



ICON/UXB Operating System Reference Manual

Volume 1A

**ICON
INTERNATIONAL**

P.O. Box 340
Orem, Utah 84059
(801) 225-6888



OPERATING SYSTEM REFERENCE MANUAL

ICON/UXB

System Commands

Volume 1A

© 1988 Icon International, Inc.
All rights reserved worldwide.

Copyright © 1987 Icon International, Inc. All rights reserved. No part of this manual shall be reproduced, stored in a retrieval system, or transmitted by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission from Icon International, Inc. While every precaution has been taken in the preparation of this manual, Icon International assumes no responsibility for errors or omissions. Neither is any liability assumed for damages resulting from the use of the information contained herein.

Copyright 1979, 1980 Regents of the University of California. Permission to copy these documents or any portion thereof as necessary for licensed use of the software is granted to licensees of this software, provided this copyright notice and statement of permission are included.

The document "Writing Tools - The STYLE and DICTION Programs" is copyrighted 1979 by Bell Telephone Laboratories. Holders of a UNIX[®]/32V software license are permitted to copy this document, or any portion of it, as necessary for licensed use of the software, provided this copyright notice and statement of permission are included.

The document "The Programming Language EFL" is copyrighted 1979 by Bell Telephone Laboratories. EFL has been approved for general release, so that one may copy it subject only to the restriction of giving proper acknowledgement to Bell Telephone Laboratories.

This manual reflects system enhancements made at Berkeley and sponsored in part by NSF Grants MCS-7807291, MCS-8005144, and MCS-74-07644-A04; DOE Contract DE-AT03-76SF00034 and Project Agreement DE-AS03-79ER10358; and by Defense Advanced Research Projects Agency (DoD) ARPA Order No. 4031, Monitored by Naval Electronics Systems Command under Contract No. N00039-80-K-0649.

This manual was prepared by the Documentation Group of Icon International, Inc., P.O. Box 340, Orem, UT 84057-0340. A form for reader's comments has been provided at the back of this publication. comments are welcomed and may be sent to the above address.

Revision B

Order Number 172-022-001 (Manual Assembly)

Order Number 171-070-001 (Pages Only)

Printed in the U.S.A.

ICON is a registered trademark of Icon International, Inc.
UNIX is a registered trademark of AT&T.

Change Record Page

Manual Part No. 172-022-001

Date	Revision	Description	Pages Affected
Jan. 1987	A	Initial production release	All
Nov. 1987	B	Incorporate additions of commands included in Releases 2.16, 3.0, and 3.1 of the ICON/UXB Operating System and separate Volume 1 into two separate binders	Main cover page, titlepage, copyright page, Table of Contents, Permuted Index, Introduction, addition of and changes to the following manual pages: ci(1), co(1), ident(1), merge(1), rcs(1), rcsdiff(1), rcsmerge(1), rlog(1), tac(1), tftp(1c), and tmail(1)



TABLE OF CONTENTS

I. Permuted Index

Special Characters	xv
aardvark — apply a command to a set of arguments	xv
apply or remove an advisory lock on an open file — based on ex	xvi
basename — calendar	xvii
calendar: reminder service — check nroff/troff files	xviii
check out RCS revisions — commands by keyword lookup	xix
commands conditionally — conventional names for terminals	xx
conversion — current working directory	xxi
current working directory — description by its name	xxii
description file — display	xxiii
display call graph profile data — efficient way	xxiv
efl: Extended Fortran Language — exec, exece, environ	xxv
exec, exit, export, login, read, readonly, set,/ — fcntl: file control options	xxvi
fcvt, gcvl: output conversion — file creation mask	xxvii
file creation mode mask — files for the manuals	xxviii
files in reverse order — Fortran file	xxix
Fortran file into individual files — generate a dump	xxx
generate a fault — getw: get character or word from stream	xxxii
getwd: get current working directory pathname — host system	xxxiii
host system — inet_network, inet_ntoe, inet_makeaddr, inet_lnaof	xxxiiii
inet_ntoa, inet_makeaddr, inet_lnaof, inet_netof — intro to system maintenance	xxxv
i-numbers — len: tell about character objects	xxxvi
length — loc: return the address of an object	xxxvii
localtime, gmtime, asctime, timezone: convert date — mag tape	xxxviii
magnetic tape manipulating program — memory segment	xxxix
memory statistics — msgs: system messages and junk mail program	xl
mt: magnetic tape manipulating program — nohup: run a command at low priority	xli
nohup: run command immune to hangups — output conversion	xlii
output conversion — plot: graphics filters	xliii
plot: graphics interface — process	xliv
process — programs to implement shared strings	xlv
protocol — RCS file	xlvi
RCS file attributes — remote login server	xlvii
remote magtape protocol module — reverse order	xlviii
revisions — seekdir, rewinddir, closedir: directory operations	xliv
seekdir, rewinddir, closedir: — setenv: set variable in environment	lx
seteuid, setruid, setgid, setegid, setrgid: set — signal mask	li
signal messages — spellout: find spelling errors	lii
spline: interpolate smooth curve — string operations. strcat, strncat, strcmp	liii
strings. xstr: — tables	liv
tables for nroff or troff — text editor	lv
text editor — trekkie game	lvi
trewin, tskipf, tstate: f77 tape I/O — (unlink) files or directories	lvii
unlink: remove a directory entry — variable in environment	lviii
variables — without checking the disks	lix
word count — zone name rule file	lix

II. Introduction

INTRODUCTION 1
SECTION FORMAT 2
HOW TO GET STARTED 3
Logging In 3
Logging Out 4
HOW TO COMMUNICATE THROUGH YOUR TERMINAL 4
HOW TO RUN A PROGRAM — THE SHELLS 5
THE CURRENT DIRECTORY 6
PATH NAMES 6
WRITING AN ICON/UXB PROGRAM 7
TEXT PROCESSING 7
STATUS INQUIRIES 8
SURPRISES 8
OTHER MANUALS AND BOOKS ON THE UNIX® OPERATING SYSTEM 8

1. Commands and Application Programs

intro introduction to commands
adb debugger
addbib create or extend bibliographic database
apply apply a command to a set of arguments
apropos locate commands by keyword lookup
ar archive and library maintainer
as M68020 assembler
at execute commands at a later time
awk pattern scanning and processing language
basename strip filename affixes
bc arbitrary-precision arithmetic language
biff be notified if mail arrives and who it is from
binmail send or receive mail among users
cal print calendar
calendar reminder service
cat catenate and print
cb C program beautifier
cc C compiler
cd change working directory
checknr check nroff/troff files
chfn change finger entry
chgrp change group
chmod change mode
chsh change default login shell
ci check in RCS revisions
clear clear terminal screen
cmp compare two files
co check out RCS revisions
col filter reverse line feeds
colcrt filter nroff output for CRT previewing
colrm remove columns from a file
comm select or reject lines common to two sorted files
compact compress and uncompress files, and cat them

compress	compress and expand data
cp	copy
cpio	copy file archives in and out
crypt	encode/decode
csch	a shell (command interpreter) with C-like syntax
ctags	create a tags file
date	print and set the date
dbx	debugger
dc	desk calculator
dd	convert and copy a file
deroff	remove nroff, troff, tbl and eqn constructs
df	disk free
diction	print wordy sentences; thesaurus for diction
diff	differential file and directory comparator
diff3	3-way differential file comparison
dis	an mc68020 disassembler
dosc	connect to proc/286 system
du	summarize disk usage
echo	echo arguments
ed	text editor
efl	Extended Fortran Language
eqn	typeset mathematics
error	analyze and disperse compiler error messages
ex	text editor
expand	expand tabs to spaces, and vice versa
explain	explain, diction- print wordy sentences; thesaurus for diction
expr	evaluate arguments as an expression
eyacc	modified yacc allowing much improved error recovery
f77	Fortran 77 compiler
false	provide truth values
file	determine file type
find	find files
finger	user information lookup program
fnt	simple text formatter
fold	fold long lines for finite width output device
fp	Functional Programming language compiler/interpreter
fpr	print Fortran file
fpu	determine presence of the floating point coprocessor
from	who is my mail from?
fsplit	split a multi-routine Fortran file into individual files
ftp	file transfer program
gcore	get core images of running processes
gprof	display call graph profile data
graph	draw a graph
grep	search a file for a pattern
groups	show group memberships
head	give first few lines
hostid	set or print identifier of current host system
hostname	set or print name of current host system
ident	identify files
indent	indent and format C program source

Table of Contents

install	install binaries
iostat	report I/O statistics
ipcrm	remove a message queue, semaphore set or shared memory id
ipcs	report inter-process communication facilities status
join	relational database operator
jove	an interactive display-oriented text editor
jove_recover	jove_recover - recover JOVE buffers after a system/editor crash
kermit	kermit file transfer
kill	terminate a process with extreme prejudice
last	indicate last logins of users and teletypes
lastcomm	show last commands executed in reverse order
ld	link editor
learn	computer aided instruction about UNIX
leave	remind you when you have to leave
lex	generator of lexical analysis programs
lint	a C program verifier
lisp	lisp interpreter
liszt	compile a Franz Lisp program
ln	make links
lock	reserve a terminal
login	sign on
look	find lines in a sorted list
lookbib	build inverted index for a bibliography, find references in a bibliography
lorder	find ordering relation for an object library
lpq	spool queue examination program
lpr	off line print
lprm	remove jobs from the line printer spooling queue
ls	list contents of directory
lxref	lisp cross reference program
m4	macro processor
mail	send and receive mail
make	maintain program groups
man	find manual information by keywords; print out the manual
merge	three-way file merge
mesg	permit or deny messages
mkdir	make a directory
mkstr	create an error message file by massaging C source
more	file perusal filter for crt viewing
msgs	system messages and junk mail program
mt	magnetic tape manipulating program
mv	move or rename files
netstat	show network status
newaliases	rebuild the data base for the mail aliases file
nice	run a command at low priority (<i>sh</i> only)
nm	print name list
nroff	text formatting
od	octal, decimal, hex, ascii dump
pagesize	print system page size
passwd	change login password
pc	Pascal compiler
pdx	pascal debugger

pi	Pascal interpreter code translator
pix	Pascal interpreter and executor
plot	graphics filters
pmerge	pascal file merger
pr	print file
print	pr to the printer
printenv	print out the environment
prmail	print out mail in the post office
prof	display profile data
ps	process status
pti	phototypesetter interpreter
ptx	permuted index
pwd	working directory name
px	Pascal interpreter
pxp	Pascal execution profiler
pxref	Pascal cross-reference program
quota	display disc usage and limits
ranlib	convert archives to random libraries
ratfor	rational Fortran dialect
rcp	remote file copy
rcs	change RCS file attributes
rcsdiff	compare RCS revisions
rcsintro	introduction to RCS commands
rcsmerge	merge RCS revisions
rdist	remote file distribution program
refer	find and insert literature references in documents
reset	reset the teletype bits to a sensible state
rev	reverse lines of a file
rlog	print log messages and other information about RCS files
rlogin	remote login
rm	remove (unlink) files or directories
rmail	handle remote mail received via uucp
rmdir	remove (unlink) directories or files
roffbib	run off bibliographic database
rsh	remote shell
ruptime	show host status of local machines
rwho	who's logged in on local machines
scsstorcs	build RCS file from SCCS file
script	make typescript of terminal session
sed	stream editor
sendbug	mail a system bug report to 4bsd-bugs
sfdate	set the time/date of a file
sh	command language
size	size of an object file
sleep	suspend execution for an interval
soelim	eliminate .so's from nroff input
sort	sort or merge files
sortbib	sort bibliographic database
spell	find spelling errors
spline	interpolate smooth curve
split	split a file into pieces

Table of Contents

strings	find the printable strings in a object, or other binary, file
strip	remove symbols and relocation bits
struct	structure Fortran programs
stty	set terminal options
style	analyze surface characteristics of a document
su	substitute user id temporarily
sum	sum and count blocks in a file
symorder	rearrange name list
sysline	display system status on status line of a terminal
tabs	set terminal tabs
tac	concatenate and print files in reverse order
tail	deliver the last part of a file
talk	talk to another user
tar	tape archiver
tbl	format tables for nroff or troff
tc	phototypesetter simulator
tcopy	copy a mag tape
teachjove	TEACHJOVE - learn how to use the JOVE editor
tee	pipe fitting
telnet	user interface to the TELNET protocol
test	condition command
tftp	trivial file transfer program
time	time a command
tip	connect to a remote system
tk	paginator for the Tektronix 4014
tmail	print out mail messages, most recent first
touch	update date last modified of a file
tr	translate characters
trman	translate version 6 manual macros to version 7 macros
troff	text formatting and typesetting
true	provide truth values
tset	terminal dependent initialization
tsort	topological sort
tty	get terminal name
ul	do underlining
unifdef	remove ifdef'ed lines
uniq	report repeated lines in a file
units	conversion program
uptime	show how long system has been up
users	compact list of users who are on the system
uucp	unix to unix copy
uuencode	encode/decode a binary file for transmission via mail
uuseed	send a file to a remote host
uux	unix to unix command execution
vers	print version number of the kernel
vgrind	grind nice listings of programs
vi	screen oriented (visual) display editor based on ex
vmstat	report virtual memory statistics
w	who is on and what they are doing
wait	await completion of process
wall	write to all users

wc	word count
what	show what versions of object modules were used to construct a file
whatis	describe what a command is
whereis	locate source, binary, and or manual for program
which	locate a program file including aliases and paths (<i>csk</i> only)
who	who is on the system
whoami	print effective current user id
whodos	display information about dosc users
window	window environment
write	write to another user
xsend	secret mail
xstr	extract strings from C programs to implement shared strings
yacc	yet another compiler-compiler
yes	be repetitively affirmative

2. System Calls

intro	introduction to system calls and error numbers
accept	accept a connection on a socket
access	determine accessibility of file
acct	turn accounting on or off
bind	bind a name to a socket
brk	change data segment size
chdir	change current working directory
chmod	change mode of file
chown	change owner and group of a file
chroot	change root directory
close	delete a descriptor
connect	initiate a connection on a socket
creat	create a new file
dup	duplicate a descriptor
execve	execute a file
exit	terminate a process
fcntl	file control
flock	apply or remove an advisory lock on an open file
fork	create a new process
fsync	synchronize a file's in-core state with that on disk
getdtablesize	get descriptor table size
getgid	get group identity
getgroups	get group access list
gethostid	get/set unique identifier of current host
gethostname	get/set name of current host
getitimer	get/set value of interval timer
getpagesize	get system page size
getpeername	get name of connected peer
getpgrp	get process group
getpid	get process identification
getpriority	get/set program scheduling priority
getrlimit	control maximum system resource consumption
getrusage	get information about resource utilization
getsockname	get socket name
getsockopt	get and set options on sockets

gettimeofday	get/set date and time
getuid	get user identity
ioctl	control device
kill	send signal to a process
killpg	send signal to a process group
link	make a hard link to a file
listen	listen for connections on a socket
lseek	move read/write pointer
mkdir	make a directory file
mknod	make a special file
mount	mount or remove file system
msgctl	message control operations
msgget	get message queue
msgop	message operations
open	open a file for reading or writing, or create a new file
pipe	create an interprocess communication channel
profil	execution time profile
ptrace	process trace
quota	manipulate disk quotas
read	read input
readlink	read value of a symbolic link
reboot	reboot system or halt processor
recv	receive a message from a socket
rename	change the name of a file
rmdir	remove a directory file
select	synchronous i/o multiplexing
semctl	semaphore control operations
semget	get set of semaphores
semop	semaphore operations
send	send a message from a socket
setgroups	set group access list
setpgrp	set process group
setquota	enable/disable quotas on a file system
setregid	set real and effective group ID
setreuid	set real and effective user ID's
shmctl	shared memory control operations
shmget	get shared memory segment
shmop	shared memory operations
shutdown	shut down part of a full-duplex connection
sigblock	block signals
sigpause	atomically release blocked signals and wait for interrupt
sigsetmask	set current signal mask
sigstack	set and/or get signal stack context
sigvec	software signal facilities
socket	create an endpoint for communication
socketpair	create a pair of connected sockets
stat	get file status
swapon	specify a swap directory
symlink	make symbolic link to a file
sync	update super-block
syscall	indirect system call

Table of Contents

truncate	truncate a file to a specified length
umask	set file creation mode mask
unlink	remove directory entry
utimes	set file times
vfork	spawn new process in a virtual memory efficient way
vhangup	virtually "hangup" the current control terminal
wait	wait for process to terminate
write	write on a file

