| msgsrej | Number of messages rejected. |
| msgsdisc | Number of messages discarded. |
| Mailer | The name of the mailer. |

After this display, a line totaling the values for all of the mailers is displayed, separated from the previous information by a line containing only equals ("=") characters.

Options

The options are as follows:

- C Read the specified file instead of the default **/etc/mail/sendmail.cf** file.
- f Read the specified statistics file instead of the statistics file specified in the **/etc/mail/sendmail.cf** file.
- o Do not display the name of the mailer in the output.

To clear the statistics file, execute, as root:

```
cp/dev/null/etc/mail/sendmail.st
```

RETURN VALUE

The **mailstats** utility exits 0 on success, and >0 if an error occurs.

DIAGNOSTICS

mailstats generates error messages if the statistics file is not accessible or if the size of the statistics file has changed. Error messages are:

mailstats: file size changed

Either **/etc/mail/sendmail.st** is zero length, meaning that no mail has been transferred since it was cleared out, or its size has changed. Since the size of this file is supposed to remain constant, any change in size means that the file is invalid.

mailstats: /etc/mail/sendmail.st: No such file or directory

The statistics file does not exist.

mailstats: /etc/mail/sendmail.st: Permission denied

The statistics file's permissions are set so that you cannot read it.

EXAMPLES

Here is a typical example of **mailstats** output:

| Statistics from Thu Jul 11 14:12:40 1996 | | | | | | | |
|--|--------|------------|--------|----------|---------|----------|--------|
| M | msgsfr | bytes_from | msgsto | bytes_to | msgsrej | msgsdiss | Mailer |
| 0 | 0 | 0K | 3 | 4K | 0 | 0 | prog |
| 3 | 3 | 4K | 0 | 0K | 0 | 0 | local |
| 5 | 2 | 1K | 11 | 11K | 4 | 0 | esmtpr |
| ===== | | | | | | | |
| T | 13 | 13K | 14 | 15K | 4 | 0 | |

This example shows that mailers 0, 3 and 5 have handled the given amounts of mail traffic since Thursday, Jul 11th. Specifically, **sendmail** has sent 11 messages containing 11 kilobytes via mailer esmtpr (M 5), but has 4 messages rejected via mailer esmtpr (M 5).

AUTHOR

mailstats was developed by the University of California, Berkeley.

FILES

- /etc/mail/sendmail.st mail traffic statistics file
- /etc/mail/sendmail.cf sendmail configuration file

SEE ALSO

sendmail(1M).



NAME

mailx - interactive message processing system

SYNOPSIS**Send mode:**

mailx [-FUm] [-s *subject*] [-r *address*] [-h *number*] *address* ...

Receive mode:

mailx -e

mailx [-UHLiNn] [-u *user*]

mailx -f [-UHLiNn] [*filename*]

Obsolescent:

mailx [-f *filename*] [-UHLiNn]

DESCRIPTION

mailx provides a comfortable, flexible environment for sending and receiving messages electronically. 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 created.

Incoming mail for each user is stored in a standard file called the **system mailbox** for that user. When using **mailx** to read messages, the system mailbox is used unless an alternate mailbox file is specified by using the **-f** option with or without a specific filename. As incoming messages are read from the system mailbox, 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 *mbx* and is usually located in the user's **HOME** directory (see **MBOX** description under ENVIRONMENT VARIABLES below for a description of this file and other environment variables used by **mailx**). Messages remain in this file until specifically removed.

Command-line options start with a hyphen (-), and any other arguments are assumed to be destinations (recipients). If no recipients are specified, **mailx** attempts to read messages from the system mailbox.

Recipient addresses specified on the command line must total less than 1024 characters in length. You may declare an **alias** or **group** (see **COMMANDS** below) to specify a recipient address or list of addresses of up to 8191 characters, and use that alias or group name (though each address in the list must still be less than 1024 characters). If you wish to specify a list of recipient addresses of greater length than this, have your system administrator declare an alias or group in the system alias file **/etc/mail/aliases** and use that alias name instead.

Options

mailx recognizes the following command-line options:

- e Test for presence of mail. **mailx** prints nothing and exits with a successful return code if there is mail to read. Sometimes used in login scripts such as **\$HOME/.profile** to check for mail during login.
- f Read messages from *filename* instead of the user's system mailbox. If *filename* is not specified, the secondary *mbx* is used.
- F Record the message in a file named after the first recipient. Overrides the **record** environment variable, if set.
- h *number* The number of network "hops" made so far. This is provided for network software to prevent infinite delivery loops.
- H Print header summary only.
- L Print complete header information only.
- i Ignore interrupts. Also see the description of the **ignore** environment below.
- n Do not initialize from the system default **mailx.rc** file.
- m Do not add MIME header lines *Mime Version*, *Content Type* & *Content Encoding* to the header information while sending mails.

- N** Do not print initial header summary.
- r address** Pass *address* to network delivery software. All tilde commands are disabled.
- s subject** Set the Subject header field to *subject*.
- u user** Read *user's mailbox*. Can be used only if read access to *user's* mailbox is not read protected.
- U** Convert UUCP-style addresses to Internet standards. Overrides the **conv** environment variable.
- d** Turn on debugging output. Neither particularly interesting nor recommended.

When reading mail, **mailx** operates in **command mode**. A header summary of the first several messages is displayed, followed by a prompt indicating that **mailx** can accept regular commands (see COMMANDS below). When sending mail, **mailx** operates 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 can 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.

The behavior of **mailx** at any given time is governed by a set of **environment variables**; flags and valued parameters that are set and cleared by using the **set** and **unset** commands. See ENVIRONMENT VARIABLES below for a summary of these parameters.

Recipients listed on the command line can be of three types: login names, shell commands, or alias groups. Login names can be any network address, including mixed network addressing. If the recipient name begins with a pipe symbol (|), the rest of the name is assumed 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** (see *lp(1)*) for recording outgoing mail on paper. Alias groups are set by the **alias** command (see COMMANDS below) and are lists of recipients of any type.

Regular commands are of the form

[*command*] [*msglist*] [*arguments*]

If no command is specified in command mode, **print** is assumed. In input mode, commands are recognized by the escape character (tilde unless redefined by the **escape** environment variable), and lines not treated as commands are treated as input for the message.

Each message is assigned a sequential number, and there is always the notion of a **current message**, marked by a > in the header summary. Many commands take an optional list of messages (*msglist*) to operate on, which defaults to the current message. A *msglist* is a list of message specifications separated by spaces. The message list can include:

- n* Message number *n*.
- .** The current message.
- ^** The first undeleted message.
- \$** The last message.
- *** All messages.
- n-m* An inclusive range of message numbers, *n* through *m*, where *n* is less than *m*.
- user* All messages from *user*.
- /string* All messages with *string* in the subject line (uppercase-lowercase differences are ignored).
- :c** All messages of type *c*, where *c* is one of:
 - 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. File names, where expected, are expanded using normal shell conventions (see *sh(1)*). Special characters are recognized by certain commands, and are documented with the commands below.

At start-up time, **mailx** reads commands from a system-wide file (`/usr/share/lib/mailx.rc`) to initialize certain parameters, then from a private start-up file (`$HOME/.mailrc`) for personalized variables. Most regular commands are legal inside start-up files, the most common use being 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**. Any errors in the start-up file cause the remaining lines in the file to be ignored.

COMMANDS

The following is a complete list of **mailx** commands:

| | |
|--|---|
| ! <i>command</i> | Escape to the shell. See the description of the SHELL environment variable below. |
| # <i>comment</i> | Null command (comment). Useful in .mailrc files. |
| = | Print the current message number. |
| ? | Print a summary of commands. |
| <new-line> | Advance to next message and print . If this is the first command entered, the first unread message is printed. (To read the current message, use print .) |
| alias <i>alias name...</i> | |
| group <i>alias name...</i> | Declare an alias for the given names. The names are substituted when <i>alias</i> is used as a recipient. Useful in the .mailrc file. |
| alternates <i>name...</i> | Declares 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, alternates prints the current list of alternate names. See also allnet under ENVIRONMENT VARIABLES. |
| cd [<i>directory</i>] | |
| chdir [<i>directory</i>] | Change directory. If <i>directory</i> is not specified, \$HOME is used. |
| copy [<i>filename</i>] | |
| copy [<i>msglist</i>] <i>filename</i> | Copy messages to the file without marking the messages as saved. Otherwise equivalent to the save command. |
| Copy [<i>msglist</i>] | 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 [<i>msglist</i>] | Delete messages from the <i>mailbox</i> . If autoprint is set, the next message after the last one deleted is printed (see ENVIRONMENT VARIABLES). See also dp . |
| discard [<i>header-field ...</i>] | |
| ignore [<i>header-field ...</i>] | Suppresses printing of the specified header fields when displaying messages on the screen. Examples of header fields to ignore are "status" and "cc." The fields are included when the message is saved. The Print and Type commands override this command. |
| dp [<i>msglist</i>] | |
| dt [<i>msglist</i>] | 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 <i>string ...</i> | Echo the given string or strings (similar to echo – see <i>echo(1)</i>). |
| edit [<i>msglist</i>] | 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 <i>ed</i> (see <i>ed(1)</i>). |
| exit | |

| | |
|---|--|
| xit | Exit from mailx , without changing the mailbox. No messages are saved in the <i>mbox</i> (see also quit). |
| file [<i>filename</i>] | |
| folder [<i>filename</i>] | Quit from the current file of messages and read in the specified file. Several special characters are recognized when used as file names, and substitutions are made as follows: % the current mailbox. % <i>user</i> the mailbox for <i>user</i> . # the previous file. & the current <i>mbox</i> . Default file is the current mailbox. |
| folders | Print the names of the files in the directory set by the folder variable (see ENVIRONMENT VARIABLES). |
| followup [<i>message</i>] | Respond to a message and record 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 (see ENVIRONMENT VARIABLES). |
| Followup [<i>msglist</i>] | Respond to the first message in the <i>msglist</i> , sending the message to the author of each message in the <i>msglist</i> . The subject line is extracted 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 (see ENVIRONMENT VARIABLES). |
| from [<i>msglist</i>] | Print the header summary for the specified messages. |
| group <i>alias name...</i> | |
| alias <i>alias name...</i> | Declare an alias for the given names. The names are substituted when <i>alias</i> is used as a recipient. Useful in the .mailrc file. |
| headers [<i>message</i>] | Prints 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 | Prints a summary of commands. |
| hold [<i>msglist</i>] | |
| preserve [<i>msglist</i>] | Holds the specified messages in the <i>mailbox</i> . |
| if <i>s r</i> <i>mail-commands</i> else <i>mail-commands</i> endif | Conditional execution, where s executes the accompanying <i>mail-commands</i> , up to an else or endif if the program is in send mode, and r causes the accompanying <i>mail-commands</i> to be executed only in receive mode. Intended for use in .mailrc files. |
| ignore <i>header-field ...</i> discard <i>header-field ...</i> | Suppresses printing of the specified header fields when displaying messages on the screen. Examples of header fields to ignore are status and cc . All fields are included when the message is saved. The Print and Type commands override this command. |
| list | Prints all commands available. No explanation is given. |
| mail <i>name ...</i> | Mail a message to the specified users. |
| mbox [<i>msglist</i>] | Arrange for the given messages to end up in the standard <i>mbox</i> save file when mailx terminates normally. See MBOX description under ENVIRONMENT VARIABLES for a description of this file. See also the exit and quit commands. |
| next [<i>message</i>] | Go to next message matching <i>message</i> . A <i>msglist</i> can 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 interpreted as a command |

in the absence of a real command. See the discussion of *msglists* above for a description of possible message specifications.

pipe [*msglist*] [*command*]
| [*msglist*] [*command*]

Pipe messages in *msglist* through the specified *command*. Each message is treated as if it were read. If *msglist* is not specified, the current message is used. If *command* is not specified, the command specified by the current value of the **cmd** variable is used. If *msglist* is specified, *command* must also be specified. If the **page** variable is set, a form feed character is inserted after each message (see ENVIRONMENT VARIABLES).

preserve [*msglist*]
hold [*msglist*]

Preserve the specified messages in the *mailbox*.

Print [*msglist*]
Type [*msglist*]

Print the specified messages on the screen, including all header fields. Overrides suppression of fields by the **ignore** command.

print [*msglist*]
type [*msglist*]

Print the specified messages. If **crt** is set, 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** (see *pg(1)*), but many users prefer **more** (see *more(1)* – see ENVIRONMENT VARIABLES).

quit

Exit from **mailx**, storing messages that were read in *mbox* and unread messages in the user's system mailbox. Messages that have been explicitly saved in a file are deleted.

Reply [*msglist*]
Respond [*msglist*]

Send a response to the author of each message in the *msglist*. The subject line is taken from the first message. If **record** is set to a file name, the response is saved at the end of that file (see ENVIRONMENT VARIABLES).

reply [*message*]
respond [*message*]

Reply to the specified message, including all other recipients of the message. If **record** is set to a file name, the response is saved at the end of that file (see ENVIRONMENT VARIABLES).

Save [*msglist*]

Save the specified messages in a file whose name is derived from the author of the first message. The name of the file is based on the author's name with all network addressing stripped off. See also the **Copy**, **followup**, and **Followup** commands and **outfolder** (see ENVIRONMENT VARIABLES).

save [*filename*]
save [*msglist*] *filename*

Save the specified messages in the given file. The file is created if it does not exist. The message is deleted from the *mailbox* when **mailx** terminates unless **keep-save** is set (see ENVIRONMENT VARIABLES and the **exit** and **quit** commands).

set
set *name*
set *name=string*
set *name=number*

Define a variable called *name*. The variable can be given a null, string, or numeric value. **Set** by itself prints all defined variables and their values (see ENVIRONMENT VARIABLES for detailed descriptions of the **mailx** variables).

shell

Invoke an interactive shell (see **SHELL** under ENVIRONMENT VARIABLES).

size [*msglist*]

Print the size in characters of the specified messages.

source *filename*

Read commands from the given file and return to command mode.

top [*msglist*]

Print the top few lines of the specified messages. If the **toplines** variable is set, it is interpreted as the number of lines to print (see ENVIRONMENT VARIABLES). The default is 5.

touch [*msglist*]

Touch the specified messages. If any message in *msglist* is not specifically saved in a file, it is placed in the *mbox* upon normal termination. See **exit** and **quit**.

Type [*msglist*]

| | |
|---|---|
| Print [<i>msglist</i>] | Print the specified messages on the screen, including all header fields. Overrides suppression of fields by the ignore command. |
| type [<i>msglist</i>] | |
| print [<i>msglist</i>] | Print the specified messages. If crt is set, 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 <i>pg(1)</i> but many users prefer <i>more(1)</i> (see ENVIRONMENT VARIABLES). |
| unalias <i>alias</i> | Discard the specified <i>alias</i> names. |
| undelete [<i>msglist</i>] | Restore the specified deleted messages. Restores only messages that were deleted in the current mail session. If autoprint is set, the last message of those restored is printed (see ENVIRONMENT VARIABLES). |
| unset <i>name...</i> | Cause the specified variables to be erased. If the variable was a shell variable imported from the execution environment, it cannot be erased. |
| version | Prints the current version and release date. |
| visual [<i>msglist</i>] | 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 [<i>msglist</i>] <i>filename</i> | Write the given messages on the specified file, except for the header (the "From ..." line) and trailing blank line. Otherwise equivalent to the save command. |
| xit exit | Exit from mailx , without changing the <i>mailbox</i> . No messages are saved in the <i>mbox</i> (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 commands can be used only when in input mode, by beginning a line with the tilde escape character (~). See **escape** (ENVIRONMENT VARIABLES) for changing this special character.

| | |
|-------------------------------|---|
| ~! <i>command</i> | Escape to the shell. |
| ~. | Simulate end of file (terminate message input). |
| ~: <i>mail-command</i> | |
| ~_ <i>mail-command</i> | 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 <i>name ...</i> | Add <i>name</i> to the blind carbon copy (Bcc) list. |
| ~c <i>name ...</i> | Add <i>name</i> to the carbon copy (Cc) list. |
| ~d | Read in the dead.letter file. See DEAD (under ENVIRONMENT VARIABLES) for a description of this file. |
| ~e | Invoke the editor on the partial message. Also see the EDITOR environment variable description below. |
| ~f [<i>msglist</i>] | Forward the specified messages. The messages are inserted into the message without alteration. |
| ~h | Prompt for Subject line and To, Cc, and Bcc lists. If the field is displayed with an initial value, it can be edited as if you had just typed it. |
| ~i <i>string</i> | Insert the value of the named variable into the text of the message. For example, ~A is equivalent to ~i sign . |
| ~m [<i>msglist</i>] | Insert the specified messages into the letter, shifting the new text to the right one tab stop. Valid only when sending a message while reading mail. |

| | |
|-----------------------|--|
| ~p | Print the message being entered. |
| ~q | Quit (terminate) input mode by simulating an interrupt. If the body of the message is not null, the partial message is saved in dead.letter . See the description of the DEAD environment variable below for a description of this file. |
| ~R name ... | Add <i>name</i> to the Reply-To list. |
| ~r filename | |
| ~< filename | |
| ~<! command | Read in the specified file. If the argument begins with an exclamation point (!), the rest of the string is assumed to be an arbitrary shell command and is executed, with the standard output inserted into the message. |
| ~s string ... | Set the subject line to <i>string</i> . |
| ~t name ... | Add the given <i>names</i> to the To list. |
| ~v | Invoke a preferred screen editor on the partial message. Also see the VISUAL environment variable description below. |
| ~w filename | Write the partial message onto the given file, without the header. |
| ~x | Exit as with ~q except the message is not saved in <i>dead.letter</i> . |
| ~ command | Pipe the body of the message through the given <i>command</i> . If <i>command</i> returns a successful exit status, the output of the command replaces the message. |

ENVIRONMENT VARIABLES

The following variables are internal **mailx** program variables. They can be imported from the execution environment or set by the **set** command at any time. The **unset** command can be used to erase variables.

| | |
|------------------------|---|
| allnet | All network names whose login names match are treated as identical. This causes the <i>msglist</i> message specifications to behave similarly. Default is noallnet . See also the alternates command and the metoo variable. |
| append | Upon termination, append messages to the end of the <i>mbox</i> file instead of inserting them at the beginning of the file. Default is noappend . |
| askbcc | Prompt for the Bcc list after the message is entered. Default is noaskbcc . |
| askcc | Prompt for the Cc list after the message is entered. Default is noaskcc . |
| asksub | Prompt for a subject if it is not specified on the command line with the -s option. Enabled by default. |
| autoprint | Enable automatic printing of messages after delete and undelete commands. Default is noautoprint . |
| bang | Enable special-case treatment of exclamation points (!) in shell escape command lines as in <i>vi</i> (1). Default is nobang . |
| charset=charset | Set the default character-set. If none is specified, mailx will attempt to use the value of LANG to look up the system default for the user's locale. If that is unsuccessful, the default value of <i>us-ascii</i> will be used. |
| cmd=command | Set the default command for the pipe command. No default value. |
| conv=conversion | Convert UUCP addresses to the specified address style. The only valid conversion currently supported is <i>internet</i> , which requires a mail delivery program conforming to the RFC822 standard for electronic mail addressing. Conversion is disabled by default. See also sendmail and the -U command-line option. |
| crt=number | Pipe messages having more than <i>number</i> lines through the command specified by the value of the PAGER variable pg by default (see <i>pg</i> (1)). Disabled by default. |
| DEAD=filename | The name of the file in which to save partial letters in case of untimely interrupt or delivery errors. Default is \$HOME/dead.letter . |
| debug | Enable verbose diagnostics for debugging. Messages are not delivered. Default is nodebug . |
| dot | When processing input from a terminal, interpret an ASCII period character on a line by itself as end-of-file. Default is nodot . |

| | |
|----------------------------------|--|
| EDITOR= <i>command</i> | The command to run when the edit or ~e command is used. Default is ed (see <i>ed(1)</i>). |
| encoding= <i>encoding</i> | Set the default encoding to be used when 8-bit characters are present. Allowable values are <i>quoted-printable</i> , <i>base64</i> and <i>8bit</i> . The short-hand <i>q-p</i> is also acceptable for <i>quoted-printable</i> . The default value will be determined based upon the value of charset . A value of <i>8bit</i> means not to encode. |
| escape= <i>c</i> | Substitute <i>c</i> for the ~ escape character. |
| folder= <i>directory</i> | 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 <i>directory</i> does not start with a slash (/), \$HOME is used as a prefix. 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 system mailbox instead of putting them in the standard <i>mbox</i> save file. Default is nohold . |
| ignore | Ignore interrupts while entering messages. Useful when communicating over noisy dial-up lines. Default is noignore . |
| ignoreeof | Ignore end-of-file during message input. Input must be terminated by a period (.) on a line by itself or by the ~. command. Default is noignoreeof . See also dot above. |
| keep | When the <i>mailbox</i> is empty, truncate it to zero length instead of removing it. Disabled by default. |
| keepsave | Keep messages that have been saved in other files in the system mailbox instead of deleting them. Default is nokeepsave . |
| MBOX= <i>filename</i> | The name of the file to save messages which have been read. The xit command overrides this function, as does saving the message explicitly in another file. Default is \$HOME/mbox . |
| metoo | Usually, when a group (alias) containing the sender is expanded, the sender is removed from the expansion. Setting this option causes the sender to be included in the group. Default is no metoo . |
| mimeheader= <i>value</i> | To add or disable MIME header when sending mail. "value" can be 'yes' or 'no'. |
| LISTER= <i>command</i> | The command (and options) to use when listing contents of the folder directory. The default is <i>ls(1)</i> . |
| NOMETAMAIL= <i>value</i> | To disable the usage of metamail to read MIME messages, set the value to TRUE. By default the NOMETAMAIL variable is not set. |
| 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 (i.e., one hop away). |
| outfolder | Cause the files used to record outgoing messages to be located in the directory specified by the folder variable. Default is nooutfolder . 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. Default is no page . |
| PAGER= <i>command</i> | The command to use as a filter for paginating output. This can also be used to specify the pager command-line options (for example, set PAGER="more -c"). Default is pg (see <i>pg(1)</i>), but many users prefer more (see <i>more(1)</i>). |
| prompt= <i>string</i> | Set the <i>command mode</i> prompt to <i>string</i> . Default is ? . |

| | |
|-------------------------|---|
| quiet | Refrain from printing the opening message and version when entering mailx . Default is noquiet . |
| record=filename | Record all outgoing mail in <i>filename</i> . Disabled by default. See also outfolder above. |
| 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 | Set the number of lines in a screen-full of headers for the headers command. |
| sendmail=command | Alternate command for delivering messages. Default is mail (see <i>mail(1)</i>). |
| sendwait | Wait for background mailer to finish before returning. Default is nosendwait . |
| SHELL=command | The name of a preferred command interpreter. Default is the user's login program (see <i>passwd(4)</i> , <i>shells(4)</i> , and <i>chsh(1)</i>). Note: in the unusual case that a user's login program is a script file from which mailx is executed, rather than a shell, then mailx requires that the user explicitly set SHELL=/usr/bin/sh in his or her \$HOME/.mailrc file. |
| 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 that is inserted into the text of a message when the ~a (autograph) command is given. No default (see also ~i under ILDE ESCAPES). |
| Sign=string | The variable inserted into the text of a message when the ~A command is given. No default (see also ~i (TILDE ESCAPES)). |
| SMARTMAILER | When SMARTMAILER is set, various commands use the From: line instead of the default From line. |
| toplines=number | The number of lines of header to print with the top command. Default is 5. |
| VISUAL=command | The name of a preferred screen editor. Default is vi (see <i>vi(1)</i>). |

EXTERNAL INFLUENCES

Environment Variables

The following are environment variables taken from the execution environment and are not alterable within **mailx**.

| | |
|--------------------------------------|---|
| HOME | The user's home directory. This is usually the current directory immediately after login. |
| MAILRC | The name of the mailer start-up file. Default is \$HOME/.mailrc . |
| LC_COLLATE LC_CTYPE | LC_COLLATE and LC_CTYPE influence mailx when the command interpreter (see SHELL below) is invoked. To determine the behavior of LC_COLLATE and LC_CTYPE , see the corresponding shell manual entry for the applicable command interpreter |
| LC_TIME | LC_TIME determines the format and contents of the date and time strings displayed. If LC_TIME is not specified in the environment, or is set to the empty string, the value of LANG is used as a default. If LANG is not specified or is set to the empty string, a default of "C" (see <i>lang(5)</i>) is used instead of LANG . If any internationalization variable contains an invalid setting, mailx behaves as if all internationalization variables are set to "C". See <i>environ(5)</i> . |
| TMPDIR | When set, the TMPDIR environment variable specifies a directory to be used for temporary files, overriding the default directory /tmp . |

International Code Set Support

Single- and multi-byte character code sets are supported within mail text. Headers are restricted to characters from the 7-bit USASCII character code set (see *ascii(5)*).

WARNINGS

Where *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 internationalization standards need some time to settle down.

mail(1), the standard mail delivery program, treats a line consisting solely of a dot (.) as the end of the message.

Using two separate mail programs to access the same mail file simultaneously (usually inadvertently from two separate windows) can cause unpredictable results.

FILES

| | |
|--------------------------------------|---|
| <code>/var/mail/</code> | Post office directory (mode 775, group ID mail) |
| <code>/var/mail/ user</code> | System mailbox for <i>user</i> (mode 660, owned by <i>user</i> , group ID mail) |
| <code>\$HOME/.mailrc</code> | Personal start-up file |
| <code>/usr/share/lib/mailx.rc</code> | Global start-up file |
| <code>\$HOME/mbox</code> | Secondary storage file |
| <code>/tmp/R[emqsx]*</code> | Temporary files |

SEE ALSO

mail(1), *pg*(1), *ls*(1).

STANDARDS CONFORMANCE

mailx: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

make - maintain, update, and regenerate groups of programs

SYNOPSIS

make [-f *makefile*] [-bBdeiknpQrsStuw] [*macro_name=value*] [*names*]

DESCRIPTION**Makefile Structure**

A makefile can contain four different kinds of lines: target lines, shell command lines, macro definitions, and include lines.

TARGET LINES

Target lines consist of a blank-separated, non-null list of targets, followed by a colon (:) or double colon (::), followed by a (possibly null) list of prerequisite files called **dependents**. Pattern Matching Notation (see *regex*(5)) is supported for the generation of file names as dependents.

SHELL COMMAND LINES

Text following a semicolon (;) on a target line, and all following lines that begin with a tab are shell commands to be executed to update the target (see the Environment section below about SHELL). The first line that does not begin with a tab or # begins a new target definition, macro definition, or include line. Shell commands can be continued across lines by using a <backslash><new-line> sequence.

Target lines with their associated command lines are called *rules*.

MACROS

Lines of the form *string1* = *string2* are macro definitions. Macros can be defined anywhere in the makefile, but are usually grouped together at the beginning. *string1* is the macro name; *string2* is the macro value. *string2* is defined as all characters up to a comment character or an unescaped new-line. Spaces and tabs immediately to the left and right of the = are ignored. Subsequent appearances of \$(*string1*) anywhere in the makefile (except in comments) are replaced by *string2*. The parentheses are optional if a single character macro name is used and there is no substitute sequence. An optional substitute sequence, \$(*string1* [:*subst1*=[*subst2*]]) can be specified, which causes all nonoverlapping occurrences of *subst1* at the end of substrings in the value of *string1* to be replaced by *subst2*. Substrings in a macro value are delimited by blanks, tabs, new-line characters, and beginnings of lines. For example, if

```
OBJS = file1.o file2.o file3.o
```

then

```
$(OBJS:.o=.c)
```

evaluates to

```
file1.c file2.c file3.c
```

Macro values can contain references to other macros (see WARNINGS):

```
ONE = 1
```

```
TWELVE = $(ONE)2
```

The value of \$(**TWELVE**) is set to \$(**ONE**)2 but when it is used in a target, command, or include line, it is expanded to 12. If the value of **ONE** is subsequently changed by another definition further down in the makefile or on the command line, any references to \$(**TWELVE**) reflect this change.

Macro definitions can also be specified on the command line and override any definitions in the makefile.

(XPG4 only. Macros on the command line are added to the **MAKEFLAGS** environment variable. Macros defined in the **MAKEFLAGS** environment variable, but without any command line macro, adds the macro to the environment overwriting any existing environment variable of the same name.)

Certain macros are automatically defined for **make** (see Built-in Macros). See the Environment section for a discussion of the order in which macro definitions are treated.

The value assigned to a macro can be overridden by a conditional macro definition. A conditional macro definition takes on the form *target* := *string1* = *string2*. When the target line associated with target is being processed, the macro value specified in the conditional macro definition is in effect. If *string1* is previously defined, the new value of *string1* will override the previous definition. The new value of *string1* takes effect when target or any dependents of target are being processed.

INCLUDE LINES

If the string **include** appears as the first seven letters of a line in a makefile, and is followed by one or more space or tab characters, the rest of the line is assumed to be a file name and is read and processed by the current invocation of **make** as another makefile after any macros in the filename have been expanded.

The default behaviour of **make** is to use **.DEFAULT** built-in target, if target does not have explicit commands associated with it and **.DEFAULT** target is defined. See the Built-In Targets Section.

General Description

make executes commands previously placed in a makefile to update one or more target names. Target names are typically names of programs. If no **-f** option is specified, the filenames **makefile**, **Makefile**, **s.makefile**, **SCCS/s.makefile**, **s.Makefile** and **SCCS/s.Makefile** are tried in that order. If **-f** is specified, the standard input is used. More than one **-f** option can be specified. The makefile arguments are processed in the order specified. A space between the **-f** and the filename **must** be present, and multiple makefile names must each have their own **-f** option preceding them. The contents of a makefile override the built-in rules and macros if they are present.

If no target names are specified on the command line, **make** updates the first target in the (first) makefile that is not an inference rule. A target is updated only if it depends on files that are newer than the target. Missing files are deemed to be out-of-date. All dependents of a target are recursively updated, if necessary, before the target is updated. This effects a depth-first update of the dependency tree for the target.

If a target does not have any dependents specified after the separator on the target line (*explicit dependents*), any shell commands associated with that target are executed if the target is out-of-date.

A target line can have either a single or double colon between the target name or names and any explicit dependent names. A target name can appear on more than one target line, but all of those lines must be of the same (single- or double-colon) type. For the usual single-colon case, at most one of these target lines can have explicit commands associated with it. If the target is out-of-date with any of its dependents on any of the lines, the explicit commands are executed, if they are specified, or else a default rule can be executed. For the double-colon case, explicit commands can be associated with more than one of the target lines containing the target name; if the target is out-of-date with any of the dependents on a particular line, the commands for that line are executed. A built-in rule may also be executed.

Target lines and their associated shell command lines are also referred to as **rules**. Hash marks (#) and new-line characters surround comments anywhere in the makefile except in rules. Comments in the rules depend on the setting of the **SHELL** macro.