SECTION I

**ICON/UXB
OPERATING
SYSTEM
PERMUTED
INDEX**

PERMUTED INDEX

<p>lib2648: subroutines for the HP diff3: sendbug: mail a system bug report to</p> <p>abort: terminate</p> <p>abs: integer fabs, floor, ceil:</p> <p>accept:</p> <p>getgroups: get group initgroups: initialize group setgroups: set group access: determine access: determine ac: login sa, accton: system acct: execution pac: printer/plotter acct: turn</p> <p>sa, sin, cos, tan, asin, signal: change the fortune: print a random, hopefully interesting,</p> <p>adduser: procedure for swapon: specify inet_makeaddr, inet_lnaof, inet_netof: Internet loc: return the mailaddr: mail</p> <p>battlestar: a tropical flock: apply or remove an yes: be repetitively basename: strip filename learn: computer</p> <p>unalias: remove</p> <p>which: locate a program file including newaliases: rebuild the data base for the mail</p> <p>aliases: valloc: malloc, free, realloc, calloc, malloc, free, realloc, calloc, alloca: memory valloc: aligned memory eyacc: modified yacc limit: renice: else: lex: generator of lexical error: style: sigstack: set worms: rain: bcd: convert to apply:</p>	<p>@: arithmetic on shell variables. csh(1) 2648 graphics terminal. lib2648(3X) 3-way differential file comparison. diff3(1) 4bsd-bugs. sendbug(1) aardvark: yet another exploration game. aardvark(6) abort: generate a fault. abort(3) abort: terminate abruptly with memory image. abort(3F) abruptly with memory image. abort(3F) abs: integer absolute value. abs(3) absolute value. abs(3) absolute value, floor, ceiling functions. floor(3M) ac: login accounting. ac(8) accept a connection on a socket. accept(2) accept: accept a connection on a socket. accept(2) access: determine accessibility of a file. access(3F) access: determine accessibility of file. access(2) access list. getgroups(2) access list. initgroups(3X) access list. setgroups(2) accessability of a file. access(3F) accessability of file. access(2) accounting. ac(8) accounting. sa(8) accounting file. acct(5) accounting information. pac(8) accounting on or off. acct(2) acct: execution accounting file. acct(5) acct: turn accounting on or off. acct(2) accton: system accounting. sa(8) acos, atan, atan2: trigonometric functions. sin(3M) action for a signal. signal(3F) adage. fortune(6) adb: debugger. adb(1) addbib: create or extend bibliographic database. addbib(1) adding new users. adduser(8) additional device for paging and swapping. swapon(8) address manipulation routines. /inet_ntoa, inet(3n) address of an object. loc(3F) addressing description. mailaddr(7) adduser: procedure for adding new users. adduser(8) adventure: an exploration game. adventure(6) adventure game. battlestar(6) advisory lock on an open file. flock(2) affirmative. yes(1) affixes. basename(1) aided instruction about UNIX. learn(1) alarm: execute a subroutine after a specified time. alarm(3F) alarm: schedule signal after specified time. alarm(3C) alias: shell macros. csh(1) aliases. csh(1) aliases: aliases file for sendmail. aliases(5) aliases and paths (csh only). which(1) aliases file. newaliases(1) aliases file for sendmail. aliases(5) aligned memory allocator. valloc(3) alloca: memory allocator. malloc(3) allocator. malloc(3) allocator. valloc(3) allowing much improved error recovery. eyacc(1) alter per-process resource limitations. csh(1) alter priority of running processes. renice(8) alternative commands. csh(1) analysis programs. lex(1) analyze and disperse compiler error messages. error(1) analyze surface characteristics of a document. style(1) and/or get signal stack context. sigstack(2) animate worms on a display terminal. worms(6) animated raindrops display. rain(6) antique media. bcd(6) apply a command to a set of arguments. apply(1) apply: apply a command to a set of arguments. apply(1)</p>
---	--

	flock:	apply or remove an advisory lock on an open file.	flock(2)
		apropos: locate commands by keyword lookup.	apropos(1)
		ar: archive and library maintainer.	ar(1)
		ar: archive (library) file format.	ar(5)
	number: convert	Arabic numerals to English.	number(6)
	bc:	arbitrary-precision arithmetic language.	bc(1)
graphics/	plot:	arc, move, cont, point, linemod, space, closepl:	plot(3X)
		archive and library maintainer.	ar(1)
	tar: tape	archive file format.	tar(5)
		archive (library) file format.	ar(5)
	tar: tape	archiver.	tar(1)
	cpio: copy file	archives in and out.	cpio(1)
	ranlib: convert	archives to random libraries.	ranlib(1)
	glob: filename expand	argument list.	csh(1)
	shift: manipulate	argument list.	csh(1)
	varargs: variable	argument list.	varargs(3)
apply: apply a command to a set of		arguments.	apply(1)
	echo: echo	arguments.	csh(1)
	echo: echo	arguments.	echo(1)
getarg, iarg: return command line		arguments.	getarg(3F)
	expr: evaluate	arguments as an expression.	expr(1)
	traper: trap	arithmetic errors.	traper(3F)
	bc: arbitrary-precision	arithmetic language.	bc(1)
	@:	arithmetic on shell variables.	csh(1)
		arithmetic: provide drill in number facts.	arithmetic(6)
	biff: be notified if mail	arrives and who it is from.	biff(1)
	expr: evaluate arguments	as an expression.	expr(1)
		as: M68020 assembler.	as(1)
	slattach: attach serial lines	as network interfaces.	slattach(8C)
gmtime, asctime, timezone: convert date and time to		ASCII. ctime, localtime,	ctime(3)
	ascii: map of	ASCII character set.	ascii(7)
	od: octal, decimal, hex,	ascii dump.	od(1)
		ascii: map of ASCII character set.	ascii(7)
	fdate: return date and time in an	ASCII string.	fdate(3F)
	atof, atoi, atol: convert	ASCII to numbers.	atof(3)
	ctime, localtime, gmtime,	asctime, timezone: convert date and time to ASCII.	ctime(3)
	sin, cos, tan,	asin, acos, atan, atan2: trigonometric functions.	sin(3M)
	as: M68020	assembler.	as(1)
	a.out:	assembler and link editor output.	a.out(5)
		assert: program verification.	assert(3X)
	setbuf, setbuffer, setlinebuf:	assign buffering to a stream.	setbuf(3S)
shutdown: close down the system		at a given time.	shutdown(8)
	at: execute commands	at a later time.	at(1)
		at: execute commands at a later time.	at(1)
	nice, nohup: run a command	at low priority (<i>sh</i> only).	nice(1)
	sin, cos, tan, asin, acos,	atan, atan2: trigonometric functions.	sin(3M)
	sin, cos, tan, asin, acos, atan,	atan2: trigonometric functions.	sin(3M)
		atof, atoi, atol: convert ASCII to numbers.	atof(3)
	atof,	atoi, atol: convert ASCII to numbers.	atof(3)
	atoi,	atol: convert ASCII to numbers.	atof(3)
	interrupt. sigpause:	atomically release blocked signals and wait for	sigpause(2)
	slattach:	attach serial lines as network interfaces.	slattach(8C)
	rcs: change RCS file	attributes.	rcs(1)
bugfler: file bug reports in folders		automatically.	bugfler(8)
	rc: command script for	auto-reboot and daemons.	rc(8)
	wait:	await completion of process.	wait(1)
		awk: pattern scanning and processing language.	awk(1)
	bg: place job in	backgammon: the game.	backgammon(6)
	wait: wait for	background.	csh(1)
	banner: print large	background processes to complete.	csh(1)
		banner on printer.	banner(6)
		banner: print large banner on printer.	banner(6)
gettytab: terminal configuration data		base.	gettytab(5)
	hosts: host name data	base.	hosts(5)
	networks: network name data	base.	networks(5)
phones: remote host phone number data		base.	phones(5)
printcap: printer capability data		base.	printcap(5)
protocols: protocol name data		base.	protocols(5)
services: service name data		base.	services(5)
termcap: terminal capability data		base.	termcap(5)
vgrinddefs: vgrind's language definition data		base.	vgrinddefs(5)
	newaliases: rebuild the data	base for the mail aliases file.	newaliases(1)
	ttytype: data	base of terminal types by port.	ttytype(5)
fetch, store, delete, firstkey, nextkey: data		base subroutines. dbminit,	dbm(3X)
vi: screen oriented (visual) display editor		based on ex.	vi(1)

basename: strip filename affixes. basename(1)
 battlestar: a tropical adventure game. battlestar(6)
 bc: arbitrary-precision arithmetic language. bc(1)
 bed: convert to antique media. bed(6)
 bcopy, bcmp, bzero, ffs: bit and byte string operations. btring(3)
 operations. btring(3)
 cb: C program cb(1)
 j0, j1, jn, y0, y1, yn: j0(3M)
 changing/ random, srandom, initstate, setstate: beasel(3F)
 random(3)
 csh(1)
 addbib(1)
 roffbib(1)
 sortbib(1)
 lookbib(1)
 lookbib(1)
 lookbib(1)
 biff(1)
 comsat(8C)
 install(1)
 whereis(1)
 strings(1)
 uencode(1C)
 fread(3S)
 bind(2)
 bind(2)
 binmail(1)
 binstl(8)
 bstring(3)
 bit(3F)
 bit(3F)
 bload(8)
 sync(8)
 update(8)
 sigblock(2)
 sigpause(2)
 sum(1)
 boggle(6)
 boggle(6)
 ching(6)
 binstl(8)
 reboot(8)
 mille(6)
 csh(1)
 sh(1)
 csh(1)
 csh(1)
 csh(1)
 brk(2)
 fread(3S)
 intro(3S)
 setbuf(3S)
 kgmon(8)
 jove_recover(1)
 sendbug(1)
 bugfiler(8)
 bugfiler(8)
 lookbib(1)
 secstorcs(1)
 secstorcs(8)
 mknod(8)
 byteorder(3n)
 bstring(3)
 swab(3)
 bstring(3)
 cc(1)
 prec(7)
 cb(1)
 indent(1)
 lint(1)
 xstr(1)
 mkstr(1)
 hypot(3M)
 cal(1)
 dc(1)
 cal(1)
 cal(1)

syscall: indirect system	calendar: reminder service.	calendar(1)
gprof: display	call.	syscall(2)
getuid, getgid: get user or group ID of the	call graph profile data.	gprof(1)
malloc, free, realloc	caller.	getuid(3F)
intro: introduction to system	calloc, alloca: memory allocator.	malloc(3)
canfield, cfscores: the solitaire card game	calls and error numbers.	intro(2)
canfield.	canfield.	canfield(6)
printcap: printer	canfield, cfscores: the solitaire card game	canfield(6)
termcap: terminal	capability data base.	printcap(5)
canfield, cfscores: the solitaire	capability data base.	termcap(5)
cribbage: the	card game canfield.	canfield(6)
cd, eval, exec, exit, export, login,/ sh, for,	card game cribbage.	cribbage(6)
	case, if, while, :, ., ., break, continue,	sh(1)
	case: selector in switch.	csh(1)
	cat: catenate and print.	cat(1)
catman: create the	cat files for the manual.	catman(8)
uncompact, ccat: compress and uncompress files, and	cat them. compact,	compact(1)
default:	catchall clause in switch.	csh(1)
cat:	catenate and print.	cat(1)
	catman: create the cat files for the manual.	catman(8)
	cb: C program beautifier.	cb(1)
	cc: C compiler.	cc(1)
compact, uncompact,	ccat: compress and uncompress files, and cat them.	compact(1)
	cd: change directory.	csh(1)
case, if, while, :, ., ., break, continue,	cd: change working directory.	cd(1)
fabs, floor,	cd, eval, exec, exit, export, login, read,/ /for,	sh(1)
fabs, floor, ceil: absolute value, floor,	ceil: absolute value, floor, ceiling functions.	floor(3M)
canfield,	ceiling functions.	floor(3M)
chdir:	cfscores: the solitaire card game canfield.	canfield(6)
brk, sbrk:	change current working directory.	chdir(2)
chdir:	change data segment size.	brk(2)
chsh:	change default directory.	chdir(3F)
cd:	change default login shell.	chsh(1)
chdir:	change directory.	csh(1)
ioinit:	change directory.	csh(1)
chfn:	change I77 I/O initialization.	ioinit(3F)
chgrp:	change finger entry.	chfn(1)
passwd:	change group.	chgrp(1)
chmod:	change login password.	passwd(1)
chmod:	change mode.	chmod(1)
chmod:	change mode of a file.	chmod(3F)
chmod:	change mode of file.	chmod(2)
umask:	change or display file creation mask.	csh(1)
chown:	change owner.	chown(8)
chown:	change owner and group of a file.	chown(2)
rcs:	change RCS file attributes.	rcs(1)
chroot:	change root directory.	chroot(2)
signal:	change the action for a signal.	signal(3F)
rename:	change the name of a file.	rename(2)
set:	change value of shell variable.	csh(1)
cd:	change working directory.	cd(1)
ching: the book of	changes and other cookies.	ching(6)
better random number generator; routines for	changing generators. /srandom, initstate, setstate:	random(3)
pipe: create an interprocess communication	channel.	pipe(2)
ungetc: push	character back into input stream.	ungetc(3S)
isspace, ispunct, isprint, iscntrl, isascii:	character classification macros. /isdigit, isalnum,	ctype(3)
eqnchar: special	character definitions for eqn.	eqnchar(7)
getc, fgetc: get a	character from a logical unit.	getc(3F)
index, rindex, lnbk, len: tell about	character objects.	index(3F)
getc, getchar, fgetc, getw: get	character or word from stream.	getc(3S)
putc, putchar, fputc, putw: put	character or word on a stream.	putc(3S)
ascii: map of ASCII	character set.	ascii(7)
putc, fputc: write a	character to a fortran logical unit.	putc(3F)
style: analyze surface	characteristics of a document.	style(1)
tr: translate	characters.	tr(1)
prec: C precedence	chart.	prec(7)
snake, snscore: display	chase game.	snake(6)
	chdir: change current working directory.	chdir(2)
	chdir: change default directory.	chdir(3F)
	chdir: change directory.	csh(1)
dcheck: file system directory consistency	check.	dcheck(8)
icheck: file system storage consistency	check.	icheck(8)
fsck: file system consistency	check and interactive repair.	fsck(8)
ci:	check in RCS revisions.	ci(1)
checknr:	check nroff/troff files.	checknr(1)

co:	check out RCS revisions.	co(1)
eqn, neqn,	checkeq: typeset mathematics.	eqn(1)
quotacheck:	checker.	quotacheck(8)
fastboot, fasthalt:	checking the disks.	fastboot(8)
	checknr: check nroff/troff files.	checknr(1)
	chfn: change finger entry.	chfn(1)
	chgrp: change group.	chgrp(1)
	ching: the book of changes and other cookies.	ching(6)
	chmod: change mode.	chmod(1)
	chmod: change mode of a file.	chmod(3F)
	chmod: change mode of file.	chmod(2)
	chown: change owner.	chown(8)
	chown: change owner and group of a file.	chown(2)
	chroot: change root directory.	chroot(2)
	chsh: change default login shell.	chsh(1)
	ci: check in RCS revisions.	ci(1)
closepl:/ plot:	circle, arc, move, cont, point, linemod, space,	plot(3X)
ispunct, isprint, iscntrl, isascii:	classification macros. /isdigit, isalnum, isspace,	ctype(3)
default: catchall	clause in switch.	csh(1)
cleanlpd:	clean line printer daemon environment.	cleanlpd(8)
	cleanlpd: clean line printer daemon environment.	cleanlpd(8)
uuclean: uucp spool directory	clean-up.	uuclean(8C)
	clear: clear terminal screen.	clear(1)
clri:	clear i-node.	clri(8)
clear:	clear terminal screen.	clear(1)
error, feof,	clearerr, fileno: stream status inquiries.	ferror(3S)
csh: a shell (command interpreter) with	C-like syntax.	csh(1)
cron:	clock daemon.	cron(8)
	close: delete a descriptor.	close(2)
shutdown:	close down the system at a given time.	shutdown(8)
fclose, fflush:	close or flush a stream.	fclose(3S)
opendir, readdir, telldir, seekdir, rewinddir,	closedir: directory operations.	directory(3)
opendir, readdir, telldir, seekdir, rewinddir,	closedir: directory operations. directory:	directory(3X)
syslog, openlog,	closelog: control system log.	syslog(3)
circle, arc, move, cont, point, linemod, space,	closepl: graphics interface. /erase, label, line,	plot(3X)
	clri: clear i-node.	clri(8)
	cmp: compare two files.	cmp(1)
	co: check out RCS revisions.	co(1)
pi: Pascal interpreter	code translator.	pi(1)
	col: filter reverse line feeds.	col(1)
log. dmesg:	colert: filter nroff output for CRT previewing.	colert(1)
	collect system diagnostic messages to form error	dmesg(8)
colrm: remove	colrm: remove columns from a file.	colrm(1)
files.	columns from a file.	colrm(1)
exec: overlay shell with specified	comm: select or reject lines common to two sorted	comm(1)
time: time	command.	csh(1)
routines for returning a stream to a remote	command.	csh(1)
rexec: return stream to a remote	command. rcmd, rresvport, ruserok:	rcmd(3X)
system: issue a shell	command.	rexec(3X)
system: execute a UNIX	command.	system(3)
test: condition	command.	system(3F)
time: time a	command.	test(1)
nice, nohup: run a	command at low priority (<i>sh</i> only).	time(1)
switch: multi-way	command branch.	nice(1)
uux: unix to unix	command execution.	csh(1)
rehash: recompute	command hash table.	uux(1C)
unhash: discard	command hash table.	csh(1)
hashstat: print	command hashing statistics.	csh(1)
nohup: run	command immune to hangups.	csh(1)
csh: a shell	(command interpreter) with C-like syntax.	csh(1)
whatis: describe what a	command is.	whatis(1)
readonly, set, shift, times, trap, umask, wait:	command language. /exec, exit, export, login, read,	sh(1)
getarg, iargc: return	command line arguments.	getarg(3F)
repeat: execute	command repeatedly.	csh(1)
rc:	command script for auto-reboot and daemons.	rc(8)
onintr: process interrupts in	command scripts.	csh(1)
apply: apply a	command to a set of arguments.	apply(1)
goto:	command transfer.	csh(1)
else: alternative	commands.	csh(1)
intro: introduction to	commands.	intro(1)
introduction to system maintenance and operation	commands. intro:	intro(8)
resintro: introduction to RCS	commands.	resintro(1)
at: execute	commands at a later time.	at(1)
apropos: locate	commands by keyword lookup.	apropos(1)

while: repeat	commands conditionally.	csh(1)
lastcomm: show last	commands executed in reverse order.	lastcomm(1)
source: read	commands from file.	csh(1)
comm: select or reject lines	common to two sorted files.	comm(1)
socket: create an endpoint for	communication.	socket(2)
pipe: create an interprocess	communication channel.	pipe(2)
ipcs: report inter-process	communication facilities status.	ipcs(1)
talkd: remote user	communication server.	talkd(8C)
users:	compact list of users who are on the system.	users(1)
files, and cat them.	compact, uncompact, ccat: compress and uncompress	compact(1)
diff: differential file and directory	comparator.	diff(1)
resdiff:	compare RCS revisions.	resdiff(1)
cmp:	compare two files.	cmp(1)
diff3: 3-way differential file	comparison.	diff3(1)
intro: introduction to	compatibility library functions.	intro(3C)
lisp:	compile a Franz Lisp program.	lisp(1)
cc: C	compiler.	cc(1)
f77: Fortran 77	compiler.	f77(1)
pc: Pascal	compiler.	pc(1)
error: analyze and disperse	compiler error messages.	error(1)
yacc: yet another	compiler/compiler.	yacc(1)
fp: Functional Programming language	compiler/interpreter.	fp(1)
wait: wait for background processes to	complete.	csh(1)
wait: await	completion of process.	wait(1)
compress, uncompress, zcat:	compress and expand data.	compress(1)
compact, uncompact, ccat:	compress and uncompress files, and cat them.	compact(1)
data.	compress, uncompress, zcat: compress and expand	compress(1)
learn:	computer aided instruction about UNIX.	learn(1)
hangman:	Computer version of the game hangman.	hangman(6)
tac:	comsat: biff server.	comsat(8C)
test:	concatenate and print files in reverse order.	tac(1)
endif: terminate	conditional command.	test(1)
if:	conditional.	csh(1)
while: repeat commands	conditional statement.	csh(1)
gettytab: terminal	conditionally.	csh(1)
uxrc:	configuration data base.	gettytab(5)
ifconfig:	configuration file for kernel.	uxrc(8)
tip, cu:	configure network interface parameters.	ifconfig(8C)
dosc:	connect: initiate a connection on a socket.	connect(2)
getpeername: get name of	connect to a remote system.	tip(1C)
socketpair: create a pair of	connect to proc/286 system.	dosc(1)
shutdown: shut down part of a full-duplex	connected peer.	getpeername(2)
accept: accept a	connected sockets.	socketpair(2)
connect: initiate a	connection.	shutdown(2)
listen: listen for	connection on a socket.	accept(2)
dcheck: file system directory	connection on a socket.	connect(2)
icheck: file system storage	connections on a socket.	listen(2)
fack: file system	consistency check.	dcheck(8)
quotacheck: file system quota	consistency check.	icheck(8)
show what versions of object modules were used to	consistency check and interactive repair.	fack(8)
newfs:	consistency checker.	quotacheck(8)
mkproto:	construct a file. what:	what(1)
deroff: remove nroff, troff, tbl and eqn	construct a new file system.	newfs(8)
setrlimit: control maximum system resource	construct a prototype file system.	mkproto(8)
vlimit: control maximum system resource	constructs.	deroff(1)
/openpl, erase, label, line, circle, arc, move,	consumption. getrlimit,	getrlimit(2)
ls: list	consumption.	vlimit(3C)
sigstack: set and/or get signal stack	cont, point, linemod, space, closepl: graphics/	plot(3X)
sh, for, case, if, while, s, ., break,	contents of directory.	ls(1)
fentl: file	context.	sigstack(2)
ioctl:	continue, cd, eval, exec, exit, export, login,/	sh(1)
init: process	continue: cycle in loop.	csh(1)
getrlimit, setrlimit:	control.	fentl(2)
vlimit:	control device.	ioctl(2)
msgctl: message	control initialization.	init(8)
semctl: semaphore	control maximum system resource consumption.	getrlimit(2)
shmctl: shared memory	control maximum system resource consumption.	vlimit(3C)
fentl: file	control operations.	msgctl(2)
lpc: line printer	control operations.	semctl(2)
syslog, openlog, closelog:	control operations.	shmctl(2)
vhangup: virtually "hangup" the current	control options.	fentl(5)
term:	control program.	lpc(8)
	control system log.	syslog(3)
	control terminal.	vhangup(2)
	conventional names for terminals.	term(7)

ecvt, fevt, gcvt: output	conversion.	ecvt(3)
long, short: integer object	conversion.	long(3F)
printf, sprintf, sprintf: formatted output	conversion.	printf(3S)
scanf, fscanf, sscanf: formatted input	conversion.	scanf(3S)
units:	conversion program.	units(1)
dd:	convert and copy a file.	dd(1)
number:	convert Arabic numerals to English.	number(6)
ranlib:	convert archives to random libraries.	ranlib(1)
atof, atoi, atol:	convert ASCII to numbers.	atof(3)
ctime, localtime, gmtime, asctime, timezone:	convert date and time to ASCII.	ctime(3)
htable:	convert NIC standard format host tables.	htable(8)
bed:	convert to antique media.	bed(6)
htonl, htons, ntohl, ntohs:	convert values between host and network byte order.	byteorder(3n)
ching: the book of changes and other	cookies.	ching(6)
fpu: determine presence of the floating point	coprocessor.	fpu(1)
ep:	copy.	cp(1)
rep: remote file	copy.	rep(1C)
uucp, uulog: unix to unix	copy.	uucp(1C)
dd: convert and	copy a file.	dd(1)
tcopy:	copy a mag tape.	tcopy(1)
doscopyd: MPS/DOS file	copy daemon.	doscopyd(8)
cpio:	copy file archives in and out.	cpio(1)
fork: create a	copy of this process.	fork(3F)
copy: standalone	copy program.	copy(8)
	copy: standalone copy program.	copy(8)
	core: format of memory image file.	core(5)
gcore: get	core images of running processes.	gcore(1)
functions. sin,	cos, tan, asin, acos, atan, atan2: trigonometric	sin(3M)
sinh,	cosh, tanh: hyperbolic functions.	sinh(3M)
wc: word	count.	wc(1)
sum: sum and	count blocks in a file.	sum(1)
	cp: copy.	cp(1)
	cpio: copy file archives in and out.	cpio(1)
- recover JOVE buffers after a system/editor	crash. jove_recover	jove_recover(1)
	crash: what happens when the system crashes.	crash(8V)
crash: what happens when the system	crashes.	crash(8V)
	creat: create a new file.	creat(2)
fork:	create a copy of this process.	fork(3F)
creat:	create a new file.	creat(2)
open: open a file for reading or writing, or	create a new file.	open(2)
fork:	create a new process.	fork(2)
socketpair:	create a pair of connected sockets.	socketpair(2)
ctags:	create a tags file.	ctags(1)
socket:	create an endpoint for communication.	socket(2)
mkstr:	create an error message file by massaging C source.	mkstr(1)
pipe:	create an interprocess communication channel.	pipe(2)
dosdisk: program to	create and display information for MPS/DOS vdisks.	dosdisk(8)
addbib:	create or extend bibliographic database.	addbib(1)
catman:	create the cat files for the manual.	catman(8)
umask: change or display file	creation mode.	csh(1)
umask: set file	creation mode mask.	umask(2)
cribbage: the card game	cribbage.	cribbage(6)
	cribbage: the card game cribbage.	cribbage(6)
	cron: clock daemon.	cron(8)
lxref: lisp	cross reference program.	lxref(1)
pxref: Pascal	cross-reference program.	pxref(1)
colcrt: filter nroff output for	CRT previewing.	colcrt(1)
more, page: file perusal filter for	ert viewing.	more(1)
	crypt: encode/decode.	crypt(1)
	crypt, setkey, encrypt: DES encryption.	crypt(3)
syntax.	csh: a shell (command interpreter) with C-like	csh(1)
	ctags: create a tags file.	ctags(1)
convert date and time to ASCII.	ctime, localtime, gmtime, asctime, timezone:	ctime(3)
time,	ctime, ltime, gmtime: return system time.	time(3F)
tip,	cu: connect to a remote system.	tip(1C)
vhangup: virtually "hangup" the	current control terminal.	vhangup(2)
gethostid, sethostid: get/set unique identifier of	current host.	gethostid(2)
gethostname, sethostname: get/set name of	current host.	gethostname(2)
hostname: get name of	current host.	hostname(3F)
hostid: set or print identifier of	current host system.	hostid(1)
hostname: set or print name of	current host system.	hostname(1)
jobs: print	current job list.	csh(1)
sigsetmask: set	current signal mask.	sigsetmask(2)
whoami: print effective	current user id.	whoami(1)
chdir: change	current working directory.	chdir(2)

getcwd: get pathname of current working directory.	getcwd(3F)
getwd: get current working directory pathname.	getwd(3)
motion.	curses(3X)
curses: screen functions with "optimal" cursor motion.	curses(3X)
spline: interpolate smooth curve.	spline(1G)
continue: cycle in loop.	cs(1)
cron: clock daemon.	cron(8)
doscopyd: MPS/DOS file copy daemon.	doscopyd(8)
dosprint: MPS/DOS spooler daemon.	dosprint(8)
lpd: line printer daemon.	lpd(8)
routed: network routing daemon.	routed(8C)
cleanlpd: clean line printer daemon.	cleanlpd(8)
rc: command script for auto-reboot and daemons.	rc(8)
ftpd: DARPA Internet File Transfer Protocol server.	ftpd(8C)
telnetd: DARPA TELNET protocol server.	telnetd(8C)
tftpd: DARPA Trivial File Transfer Protocol server.	tftpd(8C)
compress, uncompress, zcat: compress and expand data.	compress(1)
eval: re-evaluate shell data.	cs(1)
gprof: display call graph profile data.	gprof(1)
prof: display profile data.	prof(1)
ttys: terminal initialization data.	ttys(5)
gettytab: terminal configuration data base.	gettytab(5)
hosts: host name data base.	hosts(5)
networks: network name data base.	networks(5)
phones: remote host phone number data base.	phones(5)
printcap: printer capability data base.	printcap(5)
protocols: protocol name data base.	protocols(5)
services: service name data base.	services(5)
termcap: terminal capability data base.	termcap(5)
vgrind's language definition data base.	vgrind(5)
newaliases: rebuild the data base for the mail aliases file.	newaliases(1)
ttytype: data base of terminal types by port.	ttytype(5)
dbminit, fetch, store, delete, firstkey, nextkey: data base subroutines.	dbm(3X)
brk, sbrk: change data segment size.	brk(2)
null: data sink.	null(4)
types: primitive system data types.	types(5)
addbib: create or extend bibliographic database.	addbib(1)
roffbib: run off bibliographic database.	roffbib(1)
sortbib: sort bibliographic database.	sortbib(1)
join: relational database operator.	join(1)
date: print and set the date.	date(1)
gettimeofday, settimeofday: get/set date and time.	gettimeofday(2)
time, ftime: get date and time.	time(3C)
fdate: return date and time in an ASCII string.	fdate(3F)
localtime, gmtime, asctime, timezone: convert date and time to ASCII.	ctime(3)
touch: update date last modified of a file.	touch(1)
idate, itime: return date or time in numerical form.	idate(3F)
date: print and set the date.	date(1)
file.. .PP dstrules: Daylight savings time and time zone name rule	dstrules(5)
data base subroutines.	dbm(3X)
dbx: debugger.	dbx(1)
dc: desk calculator.	dc(1)
dcheck: file system directory consistency check.	dcheck(8)
dd: convert and copy a file.	dd(1)
adb: debugger.	adb(1)
dbx: debugger.	dbx(1)
pdx: pascal debugger.	pdx(1)
od: octal, decimal, hex, ascii dump.	od(1)
default: catchall clause in switch.	cs(1)
chdir: change default directory.	chdir(3F)
chsh: change default login shell.	chsh(1)
vgrind's language definition data base.	vgrind(5)
mode.. Standalone mode: definition of this Sanyo/ICON machine operation	standalone(8)
eqnchar: special character definitions for eqn.	eqnchar(7)
stty, gtty: set and get terminal state (defunct).	stty(3C)
close: delete a descriptor.	close(2)
dbminit, fetch, store, delete, firstkey, nextkey: data base subroutines.	dbm(3X)
tail: deliver the last part of a file.	tail(1)
mesg: permit or deny messages.	mesg(1)
tset: terminal dependent initialization.	tset(1)
constructs. deroff: remove nroff, troff, tbl and eqn	deroff(1)
crypt, setkey, encrypt: DES encryption.	crypt(3)
whatis: describe what a command is.	whatis(1)
mailaddr: mail addressing description.	mailaddr(7)
getdiskbyname: get disk description by its name.	getdisk(3X)

remote: remote host	description file.	remote(5)
close: delete a	descriptor.	close(2)
dup, dup2: duplicate a	descriptor.	dup(2)
getfstype, setfsent, endfsent: get file system	descriptor file entry. /getfsspec, getfsfile,	getfsent(3X)
getdtablesize: get	descriptor table size.	getdtablesize(2)
dc:	desk calculator.	dc(1)
dosprinters:	destinations for spooled output from SLPT printers.	dosprinters(5)
access:	determine accessibility of a file.	access(3F)
access:	determine accessibility of file.	access(2)
file:	determine file type.	file(1)
coprocessor. fpu:	determine presence of the floating point	fpu(1)
fold: fold long lines for finite width output	device.	fold(1)
ioctl: control	device.	ioctl(2)
swapon: specify additional	device for paging and swapping.	swapon(8)
	df: disk free.	df(1)
fmin, fmax, frac, dfmin, dfmax,	dffrac, inmax: return extreme values.	fmin(3F)
fmin, fmax, frac, dfmin,	dfmax, dffrac, inmax: return extreme values.	fmin(3F)
values. fmin, fmax, frac,	dfmin, dfmax, dffrac, inmax: return extreme	fmin(3F)
dmesg: collect system	diagnostic messages to form error log.	dmesg(8)
print wordy sentences; thesaurus for	diction. diction,explain:	diction(1)
diction- print wordy sentences; thesaurus for	diction. explain,	explain(1)
diction. explain,	diction- print wordy sentences; thesaurus for	explain(1)
for diction.	diction,explain: print wordy sentences; thesaurus	diction(1)
	diff: differential file and directory comparator.	diff(1)
diff:	diff3: 3-way differential file comparison.	diff3(1)
diff3: 3-way	diff: differential file and directory comparator.	diff(1)
	diff3: 3-way differential file comparison.	diff3(1)
dir: format of	dir: format of directories.	dir(5)
directories.	directories.	dir(5)
rm, rmdir: remove (unlink) files or	directories.	rm(1)
rmdir, rm: remove (unlink)	directories or files.	rmdir(1)
cd: change working	directory.	cd(1)
chdir: change current working	directory.	chdir(2)
chdir: change default	directory.	chdir(3F)
chroot: change root	directory.	chroot(2)
cd: change	directory.	csh(1)
chdir: change	directory.	csh(1)
getcwd: get pathname of current working	directory.	getcwd(3F)
ls: list contents of	directory.	ls(1)
mkdir: make a	directory.	mkdir(1)
scandir: scan a	directory.	scandir(3)
swapon: specify a swap	directory.	swapon(2)
uuclean: uucp spool	directory clean-up.	uuclean(8C)
diff: differential file and	directory comparator.	diff(1)
dcheck: file system	directory consistency check.	dcheck(8)
unlink: remove	directory entry.	unlink(2)
unlink: remove a	directory entry.	unlink(3F)
mkdir: make a	directory file.	mkdir(2)
rmdir: remove a	directory file.	rmdir(2)
mklost+found: make a lost+found	directory for fsck.	mklost+found(8)
pwd: working	directory name.	pwd(1)
rewinddir, closedir: directory operations.	directory: opendir, readdir, telldir, seekdir,	directory(3X)
readdir, telldir, seekdir, rewinddir, closedir:	directory operations. opendir,	directory(3)
readdir, telldir, seekdir, rewinddir, closedir:	directory operations. directory: opendir,	directory(3X)
getwd: get current working	directory pathname.	getwd(3)
popd: pop shell	directory stack.	csh(1)
pushd: push shell	directory stack.	csh(1)
	dis: an mc68020 disassembler.	dis(1)
dis: an mc68020	disassembler.	dis(1)
quota: display	disc usage and limits.	quota(1)
unhash:	discard command hash table.	csh(1)
unset:	discard shell variables.	csh(1)
binstl: program to install bootloader on	disk.	binstl(8)
synchronize a file's in-core state with that on	disk. fsync:	fsync(2)
getdiskbyname: get	disk description by its name.	getdisk(3X)
dkfmt: standalone	disk formatter.	dkfmt(8)
df:	disk free.	df(1)
park: program to park the hard	disk heads.	park(8)
quota: manipulate	disk quotas.	quota(2)
du: summarize	disk usage.	du(1)
dosdisks: list of MPS/DOS virtual	disks.	dosdisks(5)
reboot/halt the system without checking the	disks. fastboot, fasthalt:	fastboot(8)
mount, umount: mount and	dismount file system.	mount(8)
error: analyze and	disperse compiler error messages.	error(1)
rain: animated raindrops	display.	rain(6)