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**:

```
OBJS = a.o b.o
pgm: $(OBJS)
    cc $(OBJS) -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. Each command line can have one or more of the following prefixes: **-**, **@**, or **+**. These prefixes are explained below.

Commands returning non-zero status normally terminate **make**. The **-i** option or the presence of the special target **.IGNORE** in the makefile cause **make** to continue executing the makefile regardless of how many command lines cause errors, although the error messages are still printed on standard output. If **-** is present at the beginning of a command line, any error returned by that line is printed to standard output but **make** does not terminate. The prefix **-** can be used to selectively ignore errors in a makefile. If the **-k** option is specified and a command line returns an error status, work is abandoned on the current target, but continues on other branches that do not depend on that target. If the **-k** option is present in the **MAKEFLAGS** environment variable, processing can be returned to the default by specifying the **-S** option.

The **-n** option specifies printing of a command line without execution. However, if the command line has the string **\$(MAKE)** or **\${MAKE}** in it or **+** as a prefix, the line is always executed (see discussion of the **MAKEFLAGS** macro under Environment). The **-t** (touch) option updates the modified date of a file without executing any commands.

A command line is normally printed before it is executed, but if the line has a @ at the beginning, printing is suppressed. The **-s** option or the presence of the special target **.SILENT:** in the makefile suppresses printing of all command lines. The @ can be used to selectively turn off printing. Everything printed by **make** (except the initial tab) is passed directly to the shell without alteration. Thus,

```
echo a\  
b
```

produces

```
ab
```

just as the shell would.

The **-b** option allows old makefiles (those written for the old version of **make**) to run without errors. The old version of **make** assumed that if a target did not have any explicit commands associated with it, the user intended the command to be null, and would not execute any **.DEFAULT** rule that might have been defined. The current version of **make** operates in this mode by default. However, the current version of **make** provides a **-B** option which turns this mode off so that if a target does not have explicit commands associated with it and a **.DEFAULT** rule is defined, the **.DEFAULT** rule is executed. Note that the **-b** and **-B** options have no effect on the search and possible location and execution of an appropriate inference rule for the target. The search for a built-in inference rule other than **.DEFAULT** is always performed.

The signals **SIGINT**, **SIGQUIT**, **SIGHUP**, and **SIGTERM** (see *signal(5)*) cause the target to be deleted unless the target depends on the special name **.PRECIOUS**.

Options

The following is a brief description of all options and some special names. Options can occur in any order. They can be specified separately, or together with one -, except for the **-f** option.

- b** Compatibility mode for old (Version 7) makefiles. This option is turned on by default.
- B** Turn off compatibility mode for old (Version 7) makefiles.
- d** Debug mode. Print out detailed information on files and times examined. (This is very verbose and is intended for debugging the **make** command itself.)
- e** Environment variables override assignments within makefiles .
- f *makefile*** Description file name, referred to as the makefile. A file name of - denotes the standard input. The contents of the makefile override the built-in rules and macros if they are present. Note that the space between **-f** and *makefile* must be present. Multiple instances of this option are allowable (except for **-f -**), and are processed in the order specified.
- i** Ignore error codes returned by invoked commands. This mode is also entered if the special target name **.IGNORE** appears in the makefile.
- k** When a command returns nonzero status, abandon work on the current entry, but continue on other branches that do not depend on that target. This is the opposite of **-S**. If both **-k** and **-S** are specified, the last one specified is used.
- n** No execute mode. Print commands, but do not execute them. Even lines beginning with an @ are printed. However, lines that contain the string **\$(MAKE)** or **\${MAKE}** or that have + as a prefix to the command are executed.
- p** Write to standard output the complete set of macro definitions and target descriptions.
- P** Update in parallel more than one target at a time. The number of targets updated concurrently is determined by the environment variable **PARALLEL** and the presence of **.MUTEX** directives in make file.
- q** Question. The **make** command returns a zero or non-zero status code, depending on whether the target file is or is not up-to-date. Targets are not updated with this option.
- r** Clear suffix list and do not use the built-in rules.
- s** Silent mode. Command lines are not printed to standard output before their execution. This mode is also entered if the special target name **.SILENT** appears in the makefile.
- S** Terminate if an error occurs while executing the commands to bring a target up-to-date. This is the default and the opposite of **-k**. If both **-k** and **-S** are specified, the last one

- given is used. This enables overriding the presence of the **k** flag in the **MAKEFLAGS** environment variable.
- t** Touch the target files (causing them to be up-to-date) rather than issue the usual commands.
 - u** Unconditionally **make** the target, ignoring all timestamps.
 - w** Suppress warning messages. Fatal messages will not be affected.
- [*macro_name=value*]
Zero or more command line macro definitions can be specified. See the Macros section.
- [*names*]
Zero or more target names that appear in the makefile. Each target so specified is updated by **make**. If no names are specified, **make** updates the first target in the makefile that is not an inference rule.

Parallel Make

If **make** is invoked with the **-P** option, it tries to build more than one target at a time, in parallel. (This is done by using the standard UNIX system process mechanism which enables multiple processes to run simultaneously.) For the makefile shown in the example in the previous section, it would create processes to build **a.o** and **b.o** in parallel. After these processes were complete, it would build **pgm**.

The number of targets **make** will try to build in parallel is determined by the value of the environment variable **PARALLEL**. If **-P** is invoked, but **PARALLEL** is not set, then **make** will try to build no more than two targets in parallel.

You can use the **.MUTEX** directive to serialize the updating of some specified targets. This is useful when two or more targets modify a common output file, such as when inserting modules into an archive or when creating an intermediate file with the same name, as is done by **lex** and **yacc**. If the makefile in the previous section contained a **.MUTEX** directive of the form

```
.MUTEX: a.o b.o
```

it would prevent **make** from building **a.o** and **b.o** in parallel.

Environment

All variables defined in the environment (see *environ(5)*) are read by **make** and are treated and processed as macro definitions, with the exception of the **SHELL** environment variable which is always ignored. The value of the **SHELL** environment variable will not be used as a macro and will not be modified by defining the **SHELL** macro in a makefile or on the command line. Variables with no definition or empty string definitions are included by **make**.

There are four possible sources of macro definitions which are read in the following order: internal (default), current environment, the makefile(s), and command line. Because of this order of processing, macro assignments in a makefile override environment variables. The **-e** option allows the environment to override the macro assignments in a makefile. Command-line macro definitions always override any other definitions.

The **MAKEFLAGS** environment variable is processed by **make** on the assumption that it contains any legal input option (except **-f**, **-p**, and **-d**) defined for the command line. The **MAKEFLAGS** variable can also be specified in the makefile.

(XPG4 only. **MAKEFLAGS** in the makefile replaces the **MAKEFLAGS** environment variable. Command line options have precedence over **MAKEFLAGS** environment variable.)

If **MAKEFLAGS** is not defined in either of these places, **make** constructs the variable for itself, puts the options specified on the command line and any default options into it, and passes it on to invocations of commands. Thus, **MAKEFLAGS** always contains the current input options. This proves very useful for recursive **makes**. Even when the **-n** option is specified, command lines containing the string **\$(MAKE)** or **\${MAKE}** are executed; hence, one can perform a **make -n** recursively on an entire software system to see what would have been executed. This is possible because the **-n** is put into **MAKEFLAGS** and passed to the recursive invocations of **\$(MAKE)** or **\${MAKE}**. This is one way of debugging all of the makefiles for a software project without actually executing any of the commands.

Each of the commands in the rules is given to a shell to be executed. The shell used is the shell command interpreter (see *sh(1)*), or the one specified in the makefile by the **SHELL** macro. To ensure the same shell is used each time a makefile is executed, the line:

```
SHELL=/usr/bin/sh
```

or a suitable equivalent should be put in the macro definition section of the makefile.

Suffixes

Target and/or dependent names often have suffixes. Knowledge about certain suffixes is built into **make** and used to identify appropriate inference rules to be applied to update a target (see the section on Inference Rules). The current default list of suffixes is:

```
.o .c .c~ .C .C~ .cxx .cxx~ .cpp .cpp~ .cc .cc~
.y .y~ .l .l~ .L .L~ .Y .Y~ .s .s~ .sh .sh~
.h .h~ .H .H~ .p .p~ .f .f~ .r .r~
```

These suffixes are defined as the dependents of the special built-in target **.SUFFIXES**. This is done automatically by **make**.

Additional suffixes can be specified in a makefile as the dependents list for **.SUFFIXES**. These additional values are added to the default values. Multiple suffix lists accumulate. The order of the suffix list is significant (see the Inference Rules section). If the user wishes to change the order of the suffixes, he must first define **.SUFFIXES** with a null dependent list, which clears the current value for **.SUFFIXES**, and then define **.SUFFIXES** with the suffixes in the desired order. The list of suffixes built into **make** on any machine can be displayed by:

```
make -fp - 2>/dev/null </dev/null
```

The list of built-in suffixes incorporated with the definitions in a given makefile called **mymakefile** can be displayed by:

```
make -fp mymakefile 2>/dev/null </dev/null
```

Inference Rules

Certain target or dependent names (such as those ending with **.o**) have inferable dependents such as **.c** and **.s**, etc. If no update commands for such a name appear in the makefile, and if an inferable dependent file exists, that dependent file is compiled to update 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. There are currently default inference rules defined for:

Single Suffix Rules

```
.c .c~
.C .C~ .cxx .cxx~ .cpp .cpp~ .cc .cc~
.sh .sh~
.p .p~
.f .f~
.r .r~
```

Double Suffix Rules

```
.c.o .c~.o .c~.c .c.a .c~.a
.C.o .C~.o .C~.C .C.a .C~.a
.cxx.o .cxx~.o .cxx~.cxx .cxx.a .cxx~.a
.cpp.o .cpp~.o .cpp~.cpp .cpp.a .cpp~.a
.cc.o .cc~.o .cc~.cc .cc.a .cc~.a
.s.o .s~.o .s~.a
.p.o .p~.o .p~.p .p.a .p~.a
.f.o .f~.o .f~.f .f.a .f~.a
.r.o .r~.o .r~.r .r.a .r~.a
.y.o .y~.o .y~.c .y~.c
.l.o .l~.o .l.c
.h~.h .H~.H .hxx~.hxx .hpp~.hpp
.C.o .C~.o .C.a .C~.a
.L.C .L.o .L~.C .L~.L .L~.o
.Y.C .Y.o .Y~.C .Y~.Y .Y~.o
```

Double suffix inference rules (**.c.o**) define how to build **x.o** from **x.c**. Single suffix inference rules (**.c**) define how to build **x** from **x.c**. In effect, the first suffix is null. Single suffix rules are useful for building targets from only one source file; e.g., shell procedures and simple C programs.

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`). Since the `s.` of the SCCS files is a prefix, it is incompatible with **make**'s suffix point-of-view. Hence, the tilde is a way of changing any file reference into an SCCS file reference.

A 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 name that has no slashes in it and starts with a dot is identified as an inference (**implicit**) rule instead of a target (**explicit**) rule. Targets with one dot are single suffix inference rules; targets with two dots are double suffix inference rules. Users can, in a makefile, define additional inference rules and either redefine or cancel default inference rules.

The default inference rule for changing a `.c` file into a `.o` file might look like this:

```
.c.o:
    $(CC) $(CFLAGS) -c $<
```

and the default inference rule for changing a `yacc` file to a C object file might look like this:

```
.y.o:
    $(YACC) $(YFLAGS) $<
    $(CC) $(CFLAGS) -c y.tab.c
    rm y.tab.c
    mv y.tab.o $@
```

Certain macros are used in the default inference rules to permit the inclusion of optional matter in any resulting commands. For example, `CFLAGS`, `LDFLAGS`, and `YFLAGS` are used for compiler options to `cc(1)`, `lex(1)`, and `yacc(1)`, respectively. `LDFLAGS` is commonly used to designate linker/loader options. These macros are automatically defined by **make** but can be redefined by the user in the makefile.

The macro `LIBS` is, by convention, used to specify the order of inclusion of any special libraries during the linking phase of compilation. To specify a particular order of inclusion for a particular set of libraries, the existing single suffix rule for a `.c` file,

```
$(CC) $(CFLAGS) $< $(LDFLAGS) -o $@
```

can be redefined as

```
$(CC) $(CFLAGS) $< $(LDFLAGS) -o $@ $(LIBS)
```

as well as defining `LIBS` in the makefile.

There are also some special built-in macros used in the inference rules (`@`, `<`). See the Built-in Macros section.

If a target does not have explicit dependents, or if a dependent does not also have a target that matches it with associated explicit rules, **make** looks for the first inference rule that matches both the target's (dependent's) suffix (which may be null) and a file which matches the other suffix of the rule. Since it conducts this search by going through the list of `.SUFFIXES` values front to back, the order in which `.SUFFIXES` is defined is significant.

To print out the rules compiled into the **make** on any machine, type:

```
make -fp - 2>/dev/null </dev/null
```

Since **make** defines an inference rule `.c.o`, the example in the General Description section can be rewritten more simply:

```
OBJS = a.o b.o
pgm: $(OBJS)
    cc $(OBJS) -o pgm
$(OBJS): incl.h
```

Libraries

If a target or dependent 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 `$(LIB)(file.o)` both refer to an archive library that contains `file.o` (this assumes the `LIB` macro has been previously defined). The expression `$(LIB)(file1.o file2.o)` is not valid. Rules pertaining to archive libraries have the form `.xx.a` where `xx` is the suffix from which the archive member is to be made. An unfortunate byproduct 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 rv $@ $*.o
    rm -f $*.o
```

(See the section on Built-in Macros for an explanation of the <, @, and * symbols.) In fact, the .c.a rule listed above is built into **make** and is unnecessary in this example. This rule is applied to each dependent of **lib** in turn. The following example accomplishes this more efficiently:

```
lib: lib(file1.o) lib(file2.o) lib(file3.o)
    $(CC) -c $(CFLAGS) $(?:.o=.c)
    ar rv lib $?
    rm $?
    @echo lib is now up-to-date
.c.a;;
```

Here substitution in the macros is used. The \$? list is defined to be the set of object file names (inside **lib**) whose C source files are out-of-date. The substitution sequence translates the .o to .c. (Unfortunately, one cannot as yet transform to .c~; however, this may become possible in the future.) Note also, the disabling of the .c.a rule, which would have created and archived each object file, one by one. This particular construct speeds up archive library maintenance considerably, but becomes very cumbersome if the archive library contains a mix of assembly programs and C programs.

Archive members containing the definition of a symbol are designated by double parentheses around the symbol name, **lib((entry_name))**, but are otherwise handled as described above.

Built-In Targets

make has knowledge about some special targets. These must be specified in the makefile to take effect (with the exception of **.SUFFIXES**, which is automatically set by **make** but can be changed by the user).

- .DEFAULT** If a file must be made but there are no explicit commands or relevant built-in rules for it, the commands associated with the target name **.DEFAULT** are used if **.DEFAULT** has been defined in the makefile. **.DEFAULT** does not have any explicit dependents.
- .PRECIOUS** Dependents of this target are not removed when **QUIT**, **INTERRUPT**, **TERMINATE**, or **HANGUP** are received.
- .SILENT** Same effect as the **-s** option. No dependents or explicit commands need to be specified.
- .IGNORE** Same effect as the **-i** option. No dependents or explicit commands need to be specified.
- .SUFFIXES** The explicit dependents of **.SUFFIXES** are added to the built-in list of known suffixes and are used in conjunction with the inference rules. If **.SUFFIXES** does not have any dependents, the list of known suffixes is cleared. There are no commands associated with **.SUFFIXES**.
- .MUTEX** Serialize the updating of specified targets (See the *Parallel Make* Section).

Built-in Macros

There are five internally maintained macros that are useful for writing rules for building targets. In order to clearly define the meaning of these macros, some clarification of the terms **target** and **dependent** is necessary. When **make** updates a target, it may actually generate a series of targets to update. Before any rule (either explicit or implicit) is applied to the target to update it, recursion takes place on each dependent of the target. The dependent, upon recursion, becomes a target itself, and may have or generate its own dependents, which in turn are recursed upon until a target is found that has no dependents, at which point the recursion stops. Not all targets processed by **make** appear as explicit targets in the makefile; some of them are explicit dependents from the makefile while others are implicit dependents generated as **make** recursively updates the target. For instance, when the following makefile is executed:

```
pgm: a.o b.o
    cc a.o b.o -o pgm
```

the following series of targets to be made is generated:

```

--- pgm
    with two dependents and an explicit rule to follow
--- a.o
    (recursively) with an implicit dependent of a.c which matches the implicit rule .c.o
--- a.c
    (recursively) with no implicit dependents and no implicit rules. This stops the recursion
    and simply returns the last modification time of the file a.c.
--- b.o
    (recursively) with an implicit dependent of b.c which matches the implicit rule .c.o
--- b.c
    (recursively) with no implicit dependents and no implicit rules. This stops the recursion
    and merely returns the last modification time of the file b.c.
```

In the definitions below, the word *target* refers to a target specified in the makefile, an explicit dependent specified in the makefile which becomes the target when **make** recurses on it, or an implicit dependent (generated as a result of locating an inference rule and file that match the suffix of the target) which becomes the target when **make** recurses on it. The word dependent refers to an explicit dependent specified in the makefile for a particular target, or an implicit dependent generated as a result of locating an appropriate inference rule and corresponding file that matches the suffix of the target.

It may be helpful to think of target rules as user specified rules for a particular target name, and inference rules as user or **make** specified rules for a particular class of target names. It may also be helpful to remember that the value of the target name and its corresponding dependent names change as **make** recurses on both explicit and implicit dependents, and that inference rules are only applied to implicit dependents or to explicit dependents which do not have target rules defined for them in the makefile.

\$@ The **\$@** macro is the full target name of the current target, or the archive filename part of a library archive target. It is evaluated for both target and inference rules.

\$\$ The **\$\$** macro is only evaluated when the current target is an archive library member of the form **libname(member.o)** or **libname(entry)**. In these cases, **\$@** evaluates to **libname** and **\$\$** evaluates to **member.o** or the object file containing the symbol **entry**. **\$\$** is evaluated for both target and inference rules.

\$? The **\$?** macro is the list of dependents that are out-of-date with respect to the current target; essentially, those modules that have been rebuilt. It is evaluated for both target and inference rules, but is usually only used in target rules. **\$?** evaluates to one name only in an inference rule, but may evaluate to more than one name in a target rule.

\$< In an inference rule, **\$<** evaluates to the source file name that corresponds to the implicit rule which matches the suffix of the target being made. In other words, it is the file that is out-of-date with respect to the target. In the **.DEFAULT** rule, the **\$<** macro evaluates to the current target name. **\$<** is evaluated only for inference rules. 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 the current target name with the suffix deleted. It is evaluated only for inference rules.