gprof:	display call graph profile data.	gprof(1)
snake, snscore:	display chase game.	snake(6)
quota:	display disc usage and limits.	quota(1)
vi: screen oriented (visual)	display editor based on ex.	vi(1)
umask: change or	display file creation mask.	csh(1)
whodos:	display information about dosc users.	whodos(1)
dosdisk: program to create and	display information for MPS/DOS vdisks.	dosdisk(8)
prof:	display profile data.	prof(1)
sysline:	display system status on status line of a terminal.	sysline(1)
worms: animate worms on a	display terminal.	worms(6)
jove: an interactive	display-oriented text editor.	jove(1)
hypot, cabs: Euclidean	distance.	hypot(3M)
rdist: remote file	distribution program.	rdist(1)
	dkfmt: standalone disk formatter.	dkfmt(8)
	dmem, kmem: main memory.	mem(4)
error log.	dmesg: collect system diagnostic messages to form	dmesg(8)
	doctor: interact with a psychoanalyst.	doctor(6)
style: analyze surface characteristics of a	document.	style(1)
refer: find and insert literature references in	documents.	refer(1)
w: who is on and what they are	doing.	w(1)
rogue: Exploring The Dungeons of	Doom.	rogue(6)
	dosc: connect to proc/286 system.	dosc(1)
whodos: display information about	dosc users.	whodos(1)
	doscopyd: MPS/DOS file copy daemon.	doscopyd(8)
for MPS/DOS vdisks.	dosdisk: program to create and display information	dosdisk(8)
	dosdisks: list of MPS/DOS virtual disks.	dosdisks(5)
	dosprint: MPS/DOS spooler daemon.	dosprint(8)
SLPT printers.	dosprinters: destinations for spooled output from	dosprinters(5)
shutdown: shut	down part of a full-duplex connection.	shutdown(2)
shutdown: close	down the system at a given time.	shutdown(8)
rand,	drand, irand: return random values.	rand(3F)
graph:	draw a graph.	graph(1G)
arithmetic: provide	drill in number facts.	arithmetic(6)
rule file.. _PP	dstrules: Daylight savings time and time zone name	dstrules(5)
etime,	dtime: return elapsed execution time.	etime(3F)
	du: summarize disk usage.	du(1)
dump: incremental file system	dump.	dump(8)
od: octal, decimal, hex, ascii	dump.	od(1)
rdump: file system	dump across the network.	rdump(8C)
rrestore: restore a file system	dump across the network.	rrestore(8C)
	dump, dumpdates: incremental dump format.	dump(5)
dumpfs:	dump file system information.	dumpfs(8)
dump, dumpdates: incremental	dump format.	dump(5)
	dump: incremental file system dump.	dump(8)
kgmon: generate a	dump of the operating system's profile buffers.	kgmon(8)
dump,	dumpdates: incremental dump format.	dump(5)
	dumpfs: dump file system information.	dumpfs(8)
rogue: Exploring The	Dungeons of Doom.	rogue(6)
	dup, dup2: duplicate a descriptor.	dup(2)
dup,	dup2: duplicate a descriptor.	dup(2)
dup, dup2:	duplicate a descriptor.	dup(2)
echo:	echo arguments.	csh(1)
echo:	echo arguments.	echo(1)
	echo: echo arguments.	csh(1)
	echo: echo arguments.	echo(1)
ping: send ICMP	ECHO_REQUEST packets to network hosts.	ping(8)
	ecvt, fcvt, gcvt: output conversion.	ecvt(3)
	ed: text editor.	ed(1)
end, etext,	edata: last locations in program.	end(3)
ex,	edit: text editor.	ex(1)
vipw:	edit the password file.	vipw(8)
edquota:	edit user quotas.	edquota(8)
ed: text	editor.	ed(1)
ex, edit: text	editor.	ex(1)
jove: an interactive display-oriented text	editor.	jove(1)
ld: link	editor.	ld(1)
sed: stream	editor.	sed(1)
TEACHJOVE - learn how to use the JOVE	editor.	teachjove(1)
vi: screen oriented (visual) display	editor based on ex.	vi(1)
a.out: assembler and link	editor output.	a.out(5)
	edquota: edit user quotas.	edquota(8)
whoami: print	effective current user id.	whoami(1)
setregid: set real and	effective group ID.	setregid(2)
setreuid: set real and	effective user ID's.	setreuid(2)
vfork: spawn new process in a virtual memory	efficient way.	vfork(2)

grep,	ef: Extended Fortran Language.	ef(1)
etime, dtime: return	egrep, fgrep: search a file for a pattern.	grep(1)
insque, remque: insert/remove	elapsed execution time.	etime(3F)
soelim:	element from a queue.	insque(3)
setquota:	eliminate .so's from nroff input.	soelim(1)
uuencode: format of an	else: alternative commands.	csh(1)
crypt:	enable/disable quotas on a file system.	setquota(2)
mail. uuencode,uudecode:	encoded uuencode file.	uuencode(5)
crypt, setkey, encrypt: DES	encode/decode.	crypt(1)
makekey: generate	encode/decode a binary file for transmission via	uuencode(1C)
logout:	encrypt: DES encryption.	crypt(3)
/getfspec, getfsfile, getfstype, setfsent,	encryption.	crypt(3)
getgrent, getgrgid, getgrnam, setgrent,	encryption key.	makekey(8)
gethostbyaddr, gethostbyname, sethostent,	end, etext, edata: last locations in program.	end(3)
getnetent, getnetbyaddr, getnetbyname, setnetent,	end session.	csh(1)
socket: create an	end: terminate loop.	csh(1)
getprotobyname, getprotobynumber, getprotoent,	endfsent: get file system descriptor file entry.	getfsent(3X)
getpwent, getpwuid, getpwnam, setpwent,	endgrent: get group file entry.	getgrent(3)
getservbyport, getservbyname, setservent,	endhostent: get network host entry. gethostent,	gethostent(3n)
number: convert Arabic numerals to	endif: terminate conditional.	csh(1)
xsend, xget,	endnetent: get network entry.	getnetent(3n)
nlist: get	endpoint for communication.	socket(2)
chfn: change finger	endprotoent: get protocol entry. getprotoent,	getprotoent(3n)
setfsent, endfsent: get file system descriptor file	endpwent: get password file entry.	getpwent(3)
getgrnam, setgrent, endgrent: get group file	endservent: get service entry. getservent,	getservent(3n)
sethostent, endhostent: get network host	endsw: terminate switch.	csh(1)
getnetbyname, setnetent, endnetent: get network	English.	number(6)
setprotoent, endprotoent: get protocol	enroll: secret mail.	xsend(1)
getpwnam, setpwent, endpwent: get password file	entries from name list.	nlist(3)
getservbyname, setservent, endservent: get service	entry.	chfn(1)
unlink: remove directory	entry. getfsent, getfspec, getfsfile, getfstype,	getfsent(3X)
unlink: remove a directory	entry. getgrent, getgrgid,	getgrent(3)
execl, execv, execl, execlp, execvp, exec, exece,	entry. gethostent, gethostbyaddr, gethostbyname,	gethostent(3n)
cleanlpd: clean line printer daemon	entry. getnetent, getnetbyaddr,	getnetent(3n)
setenv: set variable in	entry. /getprotobyname, getprotobyname,	getprotoent(3n)
environ: user	entry. getpwent, getpwuid,	getpwent(3)
printenv: print out the	entry. getservent, getservbyport,	getservent(3n)
window: window	entry.	unlink(2)
getenv: value for	environ: execute a file.	unlink(3F)
unsetenv: remove	environ: user environment.	execl(3)
getenv: get value of	environment.	environ(7)
eqnchar: special character definitions for	environment.	cleanlpd(8)
deroff: remove nroff, troff, tbl and	environment.	csh(1)
linemod, space, closepl: graphics/ plot: openpl,	environment.	environ(7)
messages.	environment.	printenv(1)
dmesg: collect system diagnostic messages to form	environment.	window(1)
mkstr: create an	environment name.	getenv(3)
error: analyze and disperse compiler	environment variables.	csh(1)
perror, sys_errlist, sys_nerr: system	environment variables.	getenv(3F)
perror, gerror, ierrno: get system	eqn.	eqnchar(7)
intro: introduction to system calls and	eqn constructs.	deroff(1)
eyacc: modified yacc allowing much improved	eqn, neqn, checkeq: typeset mathematics.	eqn(1)
spell, spellin, spellout: find spelling	eqnchar: special character definitions for eqn.	eqnchar(7)
traper: trap arithmetic	erase, label, line, circle, arc, move, cont, point,	plot(3X)
end,	error: analyze and disperse compiler error	error(1)
hypot, cabs:	error log.	dmesg(8)
/if, while, , . . , break, continue, cd,	error message file by massaging C source.	mkstr(1)
expr:	error messages.	error(1)
history: print history	error messages.	perror(3)
screen oriented (visual) display editor based on	error messages.	perror(3F)
lpq: spool queue	error numbers.	intro(2)
execl, execv, execl, execlp, execvp,	error recovery.	eyacc(1)
	errors.	spell(1)
	errors.	traper(3F)
	etext, edata: last locations in program.	end(3)
	etime, dtime: return elapsed execution time.	etime(3F)
	Euclidean distance.	hypot(3M)
	eval, exec, exit, export, login, read, readonly,/	sh(1)
	eval: re-evaluate shell data.	csh(1)
	evaluate arguments as an expression.	expr(1)
	event list.	csh(1)
	ex. vi:	vi(1)
	ex, edit: text editor.	ex(1)
	examination program.	lpq(1)
	exec, exece, environ: execute a file.	execl(3)

/while, s, ., break, continue, cd, eval, exec, exit, export, login, read, readonly, set, / sh(1)
 exec: overlay shell with specified command. csh(1)
 exece, environ: execute a file. execl(3)
 execl, execev, execele, execlp, execev, exec, exece, execl(3)
 environ: execute a file. execl(3)
 execute a file. execl, execev, execele, execlp, execev, exec, exece, environ: execl(3)
 file. execl, execev, execele, execlp, execev, exec, exece, environ: execute a execl(3)
 sticky: executable files with persistent text. sticky(8)
 execele, execlp, execev, exec, exece, environ: execute a file. execl, execev, execl(3)
 execev: execute a file. execev(2)
 alarm: execute a subroutine after a specified time. alarm(3F)
 system: execute a UNIX command. system(3F)
 repeat: execute command repeatedly. csh(1)
 at: execute commands at a later time. at(1)
 lastcomm: show last commands lastcomm(1)
 uux: unix to unix command uux(1C)
 acct: execution accounting file. acct(5)
 sleep: suspend execution for an interval. sleep(1)
 sleep: suspend execution for an interval. sleep(3F)
 sleep: suspend execution for interval. sleep(3)
 monitor, monstartup, moncontrol: prepare execution profile. monitor(3)
 ppx: Pascal execution profiler. ppx(1)
 rexecd: remote execution server. rexecd(8C)
 etime, dtime: return elapsed execution time. etime(3F)
 profil: execution time profile. profil(2)
 pix: Pascal interpreter and executor. pix(1)
 environ: execute a file. execl, execev, execele, execlp, execev, exec, exece, execl(3)
 execev: execute a file. execev(2)
 execlp, exec, exece, environ: execute a file. execl(3)
 existing file. link(3F)
 existing file system. tunefs(8)
 /: s, ., break, continue, cd, eval, exec, exit, export, login, read, readonly, set, shift, / sh(1)
 exit from switch. csh(1)
 exit: leave shell. csh(1)
 _exit: terminate a process. exit(2)
 exit: terminate a process after flushing any exit(3)
 exit: terminate process with status. exit(3F)
 pending output. csh(1)
 break: exit while/foreach loop. csh(1)
 power, square root. exp, log, log10, pow, sqrt: exponential, logarithm, exp(3M)
 glob: filename expand argument list. csh(1)
 compress, uncompress, zcat: compress and expand data. compress(1)
 expand, unexpand: expand tabs to spaces, and vice versa. expand(1)
 expand, unexpand: expand tabs to spaces, and vice versa. expand(1)
 for diction. explain, diction- print wordy sentences; thesaurus explain(1)
 diction. diction, explain: print wordy sentences; thesaurus for diction(1)
 aardvark: yet another exploration game. aardvark(6)
 adventure: an exploration game. adventure(6)
 rogue: Exploring The Dungeons of Doom. rogue(6)
 frexp, ldexp, modf: split into mantissa and exponent. frexp(3)
 exp, log, log10, pow, sqrt: exponential, logarithm, power, square root. exp(3M)
 /: s, ., break, continue, cd, eval, exec, exit, export, login, read, readonly, set, shift, times, / sh(1)
 expr: evaluate arguments as an expression. expr(1)
 re_comp, re_exec: regular expression. expr(1)
 addbib: create or expression handler. regex(3)
 efi: extend bibliographic database. addbib(1)
 strings. xstr: Extended Fortran Language. efl(1)
 recovery. extract strings from C programs to implement shared xstr(1)
 ioinit: change eyacc: modified yacc allowing much improved error eyacc(1)
 tclose, tread, twrite, trewin, tskipf, tstate: f77: Fortran 77 compiler. f77(1)
 functions. f77 I/O initialization. ioinit(3F)
 signal: simplified software signal f77 tape I/O. topen, topen(3F)
 sigvec: software signal fabs, floor, ceil: absolute value, floor, ceiling floor(3M)
 ipc: report inter-process communication facilities. signal(3C)
 true, facilities. sigvec(2)
 checking the disks. facilities status. ipes(1)
 the disks. fastboot, fasthalt: reboot/halt the system without true(1)
 abort: generate a false, true: provide truth values. false(1)
 trpfp, fpecnt: trap and repair floating point fastboot, fasthalt: reboot/halt the system without fastboot(8)
 export, login, / sh, for, case, if, while, s, fasthalt: reboot/halt the system without checking fastboot(8)
 exit, export, login, / sh, for, case, if, while, fault. abort(3)
 faults. trpfp(3F)
 ., break, continue, cd, eval, exec, exit, sh(1)
 s, ., break, continue, cd, eval, exec, sh(1)
 fclose, fflush: close or flush a stream. fclose(3S)
 fcntl: file control. fcntl(2)
 fcntl: file control options. fcntl(5)

ecvt,	fevt, gcvt: output conversion.	ecvt(3)
	fddate: return date and time in an ASCII string.	fddate(3F)
fopen, freopen,	fdopen: open a stream.	fopen(3S)
ferror,	feof, clearerr, fileno: stream status inquiries.	ferror(3S)
inquiries.	ferror, feof, clearerr, fileno: stream status	ferror(3S)
subroutines. dbminit,	fetch, store, delete, firstkey, nextkey: data base	dbm(3X)
head: give first	few lines.	head(1)
fclose,	flush: close or flush a stream.	fclose(3S)
extreme values. fmin, fmax,	frac, dflmin, dflmax, dfrac, inmax: return	fmin(3F)
bcopy, bcmp, bzero,	fs: bit and byte string operations.	bstring(3)
	fg: bring job into foreground.	cash(1)
getc,	fgetc: get a character from a logical unit.	getc(3F)
getc, getchar,	fgetc, getw: get character or word from stream.	getc(3S)
gets,	fgets: get a string from a stream.	gets(3S)
grep, egrep,	fgrep: search a file for a pattern.	grep(1)
locate a program file including aliases and paths	(csh only). which:	which(1)
robots:	fight off villainous robots.	robots(6)
access: determine accessibility of	file.	access(2)
access: determine accessability of a	file.	access(3F)
acct: execution accounting	file.	acct(5)
chmod: change mode of	file.	chmod(2)
chmod: change mode of a	file.	chmod(3F)
chown: change owner and group of a	file.	chown(2)
colrm: remove columns from a	file.	colrm(1)
core: format of memory image	file.	core(5)
creat: create a new	file.	creat(2)
source: read commands from	file.	cash(1)
ctags: create a tags	file.	ctags(1)
dd: convert and copy a	file.	dd(1)
Daylight savings time and time zone name rule	file..PP dstrules:	dstrules(5)
execlp, execvp, exec, exece, environ: execute a	file. execl, execlv, execlv,	execl(3)
execve: execute a	file.	execve(2)
flock: apply or remove an advisory lock on an open	file.	flock(2)
fpr: print Fortran	file.	fpr(1)
group: group	file.	group(5)
link: make a hard link to a	file.	link(2)
link: make a link to an existing	file.	link(3F)
mkdir: make a directory	file.	mkdir(2)
mknod: make a special	file.	mknod(2)
mknod: build special	file.	mknod(8)
rebuild the data base for the mail aliases	file. newaliases:	newaliases(1)
open a file for reading or writing, or create a new	file. open:	open(2)
passwd: password	file.	passwd(5)
pr: print	file.	pr(1)
rcsfile: format of RCS	file.	rcsfile(5)
remote: remote host description	file.	remote(5)
rename: change the name of a	file.	rename(2)
rename: rename a	file.	rename(3F)
rev: reverse lines of a	file.	rev(1)
rmdir: remove a directory	file.	rmdir(2)
scstores: build RCS file from SCCS	file.	scstores(1)
scstores: build RCS file from SCCS	file.	scstores(8)
sfdate: set the time/date of a	file.	sfdate(1)
size: size of an object	file.	size(1)
the printable strings in a object, or other binary,	file. strings: find	strings(1)
sum: sum and count blocks in a	file.	sum(1)
symlink: make symbolic link to a	file.	symlink(2)
tail: deliver the last part of a	file.	tail(1)
touch: update date last modified of a	file.	touch(1)
uniq: report repeated lines in a	file.	uniq(1)
uuencode: format of an encoded uuencode	file.	uuencode(5)
vipw: edit the password	file.	vipw(8)
versions of object modules were used to construct a	file. what: show what	what(1)
write, writev: write on a	file.	write(2)
diff: differential	file and directory comparator.	diff(1)
cpio: copy	file archives in and out.	cpio(1)
rcs: change RCS	file attributes.	rcs(1)
bugfler:	file bug reports in folders automatically.	bugfler(8)
mkstr: create an error message	file by massaging C source.	mkstr(1)
diff3: 3-way differential	file comparison.	diff3(1)
fcntl:	file control.	fcntl(2)
fcntl:	file control options.	fcntl(5)
rcp: remote	file copy.	rcp(1C)
doscopyd: MPS/DOS	file copy daemon.	doscopyd(8)
umask: change or display	file creation mask.	cash(1)

umask: set	file creation mode mask.	umask(2)
	file: determine file type.	file(1)
	file distribution program.	rdist(1)
setfsent, endfsent: get file system descriptor	file entry. /getfsspec, getfsfile, getfstype,	getfsent(3X)
getgrgid, getgrnam, setgrent, endgrent: get group	file entry. getgrent,	getgrent(3)
getpwnam, setpwent, endpwent: get password	file entry. getpwent, getpwuid,	getpwent(3)
grep, egrep, fgrep: search a	file for a pattern.	grep(1)
uxrc: configuration	file for kernel.	uxrc(8)
open: open a	file for reading or writing, or create a new file.	open(2)
aliases: aliases	file for sendmail.	aliases(5)
uuencode, uudecode: encode/decode a binary	file for transmission via mail.	uuencode(1C)
ar: archive (library)	file format.	ar(5)
tar: tape archive	file format.	tar(5)
scstorcs: build RCS	file from SCCS file.	scstorcs(1)
scstorcs: build RCS	file from SCCS file.	scstorcs(8)
which: locate a program	file including aliases and paths (csh only).	which(1)
fsplit: split a multi-routine Fortran	file into individual files.	fsplit(1)
split: split a	file into pieces.	split(1)
merge: three-way	file merge.	merge(1)
pmerge: pascal	file merger.	pmerge(1)
mktemp: make a unique	file name.	mktemp(3)
fseek, ftell: reposition a	file on a logical unit.	fseek(3F)
more, page:	file perusal filter for crt viewing.	more(1)
stat, lstat, fstat: get	file status.	stat(2)
stat, lstat, fstat: get	file status.	stat(3F)
mkproto: construct a prototype	file system.	mkproto(8)
mount, umount: mount or remove	file system.	mount(2)
mount, umount: mount and dismount	file system.	mount(8)
newfs: construct a new	file system.	newfs(8)
repquota: summarize quotas for a	file system.	repquota(8)
setquota: enable/disable quotas on a	file system.	setquota(2)
tunefs: tune up an existing	file system.	tunefs(8)
repair. fsck:	file system consistency check and interactive	fsck(8)
getfsfile, getfstype, setfsent, endfsent: get	file system descriptor file entry. /getfsspec,	getfsent(3X)
dcheck:	file system directory consistency check.	dcheck(8)
dump: incremental	file system dump.	dump(8)
rdump:	file system dump across the network.	rdump(8C)
rrestore: restore a	file system dump across the network.	rrestore(8C)
hier:	file system hierarchy.	hier(7)
dumpfs: dump	file system information.	dumpfs(8)
quot: summarize	file system ownership.	quot(8)
quotacheck:	file system quota consistency checker.	quotacheck(8)
quotaon, quotaoff: turn	file system quotas on and off.	quotaon(8)
restore: incremental	file system restore.	restore(8)
icheck:	file system storage consistency check.	icheck(8)
mtab: mounted	file system table.	mtab(5)
fs, inode: format of	file system volume.	fs(5)
mkfs: program to make UNIX	file systems.	mkfs(8)
utime: set	file times.	utime(3C)
utimes: set	file times.	utimes(2)
uuse: send a	file to a remote host.	uuse(1C)
truncate: truncate a	file to a specified length.	truncate(2)
kermit: kermit	file transfer.	kermit(1)
ftp:	file transfer program.	ftp(1C)
tftp: trivial	file transfer program.	tftp(1C)
ftpd: DARPA Internet	File Transfer Protocol server.	ftpd(8C)
tftpd: DARPA Trivial	File Transfer Protocol server.	tftpd(8C)
file: determine	file type.	file(1)
basename: strip	filename affixes.	basename(1)
glob:	filename expand argument list.	csh(1)
ferror, feof, clearerr,	fileno: stream status inquiries.	ferror(3S)
checknr: check nroff/troff	files.	checknr(1)
cmp: compare two	files.	cmp(1)
comm: select or reject lines common to two sorted	files.	comm(1)
find: find	files.	find(1)
split a multi-routine Fortran file into individual	files. fsplit:	fsplit(1)
ident: identify	files.	ident(1)
lockf: record locking on	files.	lockf(3C)
makedev: make system special	files.	makedev(8)
mv: move or rename	files.	mv(1)
print log messages and other information about RCS	files. rlog:	rlog(1)
rmdir, rm: remove (unlink) directories or	files.	rmdir(1)
sort: sort or merge	files.	sort(1)
compact, uncompact, ccat: compress and uncompress	files, and cat them.	compact(1)
catman: create the cat	files for the manual.	catman(8)

tac: concatenate and print	files in reverse order.	tac(1)
fsync: synchronize a	file's in-core state with that on disk.	fsync(2)
rm, rmdir: remove (unlink)	files or directories.	rm(1)
sticky: executable	files with persistent text.	sticky(8)
fstab: static information about the	filesystems.	fstab(5)
more, page: file perusal	filter for crt viewing.	more(1)
colert:	filter nroff output for CRT previewing.	colert(1)
col:	filter reverse line feeds.	col(1)
plot: graphics	filters.	plot(1G)
refer:	find and insert literature references in documents.	refer(1)
find:	find files.	find(1)
	find: find files.	find(1)
look:	find lines in a sorted list.	look(1)
manual. man:	find manual information by keywords; print out the	man(1)
ttynam, isatty, ttyslot:	find name of a terminal.	ttynam(3)
ttynam, isatty:	find name of a terminal port.	ttynam(3F)
lorder:	find ordering relation for an object library.	lorder(1)
lookbib: build inverted index for a bibliography,	find references in a bibliography. indxbib,	lookbib(1)
spell, spellin, spellout:	find spelling errors.	spell(1)
binary, file. strings:	find the printable strings in a object, or other	strings(1)
chfn: change	finger entry.	chfn(1)
	finger: user information lookup program.	finger(1)
fold: fold long lines for	finite width output device.	fold(1)
tmail: print out mail messages, most recent	first.	tmail(1)
head: give	first few lines.	head(1)
dbminit, fetch, store, delete,	firstkey, nextkey: data base subroutines.	dbm(3X)
fish: play "Go	Fish".	fish(6)
	fish: play "Go Fish".	fish(6)
nice, nohup: run a command at low priority	(#h only).	nice(1)
extreme values. fmin,	fmax, frac, dffmin, dffmax, dffrac. inmax: return	fmin(3F)
return extreme values.	fmin, fmax, frac, dffmin, dffmax, dffrac. inmax:	fmin(3F)
fpu: determine presence of the	floating point coprocessor.	fpu(1)
trpfpe, fpecnt: trap and repair	floating point faults.	trpfpe(3F)
trapov: trap and repair	floating point overflow.	trapov(3F)
file.	flock: apply or remove an advisory lock on an open	flock(2)
functions. fabs,	floor, ceil: absolute value, floor, ceiling	floor(3M)
fabs, floor, ceil: absolute value,	floor, ceiling functions.	floor(3M)
fclose, fflush: close or	flush a stream.	fclose(3S)
	flush: flush output to a logical unit.	flush(3F)
flush:	flush output to a logical unit.	flush(3F)
exit: terminate a process after	flushing any pending output.	exit(3)
	fmt: simple text formatter.	fmt(1)
device.	fold: fold long lines for finite width output	fold(1)
fold:	fold long lines for finite width output device.	fold(1)
bugfiler: file bug reports in	folders automatically.	bugfiler(8)
	fopen, freopen, fdopen: open a stream.	fopen(3S)
	foreach: loop over list of names.	csh(1)
fg: bring job into	foreground.	csh(1)
	fork: create a copy of this process.	fork(3F)
	fork: create a new process.	fork(2)
idate, itime: return date or time in numerical	form.	idate(3F)
dmesg: collect system diagnostic messages to	form error log.	dmesg(8)
ar: archive (library) file	format.	ar(5)
dump, dumpdates: incremental dump	format.	dump(5)
tar: tape archive file	format.	tar(5)
indent: indent and	format C program source.	indent(1)
htable: convert NIC standard	format host tables.	htable(8)
gettable: get NIC	format host tables from a host.	gettable(8C)
uuencode:	format of an encoded uuencode file.	uuencode(5)
dir:	format of directories.	dir(5)
fs, inode:	format of file system volume.	fs(5)
core:	format of memory image file.	core(5)
rcsfile:	format of RCS file.	rcsfile(5)
tbl:	format tables for nroff or troff.	tbl(1)
scanf, fscanf, sscanf:	formatted input conversion.	scanf(3S)
printf, sprintf, sprintf:	formatted output conversion.	printf(3S)
dkfmt: standalone disk	formatter.	dkfmt(8)
fmt: simple text	formatter.	fmt(1)
nroff: text	formatting.	nroff(1)
troff, nroff: text	formatting and typesetting.	troff(1)
ms: text	formatting macros.	ms(7)
me: macros for	formatting papers.	me(7)
f77:	Fortran 77 compiler.	f77(1)
ratfor: rational	Fortran dialect.	ratfor(1)
fpr: print	Fortran file.	fpr(1)

fsplit: split a multi-routine Fortran file into individual files.	fsplit(1)
efl: Extended Fortran Language.	efl(1)
intro: introduction to FORTRAN library functions.	intro(3F)
putc, fputc: write a character to a fortran logical unit.	putc(3F)
struct: structure Fortran programs.	struct(1)
adage. fortune: print a random, hopefully interesting, , break, continue, cd, eval, exec, exit, export,	fortune(6)
login,/ sh, for, case, if, while, t, . , . , break, continue, cd, eval, exec,	sh(1)
exit, export,/ sh, for, case, if, while, : compiler/interpreter.	sh(1)
trpfpe, fp: Functional Programming language	fp(1)
printf, sprintf: formatted output conversion.	trpfpe(3F)
coprocessor. fpu: determine presence of the floating point	fpr(1)
putc, putchar, fputc, putw: put character or word on a stream.	printf(3S)
putc, fputc: write a character to a fortran logical unit.	fpu(1)
puts, fputs: put a string on a stream.	putc(3S)
liszt: compile a Franz Lisp program.	putc(3F)
df: disk fread, fwrite: buffered binary input/output.	puts(3S)
malloc, free.	liszt(1)
fopen, freopen, fdopen: open a stream.	fread(3S)
exponent. frexp, ldexp, modf: split into mantissa and	df(1)
from: who is my mail from!	malloc(3)
scanf, fscanf, sscanf: formatted input conversion.	fopen(3S)
mklost+found: make a lost+found directory for	frexp(3)
repair. fsck.	from(1)
individual files. fsck: file system consistency check and interactive	fs(5)
stat, lstat, fseek, ftell: reposition a file on a logical unit.	scanf(3S)
stat, lstat, fseek, ftell, rewind: reposition a stream.	mklost+found(8)
on disk. fsplit: split a multi-routine Fortran file into	fseek(8)
fseek, fstat: static information about the filesystems.	fseek(3F)
time, fstat: get file status.	fseek(3S)
shutdown: shut down part of a gamma: log gamma	fsplit(1)
compiler/interpreter. fp: Functional Programming language	fstat(5)
bit: and, or, xor, not, rshift, lshift bitwise	stat(2)
fabs, floor, ceil: absolute value, floor, ceiling	stat(3F)
intro: introduction to library functions.	fsync(2)
intro: introduction to compatibility library functions.	fseek(3F)
intro: introduction to FORTRAN library functions.	fseek(3S)
intro: introduction to mathematical library functions.	time(3C)
intro: introduction to network library functions.	ftp(1C)
intro: introduction to miscellaneous library functions.	ftpd(8C)
j0, j1, jn, y0, y1, yn: bessel functions.	shutdown(2)
cos, tan, asin, acos, atan, atan2: trigonometric functions.	gamma(3M)
sinh, cosh, tanh: hyperbolic functions.	fp(1)
bessel functions: of two kinds for integer orders.	bit(3F)
curves: screen functions with "optimal" cursor motion.	floor(3M)
fread, fwrite: buffered binary input/output.	intro(3)
aardvark: yet another exploration game.	intro(3C)
adventure: an exploration game.	intro(3F)
backgammon: the game.	intro(3M)
battlestar: a tropical adventure game.	intro(3n)
hunt: a multi-player multi-terminal game.	intro(3X)
monop: Monopoly game.	j0(3M)
snake, snscore: display chase game.	sin(3M)
trek: trekkie game.	sinh(3M)
worm: Play the growing worm game.	bessel(3F)
canfield, cfscores: the solitaire card game canfield.	curves(3X)
cribbage: the card game cribbage.	fread(3S)
hangman: Computer version of the game hangman.	aardvark(6)
boggle: play the game of boggle.	adventure(6)
wump: the game of hunt-the-wumpus.	backgammon(6)
gamma: log gamma function.	battlestar(6)
gamma: log gamma function.	hunt(6)
gcov: get core images of running processes.	monop(6)
gevt: output conversion.	snake(6)
kgmon: generate a dump of the operating system's profile	trek(6)
	worm(6)
	canfield(6)
	cribbage(6)
	hangman(6)
	boggle(6)
	wump(6)
	gamma(3M)
	gamma(3M)
	gcov(1)
	ecvt(3)
	kgmon(8)

abort:	generate a fault.	abort(3)
makekey:	generate encryption key.	makekey(8)
ncheck:	generate names from i-numbers.	ncheck(8)
rand, srand:	random number generator.	rand(3C)
lex:	generator of lexical analysis programs.	lex(1)
/arandom, initstate, setstate:	better random number generator; routines for changing	random(3)
random number generator;	error, ierrno: get system error messages.	random(3)
error,	getarg, iargc: return command line arguments.	error(3F)
	getc, fgetc: get a character from a logical unit.	getarg(3F)
	getc, getchar, fgetc, getw: get character or word	getc(3F)
from stream.	getchar, fgetc, getw: get character or word from	getc(3S)
stream. getc,	getcwd: get pathname of current working directory.	getc(3S)
	getdiskbyname: get disk description by its name.	getcwd(3F)
	getdtablesize: get descriptor table size.	getdisk(3X)
	getgid, getegid: get group identity.	getdtablesize(2)
	getenv: get value of environment variables.	getgid(2)
	getenv: value for environment name.	getenv(3F)
	getuid, geteuid: get user identity.	getenv(3)
setfsent, endfsent:	getfsent, getfspec, getfsfile, getfstype,	getuid(2)
get file system descriptor/	getfsfile, getfstype, setfsent, endfsent: get file	getfsent(3X)
system descriptor file entry. getfsent, getfspec,	getfspec, getfsfile, getfstype, setfsent,	getfsent(3X)
endfsent: get file system descriptor/	getfstype, setfsent, endfsent: get file system	getfsent(3X)
descriptor file/	getgid: get user or group ID of the caller.	getfsent(3X)
getfsent, getfsspec, getfsfile,	getgid, getegid: get group identity.	getuid(3F)
getuid,	getgrnt, getgrgid, getgrnam, setgrnt, endgrnt:	getgid(2)
	getgrgid, getgrnam, setgrnt, endgrnt: get group	getgrnt(3)
get group file entry.	getgrnam, setgrnt, endgrnt: get group file entry.	getgrnt(3)
file entry. getgrnt,	getgroups: get group access list.	getgrnt(3)
getgrnt, getgrgid,	gethostbyaddr, gethostbyname, sethostent,	getgrnt(3)
	gethostbyname, sethostent, endhostent: get network	getgroups(2)
endhostent: get network host entry. gethostent,	gethostent, gethostbyaddr, gethostbyname,	gethostent(3n)
host entry. gethostent, gethostbyaddr,	gethostent, gethostbyaddr, gethostbyname,	gethostent(3n)
sethostent, endhostent: get network host entry.	gethostid, sethostid: get/set unique identifier of	gethostent(3n)
current host.	gethostname, sethostname: get/set name of current	gethostid(2)
host.	getitimer, setitimer: get/set value of interval	gethostname(2)
timer.	getlog: get user's login name.	getitimer(2)
	getlogin: get login name.	getlog(3F)
	getnetbyaddr, getnetbyname, setnetent, endnetent:	getlogin(3)
get network entry. getnetent,	getnetbyname, setnetent, endnetent: get network	getnetent(3n)
entry. getnetent, getnetbyaddr,	getnetent, getnetbyaddr, getnetbyname, setnetent,	getnetent(3n)
endnetent: get network entry.	getpagesize: get system page size.	getnetent(3n)
	getpass: read a password.	getpagesize(2)
	getpeername: get name of connected peer.	getpass(3)
	getpgrp: get process group.	getpeername(2)
	getpid: get process id.	getpgrp(2)
	getpid, getppid: get process identification.	getpid(3F)
	getppid: get process identification.	getpid(2)
getpid,	getpriority, setpriority: get/set program	getpid(2)
scheduling priority.	getprotobyname, setprotoent, endprotoent: get	getpriority(2)
protocol entry. getprotoent, getprotobyname,	getprotoent, getprotobyname, setprotoent,	getprotoent(3n)
endprotoent: get protocol entry. getprotoent,	getprotoent, getprotobyname, getprotobyname,	getprotoent(3n)
setprotoent, endprotoent: get protocol entry.	getpw: get name from uid.	getprotoent(3n)
	getpwent, getpwuid, getpwnam, setpwent, endpwent:	getpw(3C)
	getpwnam, setpwent, endpwent: get password file	getpwent(3)
get password file entry.	getpwuid, getpwnam, setpwent, endpwent: get	getpwent(3)
entry. getpwent, getpwuid,	getrlimit, setrlimit: control maximum system	getpwent(3)
password file entry. getpwent,	getrusage: get information about resource	getrlimit(2)
resource consumption.	gets, fgets: get a string from a stream.	getrusage(2)
utilization.	getservbyname, setservent, endservent: get service	gets(3S)
	getservbyport, getservbyname, setservent,	getservent(3n)
entry. getservent, getservbyport,	getservent, getservbyport, getservbyname,	getservent(3n)
endservent: get service entry. getservent,	get/set date and time.	getservent(3n)
setservent, endservent: get service entry.	get/set name of current host.	gettimeofday(2)
gettimeofday, settimeofday:	get/set program scheduling priority.	gethostname(2)
gethostname, sethostname:	get/set unique identifier of current host.	getpriority(2)
getpriority, setpriority:	get/set value of interval timer.	gethostid(2)
gethostid, sethostid:	getsockname: get socket name.	getitimer(2)
getitimer, setitimer:	getsockopt, setsockopt: get and set options on	getsockname(2)
	gettable: get NIC format host tables from a host.	getsockopt(2)
sockets.	gettimeofday, settimeofday: get/set date and time.	gettable(8C)
	getty: set terminal mode.	gettimeofday(2)
	gettytab: terminal configuration data base.	getty(8)
	getuid, geteuid: get user identity.	gettytab(5)
	getuid, getgid: get user or group ID of the caller.	getuid(2)
	getw: get character or word from stream.	getuid(3F)
getc, getchar, fgetc,		getc(3S)