These five macros can have alternative forms. When an uppercase **D** or **F** is appended to any of the five 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. When the **\$?** macro contains more than one dependent name, the **\$(?D)** expands to a list of directory name parts and the **\$(?F)** expands to a list of the filename parts.

In addition to the built-in macros listed above, other commonly used macros are defined by **make**. These macros are used in the default inference rules, and can be displayed with the **-p** option. These macros can

be used in target rules in the makefile. They can also be redefined in the makefile.

\$\$@

The **\$\$@** macro has meaning *only* on dependency lines. Macros of this form are called **dynamic dependencies** because they are evaluated at the time the dependency is actually processed. **\$\$@** evaluates to exactly the same thing as **\$@** does on a command line; i.e., the current target name. This macro is useful for building large numbers of executable files, each of which has only one source file. For instance, the following HP-UX commands could all be built using the same rule:

```
CMDS = cat echo cmp chown
$(CMD) : $$@.c
$(CC) -O $? -o $@
```

If this makefile is invoked with **make cat echo cmp chown**, **make** builds each target in turn using the generic rule, with **\$\$@** evaluating to **cat** while **cat** is the target, to **echo** when the target is **echo**, and so forth.

The dynamic dependency macro can also take the F form, **\$\$(@F)** which represents the filename part of **\$\$@**. This is useful if the targets contain pathnames. For example:

```
INCDIR = /usr/include
INCLUDES = $(INCDIR)/stdio.h \
           $(INCDIR)/pwd.h \
           $(INCDIR)/dir.h \
           $(INCDIR)/a.out.h
$(INCLUDES) : $$(@F)
cp $? $@
chmod 0444 $@
```

Special Macros

The **VPATH** macro allows **make** to search a colon separated list of directories for dependents. Lines of the form **VPATH= path1:path2 ...** causes **make** to first search the current directory for a dependent and if the dependent is not found, **make** searches *path1* and continues until the directories specified in the **VPATH** macro are exhausted.

EXTERNAL INFLUENCES

Environment Variables

LANG provides a default value for the internationalization variables that are unset or null. If **LANG** is unset or null, the default value of "C" (see *lang(5)*) is used. If any of the internationalization variables contains an invalid setting, **make** will behave as if all internationalization variables are set to "C". See *environ(5)*.

LC_ALL If set to a non-empty string value, overrides the values of all the other internationalization variables.

LC_CTYPE determines the interpretation of text as single and/or multi-byte characters, the classification of characters as printable, and the characters matched by character class expressions in regular expressions.

LC_MESSAGES determines the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.

NLSPATH determines the location of message catalogues for the processing of **LC_MESSAGES**.

PROJECTDIR provides a directory to be used to search for SCCS files not found in the current directory. In all of the following cases, the search for SCCS files will be made in the directory **SCCS** in the identified directory. If the value of **PROJECTDIR** begins with a slash, it is considered an absolute pathname; otherwise, the home directory of a user of that name is examined for a subdirectory **src** or **source**. If such a directory is found, it is used. Otherwise, the value is used as a relative pathname.

If **PROJECTDIR** is not set or has a null value, the current directory is searched first, followed by a search in the **SCCS** directory in the current directory.

The setting of **PROJECTDIR** affects all files listed in the remainder of this utility description for files with a component named **SCCS**.

International Code Set Support

Single and multi-byte character code sets are supported.

RETURN VALUES

make returns a 0 upon successful completion or a value greater than 0 if an error occurred. If the **-q** option is specified, **make** returns 0 if the target was up-to-date and a value greater than 0 if the target was not up-to-date.

EXAMPLES

The following example creates an executable file from a C source code file without a makefile, if program.c exists in the current directory:

```
make program
```

The following example shows more than one makefile specified and some command line macros defined, and updates the first target in module1:

```
make -f module1 -f module2 RELEASE=1.0 CFLAGS=-g
```

The following example updates two targets in a default makefile currently residing in the current directory:

```
make clobber prog
```

The following example updates the prog target in a specified makefile, allows environment variables to override any common variables in the makefile, clears the built-in suffix list and ignore the built-in rules, and outputs exhaustive debugging information:

```
make -erd -f module1 prog
```

WARNINGS

Be wary of any file (such as an include file) whose access, modification, and last change times cannot be altered by the **make**-ing process. For example, if a program depends on an include file that in turn depends on another include file, and if one or both of these files are out-of-date, **make** tries to update these files each time it is run, thus unnecessarily re-**make**ing up-to-date files that are dependent on the include file. The solution is to manually update these files with the **touch** command before running **make** (see *touch(1)*).

Some commands return non-zero status inappropriately; use **-i** to overcome the difficulty.

File names with the characters = : @ \$ do not work.

Built-in commands that are directly executed by the shell such as **cd** (see *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**.

The macro **\$(a:.o=.c~)** does not work.

Expanded target lines cannot contain more than 16384 characters, including the terminating new-line.

If no makefile exists in the current directory, typing

```
make filename
```

results in **make** attempting to build **filename** from **filename.c**

If **make** is invoked in a shell script with a quoted argument that evaluates to NULL (such as **\$@**), **make** fails.

DEPENDENCIES**NFS Warning:**

When comparing modification times of files located on different NFS servers, **make** behaves unpredictably if the clocks on the servers are not synchronized.

FILES

```
[Mm]akefile
s.[Mm]akefile
SCCS/s.[Mm]akefile
```

SEE ALSO

cc_bundled(1), cd(1), lex(1), mkmf(1), sh(1), environ(5), lang(5), regexp(5).

A Nutshell Handbook, Managing Projects With Make by Steve Talbot, Second Edition, O'Reilly & Associates, Inc., 1986.

STANDARDS CONFORMANCE

make: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

makekey - generate encryption key

SYNOPSIS

`/usr/sbin/makekey`

DESCRIPTION

makekey improves the usefulness of encryption schemes depending on a key by increasing the amount of time required to search the key space. It reads 10 bytes from its standard input and writes 13 bytes on its standard output. The output depends on the input in a way intended to be difficult to compute (i.e., to require a substantial fraction of a second).

The first eight input bytes (the *input key*) can be arbitrary ASCII characters. The last two (the *salt*) are best chosen from the set of digits, ., /, and uppercase and lowercase letters. The salt characters are repeated as the first two characters of the output. The remaining 11 output characters are chosen from the same set as the salt and constitute the *output key*.

The transformation performed is essentially the following: the salt is used to select one of 4,096 cryptographic machines all based on the National Bureau of Standards DES algorithm, but broken in 4,096 different ways. Using the *input key* as key, a constant string is fed into the machine and recirculated a number of times. The 64 bits that come out are distributed into the 66 *output key* bits in the result.

makekey is intended for programs that perform encryption (e.g., *ed*(1) and *crypt*(1)). Usually, its input and output will be pipes.

SEE ALSO

crypt(1), *ed*(1), *passwd*(4).

NAME

man - find manual information by keywords; print out a manual entry

SYNOPSIS

man [-**M** path] -**k** keyword...

man [-**M** path] -**f** file...

man [-] [-**M** path] [-**T** macro-package] [section [subsection]] entry_name...

DESCRIPTION

man accesses information from the HP-UX manual pages. It can be used to:

- List all manual entries whose one-line description contains any of a specified set of keywords.
- Display or print one-line descriptions of entries specified by name.
- Search on-line manual directories by entry name and display or print the specified entry or entries.
- Search a specified on-line manual section (directory) and display or print the specified entry or entries in that section.

Searching for Entry Names by Keyword (first form)

The first form above searches the one-line descriptions of individual entries for specified keywords. Arguments are as follows:

-**k** *keyword* -**k** followed by one or more keywords causes **man** to print the one-line description of each manual entry whose one-line description contains text matching one or more of the specified keywords (similar to the behavior of *grep*(1)). Keywords are separated by blanks (space or tab).

Before this option can be used, file `/usr/share/lib/whatis` must exist. `/usr/share/lib/whatis` can be created by running *catman*(1M).

Obtaining One-Line Description of an Entry (second form)

The second form above finds and displays or prints the one-line descriptions of specified individual entries. Arguments are as follows:

-**f** *file* -**f** followed by one or more file names causes **man** to print the one-line description of each manual entry found whose name matches *file*. When specifying two or more files, *file* arguments are separated by blanks (space or tab). If entry names matching *file* exist in two or more sections, the one-line description of each matched file name is output.

Before this option can be used, file `/usr/share/lib/whatis` must exist. `/usr/share/lib/whatis` can be created by running *catman*(1M).

Viewing Individual Manual Entries (third form)

The third form shown above is used for viewing one or more individual manual entries. **man** in this form recognizes the following arguments:

- (optional) When the - argument is present, **man** sends the formatted manual entry directly to standard output without processing it through the output filter specified by the **PAGER** environment variable.

-**M** *path* Change the search path for manual pages. *path* is a colon-separated list of directories that contain manual page directory subtrees. When used with the -**k** or -**f** options, the -**M** option must appear first.

-**T** *macro-package*

man uses macro-package rather than the standard -man macros defined in `/usr/share/lib/tmac/tmac.an` for formatting manual pages.

When specifying the -**T** option to **man**, the full path must be given. For example:

man -**T** `/usr/share/lib/tmac/tmac.s ls`

section[*subsection*]

(optional) Search in the specified section for the given *entry_name*. *section* specifies a single section number or one of the words **local**, **new**, **old**, or **public** to search

for one or more of the entries indicated. *section* corresponds to the section number where the entry appears in the *HP-UX Reference*. It can be followed by an optional uppercase/lowercase subsection identifier such as **3C** which would indicate a library routine in Section 3. **3**, **3c**, and **3C** are interpreted as equivalent, since all Section 3 manual entries are stored in the same or in related directories (such as **/usr/share/man/man3.Z** and **/usr/share/man/man3**). However, if an entry is in Section 1M, *section* must be specified as **1m** or **1M**.

entry_name Search for a specific entry name where *entry_name* is the name of the manual entry without its section-number suffix. Except for names exceeding 11 characters, *entry_name* is identical to the name of the manual entry as listed at the top of each page, or is the same as one of the keywords in the left-hand part of the one-line description in the corresponding manual entry.

If *entry_name* is longer than 11 characters, **man** first searches for the full-length *entry_name*. If not found, *entry_name* is truncated to 11 characters to ensure that there is room for the *section* suffix in 14-character source file names. Files in the **/usr/share/man/*** directories are normally installed with the filename truncated to 11 characters where necessary so that the name plus a three-character section suffix does not exceed the maximum filename length on short filename systems.

If *section* is not specified (see previous argument description), **man** searches all sections of the manual in order: **man1**, **man2**, **man1M**, **man3**, **man4**, **man5**, **man6**, **man7**, **man8**, **man9**, **manlocal**, **mannew**, **manold**, then **manpublic**; and printing the first matching entry it encounters.

If there is more than one manual entry among the sections, the first manual entry is displayed. For example, **man intro** will display only **intro(1)**. **man 4 intro** will display **intro(4)**.

If the standard output is a teletype, and if the **-t** flag is not given, **man** pipes its output through **more** (see **more(1)**), with the **-s** option, to eliminate multiple blank lines and stop after each screenful. This default behavior can be changed by setting the **PAGER** variable in the user's environment. The value of **PAGER** must be a string that names an output filter (such as **pg(1)**), along with the desired options.

File Search Conventions

man searches in several directories, as appropriate, for the specified manual entry. The search continues until either the entry is found or all candidate directories are searched. The first three directories searched, in order, are: **/usr/share/man**, **/usr/contrib/man**, and **/usr/local/man**.

The **MANPATH** environment variable can be used to specify directories to be searched, and, if set, overrides the default paths given above. Upon logging in, **/etc/profile** (or **/etc/csh.login**) sets the **MANPATH** environment variable to default settings. If the file **/etc/MANPATH** exists, the default settings are taken from this file. The **MANPATH** variable follows the same form as the **PATH** variable (see **environ(5)**).

Within each of these directories, **man** searches in the **cat*.Z** subdirectories, the **man*.Z** subdirectories, the **cat*** subdirectories, and the **man*** subdirectories. **man*.Z** and **man*** directories contain **nroff(1)**-compatible source text for the entries. **cat*.Z** and **cat*** directories contain the formatted versions of the entries. **man*.Z** and **cat.Z** directories contain entries in compressed form. Files in these directories are uncompressed by **uncompress** (see **compress(1)**) before being processed for printing or display.

If the **LANG** environment variable is set to any valid language name defined by **lang(5)**, and the **MANPATH** variable is not set, or is set to the default directories, **man** searches in three additional directories for the manual entry before searching in **/usr/share/man**. First, **man** searches in **/usr/share/man/\$LANG**, then in **/usr/contrib/man/\$LANG**, then in **/usr/local/man/\$LANG**. Thus, native-language manual entries are displayed if they are present and installed properly in the system.

If the **MANPATH** environment variable is set to anything other than the default, the above directories with **\$LANG** as part of the path are not automatically searched. All directories must be explicitly given in **MANPATH**. The **%L**, **%l**, **%t**, and **%c** specifiers can be used as path components to cause locale-specific directories to be searched. See **environ(5)** for a complete description of **MANPATH**.

man uses the most recent version that it finds in the subdirectories searched. If the most recent version is in:

| | |
|---------------|---|
| man*.Z | The entry is uncompressed, formatted, and displayed. If the cat*.Z directory exists, the formatted entry is compressed and installed in cat*.Z . If the cat* directory exists, the formatted entry is installed in cat* . |
| cat*.Z | The entry is uncompressed and displayed. |
| man* | The entry is formatted, and displayed. If the cat*.Z directory exists, it is compressed, and installed in cat*.Z . If the cat* directory exists, the formatted entry is installed in cat* . |
| cat* | The entry is displayed. |

If only the **cat*** or **cat*.Z** subdirectory is present and/or *nroff*(1) is not installed, only entries that are already formatted can be displayed.

If you choose to have the formatted entries on your system, run *catman*(1M) with the default, which creates the **cat*.Z** directories (after removing any **cat*** directories that exist on your system) and also creates the file `/usr/share/lib/whatIs` used by the **man -k** option. If you choose to have the **cat*** directories, it would be space-saving to remove any **cat*.Z** directories that may exist on your system. Beware that **man** updates both directories (**cat*** and **cat*.Z**) if they both exist.

Special Manual Entries

Some situations may require creation of manual entries for local use or distribution by third-party software suppliers. The manual formatting macros have been structured to redefine page footers so that manual entries not originating from Hewlett-Packard Company do not show the **HP** name in the footer. For more information about this change and a description of the manual formatting macros used with **nroff** or **troff**, see *man*(5).

EXTERNAL INFLUENCES

Environment Variables

LANG determines the language in which messages are displayed. **LANG** is also used to determine the search path (as described above).

If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang*(5)) is used instead of **LANG** for messages, but not for the search path.

If any internationalization variable contains an invalid setting, **man** behaves as if all internationalization variables are set to "C". See *environ*(5).

MANPATH, if set, gives a list of directories to be searched for the given entry, replacing the default paths.

PAGER, if set, defines an output filter to be used instead of *more*(1) to paginate output.

International Code Set Support

Single- and multi-byte character code sets are supported.

EXAMPLES

List the manual entries that contain the word **grep** in their respective one-line description (NAME) lines:

```
man -k grep
```

The output is:

```
grep, egrep, fgrep (1) - search a file for a pattern
zgrep(1)                - search possibly compressed files for a
                        regular expression
```

Print the one-line description of the *grep*(1) manual entry:

```
man -f grep
```

Print the entire *grep*(1) manual entry:

```
man grep
```

Set a search path that includes a path directly below the current directory. The manual entry, **mypage** is assumed to exist in the directory `./man1` (or `./man1.Z`, `cat1`, or `cat1.Z`).

```
MANPATH=./usr/share/man:/usr/contrib/man:/usr/local/man
export MANPATH
man mypage
```

Display the manual entry for *id*(1), with the output piped through `pg -c`:

```
PAGER="pg -c"
export PAGER
man id
```

List all printed manuals available for the current system (see *manuals*(5)):

```
man manuals
```

Display *intro*(4) and *intro*(3):

```
man 4 intro
man 3 intro
```

WARNINGS

Manual entries are structured such that they can be printed on a phototypesetter, conventional line printer, and screen display devices. However, due to line printer and display device limitations, some information may be lost in certain situations.