	getwd: get current working directory pathname.	getwd(3)
head:	give first few lines.	head(1)
shutdown: close down the system at a	given time.	shutdown(8)
	glob: filename expand argument list.	glob(1)
ASCII. ctime, localtime,	gmtime, asctime, timezone: convert date and time to	ctime(3)
time, ctime, ltime,	gmtime: return system time.	time(3F)
fish: play	"Go Fish".	fish(6)
setjmp, longjmp: non-local	goto.	setjmp(3)
	goto: command transfer.	goto(1)
	gprof: display call graph profile data.	gprof(1)
graph: draw a	graph.	graph(1G)
	graph: draw a graph.	graph(1G)
gprof: display call	graph profile data.	gprof(1)
plot:	graphics filters.	plot(1G)
arc, move, cont, point, linemod, space, closepl:	graphics interface. /erase, label, line, circle,	plot(3X)
plot:	graphics interface.	plot(5)
lib2648: subroutines for the HP 2648	graphics terminal.	lib2648(3X)
	grep, egrep, fgrep: search a file for a pattern.	grep(1)
	grind nice listings of programs.	grind(1)
vgrind:	group.	chgrp(1)
chgrp: change	group.	getpgrp(2)
getpgrp: get process	group.	killpg(2)
killpg: send signal to a process	group.	setpgrp(2)
setpgrp: set process	group access list.	getgroups(2)
getgroups: get	group access list.	initgroups(3X)
initgroups: initialize	group access list.	setgroups(2)
setgroups: set	group file.	group(5)
group:	group file entry. getgrent,	getgrent(3)
getgrgid, getgrnam, setgrent, endgrent: get	group: group file.	group(5)
	group ID.	setregid(2)
setregid: set real and effective	group ID. setuid, seteuid,	setuid(3)
setruid, setgid, setegid, setrgid: set user and	group ID of the caller.	getuid(3F)
getuid, getgid: get user or	group identity.	getgid(2)
getgid, getegid: get	group memberships.	groups(1)
groups: show	group of a file.	chown(2)
chown: change owner and	groups.	make(1)
make: maintain program	groups: show group memberships.	groups(1)
	growing worm game.	worm(6)
worm: Play the	gtty: set and get terminal state (defunct).	stty(3C)
stty,	halt a job or process.	cs(1)
stop:	halt processor.	reboot(2)
reboot: reboot system or	halt: stop the processor.	halt(8)
	handle remote mail received via uucp.	rmail(1)
rmail:	handler.	regex(3)
re_comp, re_exec: regular expression	hangman.	hangman(6)
hangman: Computer version of the game	hangman: Computer version of the game hangman.	hangman(6)
	"hangup" the current control terminal.	vhangup(2)
vhangup: virtually	hangups.	cs(1)
nohup: run command immune to	happens when the system crashes.	crash(8V)
crash: what	hard disk heads.	park(8)
park: program to park the	hard link to a file.	link(2)
link: make a	hash table.	cs(1)
rehash: recompute command	hash table.	cs(1)
unhash: discard command	hashing statistics.	cs(1)
hashstat: print command	hashstat: print command hashing statistics.	cs(1)
	have to leave.	leave(1)
leave: remind you when you	heads.	park(8)
park: program to park the hard disk	hex, ascii dump.	od(1)
od: octal, decimal,	hier: file system hierarchy.	hier(7)
	hierarchy.	hier(7)
hier: file system	history event list.	cs(1)
history: print	history: print history event list.	cs(1)
	hopefully interesting. adage.	fortune(6)
fortune: print a random,	host. gethostid,	gethostid(2)
sethostid: get/set unique identifier of current	host.	gethostname(2)
gethostname, sethostname: get/set name of current	host.	gettable(8C)
gettable: get NIC format host tables from a	host.	hostname(3F)
hostname: get name of current	host.	uucsend(1C)
uucsend: send a file to a remote	host and network byte order.	byteorder(3n)
htonl, htons, ntohl, ntohs: convert values between	host description file.	remote(5)
remote: remote	host entry. gethostent, gethostbyaddr,	gethostent(3n)
gethostbyname, sethostent, endhostent: get network	host name data base.	hosts(5)
hosts:	host phone number data base.	phones(5)
phones: remote	host status of local machines.	ruptime(1C)
ruptime: show	host system.	hostid(1)
hostid: set or print identifier of current		

hostname: set or print name of current host system. hostname(1)
 htable: convert NIC standard format host tables. htable(8)
 gettable: get NIC format host tables from a host. gettable(8C)
 hostid: set or print identifier of current host hostid(1)
 hostname: set or print name of current host system. hostname(1)
 hostnm: get name of current host. hostnm(3F)
 hosts. ping(8)
 hosts: host name data base. hosts(5)
 how long system has been up. uptime(1)
 how to use the JOVE editor. teachjove(1)
 HP 2648 graphics terminal. lib2648(3X)
 htable: convert NIC standard format host tables. htable(8)
 htonl, htols, ntohl, ntohs: convert values between byteorder(3n)
 htols, ntohl, ntohs: convert values between host byteorder(3n)
 hunt: a multi-player multi-terminal game. hunt(6)
 hunt-the-wumpus. wump(6)
 hyperbolic functions. sinh(3M)
 hypot, cabs: Euclidean distance. hypot(3M)
 iargc: return command line arguments. getarg(3F)
 icode: file system storage consistency check. icode(8)
 ICMP ECHO_REQUEST packets to network hosts. ping(8)
 id. getpid(3F)
 id. ipcrm: remove ipcrm(1)
 ID. setregid(2)
 ID. setuid, seteuid, setruid, setuid(3)
 id. whoami(1)
 ID of the caller. getuid(3F)
 id temporarily. su(1)
 idate, itime: return date or time in numerical idate(3F)
 ident: identify files. ident(1)
 identification. getpid(2)
 identifier of current host. gethostid(2)
 identifier of current host system. hostid(1)
 identify files. ident(1)
 identity. getgid(2)
 identity. getuid(2)
 ID's. setreuid(2)
 ierrno: get system error messages. perror(3F)
 if: conditional statement. csh(1)
 if mail arrives and who it is from. biff(1)
 if, while, :, . . . , break, continue, cd, sh(1)
 ifconfig: configure network interface parameters. ifconfig(8C)
 ifdef'ed lines. undef(1)
 image. abort(3F)
 image file. core(5)
 images of running processes. gcore(1)
 immediate notification. csh(1)
 immune to hangups. csh(1)
 IMP log interpreter. implog(8C)
 IMP logger process. implogd(8C)
 implement shared strings. xstr(1)
 implog: IMP log interpreter. implog(8C)
 implogd: IMP logger process. implogd(8C)
 improved error recovery. eyacc(1)
 including aliases and paths (csh only). which(1)
 in-core state with that on disk. fsync(2)
 incremental dump format. dump(5)
 incremental file system dump. dump(8)
 incremental file system restore. restore(8)
 indent and format C program source. indent(1)
 indent: indent and format C program source. indent(1)
 independent operation routines. tgetent, termcap(3X)
 index. ptx(1)
 index for a bibliography, find references in a lookbib(1)
 index, rindex, lnblink, len: tell about character index(3F)
 index, rindex: string operations. strcat, string(3)
 indicate last logins of users and teletypes. last(1)
 indirect system call. syscall(2)
 individual files. fsplit(1)
 indexbib, lookbib: build inverted index for a lookbib(1)
 inet_addr, inet_network, inet_ntoa, inet_makeaddr, inet(3n)
 inet_addr, inet_network, inet_ntof: Internet address/ inet(3n)
 inet_makeaddr, inet_ntof: Internet inet(3n)
 inet_netof: Internet address manipulation routines. inet(3n)
 inet_network, inet_ntoa, inet_makeaddr, inet_addr, inet(3n)

Internet address/ inet_addr, inet_network,	inet_ntoa, inet_makeaddr, inet_lnaof, inet_netof:	inet(3n)
dumpfs: dump file system	information.	dumpfs(8)
mttys: Multi-Link partition	information.	mttys(5)
pac: printer/plotter accounting	information.	pac(8)
whodos: display	information about dosc users.	whodos(1)
rlog: print log messages and other	information about RCS files.	rlog(1)
getrusage: get	information about resource utilization.	getrusage(2)
vtimes: get	information about resource utilization.	vtimes(3C)
fstab: static	information about the filesystems.	fstab(5)
man: find manual	information by keywords; print out the manual.	man(1)
dosdisk: program to create and display	information for MPS/DOS vdisks.	dosdisk(8)
finger: user	information lookup program.	finger(1)
miscellaneous: miscellaneous useful	information pages.	intro(7)
	init: process control initialization.	init(8)
	initgroups: initialize group access list.	initgroups(3X)
	initialization.	init(8)
ioinit: change /77 I/O	initialization.	ioinit(3F)
tset: terminal dependent	initialization.	tset(1)
ttys: terminal	initialization data.	ttys(5)
initgroups:	initialize group access list.	initgroups(3X)
connect:	initiate a connection on a socket.	connect(2)
popen, pclose:	initiate I/O to/from a process.	popen(3)
generator; routines for changing/ random, srandom,	initstate, setstate: better random number	random(3)
fimin, fimax, frac, dfimin, dfimax, dfrac,	inmax: return extreme values.	fimin(3F)
cli: clear	i-node.	cli(8)
fs,	inode: format of file system volume.	fs(5)
read, readv: read	input.	read(2)
soelim: eliminate .so's from nroff	input.	soelim(1)
scanf, fscanf, sscanf: formatted	input conversion.	scanf(3S)
ungetc: push character back into	input stream.	ungetc(3S)
fread, fwrite: buffered binary	input/output.	fread(3S)
stdio: standard buffered	input/output package.	intro(3S)
error, feof, clearerr, fileno: stream status	inquiries.	error(3S)
refer: find and	insert literature references in documents.	refer(1)
insque, remque:	insert/remove element from a queue.	insque(3)
	insque, remque: insert/remove element from a queue.	insque(3)
	install binaries.	install(1)
install:	install bootloader on disk.	binstl(8)
binstl: program to	install: install binaries.	install(1)
	instruction about UNIX.	learn(1)
learn: computer aided	interact with a psychoanalyst.	doctor(6)
doctor:	interactive display-oriented text editor.	jove(1)
jove: an	interactive repair.	fack(8)
fsck: file system consistency check and	interesting, adage.	fortune(6)
fortune: print a random, hopefully	interface. /erase, label, line, circle, arc, move,	plot(3X)
cont, point, linemod, space, closepl: graphics	interface.	plot(5)
plot: graphics	interface.	tty(4)
tty: general terminal	interface parameters.	ifconfig(8C)
ifconfig: configure network	interface to the TELNET protocol.	telnet(1C)
telnet: user	interfaces.	slattach(8C)
slattach: attach serial lines as network	internet.	sendmail(8)
sendmail: send mail over the	Internet address manipulation routines.	inet(3n)
/inet_ntoa, inet_makeaddr, inet_lnaof, inet_netof:	Internet File Transfer Protocol server.	ftpd(8C)
ftpd: DARPA	interpolate smooth curve.	spline(1G)
spline:	interpreter.	implog(8C)
implog: IMP log	interpreter.	lisp(1)
lisp: lisp	interpreter.	pti(1)
pti: phototypesetter	interpreter.	px(1)
px: Pascal	interpreter and executor.	pix(1)
pix: Pascal	interpreter code translator.	pi(1)
pi: Pascal	interpreter) with C-like syntax.	csh(1)
csh: a shell (command	interprocess communication channel.	pipe(2)
pipe: create an	inter-process communication facilities status.	ipcs(1)
ipes: report	interrupt. sigpause:	sigpause(2)
atomically release blocked signals and wait for	interrupts in command scripts.	csh(1)
onintr: process	introduction to commands.	intro(1)
intro:	introduction to compatibility library functions.	intro(3C)
intro:	introduction to FORTRAN library functions.	intro(3F)
intro:	introduction to library functions.	intro(3)
intro:	introduction to mathematical library functions.	intro(3M)
intro:	introduction to miscellaneous library functions.	intro(3X)
intro:	introduction to network library functions.	intro(3n)
rcsintro:	introduction to RCS commands.	rcsintro(1)
intro:	introduction to system calls and error numbers.	intro(2)
commands. intro:	introduction to system maintenance and operation	intro(8)

ncheck: generate names from
in a bibliography. indxbib, lookbib: build
tread, twrite, trewin, tskipf, ttate: f77 tape
ioinit: change f77
select: synchronous
iostat: report
popen, pclose: initiate

shared memory id.
status.
rand, drand,
isascii:/ isalpha, isupper, islower, isdigit,
isspace, ispunct, isprint, iscntrl, isascii:/
isalnum, isspace, ispunct, isprint, iscntrl,
ttynam,
ttypename,
/isdigit, isalnum, isspace, ispunct, isprint,
iscntrl, isascii:/ isalpha, isupper, islower,
isprint, iscntrl, isascii:/ isalpha, isupper,
/islower, isdigit, isalnum, isspace, ispunct,
/isupper, islower, isdigit, isalnum, isspace,
isalpha, isupper, islower, isdigit, isalnum,
system:
ispunct, isprint, iscntrl, isascii:/ isalpha,
idate,
j0,
j0, j1,
bg: place
fg: bring
jobs: print current
stop: halt a
kill: kill
lprm: remove

jove_recover - recover
TEACHJOVE - learn how to use the
system/editor crash.
msgs: system messages and
kermit:

uxrc: configuration file for
vers: print version number of the
makekey: generate encryption
apropos: locate commands by
man: find manual information by
profile buffers.
kill:

bessel functions: of two
dmem,
linemod, space, closepl:/ plot: openpl, erase,
awk: pattern scanning and processing
bc: arbitrary-precision arithmetic
efl: Extended Fortran
set, shift, times, trap, umask, wait: command
fp: Functional Programming
vgrindfs: vgrind's
order.
frexp,
TEACHJOVE -
leave: remind you when you have to
exit:
index, rindex, lnbk,

i-numbers.
inverted index for a bibliography, find references
I/O. topen, tclose,
I/O initialization.
i/o multiplexing.
I/O statistics.
I/O to/from a process.
ioctl: control device.
ioinit: change f77 I/O initialization.
iostat: report I/O statistics.
ipcrm: remove a message queue, semaphore set or
ipcs: report inter-process communication facilities
irand: return random values.
isalnum, isspace, ispunct, isprint, iscntrl,
isalpha, isupper, islower, isdigit, isalnum,
isascii: character classification macros. /isdigit,
isatty: find name of a terminal port.
isatty, ttyslot: find name of a terminal.
iscntrl, isascii: character classification macros.
isdigit, isalnum, isspace, ispunct, isprint,
islower, isdigit, isalnum, isspace, ispunct,
isprint, iscntrl, isascii: character classification/
ispunct, isprint, iscntrl, isascii: character/
isspace, ispunct, isprint, iscntrl, isascii:/
issue a shell command.
isupper, islower, isdigit, isalnum, isspace,
itime: return date or time in numerical form.
j0, j1, jn, y0, y1, yn: bessel functions.
j1, jn, y0, y1, yn: bessel functions.
jn, y0, y1, yn: bessel functions.
job in background.
job into foreground.
job list.
job or process.
jobs and processes.
jobs from the line printer spooling queue.
jobs: print current job list.
join: relational database operator.
jove: an interactive display-oriented text editor.
JOVE buffers after a system/editor crash.
JOVE editor.
jove_recover - recover JOVE buffers after a
junk mail program.
kermit file transfer.
kermit: kermit file transfer.
kernel.
kernel.
key.
keyword lookup.
keywords; print out the manual.
kgmon: generate a dump of the operating system's
kill jobs and processes.
kill: kill jobs and processes.
kill: send a signal to a process.
kill: send signal to a process.
kill: terminate a process with extreme prejudice.
killpg: send signal to a process group.
kinds for integer orders.
kmem: main memory.
label, line, circle, arc, move, cont, point,
language.
language.
Language.
language. /exit, export, login, read, readonly,
language compiler/interpreter.
language definition data base.
lastcomm: show last commands executed in reverse
ld: link editor.
ldexp, modf: split into mantissa and exponent.
learn: computer aided instruction about UNIX.
learn how to use the JOVE editor.
leave.
leave: remind you when you have to leave.
leave shell.
len: tell about character objects.

ncheck(8)
lookbib(1)
topen(3F)
ioinit(3F)
select(2)
iostat(1)
popen(3)
ioctl(2)
ioinit(3F)
iostat(1)
ipcrm(1)
ipcs(1)
rand(3F)
ctype(3)
ctype(3)
ctype(3)
ttynam(3F)
ttypename(3)
ctype(3)
ctype(3)
ctype(3)
ctype(3)
ctype(3)
system(3)
ctype(3)
idate(3F)
j0(3M)
j0(3M)
j0(3M)
csh(1)
csh(1)
csh(1)
csh(1)
lprm(1)
csh(1)
join(1)
jove(1)
jove_recover(1)
teachjove(1)
jove_recover(1)
msgs(1)
kermit(1)
kermit(1)
uxrc(8)
vers(1)
makekey(8)
apropos(1)
man(1)
kgmon(8)
csh(1)
csh(1)
kill(3F)
kill(2)
kill(1)
killpg(2)
bessel(3F)
mem(4)
plot(3X)
awk(1)
bc(1)
efl(1)
sh(1)
fp(1)
vgrindfs(5)
lastcomm(1)
ld(1)
frexp(3)
learn(1)
teachjove(1)
leave(1)
leave(1)
csh(1)
index(3F)

truncate: truncate a file to a specified length.	truncate(2)
lex: generator of lexical analysis programs.	lex(1)
lexical analysis programs.	lex(1)
lib2648: subroutines for the HP 2648 graphics libraries.	lib2648(3X)
libraries.	ranlib(1)
library.	lorder(1)
(library) file format.	ar(5)
library functions.	intro(3)
library functions.	intro(3C)
library functions.	intro(3F)
library functions.	intro(3M)
library functions.	intro(3n)
library functions.	intro(3X)
library maintainer.	ar(1)
limit: alter per-process resource limitations.	csh(1)
limitations.	csh(1)
limitations.	csh(1)
limits.	quota(1)
line arguments.	getarg(3F)
line, circle, arc, move, cont, point, linemod,	plot(3X)
line feeds.	col(1)
line of a terminal.	sysline(1)
line print.	lpr(1)
line printer control program.	lpc(8)
line printer daemon.	lpd(8)
line printer daemon environment.	cleanlpd(8)
line printer spooling queue.	lprm(1)
linemod, space, closepl: graphics interface.	plot(3X)
lines.	head(1)
lines.	unifdef(1)
lines as network interfaces.	slattach(8C)
lines common to two sorted files.	comm(1)
lines for finite width output device.	fold(1)
lines in a file.	uniq(1)
lines in a sorted list.	look(1)
lines of a file.	rev(1)
link.	readlink(2)
link editor.	ld(1)
link editor output.	a.out(5)
link: make a hard link to a file.	link(2)
link: make a link to an existing file.	link(3F)
link to a file.	link(2)
link to a file.	symlink(2)
link to an existing file.	link(3F)
links.	ln(1)
lint: a C program verifier.	lint(1)
lisp cross reference program.	lxref(1)
lisp interpreter.	lisp(1)
lisp: lisp interpreter.	lisp(1)
Lisp program.	liszt(1)
list.	csh(1)
list.	csh(1)
list.	csh(1)
list.	csh(1)
list.	getgroups(2)
list.	initgroups(3X)
list.	look(1)
list.	nlist(3)
list.	nm(1)
list.	setgroups(2)
list.	symorder(1)
list.	varargs(3)
list contents of directory.	ls(1)
list of MPS/DOS virtual disks.	dosdisks(5)
list of names.	csh(1)
list of users who are on the system.	users(1)
listen: listen for connections on a socket.	listen(2)
listen: listen for connections on a socket.	listen(2)
listings of programs.	vgrind(1)
liszt: compile a Franz Lisp program.	liszt(1)
literature references in documents.	refer(1)
ln: make links.	ln(1)
lnblnk, len: tell about character objects.	index(3F)
load standalone programs.	load(8)
loc: return the address of an object.	loc(3F)
lorder: find ordering relation for an object	
ar: archive	
intro: introduction to	
intro: introduction to compatibility	
intro: introduction to FORTRAN	
intro: introduction to mathematical	
intro: introduction to network	
intro: introduction to miscellaneous	
ar: archive and	
limit: alter per-process resource	
unlimit: remove resource	
quota: display disc usage and	
getarg, larg: return command	
space, closepl:/ plot: openpl, erase, label,	
col: filter reverse	
sysline: display system status on status	
lpr: off	
lpc:	
lpd:	
cleanlpd: clean	
lprm: remove jobs from the	
/erase, label, line, circle, arc, move, cont, point,	
head: give first few	
unifdef: remove ifdef'd	
slattach: attach serial	
comm: select or reject	
fold: fold long	
uniq: report repeated	
look: find	
rev: reverse	
readlink: read value of a symbolic	
ld:	
a.out: assembler and	
link: make a hard	
symlink: make symbolic	
link: make a	
ln: make	
lxref:	
lisp:	
liszt: compile a Franz	
glob: filename expand argument	
history: print history event	
jobs: print current job	
shift: manipulate argument	
getgroups: get group access	
initgroups: initialize group access	
look: find lines in a sorted	
nlist: get entries from name	
nm: print name	
setgroups: set group access	
symorder: rearrange name	
varargs: variable argument	
ls:	
dosdisks:	
foreach: loop over	
users: compact	
listen:	
vgrind: grind nice	
refer: find and insert	
index, rindex,	
load: program to	

and time to ASCII. <i>ctime</i> , (<i>csh</i> only). which:	localtime, gmtime, asctime, timezone: convert date	<i>ctime</i> (3)
apropos:	locate a program file including aliases and paths	which(1)
whereis:	locate commands by keyword lookup.	apropos(1)
end, <i>etext</i> , <i>edata</i> : last	locate source, binary, and or manual for program.	whereis(1)
<i>flock</i> : apply or remove an advisory	locations in program.	end(3)
	lock on an open file.	<i>flock</i> (2)
	lock: reserve a terminal.	lock(1)
	lockf: record locking on files.	lockf(3C)
lockf: record	locking on files.	lockf(3C)
collect system diagnostic messages to form error	log. <i>dmesg</i> :	<i>dmesg</i> (8)
syslog, <i>openlog</i> , <i>closelog</i> : control system	log.	syslog(3)
gamma:	log gamma function.	gamma(3M)
<i>implog</i> : IMP	log interpreter.	<i>implog</i> (8C)
power, square root. <i>exp</i> ,	log, <i>log10</i> , <i>pow</i> , <i>sqrt</i> : exponential, logarithm,	<i>exp</i> (3M)
<i>rlog</i> : print	log messages and other information about RCS files.	<i>rlog</i> (1)
syslog:	log systems messages.	syslog(8)
square root. <i>exp</i> , <i>log</i> ,	<i>log10</i> , <i>pow</i> , <i>sqrt</i> : exponential, logarithm, power,	<i>exp</i> (3M)
<i>exp</i> , <i>log</i> , <i>log10</i> , <i>pow</i> , <i>sqrt</i> : exponential,	logarithm, power, square root.	<i>exp</i> (3M)
<i>rwho</i> : who's	logged in on local machines.	<i>rwho</i> (1C)
<i>implogd</i> : IMP	logger process.	<i>implogd</i> (8C)
<i>flush</i> : flush output to a	logical unit.	<i>flush</i> (3F)
<i>fseek</i> , <i>ftell</i> : reposition a file on a	logical unit.	<i>fseek</i> (3F)
<i>getc</i> , <i>fgetc</i> : get a character from a	logical unit.	<i>getc</i> (3F)
<i>putc</i> , <i>fputc</i> : write a character to a fortran	logical unit.	<i>putc</i> (3F)
<i>rlogin</i> : remote	<i>login</i>	<i>rlogin</i> (1C)
<i>ac</i> :	<i>login</i> accounting.	<i>ac</i> (8)
	<i>login</i> : <i>login</i> new user.	<i>csh</i> (1)
<i>getlog</i> : get user's	<i>login</i> name.	<i>getlog</i> (3F)
<i>getlogin</i> : get	<i>login</i> name.	<i>getlogin</i> (3)
<i>login</i> :	<i>login</i> new user.	<i>csh</i> (1)
<i>passwd</i> : change	<i>login</i> password.	<i>passwd</i> (1)
/break, continue, <i>cd</i> , <i>eval</i> , <i>exec</i> , <i>exit</i> , <i>export</i> ,	<i>login</i> , read, <i>readonly</i> , <i>set</i> , <i>shift</i> , <i>times</i> , <i>trap</i> ,/	<i>sh</i> (1)
<i>utmp</i> , <i>wtmp</i> :	<i>login</i> records.	<i>utmp</i> (5)
<i>rlogind</i> : remote	<i>login</i> server.	<i>rlogind</i> (8C)
<i>chsh</i> : change default	<i>login</i> shell.	<i>chsh</i> (1)
	<i>login</i> : sign on.	<i>login</i> (1)
<i>last</i> : indicate last	logins of users and teletypes.	<i>last</i> (1)
	logout: end session.	<i>csh</i> (1)
<i>setjmp</i> ,	longjmp: non-local goto.	<i>setjmp</i> (3)
find references in a bibliography. <i>indxbib</i> ,	look: find lines in a sorted list.	look(1)
<i>apropos</i> : locate commands by keyword	lookbib: build inverted index for a bibliography,	lookbib(1)
<i>finger</i> : user information	lookup.	apropos(1)
break: exit while/foreach	lookup program.	<i>finger</i> (1)
continue: cycle in	loop.	<i>csh</i> (1)
end: terminate	loop.	<i>csh</i> (1)
foreach:	loop.	<i>csh</i> (1)
library.	loop over list of names.	<i>csh</i> (1)
<i>mklost+found</i> : make a	lorder: find ordering relation for an object	lorder(1)
	lost+found directory for <i>fsck</i>	<i>mklost+found</i> (8)
	<i>lpc</i> : line printer control program.	<i>lpc</i> (8)
	<i>lpd</i> : line printer daemon.	<i>lpd</i> (8)
	<i>lpq</i> : spool queue examination program.	<i>lpq</i> (1)
	<i>lpr</i> : off line print.	<i>lpr</i> (1)
queue.	<i>lprm</i> : remove jobs from the line printer spooling	<i>lprm</i> (1)
	<i>ls</i> : list contents of directory.	<i>ls</i> (1)
bit: and, or, xor, not, <i>rshift</i> ,	<i>lseek</i> : move read/write pointer.	<i>lseek</i> (2)
<i>stat</i> ,	<i>lshift</i> bitwise functions.	bit(3F)
<i>stat</i> ,	<i>lstat</i> , <i>fstat</i> : get file status.	<i>stat</i> (2)
<i>time</i> , <i>ctime</i> ,	<i>lstat</i> , <i>fstat</i> : get file status.	<i>stat</i> (3F)
	<i>ltime</i> , <i>gmtime</i> : return system time.	<i>time</i> (3F)
	<i>bxref</i> : lisp cross reference program.	<i>bxref</i> (1)
	<i>m4</i> : macro processor.	<i>m4</i> (1)
<i>as</i> :	<i>M68020</i> assembler.	<i>as</i> (1)
Standalone mode: definition of this Sanyo/ICON	machine operation mode..	standalone(8)
<i>ruptime</i> : show host status of local	machines.	<i>ruptime</i> (1C)
<i>rwho</i> : who's logged in on local	machines.	<i>rwho</i> (1C)
<i>m4</i> :	macro processor.	<i>m4</i> (1)
alias: shell	macros.	<i>csh</i> (1)
<i>isprint</i> , <i>isctrl</i> , <i>isascii</i> : character classification	macros. / <i>isdigit</i> , <i>isalnum</i> , <i>isspace</i> , <i>ispunct</i> ,	<i>ctype</i> (3)
<i>ms</i> : text formatting	macros.	<i>ms</i> (7)
translate version 6 manual macros to version 7	macros. <i>trman</i> :	<i>trman</i> (1)
<i>me</i> :	macros for formatting papers.	<i>me</i> (7)
<i>man</i> :	macros to typeset manual.	<i>man</i> (7)
<i>trman</i> : translate version 6 manual	macros to version 7 macros.	<i>trman</i> (1)
<i>tecopy</i> : copy a	mag tape.	<i>tecopy</i> (1)

	mt:	magnetic tape manipulating program.	mt(1)
	rmt:	remote magtape protocol module.	rmt(8C)
	mail:	send and receive	mail(1)
encode/decode a binary file for transmission via	mail.	uuencode,uudecode:	uuencode(1C)
xsend, xget, enroll: secret	mail.	xsend(1)
	sendbug:	mail a system bug report to 4bsd-bugs.	sendbug(1)
	mailaddr:	mail addressing description.	mailaddr(7)
newaliases: rebuild the data base for the	mail	aliases file.	newaliases(1)
binmail: send or receive	mail	among users.	binmail(1)
biff: be notified if	mail	arrives and who it is from.	biff(1)
from: who is my	mail	from?.	from(1)
prmail: print out	mail	in the post office.	prmail(1)
tmail: print out	mail	messages, most recent first.	tmail(1)
sendmail: send	mail	over the internet.	sendmail(8)
msgs: system messages and junk	mail	program.	msgs(1)
rmail: handle remote	mail	received via uucp.	rmail(1)
	mail:	send and receive mail.	mail(1)
	mailaddr:	mail addressing description.	mailaddr(7)
	dmem, kmem:	main memory.	mem(4)
	make:	maintain program groups.	make(1)
ar: archive and library	maintainer.	ar(1)
intro: introduction to system	maintenace	and operation commands.	intro(8)
mkdir:	make	a directory.	mkdir(1)
mkdir:	make	a directory file.	mkdir(2)
link:	make	a hard link to a file.	link(2)
link:	make	a link to an existing file.	link(3F)
mklost+found:	make	a lost+found directory for fsck.	mklost+found(8)
mknod:	make	a special file.	mknod(2)
mktemp:	make	a unique file name.	mktemp(3)
ln:	make	links.	ln(1)
	make:	maintain program groups.	make(1)
	symlink:	make symbolic link to a file.	symlink(2)
	makedev:	make system special files.	makedev(8)
	script:	make typescript of terminal session.	script(1)
mkfs: program to	make	UNIX file systems.	mkfs(8)
	makedev:	make system special files.	makedev(8)
	makekey:	generate encryption key.	makekey(8)
allocator.	malloc, free, realloc, calloc, alloca:	memory	malloc(3)
the manual.	man:	find manual information by keywords; print out	man(1)
	man:	macros to typeset manual.	man(7)
	shift:	manipulate argument list.	sh(1)
	quota:	manipulate disk quotas.	quota(2)
	route:	manipulate the routing tables.	route(8C)
	mt:	magnetic tape manipulating program.	mt(1)
inet_1naof, inet_netof: Internet address	manipulation	routines. /inet_ntoa, inet_makeaddr,	inet(3n)
frexp, ldexp, modf: split into	mantissa	and exponent.	frexp(3)
catman: create the cat files for the	manual.	catman(8)
find manual information by keywords; print out the	manual.	man:	man(1)
man: macros to typeset	manual.	man(7)
whereis: locate source, binary, and or	manual	for program.	whereis(1)
manual. man: find	manual	information by keywords; print out the	man(1)
trman: translate version 6	manual	macros to version 7 macros.	trman(1)
route:	manually	manipulate the routing tables.	route(8C)
umask: change or display file creation	mask.	sh(1)
sigsetmask: set current signal	mask.	sigsetmask(2)
umask: set file creation mode	mask.	umask(2)
mkstr: create an error message file by	massaging	C source.	mkstr(1)
intro: introduction to	mathematical	library functions.	intro(3M)
eqn, neqn, checkeq: typeset	mathematics.	eqn(1)
getrlimit, setrlimit: control	maximum	system resource consumption.	getrlimit(2)
vlimit: control	maximum	system resource consumption.	vlimit(3C)
dis: an	mc68020	disassembler.	dis(1)
	me:	macros for formatting papers.	me(7)
	media.	bcd(6)
bcd: convert to antique	memberships.	groups(1)
groups: show group	memory.	mem(4)
dmem, kmem: main	memory	allocator.	malloc(3)
malloc, free, realloc, calloc, alloca:	memory	aligned	valloc(3)
valloc: aligned	memory	control operations.	shmctl(2)
shmctl: shared	memory	efficient way.	vfork(2)
vfork: spawn new process in a virtual	memory	id. ipcrm:	ipcrm(1)
remove a message queue, semaphore set or shared	memory	image.	abort(3F)
abort: terminate abruptly with	memory	image file.	core(5)
core: format of	memory	operations.	shmop(2)
shmop: shared	memory	segment.	shmget(2)
shmget: get shared			