FILES

| | |
|--|--|
| <code>/usr/share/lib/whatis</code> | keyword database |
| <code>/usr/share/man/cat* [.Z]/*</code> | formatted manual entries [compressed] |
| <code>/usr/share/man/man* [.Z]/*</code> | raw (<i>nroff</i> (1) source) manual entries [compressed] |
| <code>/usr/contrib/man/cat* [.Z]/*</code> | |
| <code>/usr/contrib/man/man* [.Z]/*</code> | |
| <code>/usr/local/man/cat* [.Z]/*</code> | |
| <code>/usr/local/man/man* [.Z]/*</code> | |
| <code>/usr/share/man/\$LANG/cat* [.Z]/*</code> | formatted native-language manual entries [compressed] |
| <code>/usr/share/man/\$LANG/man* [.Z]/*</code> | raw (<i>nroff</i> (1) source) native-language manual entries [compressed] |
| <code>/usr/contrib/man/\$LANG/cat* [.Z]/*</code> | |
| <code>/usr/contrib/man/\$LANG/man* [.Z]/*</code> | |
| <code>/usr/local/man/\$LANG/cat* [.Z]/*</code> | |
| <code>/usr/local/man/\$LANG/man* [.Z]/*</code> | |

SEE ALSO

`col`(1), `compress`(1), `grep`(1), `more`(1), `catman`(1M), `fixman`(1M), `environ`(5), `intro`(1), `intro`(1M), `intro`(2), `intro`(3), `intro`(4), `intro`(5), `intro`(7), `intro`(9), `introduction`(9), `man`(5), `manuals`(5).

STANDARDS CONFORMANCE

`man`: XPG4

NAME

mediainit - initialize disk or partition DDS tape

SYNOPSIS

mediainit [-**vr**] [-**f** *fmt_optn*] [-**i** *interleave*] [-**p** *size*] *pathname*

DESCRIPTION

mediainit initializes mass storage media by formatting the media, writing and reading test patterns to verify media integrity, then sparing any defective blocks found. This process prepares the disk or tape for error-free operation. Initialization destroys all existing user data in the area being initialized.

mediainit can also be used for partitioning DDS tape media. See the **-p** option below for further details.

Options

The following command options are recognized. They can be specified in any order, but all must precede the *pathname*. Options without parameters can be listed individually or grouped together. Options with parameters must be listed individually, but white space between the option and its parameter is discretionary.

- v** Normally, **mediainit** provides only fatal error messages which are directed to standard error. The **-v** (verbose) option sends device-specific information related to low-level operation of **mediainit** to standard output (stdout). This option is most useful to trained service personnel because it usually requires detailed knowledge of device operation before the information can be interpreted correctly.
- r** (re-certify) This option forces a complete tape certification whether or not the tape has been certified previously. All record of any previously spared blocks is discarded, so any bad blocks will have to be rediscovered. This option should be used only if:
 - It is suspected that numerous blocks on the tape have been spared which should not have been, or
 - It is necessary to destroy (overwrite) all previous data on the tape.
- f *fmt_optn*** The format option is a device-specific number in the range 0 through 239. It is intended solely for use with certain SS/80 devices that support multiple media formats (independent from interleave factor). For example, certain microfloppy drives support 256-, 512-, and 1024-byte sectors. **mediainit** passes any supplied format option directly through to the device. The device then either accepts the format option if it is supported, or rejects it if it is not supported. Refer to device operating manuals for additional information. The default format option is 0.
- i *interleave*** The interleave factor, *interleave*, refers to the relationship between sequential logical records and sequential physical records. It defines the number of physical records on the media that lie between the beginning points of two consecutively numbered logical records. The choice of interleave factor can have a substantial impact on disk performance.
- p *size*** Partition DDS cartridge media into two logical separate volumes: partition 0 and partition 1:
 - *size* specifies the minimum size of partition 1 (in Mbytes). The maximum allowed value is 1200.
 - Partition 0 is the remainder of the tape (partition 0 physically follows partition 1 on the tape).

The actual size of partition 1 is somewhat larger than the requested size to allow for tape media errors during writing. Thus, a *size* of 400 formats the DDS tape into two partitions where partition 1 holds at least 400 Megabytes of data, and the remainder of the tape is used for partition 0 (for a 1300 Mbyte DDS cartridge, this means that partition 0 has a size somewhat less than 900 Mbytes).

Note that it is unnecessary to format a DDS tape before use unless the tape is being partitioned. Unformatted DDS media does not require initialization when used as a single partition tape. Accessing partition 1 on a single-partition tape produces an error. To change a two-partition tape to single-partition, use **mediainit** with 0 specified as the *size*.

pathname *pathname* is the path name to the character (raw) device special file associated with the device unit or volume that is to be initialized. **mediainit** aborts if you lack either read or write permission to the device special file, or if the device is currently open for any other process. This prevents accidental initialization of the root device or any mounted volume. **mediainit** locks the unit or volume being initialized so that no other processes can access it.

Except for SCSI devices, *pathname* must be a device special file whose minor number of the device being initialized has the diagnostic bit set. For device special files with the diagnostic bit set, the section number is meaningless. The entire device is accessed.

When a given unit contains multiple volumes as defined by the drive controller, any available unit or volume associated with that controller can be initialized, independent of other units and volumes that share the same controller. Thus, you can initialize one unit or volume to any format or interleave factor without affecting formats or data on companion units or volumes. However, be aware that the entire unit or volume (as defined by the drive controller) is initialized without considering the possibility that it may be subdivided into smaller structures by the the operating software. When such structures exist, unexpected loss of data is possible.

mediainit dominates controller resources and limits access by competing processes to other units or volumes sharing the same controller. If other simultaneous processes need access to the same controller, some access degradation can be expected until initialization is complete; especially if you are initializing a tape cartridge in a drive that shares the root disk controller.

In general, **mediainit** attempts to carefully check any **-f** (format option) or **-i** (interleave options) supplied, and aborts if an option is out of range or inappropriate for the media being initialized. Specifying an interleave factor or format option value of **0** has the same effect as not specifying the option at all.

For disks that support interleave factors, the acceptable range is usually **1** (no interleave) through $n-1$, where n is the number of sectors per track. Refer to the appropriate device operating manual for recommended values.

If a disk being initialized requires an interleave factor but none is specified, **mediainit** provides an appropriate, though not necessarily optimum default.

When a given device supports format options, the allowable range of interleave factors may be related to the specified format option. In such instances, **mediainit** cannot check the interleave factor if one is specified.

Notes

Most types of mass storage media must be initialized before they can be used. HP hard disks, flexible disks, and cartridge tapes require some form of initialization, but 9-track tapes do not. Initialization usually involves formatting the media, writing and reading test patterns, then sparing any defective blocks. Depending upon the media and device type, none, some, or all of the initialization process may have been performed at the factory. **mediainit** completes whatever steps are appropriate to prepare the media for error-free operation.

Most HP hard disks are formatted and exhaustively tested at the factory by use of a process more thorough but also more time-consuming than appropriate for **mediainit**. However, **mediainit** is still valuable for ensuring the integrity of the media after factory shipment, formatting with the correct interleave factor, and sparing any blocks which may have become defective since original factory testing was performed.

HP flexible disks are not usually formatted prior to shipment, so they must undergo the entire initialization process before they can be used.

When a tape is certified, it is thoroughly tested and defective blocks are spared. **mediainit** usually certifies a tape only if it has not been certified previously. If the tape has been previously certified and spared, **mediainit** usually reorganizes the tape's spare block table, retaining any previous spares, and optimizing their assignment for maximum performance under sequential access. Reorganizing the spare block table takes only a few seconds, whereas complete certification takes about a half-hour for 150-foot tapes, and over an hour for 600-foot tapes.

Reorganization of a tape's spare block table technically renders any existing data undefined, but the data is not usually destroyed by overwriting. To ensure that old tape data is destroyed, which is useful for security, complete tape re-certification can be forced with the **-r** option.

Some applications may require that a file system be placed on the media before use. **mediainit** does not create a file system; it only prepares media for writing and reading. If such a file system is required, other utilities such as **newfs**, **lifinit**, or **mkfs** must be invoked after running **mediainit** (see *newfs(1M)*, *lifinit(1)*, and *mkfs(1M)*).

RETURN VALUE

mediainit returns one of the following values:

- 0** Successful completion.
- 1** A device-related error occurred.
- 2** A syntax-related error was encountered.

ERRORS

Appropriate error messages are printed on standard error during execution of **mediainit**.

WARNINGS

For a device that contains multiple units on a single controller, each unit can be initialized independently from any other unit. It should be noted, however, that **mediainit** requires that there be no other processes accessing the device before initialization begins, regardless of which unit is being initialized. If there are accesses currently in progress, **mediainit** aborts.

Aborting **mediainit** is likely to leave the medium in a corrupt state, even if it was previously initialized. To recover, the initialization must be restarted.

During the initialization process, **open()** rejects all other accesses to the device being initialized, producing the error EACCES (see *open(2)*).

DEPENDENCIES**Series 800**

Partitioning of DDS tape media (**-p** option) is not supported.

AUTHOR

mediainit was developed by HP.

SEE ALSO

lifinit(1), *mkfs(1M)*, *newfs(1M)*.

NAME

merge - three-way file merge

SYNOPSIS

merge [-p] *file1 file2 file3*

DESCRIPTION

merge combines two files that are revisions of a single original file. The original file is *file2*, and the revised files are *file1* and *file3*. **merge** identifies all changes that lead from *file2* to *file3* and from *file2* to *file1*, then deposits the merged text into *file1*. If the **-p** option is used, the result goes to standard output instead of *file1*.

An overlap occurs if both *file1* and *file3* have changes in the same place. **merge** prints how many overlaps occurred, and includes both alternatives in the result. The alternatives are delimited as follows:

```
<<<<<<< file1
lines in file1
=====
lines in file3
>>>>>>> file3
```

If there are overlaps, edit the result in *file1* and delete one of the alternatives.

This command is particularly useful for revision control, especially if *file1* and *file3* are the ends of two branches that have *file2* as a common ancestor.

EXAMPLES

A typical use for **merge** is as follows:

1. To merge an RCS branch into the trunk, first check out the three different versions from RCS (see *co(1)*) and rename them for their revision numbers: 5.2, 5.11, and 5.2.3.3. File 5.2.3.3 is the end of an RCS branch that split off the trunk at file 5.2.
2. For this example, assume file 5.11 is the latest version on the trunk, and is also a revision of the "original" file, 5.2. Merge the branch into the trunk with the command:

```
merge 5.11 5.2 5.2.3.3
```
3. File 5.11 now contains all changes made on the branch and the trunk, and has markings in the file to show all overlapping changes.
4. Edit file 5.11 to correct the overlaps, then use the **ci** command to check the file back in (see *ci(1)*).

WARNINGS

merge uses the *ed(1)* system editor. Therefore, the file size limits of *ed(1)* apply to **merge**.

AUTHOR

merge was developed by Walter F. Tichy.

SEE ALSO

diff3(1), *diff(1)*, *rcsmerge(1)*, *co(1)*.

m

NAME

mesg - permit or deny messages to terminal

SYNOPSIS

```
mesg [[-] g] [[-] y] [[-] n]  
mesg
```

DESCRIPTION

The command form **mesg** [-] *n* forbids messages via **write** by revoking write permission to users without appropriate privilege on the user's terminal (see *write(1)*). The command form **mesg** [-] *g* reinstates permission so that only legitimate commands (such as *write(1)*) can be used by other users to send messages. **mesg** [-] *y* allows applications such as **write** or **talk** to send messages to the user's terminal (that is, without restrictions). **mesg** without any other argument reports the current state without changing it.

RETURN VALUE

mesg returns the following values:

- 0** Messages are receivable.
- 1** Messages are not receivable.
- 2** An error occurred.

EXTERNAL INFLUENCES**Environment Variables**

LC_MESSAGES determines the language in which messages are displayed.

If **LC_MESSAGES** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**.

If any internationalization variable contains an invalid setting, **mesg** behaves as if all internationalization variables are set to "C". See *environ(5)*.

FILES

/dev/tty*

SEE ALSO

write(1).

STANDARDS CONFORMANCE

mesg: SVID2, SVID3, XPG2, XPG3, XPG4

NAME

mkdir - make a directory

SYNOPSIS

mkdir [-p] [-m *mode*] *dirname* ...

DESCRIPTION

mkdir creates specified directories in mode 0777 (possibly altered by **umask** unless specified otherwise by a **-m mode** option (see *umask(1)*). Standard entries, **.** (for the directory itself) and **..** (for its parent) are created automatically. If *dirname* already exists, **mkdir** exits with a diagnostic message, and the directory is not changed.

Options

mkdir recognizes the following command-line options:

-m mode After creating the directory as specified, the file permissions are set to *mode*, which is a symbolic mode string as defined for **chmod** (see *chmod(1)*). The *umask(1)* has precedence over **-m**.

-p Intermediate directories are created as necessary. Otherwise, the full path prefix of *dirname* must already exist. **mkdir** requires write permission in the parent directory.

For each directory name in the pathname prefix of the *dirname* argument that is not the name of an existing directory, the specified directory is created using the current **umask** setting, except that the equivalent of **chmod u+wx** is done on each component to ensure that **mkdir** can create lower directories regardless of the setting of **umask**. Each directory name in the pathname prefix of the *dirname* argument that matches an existing directory is ignored without error. If an intermediate path component exists, but has permissions set to prevent writing or searching, **mkdir** fails with an error message.

If the **-m** option is used, the directory specified by *dirname* (excluding directories in the pathname prefix) is created with the permissions specified by *mode*.

Only **LINK_MAX** subdirectories can be created (see *limits(5)*).

Access Control Lists - JFS File Systems Only

If the parent directory has an access control list (ACL, see *acl(5)*), and that ACL contains default entries, an ACL is created for the new directory, and the parent directory's default entries are applied to the new directory's ACL, both as regular entries and as default entries.

EXTERNAL INFLUENCES**Environment Variables**

LANG provides a default value for the internationalization variables that are unset or null. If **LANG** is unset or null, the default value of "C" (see *lang(5)*) is used. If any of the internationalization variables contains an invalid setting, **mkdir** will behave as if all internationalization variables are set to "C". See *environ(5)*.

LC_ALL If set to a non-empty string value, overrides the values of all the other internationalization variables.

LC_CTYPE determines the interpretation of text as single and/or multi-byte characters, the classification of characters as printable, and the characters matched by character class expressions in regular expressions.

LC_MESSAGES determines the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output.

NLSPATH determines the location of message catalogues for the processing of **LC_MESSAGES**.

International Code Set Support

Single- and multi-byte character code sets are supported.

DIAGNOSTICS

mkdir returns exit code 0 if all directories were created successfully. Otherwise, it prints a diagnostic and returns non-zero.

mkdir returns exit code 0 if the **-p** option was specified, and all the specified directories now exist. If any of the intermediate directories do not have search or write permission (with the **-p** option), **mkdir** prints a diagnostic and returns non-zero.

EXAMPLES

Create directory **gem** beneath existing directory **raw** in the current directory:

```
mkdir raw/gem
```

Create directory path **raw/gem/diamond** underneath the current directory and set permissions on directory **diamond** to read-only for all users (**a=r**):

```
mkdir -p -m "a=r" raw/gem/diamond
```

which is equivalent to (see *chmod(1)*):

```
mkdir -p -m 444 raw/gem/diamond
```

If directories **raw** or **raw** and **gem** already exist, only the missing directories in the specified path are created.

SEE ALSO

rm(1), *setacl(1)*, *sh(1)*, *umask(1)*, *acly(5)*.

STANDARDS CONFORMANCE

mkdir: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2

NAME

mkfifo - make FIFO (named pipe) special files

SYNOPSIS

mkfifo [-p] [-m *mode*] *filename* ...

DESCRIPTION

mkfifo creates the FIFO special files named by its operand list. The operands are taken sequentially in the order specified and, if the user has write permission in the appropriate directory, the FIFO is created with permissions 0666 modified by the user's file mode creation mask (see *umask*(2)).

The specific actions performed are equivalent to calling

```
mkfifo( filename, 0666 )
```

for each filename in the operand list (see *mkfifo*(2)).

Options

mkfifo recognizes the following command-line options:

- m *mode* After creating the FIFO special file, set the permission bits of the new file to the specified *mode* value. The *mode* option argument is a symbolic mode string as defined in *chmod*(1).
(XPG4 Only.) In the symbolic mode strings, the operators + and - will be interpreted relative to an initial mode of a=rw.
- p Create any missing intermediate path name components.

EXTERNAL INFLUENCES**Environment Variables**

LANG determines the locale to use for the locale categories when both **LC_ALL** and the corresponding environment variable (beginning with **LC_**) do not specify a locale. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang*(5)) is used.

LC_ALL determines the locale to use to override any values for locale categories specified by the settings of **LANG** or any environment variables beginning with **LC_**.

LC_CTYPE determines the locale for the interpretation of sequences of bytes of text data as characters (e.g., single- versus multibyte characters in arguments).

If any internationalization variable contains an invalid setting, **mkfifo** behaves as if all internationalization variables are set to "C". See *environ*(5).

International Code Set Support

Single-byte character code sets are supported.

RETURN VALUE

mkfifo returns zero if invoked with at least one operand and if all FIFO special files were created successfully. Otherwise, it prints a diagnostic message and returns non-zero.

EXAMPLES

The following command creates a FIFO special file named **peacepipe** in the current directory:

```
mkfifo peacepipe
```

SEE ALSO

chmod(1), *umask*(1), *mknod*(1M), *mkfifo*(3C).

STANDARDS CONFORMANCE

mkfifo: XPG3, XPG4, POSIX.2

NAME

mkmf - make a makefile

SYNOPSIS

mkmf [-**acdeil**] [-**f** *makefile*] [-**F** *template*] [-**M** *language*] [*macroname=value ...*]

DESCRIPTION

The **mkmf** command creates a makefile that informs the **make** command how to construct and maintain programs and libraries (see *make(1)*). After gathering up all source code file names in the current working directory and inserting them into the makefile, **mkmf** scans source code files for included files and generates dependency information that is appended to the makefile. Source code files are identified by their file name suffixes. **mkmf** recognizes the following suffixes:

| | |
|-----------|----------------------|
| .c | C |
| .C | C++ |
| .f | FORTRAN |
| .h | Include files |
| .i | Pascal include files |
| .l | Lex or Lisp |
| .o | Object files |
| .p | Pascal |
| .r | Ratfor |
| .s | Assembler |
| .y | Yacc |

The **mkmf** command checks for an existing makefile before creating one. If no **-f** option is present, **mkmf** tries the makefiles **makefile** and **Makefile**, respectively.

After the makefile has been created, arbitrary changes can be made using a text editor. **mkmf** can also be used to re-edit the macro definitions in the makefile, regardless of changes that may have been made since it was created.

By default, **mkmf** creates a program makefile. To create a makefile that handles libraries, the **-l** option must be used.

Make Requests

Given a makefile created by **mkmf**, **make** recognizes the following requests:

| | |
|----------------|--|
| all | Compile and load a program or library. |
| clean | Remove all object and core files. |
| clobber | Remove all files that can be regenerated. |
| depend | Update included file dependencies in a makefile. |
| echo | List the names of the source code files on standard output. |
| extract | Extract all object files from the library and place them in the same directory as the source code files. The library is not altered. |
| index | Print an index of functions on standard output. |
| install | Compile and load the program or library and move it to its destination directory. |
| print | Print source code files on standard output. |
| tags | Create a tags file for the ex editor (see <i>ex(1)</i> and <i>ctags(1)</i>), for C, Pascal, and Fortran source code files. |
| update | Recompile only if there are source code files that are newer than the program or library, link and install the program or library. |

Several requests can be given simultaneously. For example, to (1) compile and link a program, (2) move the program to its destination directory, and (3) remove any unnecessary object files, use:

```
make install clean
```

Macro Definitions

mkmf understands the following macro definitions:

| | |
|-----------------------|--|
| CFLAGS | C compiler flags. After searching for included files in the directory currently being processed, mkmf searches in directories named in -I compiler options and then in the /usr/include directory. |
| COMPILESYSTYPE | Location of /usr/include . If the COMPILESYSTYPE macro or environment variable is defined, mkmf searches for included files in /\$COMPILESYSTYPE/usr/include instead of /usr/include . |
| CXXFLAGS | C++ compiler flags. After searching for included files in the directory currently being processed, mkmf searches in directories named in -I compiler options and then in the /usr/include/CC directory, followed by the /usr/include directory. |
| DEST | Directory where the program or library is to be installed. |
| EXTHDRS | List of included files external to the current directory. mkmf automatically updates this macro definition in the makefile if dependency information is being generated. |
| FFLAGS | Fortran compiler flags. After searching for included files in the directory currently being processed, mkmf searches in directories named in -I compiler options, then in the /usr/include directory. |
| HDRS | List of included files in the current directory. mkmf automatically updates this macro definition in the makefile. |
| INSTALL | Installation program name. |
| LD | Link editor name. |
| LDFLAGS | Link editor flags. |
| LIBRARY | Library name. This macro also implies the -l option. |
| LIBS | List of libraries needed by the link editor to resolve external references. |
| MAKEFILE | Makefile name. |
| OBJS | List of object files. mkmf automatically updates this macro definition in the makefile. |
| PROGRAM | Program name. |
| SRCS | List of source code files. mkmf automatically updates this macro definition in the makefile. |
| SUFFIX | List of additional file name suffixes for mkmf to know about. |
| SYSHDRS | List of included files found in the /usr/include directory hierarchy. mkmf automatically updates this macro definition in the makefile if dependency information is being generated. If SYSHDRS is omitted from the makefile, mkmf does not generate /usr/include dependencies. |

Both these and any other macro definitions already within the makefile can be replaced by definitions on the command line in the form *macroname=value*. For example, to change the C compiler flags and the program name, type the following line:

```
mkmf "CFLAGS=-I../include -O" PROGRAM=mkmf
```

Note that macro definitions such as **CFLAGS** with blanks in them must be enclosed in double quote (") marks.

Environment

The environment is read by **mkmf**. All variables are assumed to be macro definitions with the exception of **HDRS**, **EXTHDRS**, **SRCS**, and **OBJS**. Environment variables are processed after command line macro definitions and the macro definitions in a *makefile*. The **-e** option forces the environment to override the macro definitions in a *makefile*.

File Name Suffixes

mkmf can recognize additional file name suffixes, or ignore ones that it already recognizes, by specifying suffix descriptions in the **SUFFIX** macro definition. Each suffix description takes the form *.suffix:tI* where *t* is a character indicating the contents of the file (**s** = source file, **o** = object file, **h** = header file, **x** = executable file) and *I* is an optional character indicating the include syntax for header files (**C** = C syntax, **C++** = C syntax plus the addition of **/usr/include/CC** as a standard search directory, **F** = Fortran and

Ratfor syntax, P = Pascal syntax). The following list shows the default configuration for **mkmf**:

```
.c:sC      C
.C:sC++    C++
.f:sF      Fortran
.h:h       Include files
.i:h       Pascal include files
.l:sC      Lex or Lisp
.o:o       Object files
.p:sP      Pascal
.r:sF      Ratfor
.s:s       Assembler
.y:sC      Yacc
```

For example, to change the object file suffix to **.obj**, undefine the Pascal include file suffix, and prevent Fortran files from being scanned for included files, the **SUFFIX** macro definition could be:

```
SUFFIX = .obj:o .i: .f:s
```

Include Statement Syntax

The syntax of include statements for C, C++, Fortran, and Pascal source code are of the form:

C/C++:

```
#include "filename"
#include filename
```

where # must be the first character in the line.

Fortran:

```
$include ' filename '$
$INCLUDE ' filename '$
```

where \$ must be the first character in the line. Alternatively, the \$ can be omitted if the include statement starts in column 7. In either case the trailing \$ can be omitted.

Pascal:

```
$include ' filename '$
$INCLUDE ' filename '$
```

where \$ must be the first character in the line and the trailing \$ is optional.

User-defined Templates

If **mkmf** cannot find a makefile within the current directory, it normally uses one of the standard makefile templates, **C.p** or **C.l**, in **/usr/ccs/lib/mf** unless the user has alternative **C.p** or **C.l** template files in a directory **\$PROJECT/lib/mf** where **\$PROJECT** is the absolute path name of the directory assigned to the **PROJECT** environment variable.

Options

mkmf recognizes the following options:

- a Include source files beginning with a . in the makefile.
- c Suppress "creating makefile from ..." message.
- d Turn off scanning of source code for **include** files. Old dependency information is left untouched in the makefile.
- e Environment variables override macro definitions within *makefiles*.
- f *makefile* Specify an alternative *makefile* file name. The default file name is **Makefile**.
- i Prompt the user for the name of the program or library and the directory where it is to be installed. If a carriage-return is typed in response to each of these queries, **mkmf** assumes that the default program name is **a.out** or the default library name is **lib.a**, and the destination directory is the current directory.
- l Force the makefile to be a library makefile.

- F *template* Specify an alternative makefile template path name. The path name can be relative or absolute.
- M *language* Specify an alternative *language*-specific makefile template. The default language is C and the corresponding program and library makefile templates are C.p and C.l, respectively. **mkmf** looks for these templates in `/usr/ccs/lib/mf` or `$PROJECT/lib/mf`.

DIAGNOSTICS

Exit status 0 is normal. Exit status 1 indicates an error.

WARNINGS

The name of the makefile is included as a macro definition within the makefile and must be changed if the makefile is renamed.

Since executable files are dependent on libraries, standard library abbreviations must be expanded to full path names within the **LIBS** macro definition in the makefile.

Generated dependency information appears after a line in the makefile beginning with **###**. This line must not be removed, nor must any other information be inserted in the makefile below this line.

The name of a program or library must not conflict with any predefined target names in a makefile. It is especially important to avoid the the name **update** to prevent **make** from recursively executing itself an infinite number of times.

AUTHOR

mkmf was developed by the University of California, Berkeley.

FILES

| | |
|-----------------------------------|--|
| <code>/usr/ccs/lib/mf/C.p</code> | Standard program makefile template |
| <code>/usr/ccs/lib/mf/C.l</code> | Standard library makefile template |
| <code>\$PROJECT/lib/mf/C.p</code> | User-defined program makefile template |
| <code>\$PROJECT/lib/mf/C.l</code> | User-defined library makefile template |

SEE ALSO

`ar(1)`, `ctags(1)`, `ld(1)`, `make(1)`.

"Automatic Generation of Make Dependencies", *Software-Practice and Experience*, Walden, K., vol. 14, no. 6, pp. 575-585, June 1984.

NAME

mkmsgs - create message files for use by gettext()

SYNOPSIS

mkmsgs [-o] [-i *locale*] *textfile* *msgfile*

DESCRIPTION

The **mkmsgs** command takes as input a file of localized text strings and generates a message file that can be accessed by the *gettext*(3C) routine. *textfile* is the name of the file that contains the text strings. *msgfile* is the name of the output message file. **mkmsgs** appends the suffix *.cat* to the message file name. The combined length of the file name should be less than 14 bytes for short file name file system. The *msgfile* file should not contain a colon since it will confuse the formatting routines.

The *textfile* file contains the localized text strings. The text strings are separated by a newline character. The text strings are processed sequentially and copied to the *msgfile* message file. An empty line in the input results in a corresponding empty message written to the *msgfile* message file.

Options

The **mkmsgs** command supports the following options:

- o Overwrite the *msgfile* message file if it exists.
- i *locale* The *msgfile* message file is installed in the system-wide localization directory corresponding to the specified *locale*. Only a user with the appropriate privileges can create or overwrite the message file in that directory. The directory will be created if it does not exist.

EXTERNAL INFLUENCES**Environment Variables**

LC_CTYPE determines the interpretation of messages as single- and/or multibyte characters.

Messages are issued in **LANG** if it is set to a valid language and **LANG** messages are available. Otherwise "C" locale messages are issued.

If **LC_CTYPE** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang*(5)) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **mkmsgs** behaves as if all internationalization variables are set to "C". See *environ*(5).

International Code Set Support

Single- and multibyte character code sets are supported.

EXAMPLES

The following example shows the format of the input text strings:

```
global %s not found\n
\n\n<press return to continue>\n\n
\t%s, %d, %d, typ = %d, disp = '%s'\n
```

WARNINGS

mkmsgs is provided for SVID3 compatibility only. The user is encouraged to use the NLS mechanism developed by HP and the X/Open Company, Ltd.

SEE ALSO

gencat(1), gettext(3C), setlocale(3C).

STANDARDS COMPLIANCE

mkmsgs: SVID3

NAME

mkstr - extract error messages from C source into a file

SYNOPSIS

mkstr [-] *messagefile prefix file* ...

DESCRIPTION

mkstr examines a C program and creates a file containing error message strings used by the program. Programs with many error diagnostics can be made much smaller by referring to places in the file, and reduce system overhead in running the program.

mkstr processes each of the specified *files*, placing a revised version of each in a file whose name consists of the specified *prefix* concatenated in front of the original name. A typical usage of *mkstr* would be

```
mkstr mystrings xx *.c
```

This command would cause all the error messages from the C source files in the current directory to be placed in the file *mystrings* and revised copies of the source for these files to be placed in files whose names are prefixed with *xx*.

When processing the error messages in the source for transfer to the message file, **mkstr** searches for the string **error(** in the input file. Each time it is encountered, the C string starting after the leading quote is placed in the message file, followed by a null character and a new-line character. The null character terminates the message so that it can be easily used when retrieved, and the new-line character makes it possible to conveniently list the error message file (using **cat**, **more**, etc. — see *cat(1)* and *more(1)*) to review its contents.

The modified copy of the input file is identical to the original, except that each occurrence of any string that was moved to the error message file is replaced by an offset pointer usable by **lseek** to retrieve the message.

If the command line includes the optional **-**, extracted error messages are placed at the end of the specified message file (append) instead of overwriting it. This enables you to process individual files that are part of larger programs that have been previously processed by **mkstr** without reprocessing all the files.

All functions used by the original program whose names end in "error" that also can take a constant string as their first argument should be rewritten so that they search for the string in the error message file.

For example, a program based on the previous example usage would resemble the following:

```
#include <stdio.h>
#include <sys/types.h>
#include <fcntl.h>

char errfile[] = "mystrings" ;

error(offset, a2, a3, a4)
int offset, a1, a2, a3;
{
    char msg[256];
    static int fd = -1;

    if (fd < 0) {
        fd = open(errfile, O_RDONLY);
        if (fd < 0) {
            perror(errfile);
            exit(1);
        }
    }

    if (lseek(fd, (off_t) offset, 0) || read(fd, msg, 256) <= 0) {
        printf("? Can't find error message in %s:\n", errfile);
        perror(errfile);
        exit(1);
    }

    printf(msg, a1, a2, a3);
}
```

EXTERNAL INFLUENCES**Environment Variables**

LC_CTYPE determines the interpretation of comments and string literals as single- and/or multi-byte characters.

If **LC_CTYPE** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of "C" (see *lang(5)*) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **mkstr** behaves as if all internationalization variables are set to "C". See *environ(5)*.

International Code Set Support

Single- and multi-byte character code sets are supported within file names, comments, and string literals.

SEE ALSO

lseek(2), *perror(3C)*, *xstr(1)*.

BUGS

Strings in calls to functions whose names end in **error**, notably **perror()**, may be replaced with offsets by **mkstr**.

Calls to error functions whose first argument is not a string constant are left unmodified without warning.

NAME

mktemp - make a name for a temporary file

SYNOPSIS

mktemp [-c] [-d *directory_name*] [-p *prefix*]

DESCRIPTION

mktemp makes a name that is suitable for use as the pathname of a temporary file, and writes that name to the standard output. The name is chosen such that it does not duplicate the name of an existing file. If the **-c** option is specified, a zero-length file is created with the generated name.

The name generated by **mktemp** is the concatenation of a directory name, a slash (/), the value of the **LOGNAME** environment variable truncated to **{NAME_MAX}** - 6 characters, and the process ID of the invoking process.

The directory name is chosen as follows:

1. If the **-d** option is specified, *directory_name* is used.
2. Otherwise, if the **TMPDIR** environment variable is set and a string that would yield a unique name can be obtained by using the value of that variable as a directory name, this value is used.
3. Otherwise, if a string that would yield a unique name can be obtained using **/tmp** as the directory, **/tmp** is used.
4. Otherwise, **.** (current directory) is used.

If the **-p** option is specified, *prefix* is used instead of the value of the **LOGNAME** environment variable for name generation.

RETURN VALUE

mktemp returns zero on successful completion and non-zero if syntax, file access, or file creation errors were encountered or a unique pathname could not be generated.

SEE ALSO

mktemp(3C), umask(1).


m

NAME

mm, osdd - print documents formatted with the mm macros

SYNOPSIS

mm [*options*] [*files*]

osdd [*options*] [*files*]

DESCRIPTION

mm can be used to format and print documents using **nroff** and the **mm** text-formatting macro package (see *nroff(1)*). It has options to specify preprocessing by **tbl** and/or **neqn**, (see *tbl(1)* and *neqn(1)*), and postprocessing by various terminal-oriented output filters. The proper pipelines and the required arguments and flags for **nroff** and **mm** are generated, depending on the options selected.

osdd is equivalent to the command **mm -mosd**.

Options

mm recognizes the following *options* and command-line arguments. Any other arguments or options (such as **-rC3**) are passed to **nroff** or to **mm**, as appropriate. Such options can occur in any order, but they must appear before the *files* arguments. If no arguments are given, **mm** prints a list of its options.

- Tterm** Specifies the type of output terminal; for a list of recognized values for *term*, type **help term2**. If this option is *not* used, **mm** uses the value of the shell variable **\$TERM** from the environment (see *profile(4)* and *environ(5)*) as the value of *term* if **\$TERM** is set; otherwise, **mm** uses **450** as the value of *term*. If several terminal types are specified, the last one is used.
- 12** Indicates that the document is to be produced in 12-pitch. Can be used when **\$TERM** is set to one of **300**, **300s**, **450**, and **1620**. (The pitch switch on the DASI 300 and 300s terminals must be manually set to **12** if this option is used.)
- c** Causes **mm** to invoke *col(1)*; note that *col(1)* is invoked automatically by **mm** unless *term* is one of **300**, **300s**, **450**, **37**, **4000a**, **382**, **4014**, **tek**, **1620**, and **x**.
- e** Causes **mm** to invoke **neqn**.
- t** Causes **mm** to invoke **tbl**.
- E** Invokes the **-e** option of **nroff**.

DIAGNOSTICS

mm sends the message **mm: no input file** if none of the arguments is a readable file and **mm** is not used as a filter.

EXAMPLES

Assuming that the shell variable **\$TERM** is set in the environment to **450**, the two command lines below are equivalent:

```
mm -t -rC3 -12 ghh*
tbl ghh* | nroff -cm -T450-12 -h -rC3
```

mm reads the standard input when **-** is specified instead of any file names (mentioning other files along with **-** leads to disaster). This option allows **mm** to be used as a filter, as in this example:

```
cat dws | mm -
```

Hints

- mm** invokes **nroff** with the **-h** option. With this option, **nroff** assumes that the terminal has tabs set every 8 character positions.
- Use the **-olist** option of **nroff** to specify ranges of pages to be output. Note, however, that **mm**, if invoked with one or more of the **-e**, **-t**, and **-** options, *together* with the **-olist** option of **nroff** may cause a harmless “broken pipe” diagnostic if the last page of the document is not specified in *list*.
- If you use the **-s** option of **nroff** (to stop between pages of output), use line-feed (rather than return or new-line) to restart the output. The **-s** option of **nroff** does not work with the **-c** option of **mm**, or if **mm** automatically invokes **col** (see **-c** option above and *col(1)*).

- If you specify an incorrect output terminal type, **mm** produces (often subtle) unpredictable results. However, if you are redirecting output into a file, use the **-T37** option, then use the appropriate terminal filter when actually printing the formatted file.

SEE ALSO

col(1), **env(1)**, **nroff(1)**, **tbl(1)**, **profile(4)**, **term(4)**, **mm(5)**.

mm section in *Text Formatting: User's Guide*.

NAME

model - print hardware model information

SYNOPSIS

model

DESCRIPTION

model prints the machine hardware model. Two or more fields may be displayed: computer, model number, and sometimes the clock or an additional model number.

Its output is similar to that of **uname -m**. However, it is recommended that the **model** command or the **getconf** command be used to obtain the model name since future model names may not be compatible with **uname**.

EXAMPLES

Here are examples of what the **model** command displays.

The **model** output below indicates an HP 9000 Model 715 with a 50 MHz clock.

9000/715/50

The **model** output below indicates an HP 9000 879 K-Class model K260.

9000/879/K260

The **model** output below indicates an HP 9000 871 D-Class model D370.

9000/871/D370

The **model** output below indicates an HP 9000 Model 720.

9000/720

SEE ALSO

getconf(1), **uname(1)**.

m

NAME

more, page - file perusal filter for crt viewing

SYNOPSIS

```
more [-n] [-cdefisuvz] [-n number] [-p command] [-t tagstring] [-x tabs] [-W option]
[+linenumber] [+ / pattern] [name ...]
```

```
page [-n] [-cdefisuvz] [-n number] [-p command] [-t tagstring] [-x tabs] [-W option]
[+linenumber] [+ / pattern] [name ...]
```

REMARKS:

pg is preferred in some standards and has some added functionality, but does not support character highlighting (see *pg(1)*).

DESCRIPTION

more is a filter for examining continuous text, one screenful at a time, on a soft-copy terminal. It is quite similar to **pg**, and is retained primarily for backward compatibility. **more** normally pauses after each screenful, printing the filename at the bottom of the screen. To display one more line, press **<Return>**. To display another screenful press **<Space>**. Other possibilities are described later.

more and **page** differ only slightly. **more** scrolls the screen upward as it prints the next page. **page** clears the screen and prints a new screenful of text when it prints a new page. Both provide one line of overlap between screenfuls.

name can be a filename or **-**, specifying standard input. **more** processes file arguments in the order given.

more supports the Basic Regular Expression syntax (see *regex(5)*).

more recognizes the following command line options:

- n number** Set the number of lines in the display window to *number*, a positive decimal integer. The default is one line less than the the number of lines displayed by the terminal; on a screen that displays 24 lines, the default is 23. The **-n** flag overrides any values obtained from the environment.
- n** Same as **-n number** except that the number of lines is set to *n*.
- c** Draw each page by beginning at the top of the screen, and erase each line just before drawing on it. This avoids scrolling the screen, making it easier to read while **more** is writing. This option is ignored if the terminal has no clear-to-end-of-line capability.
- d** Prompt user with the message **Press space to continue, q to quit, h for help** at the end of each screenful. This is useful if **more** is being used as a filter in some setting, such as a training class, where many users might be unsophisticated.
- e** Exit immediately after writing the last line of the last file in the argument list
- f** Count logical lines, rather than screen lines. That is, long lines are not folded. This option is recommended if *nroff* output is being piped through *ul*, since the latter can generate escape sequences. These escape sequences contain characters that would ordinarily occupy screen positions, but which do not print when sent to the terminal as part of an escape sequence. Thus **more** might assume lines are longer than they really are, and fold lines erroneously.
- i** Perform pattern matching in searches without regard to case.
- s** Squeeze multiple blank lines from the output, producing only one blank line. Especially helpful when viewing *nroff* output, this option maximizes the useful information present on the screen.
- u** Normally, **more** handles underlining and bold such as produced by *nroff* in a manner appropriate to the particular terminal: if the terminal supports underlining or has a highlighting (usually inverse-video) mode, **more** outputs appropriate escape sequences to enable underlining, else highlighting mode, for underlined information in the source file. If the terminal supports highlighting, **more** uses that mode information that should be printed in boldface type. The **-u** option suppresses this processing, as do the "ul" and "os" terminfo flags.

- v** Do not display nonprinting characters graphically; by default, all non-ASCII and control characters (except **<Tab>**, **<Backspace>**, and **<Return>**) are displayed visibly in the form **^X** for **<Ctrl-x>**, or **M-x** for non-ASCII character **x**.
 - z** Same as not specifying **-v**, with the exception of displaying **<Backspace>** as **^H**, **<Return>** as **^M**, and **<Tab>** as **^I**.
 - p *command*** Execute the **more** command initially in the *command* argument for each file examined. If the command is a positioning command, such as a line number or a regular expression search, sets the current position to represent the final results of the command, without writing any intermediate lines of the file. If the positioning command is unsuccessful, the first line in the file is the current position.
 - t *tagstring*** Write the screenful of the file containing the tag named by the *tagstring* argument. The specified tag appears in the current position. If both **-p** and **-t** options are specified, **more** processes **-t** first; that is, the file containing the *tagstring* is selected by **-t** and then the command is executed.
 - x *tabs*** Set the tabstops every *tabs* position. The default value for the *tabs* argument is 8.
 - W *option*** Provides optional extensions to the **more** command. Currently, the following two options are supported:
 - notite** Prevents **more** from sending the terminal initialization string before displaying the file. This argument also prevents **more** from sending the terminal de-initialization string before exiting.
 - tite** Causes **more** to send the initialization and de-initialization strings. This is the default.
 - +*linenumber*** Start listing such that the current position is set to *linenumber*.
 - +/*pattern*** Start listing such that the current position is set to two lines above the line matching the regular expression *pattern*.
- Note: Unlike editors, this construct should NOT end with a **/**. If it does, the trailing slash is taken as character in the search pattern.

The number of lines available per screen is determined by the **-n** option, if present or by examining values in the environment. The actual number of lines written is one less than this number, as the last line of the screen is used to write a user prompt and user input.

The number of columns available per line is determined by examining values in the environment. **more** writes lines containing more characters than would fit into this number of columns by breaking the line into one more logical lines where each of these lines but the last contains the number of characters needed to fill the columns. The logical lines are written independently of each other; that is, commands affecting a single line affect them separately.

While determining the number of lines and the number of columns, if the methods described above do not yield any number then **more** uses terminfo descriptor files (see *term(4)*). If this also fails then the number of lines is set to 24 and the number of columns to 80.

When standard output is a terminal and **-u** is not specified, **more** treats backspace characters and carriage-return characters specially.

- A character, followed first by a backspace character, then by an underscore (**_**), causes that character to be written as underlined text, if the terminal supports that. An underscore, followed first by a backspace character, then any character, also causes that character to be written as underlined text, if the terminal supports that.
- A backspace character that appears between two identical printable characters causes the first of those two characters to be written as emboldened text, if the terminal type supports that, and the second to be discarded. Immediately subsequent occurrences of backspaces/character pairs for that same character is also discarded.
- Other backspace character sequences is written directly to the terminal, which generally causes the character preceding the backspace character to be suppressed in the display.
- A carriage-return character at the end of a line is ignored, rather than being written as a control character.

If the standard output is not a terminal device, **more** always exits when it reaches end-of-file on the last file in its argument list. Otherwise, for all files but the last, **more** prompts, with an indication that it has reached the end of file, along with the name of the next file. For the last file specified, or for the standard input if no file is specified, **more** prompts, indicating end-of-file, and accept additional commands. If the next command specifies forward scrolling, **more** will exit. If the **-e** option is specified, **more** will exit immediately after writing the last line of the last file.

more uses the environment variable **MORE** to preset any flags desired. The **MORE** variable thus sets a string containing flags and arguments, preceded with hyphens and blank-character-separated as on the command line. Any command-line flags or arguments are processed after those in the **MORE** variable, as if the command line were as follows:

```
more $MORE flags arguments
```

For example, to view files using the **-c** mode of operation, the shell command sequence

```
MORE='-c' ; export MORE
```

or the **csh** command

```
setenv MORE -c
```

causes all invocations of **more**, including invocations by programs such as *man* and *msgs*, to use this mode. The command sequence that sets up the **MORE** environment variable is usually placed in the *.profile* or *.cshrc* file.

In the following descriptions, the *current position* refers to two things:

- the position of the current line on the screen
- the line number (in the file) of the current line on the screen

The line on the screen corresponding to the current position is the third line on the screen. If this is not possible (there are fewer than three lines to display or this is the first page of the file, or it is the last page of the file), then the current position is either the first or last line on the screen.

Other sequences that can be typed when **more** pauses, and their effects, are as follows (*i* is an optional integer argument, defaulting to 1):

| | |
|-------------------|--|
| <i>i</i> <Return> | |
| <i>i</i> j | |
| <i>i</i> <Ctrl-e> | |
| <i>i</i> <Space> | Scroll forward <i>i</i> lines. The default <i>i</i> for <Space> is one screenful; for <i>j</i> and <Return> it is one line. The entire <i>i</i> lines are written, even if <i>i</i> is more than the screen size. At end-of-file, <Return> causes more to continue with the next file in the list, or exits if the current file is the last file in the list. |
| <i>i</i> d | |
| <i>i</i> <Ctrl-d> | Scroll forward <i>i</i> lines, with a default of one half of the screen size. If <i>i</i> is specified, it becomes the new default for subsequent <i>d</i> and <i>u</i> commands. |
| <i>i</i> u | |
| <i>i</i> <Ctrl-u> | Scrolls backward <i>i</i> lines, with a default of one half of the screen size. If <i>i</i> is specified, it becomes the new default for subsequent <i>d</i> and <i>u</i> commands. |
| <i>i</i> k | |
| <i>i</i> <Ctrl-y> | Scrolls backward <i>i</i> lines, with a default of one line. The entire <i>i</i> lines are written, even if <i>i</i> is more than the screen size. |
| <i>i</i> z | Display <i>i</i> more lines and sets the new window (screenful) size to <i>i</i> . |
| <i>i</i> g | Go to line <i>i</i> in the file, with a default of 1 (beginning of file). Scroll or rewrite the screen so that the line is at the current position. If <i>i</i> is not specified, then more displays the first screenful in the file. |
| <i>i</i> G | Go to line <i>i</i> in the file, with a default of the end of the file. If <i>i</i> is not specified, scrolls or rewrites screen so that the last line in the file is at the bottom of the screen. If <i>i</i> is specified, scrolls or rewrites the screen so that the line is at the current position. |
| <i>i</i> s | Skip forward <i>i</i> lines, with a default of 1, and write the next screenful beginning at that point. If <i>i</i> would cause the current position to be such that less than one screenful would be written, the last screenful in the file is written. |

| | |
|----------------------------------|--|
| <i>if</i> | |
| <i>i</i> <Ctrl-f> | Move forward <i>i</i> lines, with a default of one screenful. At end-of-file, more will continue with the next file in the list, or exit if the current file is the last file in the list. |
| <i>ib</i> | |
| <i>i</i> <Ctrl-b> | Move backward <i>i</i> lines, with a default of one screenful. If <i>i</i> is more than the screen size, only the final screenful will be written. |
| <i>q</i> | |
| <i>Q</i> | |
| <i>:q</i> | |
| <i>:Q</i> | |
| <i>ZZ</i> | Exit from more . |
| <i>=</i> | |
| <i>:f</i> | |
| <Ctrl-g> | Write the name of the file currently being examined, the number relative to the total number of files there are to examine, the current line number, the current byte number, and the total bytes to write and what percentage of the file precedes the current position. All of these items reference the first byte of the line after the last line written. |
| <i>v</i> | Invoke an editor to edit the current file being examined. The name of the editor is taken from the environment variable EDITOR , or default to vi . If EDITOR represents either vi or ex , the editor is invoked with options such that the current editor line is the physical line corresponding to the current position in more at the time of the invocation. When the editor exits, more resumes on the current file by rewriting the screen with the current line as the current position. |
| <i>h</i> | Display a description of all the more commands. |
| <i>i</i> / [!] <i>expression</i> | Search forward in the file for the <i>i</i> -th line containing the regular expression <i>expression</i> . The default value for <i>i</i> is 1. The search starts at the line following the current position. If the search is successful, the screen is modified so that the searched-for line is in the current position. The null regular expression (/<Return>) repeats the search using the previous regular expression. If the character ! is included, the lines for searching are those that do not contain <i>expression</i> . If there are less than <i>i</i> occurrences of <i>expression</i> , and the input is a file rather than a pipe, then the position in the file remains unchanged. The user's erase and kill characters can be used to edit the regular expression. Erasing back past the first column cancels the search command. |
| <i>i</i> ? [!] <i>expression</i> | Same as /, but searches backward in the file for the <i>i</i> th line containing the regular expression <i>expression</i> . Note: Unlike editors, the ? construct should NOT end with a /. If it does, the trailing slash is taken as a character in the search pattern. |
| <i>in</i> | Repeat the previous search for the <i>i</i> -th line (default 1) containing the last <i>expression</i> (or not containing the last <i>expression</i> , if the previous search was /! or ?!). |
| <i>n</i> | Repeat the search for the opposite direction of the previous search for the <i>i</i> -th line (default 1) containing the last <i>expression</i> |
| <i>''</i> | (2 apostrophes) Return to the position from which the last large movement command was executed ("large movement" is defined as any movement of more than a screenful of lines). If no such movements have been made, return to the beginning of the file. |
| <i>!command</i> | Invoke a shell with <i>command</i> . The characters % and ! in <i>command</i> are replaced with the current file name and the previous shell command, respectively. If there is no current file name, % is not expanded. The sequences \% and \! are replaced by % and ! respectively. |

| | |
|------------------------------|--|
| :e [<i>file</i>] | |
| E [<i>file</i>] | Examine a new file. If the <i>file</i> argument is not specified, the "current" file (see the :n and :p commands) from the list of files in the command line is re-examined. The filename is subjected to the process of shell word expansions. If <i>file</i> is a # (number sign) character, the previously examined file is re-examined. |
| i:n | Examine the next file. If <i>i</i> is specified, examines the <i>i</i> -th next file specified in the command line. |
| i:p | Examine the previous file. If a number <i>i</i> is specified, examines the <i>i</i> -th previous file specified in the command line. |
| :t <i>tagstring</i> | Go to the supplied <i>tagstring</i> and scroll or rewrite the screen with that line in the current position. |
| m <i>letter</i> | Mark the current position with the specified letter, where <i>letter</i> represents the name of one of the lower-case letters of the portable character set. |
| ' <i>letter</i> | Return to the position that was previously marked with the specified <i>letter</i> , making that line the current position. |
| r | |
| <Ctrl-l> | Refresh the screen. |
| R | Refresh the screen, discarding any buffered input. |
| . | Dot. Repeat the previous command. |
| ^\ <Ctrl-c> | Halt a partial display of text. more stops sending output, and displays the usual prompt. Unfortunately, some output is lost as a result. |

The commands take effect immediately; i.e., it is not necessary to press **<Return>**. Up to the time when the command character itself is given, the line-kill character can be used to cancel the numerical argument being formed.

If the standard output is not a teletype, **more** is equivalent to **cat(1)**.

more supports the **SIGWINCH** signal, and redraws the screen in response to window size changes.

EXTERNAL INFLUENCES

Environment Variables

| | |
|--------------------|---|
| COLUMNS | Overrides the system-selected horizontal screen size. |
| EDITOR | Used by the v command to select an editor. |
| LANG | Provides a default value for the internationalization variables that are unset or null. If LANG is unset or null, the default value of "C" (see <i>lang(5)</i>) is used. If any of the internationalization variables contains an invalid setting, more will behave as if all internationalization variables are set to "C". See <i>environ(5)</i> . |
| LC_ALL | If set to a non-empty string value, overrides the values of all the other internationalization variables. |
| LC_CTYPE | Determines the interpretation of text as single and/or multi-byte characters, the classification of characters as printable, and the characters matched by character class expressions in regular expressions. |
| LC_MESSAGES | Determines the locale that should be used to affect the format and contents of diagnostic messages written to standard error and informative messages written to standard output. |
| NLSPATH | Determines the location of message catalogues for the processing of LC_MESSAGES . |
| LINES | Overrides the system-selected vertical screen size, used as the number of lines in a screenful. The -n option takes precedence over the LINES variable for determining the number of lines in a screenful. |
| MORE | Determines a string containing options, preceded with hyphens and blank-character-separated as on the command line. Any command-line options are processed after those in the MORE variable. The MORE variable takes precedence over the TERM and LINES variables for determining the number of lines in a screenful. |

TERM Determines the name of the terminal type.

International Code Set Support

Single- and multi-byte character code sets are supported.

APPLICATION USAGE

When the standard output is not a terminal, none of the filter-modification options is effective. This is based on historical practice. For example, a typical implementation of **man** pipes its output through **more -s** to squeeze excess white space for terminal users. When **man** is piped to **lp**, however, it is undesirable for this squeezing to happen.

EXAMPLES

To view a simple file, use:

```
more filename
```

To preview *nroff* output, use a command resembling:

```
nroff -mm +2 doc.n | more -s
```

If the file contains tables, use:

```
tbl file | nroff -mm | col | more -s
```

To display file **stuff** in a fifteen line-window and convert multiple adjacent blank lines into a single blank line:

```
more -s -n 15 stuff
```

To examine each file with its last screenful:

```
more -p G file1 file2
```

To examine each file starting with line 100 in the current position (third line, so line 98 is the first line written):

```
more -p 100g file1 file2
```

To examine the file that contains the tagstring *tag* with line 30 in the current position:

```
more -t tag -p 30g
```

WARNINGS

Standard error, file descriptor 2, is normally used for input during interactive use and should not be redirected (see Input/Output section in the manpage of the shell in use).

FILES

*/usr/share/lib/terminfo/?/** compiled terminal capability data base

AUTHOR

more was developed by Mark Nudleman, University of California, Berkeley, OSF, and HP.

SEE ALSO

csh(1), man(1), pg(1), sh(1), term(4), terminfo(4), environ(5), lang(5), regexp(5).

STANDARDS CONFORMANCE

more: XPG4

NAME

mpsched - control the processor or locality domain on which a specific process executes

SYNOPSIS

mpsched -h

mpsched -s

mpsched [-g] [-P *policy*] [-T *policy*] [-l *locality-domain-id*] [-c *spu*] *command*

mpsched [-qu] [-P *policy*] [-l *locality-domain-id*] [-c *spu*] {-p *pid*}...

DESCRIPTION

mpsched controls the processor (*spu*), or locality domain (*locality-domain-id*) on which a process executes. It can do this by binding a process to a particular processor or locality domain, or by setting the launch policy for the process.

The command may be invoked in four manners.

- With -h, it prints a help message.
- With -s, it returns the hardware configuration of the system. This includes information about the number of locality domains and processors active in the system.
- With a *command* and its *arguments*, it applies the binding or launch policy to this command.
- With -p, it applies the binding or launch policy to the specified *pid*.

Options

The command-line options are:

-c *spu*

Bind the specified processes to the *spu* listed. This will ensure that the processes always run on the indicated processor. This option may be used with the -P, -T, and -p options.

-g Enable gang scheduling on th No other options may be used with -g.

-h Print a help message.

-l *locality-domain-id*

Bind the specified processes to the *locality-domain* listed. This will ensure that the processes always run on the indicated domain. This option may be used with the -P, -T, and -p options.

-p *pid*

Specify process ID, *pid*. To use the *pid* option, the caller must be a member of a group having PRIV_MPCTL access, be superuser, or have the same effective user ID as the *pid*. Specifying a command instead of the -p option does not require special privileges. Multiple -p options may be specified per command line, although each -p option can take only a single process ID.

-q Query the system regarding process bindings. This will return information about whether processes are bound to a processor or locality domain. It will also report on the thread and process launch policies for the processes. If this option is used in conjunction with -p then only those processes specified are queried. If this option is specified alone, then the status of all processes on the system that differ from the default settings are displayed.

-s Print the system hardware configuration. No other options may be specified.

-u Unbind the processes from any processor or locality domain bindings that may be present. This option can be used only with -p and no other options may be specified.

-P *policy*

Apply the specified *policy* to the processes. Launch policies affect the locality domain on which a process is spawned. This option may be used with the -T, -p, -c, and -l options. *policy* is one of the following values.

RR Round robin launch policy. Under this policy, successive child processes are launched in a round robin fashion across the other locality domains in the system relative to creating thread.

LL Least loaded launch policy. Under this policy, child processes are launched on the least loaded node in the system at the time of creation.

FILL Fill first launch policy. Under this policy, successive processes are launched on the same locality domain as their parent until one has been launched on each processor in the locality domain. At that point, new processes are created on the next locality domain.

PACKED Packed launch. Under this policy, successive processes are launched on the same locality domain as their parent. A different domain is never selected.

NONE No special policy. The default HP-UX launch policy is used.

-T *policy*

Apply the specified *policy* to the threads of the process. The scheduling policies are the same as for the **-P** option except that they apply to newly created threads instead of processes. Also, thread policies may only be specified on commands launched from the command line of **mpsched**. The option may be used with the **-P**, **-p**, **-l**, and **-c** options.

Operands

The command-line operands are:

command

A command including its arguments.

RETURN VALUE

mpsched returns exit status 0 if *command* is successfully scheduled or -1 if it fails.

EXAMPLES

Execute the **a.out** file on processor 2:

```
mpsched -c 2 a.out
```

Set the process launch policy for the existing process with pid 24217 to round robin:

```
mpsched -P RR -p 24217
```

Bind the processes with pids 1247 and 1842 to processor 4:

```
mpsched -c 4 -p 1247 -p 1842
```

AUTHOR

mpsched was developed by HP.

SEE ALSO

getprivgrp(1), setprivgrp(1M), fork(2), getprivgrp(2), mpctl(2), privgrp(4).

NAME

mt - magnetic tape manipulating program

SYNOPSIS

mt [-f *tapename*] *command* [*count*]

Obsolescent

mt [-t *tapename*] *command* [*count*]

DESCRIPTION

mt is used to give commands to the tape drive. If *tapename* is not specified, the environment variable **TAPE** is used; if **TAPE** is not defined, the default drive is used.

mt winds the tape in the requested direction (forward or backward), stopping after the specified *count* EOF marks or records are passed. If *count* is not specified, one is assumed. Each EOF mark counts as one record. When winding backwards, the tape always stops at the BOT marker, regardless of the number remaining in *count*.

mt accepts the following *commands*:

| | |
|---------------|---|
| eof | Write <i>count</i> EOF marks. |
| fsf | Forward space <i>count</i> files. |
| fsr | Forward space <i>count</i> records. |
| bsf | Backward space <i>count</i> files. |
| bsr | Backward space <i>count</i> records. |
| rew | Rewind tape. |
| offl | Rewind tape and go offline. |
| eod | Seek to end of data (DDS and QIC drives only). |
| smk | Write <i>count</i> setmarks (DDS drives only). |
| fss | Forward space <i>count</i> setmarks (DDS drives only). |
| bss | Backward space <i>count</i> setmarks (DDS drives only). |
| status | Print status information about the tape drive. |
| res | Reserve tape drive for sole use by the host issuing the mt command (stape driver only). |
| rel | Release tape drive from sole use by the host issuing the mt command (stape driver only). |

Spacing operations (back or forward space file or record) leave the tape positioned past the object being spaced to in the direction of motion. That is, backspacing a file leaves the the tape positioned before the file mark, forward spacing a file leaves the tape positioned after the file mark. This is consistent with all classical usage on tapes.

WARNINGS

Only raw, no-rewind Berkeley-type devices should be specified. This type of device will not reposition the tape upon close. An example of such a device is **/dev/rmt/0mnb**. Please refer to *mt(7)* for more details.

It is possible to wind the tape beyond the EOT marker and off the end of the reel.

A reservation may only be cleared with a release by the host that issued the original reserve. In the event that the host that holds the reservation is no longer available, the **st** command may be used to reclaim the device by issuing a bus device reset. Refer to *st(1M)* for more details.

The reserve/release functionality can only be issued to drives using the *stape* driver.

EXAMPLES

Rewind the tape associated with the device file **/dev/rmt/0mnb**:

```
mt -f /dev/rmt/0mnb rew
```

FILES

/dev/rmt/* Raw magnetic tape interface

`/dev/rmt/0mn` Default tape interface

AUTHOR

`mt` was developed by the University of California, Berkeley.

SEE ALSO

`dd(1)`, `mt(7)`, `st(1M)`.

NAME

mv - move or rename files and directories

SYNOPSIS

```
mv [-f|-i] [-e extarg] file1 new-file
mv [-f|-i] [-e extarg] file1 [file2 ...] dest-directory
mv [-f|-i] [-e extarg] directory1 [directory2 ...] dest-directory
```

DESCRIPTION

The **mv** command moves:

- One file (*file1*) to a new or existing file (*new-file*).
- One or more files (*file1*, [*file2*, ...]) to an existing directory (*dest-directory*).
- One or more directory subtrees (*directory1*, [*directory2*, ...]) to a new or existing directory (*dest-directory*).

Moving *file1* to *new-file* is used to rename a file within a directory or to relocate a file within a file system or across different file systems. When the destination is a directory, one or more files are moved into that directory. If two or more files are moved, the destination must be a directory. When moving a single file to a new file, if *new-file* exists, its contents are destroyed.

If the access permissions of the destination *dest-directory* or existing destination file *new-file* forbid writing, **mv** asks permission to overwrite the file. This is done by printing the mode (see *chmod(2)* and Access Control Lists below), followed by the first letters of the words *yes* and *no* in the language of the current locale, prompting for a response, and reading one line from the standard input. If the response is affirmative and the action is permissible, the operation occurs; if not, the command proceeds to the next source file, if any.

If *file1* is a file and *new-file* is a link to another file with other links, the other links remain and *new-file* becomes a new file. If *file1* is a file with links or a link to a file, the existing file or link remains intact, but the name is changed to *new-file* which may or may not be in the directory where *file1* resided, depending on directory path names used in the **mv** command. The last access and modification times of the file or files being moved remain unchanged.

Options

mv recognizes the following options:

- | | |
|------------------|--|
| -f | Perform mv commands without prompting for permission. This option is assumed when the standard input is not a terminal. |
| -i | Causes mv to write a prompt to standard output before moving a file that would overwrite an existing file. If the response from the standard input is affirmative, the file is moved if permissions allow the move. |
| -e extarg | Specifies the handling of any extent attributes of the file(s) to be moved. <i>extarg</i> can be one of the following values: |
| warn | Issue a warning message if extent attributes cannot be preserved, but move the file anyway. |
| ignore | Do not preserve extent attributes. |
| force | Do not move the file if the extent attributes cannot be preserved. |
- If multiple source files are specified with a single target directory, **mv** will move the files that either do not have extent attributes or that have extent attributes that can be preserved. **mv** will not move the files if it cannot preserve their extent attributes.

Extent attributes cannot be preserved if the files are being moved to a file system that does not support extent attributes or if that file system has a different block size than the original. If **-e** is not specified, the default value for *extarg* is **warn**.

Access Control Lists (ACLs)

If optional ACL entries are associated with *new-file*, **mv** displays a plus sign (+) after the access mode when asking permission to overwrite the file.

If *new-file* is a new file, it inherits the access control list of *file1*, altered to reflect any difference in ownership between the two files (see *acl(5)* and *aclv(5)*). In JFS file systems, new files created by **mv** do not inherit their parent directory's default ACL entries (if any), but instead retain their original ACLs. When moving files from a JFS file system to an HFS file system or vice versa, optional ACL entries are lost.

EXTERNAL INFLUENCES

Environment Variables

LC_CTYPE determines the interpretation of text as single byte and/or multibyte characters.

LANG and **LC_CTYPE** determine the local language equivalent of y (for yes/no queries).

LANG determines the language in which messages are displayed.

If **LC_CTYPE** is not specified in the environment or is set to the empty string, the value of **LANG** is used as a default for each unspecified or empty variable. If **LANG** is not specified or is set to the empty string, a default of **C** (see *lang(5)*) is used instead of **LANG**. If any internationalization variable contains an invalid setting, **mv** behaves as if all internationalization variables are set to **C**. See *environ(5)*.

International Code Set Support

Single character and multibyte character code sets are supported.

EXAMPLES

Rename a file in the current directory:

```
mv old-filename new-filename
```

Rename a directory in the current directory:

```
mv old-dirname new-dirname
```

Rename a file in the current directory whose name starts with a nonprinting control character or a character that is special to the shell, such as **-** and ***** (extra care may be required depending on the situation):

```
mv ./bad-filename new-filename
mv ./?bad-filename new-filename
mv ./*bad-filename new-filename
```

Move directory **sourcedir** and its contents to a new location (**targetdir**) in the file system (upon completion, a subdirectory named **sourcedir** resides in directory **targetdir**):

```
mv sourcedir targetdir
```

Move all files and directories (including links) in the current directory to a new location underneath **targetdir**:

```
mv * targetdir
```

Move all files and directories (including links) in **sourcedir** to a new location underneath **targetdir** (**sourcedir** and **targetdir** are in separate directory paths):

```
mv sourcedir/* targetdir
```

WARNINGS

If *file1* and *new-file* exist on different file systems, **mv** copies the file and deletes the original. In this case the mover becomes the owner and any linking relationship with other files is lost. **mv** cannot carry hard links across file systems. If *file1* is a directory, **mv** copies the entire directory structure onto the destination file system and deletes the original.

mv cannot be used to perform the following operations:

- Rename either the current working directory or its parent directory using the **.** or **..** notation.
- Rename a directory to a new name identical to the name of a file contained in the same parent directory.

DEPENDENCIES

NFS

Access control lists of networked files are summarized (as returned in *st_mode* by *stat(2)*), but not copied to the new file. When using **mv** on such files, a **+** is not printed after the mode value when asking for permission to overwrite a file.

AUTHOR


mv was developed by AT&T, the University of California, Berkeley and HP.

SEE ALSO

cp(1), cpio(1), ln(1), rm(1), link(1M), lstat(2), readlink(2), stat(2), symlink(2), symlink(4), acl(5), aclv(5).

STANDARDS CONFORMANCE

mv: SVID2, SVID3, XPG2, XPG3, XPG4, POSIX.2



m