vmstat: report virtual	memory statistics.	vmstat(1)
merge: three-way file	merge.	merge(1)
sort: sort or	merge files.	sort(1)
rcsmerge:	merge RCS revisions.	rcsmerge(1)
	merge: three-way file merge.	merge(1)
pmerge: pascal file	merger.	pmerge(1)
	msg: permit or deny messages.	msg(1)
msgctl:	message control operations.	msgctl(2)
mkstr: create an error	message file by massaging C source.	mkstr(1)
recv, recvfrom, recvmsg: receive a	message from a socket.	recv(2)
send, sendto, sendmsg: send a	message from a socket.	send(2)
msgop:	message operations.	msgop(2)
msgget: get	message queue.	msgget(2)
iperm: remove a	message queue, semaphore set or shared memory id.	iperm(1)
error: analyze and disperse compiler error	messages.	error(1)
msg: permit or deny	messages.	msg(1)
perror, sys_errlist, sys_nerr: system error	messages.	perror(3)
perror, gerror, ierrno: get system error	messages.	perror(3F)
psignal, sys_siglist: system signal	messages.	psignal(3)
syslog: log systems	messages.	syslog(8)
msgs: system	messages and junk mail program.	msg(1)
rlog: print log	messages and other information about RCS files.	rlog(1)
tmail: print out mail	messages, most recent first.	tmail(1)
dmesg: collect system diagnostic	messages to form error log.	dmesg(8)
mille: play	Mille Bournes.	mille(6)
	mille: play Mille Bournes.	mille(6)
intro: introduction to	miscellaneous library functions.	intro(3X)
pages.	miscellaneous: miscellaneous useful information	intro(7)
miscellaneous:	miscellaneous useful information pages.	intro(7)
	mkdir: make a directory.	mkdir(1)
	mkdir: make a directory file.	mkdir(2)
	mkfs: program to make UNIX file systems.	mkfs(8)
	mklost+found: make a lost+found directory for fsck.	mklost+found(8)
	mknod: build special file.	mknod(8)
	mknod: make a special file.	mknod(2)
	mkproto: construct a prototype file system.	mkproto(8)
source.	mkstr: create an error message file by massaging C	mkstr(1)
	mktemp: make a unique file name.	mktemp(3)
chmod: change	mode.	chmod(1)
getty: set terminal	mode.	getty(8)
definition of this Sanyo/ICON machine operation	mode.. Standalone mode:	standalone(8)
operation mode.. Standalone	mode: definition of this Sanyo/ICON machine	standalone(8)
umask: set file creation	mode mask.	umask(2)
chmod: change	mode of a file.	chmod(3F)
chmod: change	mode of file.	chmod(2)
frexp, ldexp,	modf: split into mantissa and exponent.	frexp(3)
touch: update date last	modified of a file.	touch(1)
recovery. eyacc:	modified yacc allowing much improved error	eyacc(1)
rmt: remote magtape protocol	module.	rmt(8C)
what: show what versions of object	modules were used to construct a file.	what(1)
monitor, monstartup,	moncontrol: prepare execution profile.	monitor(3)
profile.	monitor, monstartup, moncontrol: prepare execution	monitor(3)
	monop: Monopoly game.	monop(6)
monop:	Monopoly game.	monop(6)
monitor,	monstartup, moncontrol: prepare execution profile.	monitor(3)
	more, page: file perusal filter for crt viewing.	more(1)
tmail: print out mail messages,	most recent first.	tmail(1)
curses: screen functions with "optimal" cursor	motion.	curses(3X)
mount, umount:	mount and dismount file system.	mount(8)
mount, umount:	mount or remove file system.	mount(2)
	mount, umount: mount and dismount file system.	mount(8)
	mount, umount: mount or remove file system.	mount(2)
mtab:	mounted file system table.	mtab(5)
plot: openpl, erase, label, line, circle, arc,	move, cont, point, linemod, space, closepl:/	plot(3X)
mv:	move or rename files.	mv(1)
lseek:	move read/write pointer.	lseek(2)
doscopyd:	MPS/DOS file copy daemon.	doscopyd(8)
dosprint:	MPS/DOS spooler daemon.	dosprint(8)
program to create and display information for	MPS/DOS vdisks. dosdisk:	dosdisk(8)
dosdisks: list of	MPS/DOS virtual disks.	dosdisks(5)
	ms: text formatting macros.	ms(7)
	msgctl: message control operations.	msgctl(2)
	msgget: get message queue.	msgget(2)
	msgop: message operations.	msgop(2)
	msgs: system messages and junk mail program.	msgs(1)

	mt: magnetic tape manipulating program.	mt(1)
	mtab: mounted file system table.	mtab(5)
	mttys: Multi-Link partition information.	mttys(5)
eyacc: modified yacc allowing	much improved error recovery.	eyacc(1)
mttys:	Multi-Link partition information.	mttys(5)
hunt: a	multi-player multi-terminal game.	hunt(6)
select: synchronous i/o	multiplexing.	select(2)
fsplit: split a	multi-routine Fortran file into individual files.	fsplit(1)
hunt: a multi-player	multi-terminal game.	hunt(6)
switch:	multi-way command branch.	csh(1)
	mv: move or rename files.	mv(1)
	my mail from?	from(1)
from: who is	name.	getdisk(3X)
getdiskbyname: get disk description by its	name.	getenv(3)
getenv: value for environment	name.	getlog(3F)
getlog: get user's login	name.	getlogin(3)
getlogin: get login	name.	getsockname(2)
getsockname: get socket	name.	mktemp(3)
mktemp: make a unique file	name.	pwd(1)
pwd: working directory	name.	tty(1)
tty: get terminal	name data base.	hosts(5)
hosts: host	name data base.	networks(5)
networks: network	name data base.	protocols(5)
protocols: protocol	name data base.	services(5)
services: service	name from uid.	getpw(3C)
getpw: get	name list.	nlist(3)
nlist: get entries from	name list.	nm(1)
nm: print	name of a file.	symorder(1)
symorder: rearrange	name of a terminal.	rename(2)
rename: change the	name of a terminal port.	ttynam(3)
ttynam, isatty, ttyslot: find	name of connected peer.	ttynam(3F)
ttynam, isatty: find	name of current host.	getpeername(2)
getpeername: get	name of current host.	gethostname(2)
gethostname, sethostname: get/set	name of current host system.	hostname(1)
hostname: set or print	name rule file.	dstrules(5)
PP dstrules: Daylight savings time and time zone	name to a socket.	bind(2)
bind: bind a	names.	csh(1)
foreach: loop over list of	names for terminals.	term(7)
term: conventional	names from i-numbers.	ncheck(8)
ncheck: generate	ncheck: generate names from i-numbers.	ncheck(8)
	eqn, checkeq: typeset mathematics.	eqn(1)
	netstat: show network status.	netstat(1)
	network.	rdump(8C)
rdump: file system dump across the	network.	rrestore(8C)
rrestore: restore a file system dump across the	network byte order. htonl, htons,	byteorder(3n)
ntohl, ntohs: convert values between host and	network entry. getnetent, getnetbyaddr,	getnetent(3n)
getnetbyname, setnetent, endnetent: get	network host entry. gethostent, gethostbyaddr,	gethostent(3n)
gethostbyname, sethostent, endhostent: get	network hosts.	ping(8)
ping: send ICMP ECHO_REQUEST packets to	network interface parameters.	ifconfig(8C)
ifconfig: configure	network interfaces.	slattach(8C)
slattach: attach serial lines as	network library functions.	intro(3n)
intro: introduction to	network name data base.	networks(5)
networks:	network routing daemon.	routed(8C)
routed:	network status.	netstat(1)
netstat: show	networks: network name data base.	networks(5)
	new file.	creat(2)
creat: create a	new file. open:	open(2)
open a file for reading or writing, or create a	new file system.	newfs(8)
newfs: construct a	new process.	fork(2)
fork: create a	new process in a virtual memory efficient way.	vfork(2)
vfork: spawn	new user.	csh(1)
login: login	new users.	adduser(8)
adduser: procedure for adding	newaliases: rebuild the data base for the mail	newaliases(1)
aliases file.	newfs: construct a new file system.	newfs(8)
	nextkey: data base subroutines.	dbm(3X)
dbm, fetch, store, delete, firstkey,	NIC format host tables from a host.	gettable(8C)
gettable: get	NIC standard format host tables.	htable(8)
htable: convert	nice listings of programs.	vgrind(1)
vgrind: grind	nice, nohup: run a command at low priority	nice(1)
(sh only).	nice: run low priority process.	csh(1)
	nice: set program priority.	nice(3C)
	nlist: get entries from name list.	nlist(3)
	nm: print name list.	nm(1)
only). nice,	nohup: run a command at low priority (sh	nice(1)

setjmp, longjmp: bit: and, or, xor
 notify: request immediate
 biff: be
 soelim: eliminate .so's from
 tbl: format tables for
 colcrt: filter
 troff,
 deroff: remove
 checknr: check
 network byte order. htonl, htons,
 order. htonl, htons, ntohl,
 phones: remote host phone
 arithmetic: provide drill in
 rand, srand: random
 random, srandom, initstate, setstate: better random
 vers: print version
 atof, atoi, atol: convert ASCII to
 intro: introduction to system calls and error
 number: convert Arabic
 idate, itime: return date or time in
 loc: return the address of an
 long, short: integer
 size: size of an
 lorder: find ordering relation for an
 what: show what versions of
 strings: find the printable strings in a
 index, rindex, lnbink, len: tell about character
 od:
 prmail: print out mail in the post
 nohup: run a command at low priority (*sh*
 program file including aliases and paths (*cs*
 file. open:
 fopen, freopen, fdopen:
 flock: apply or remove an advisory lock on an
 a new file.
 closedir: directory operations.
 closedir: directory operations. directory:
 syslog,
 cont, point, linemod, space, closepl:/ plot:
 kgmon: generate a dump of the
 intro: introduction to system maintenance and
 mode: definition of this Sanyo/ICON machine
 tgetstr, tgoto, tputs: terminal independent
 bcomp, bzero, bzero, fls: bit and byte string
 telldir, seekdir, rewinddir, closedir: directory
 telldir, seekdir, rewinddir, closedir: directory
 msgctl: message control
 msgop: message
 semctl: semaphore control
 semop: semaphore
 shmctl: shared memory control
 shmop: shared memory
 strepy, strncpy, strlen, index, rindex: string
 join: relational database
 curses: screen functions with
 fentl: file control
 stty: set terminal
 getsockopt, setsockopt: get and set
 ntohs: convert values between host and network byte
 lastcomm: show last commands executed in reverse
 tac: concatenate and print files in reverse
 lorder: find
 bessel functions: of two kinds for integer
 vi: screen
 a.out: assembler and link editor
 terminate a process after flushing any pending
 ecvt, fcvt, gcvt:
 nohup: run command immune to hangups. csh(1)
 non-local goto. setjmp(3)
 not, rshift, lshift bitwise functions. bit(3F)
 notification. csh(1)
 notified if mail arrives and who it is from. biff(1)
 notify: request immediate notification. csh(1)
 nroff input. soelim(1)
 nroff or troff. tbl(1)
 nroff output for CRT previewing. colcrt(1)
 nroff: text formatting. nroff(1)
 nroff: text formatting and typesetting. troff(1)
 nroff, troff, tbl and eqn constructs. deroff(1)
 nroff/troff files. checknr(1)
 ntohl, ntohs: convert values between host and byteorder(3n)
 ntohs: convert values between host and network byte byteorder(3n)
 null: data sink. null(4)
 number: convert Arabic numerals to English. number(6)
 number data base. phones(5)
 number facts. arithmetic(6)
 number generator. rand(3C)
 number generator; routines for changing generators. random(3)
 number of the kernel. vers(1)
 numbers. atof(3)
 numbers. intro(2)
 numerals to English. number(6)
 numerical form. idate(3F)
 object. loc(3F)
 object conversion. long(3F)
 object file. size(1)
 object library. lorder(1)
 object modules were used to construct a file. what(1)
 object, or other binary, file. strings(1)
 objects. index(3F)
 octal, decimal, hex, ascii dump. od(1)
 od: octal, decimal, hex, ascii dump. od(1)
 office. prmail(1)
 onintr: process interrupts in command scripts. csh(1)
 only). nice, nice(1)
 only). which: locate a which(1)
 open a file for reading or writing, or create a new open(2)
 open a stream. fopen(3S)
 open file. flock(2)
 open: open a file for reading or writing, or create open(2)
 opendir, readdir, telldir, seekdir, rewinddir, directory(3)
 opendir, readdir, telldir, seekdir, rewinddir, directory(3X)
 openlog, closelog: control system log. syslog(3)
 openpl, erase, label, line, circle, arc, move, plot(3X)
 operating system's profile buffers. kgmon(8)
 operation commands. intro(8)
 operation mode.. Standalone standalone(8)
 operation routines. tgetent, tgetnum, tgetflag, termcap(3X)
 operations. bstring(3)
 operations. opendir, readdir, directory(3)
 operations. directory: opendir, readdir, directory(3X)
 operations. msgctl(2)
 operations. msgop(2)
 operations. semctl(2)
 operations. semop(2)
 operations. shmctl(2)
 operations. shmop(2)
 operations. strcat, strncat, strcmp, strncmp, string(3)
 operator. join(1)
 "optimal" cursor motion. curses(3X)
 options. fentl(5)
 options. stty(1)
 options on sockets. getsockopt(2)
 order. htonl, htons, ntohl, byteorder(3n)
 order. lastcomm(1)
 order. tac(1)
 ordering relation for an object library. lorder(1)
 orders. bessel(3F)
 oriented (visual) display editor based on ex. vi(1)
 output. a.out(5)
 output. exit: exit(3)
 output conversion. ecvt(3)

printf, sprintf, sprintf: formatted output conversion.	printf(3S)
fold: fold long lines for finite width output device.	fold(1)
colcrt: filter nroff output for CRT previewing.	colcrt(1)
dosprinters: destinations for spooled output from SLPT printers.	dosprinters(5)
flush: flush output to a logical unit.	flush(3F)
foreach: loop over list of names.	cash(1)
sendmail: send mail over the internet.	sendmail(8)
trapov: trap and repair floating point overflow.	trapov(3F)
exec: overlay shell with specified command.	cash(1)
chown: change owner.	chown(8)
chown: change owner and group of a file.	chown(2)
quot: summarize file system ownership.	quot(8)
pac: printer/plotter accounting information.	pac(8)
packets to network hosts.	ping(8)
page: file perusal filter for crt viewing.	more(1)
page size.	getpagesize(2)
page size.	pagesize(1)
pages.	intro(7)
pagesize: print system page size.	pagesize(1)
paginator for the Tektronix 4014.	tk(1)
paging and swapping.	swapon(8)
pair of connected sockets.	socketpair(2)
papers.	me(7)
parameters.	ifconfig(8C)
park: program to park the hard disk heads.	park(8)
park the hard disk heads.	park(8)
partition information.	mttys(5)
Pascal compiler.	pc(1)
Pascal cross-reference program.	pxref(1)
pascal debugger.	pdx(1)
Pascal execution profiler.	pxp(1)
pascal file merger.	pmerge(1)
Pascal interpreter.	px(1)
Pascal interpreter and executor.	pix(1)
Pascal interpreter code translator.	pi(1)
passwd: change login password.	passwd(1)
passwd: password file.	passwd(5)
password.	getpass(3)
password.	passwd(1)
password file.	passwd(5)
password file.	vipw(8)
password file entry. getpwent.	getpwent(3)
pathname.	getwd(3)
pathname of current working directory.	getcwd(3F)
paths (csh only).	which(1)
pattern.	grep(1)
pattern scanning and processing language.	awk(1)
pause: stop until signal.	pause(3C)
pc: Pascal compiler.	pc(1)
pclose: initiate I/O to/from a process.	popen(3)
pdx: pascal debugger.	pdx(1)
peer.	getpeername(2)
pending output.	exit(3)
permit or deny messages.	mesg(1)
permuted index.	ptx(1)
per-process resource limitations.	csh(1)
perrot, gerror, ierrno: get system error messages.	perrot(3F)
perrot, sys_errlist, sys_nerr: system error persistent text.	perrot(3)
perusal filter for crt viewing.	sticky(8)
phone number data base.	more(1)
phones: remote host phone number data base.	phones(5)
phones: remote host phone number data base.	phones(5)
phototypesetter interpreter.	pti(1)
tc: phototypesetter simulator.	tc(1)
pi: Pascal interpreter code translator.	pi(1)
hosts. ping: send ICMP ECHO_REQUEST packets to network	ping(8)
pipe: create an interprocess communication channel.	pipe(2)
tee: pipe fitting.	tee(1)
pix: Pascal interpreter and executor.	pix(1)
bg: place job in background.	csh(1)
fish: play "Go Fish".	fish(6)
mille: play Mille Bournes.	mille(6)
boggle: play the game of boggle.	boggle(6)
worm: Play the growing worm game.	worm(6)
plot: graphics filters.	plot(1G)

move, cont, point, linemod, space, closepl:/	plot: graphics interface.	plot(5)
fpu: determine presence of the floating	plot: openpl, erase, label, line, circle, arc,	plot(3X)
trpfpe, fpeent: trap and repair floating	pmerge: pascal file merger.	pmerge(1)
/erase, label, line, circle, arc, move, cont,	point coprocessor.	fpu(1)
trapov: trap and repair floating	point faults.	trpfpe(3F)
lseek: move read/write	point, linemod, space, closepl: graphics interface.	plot(3X)
popd:	point overflow.	trapov(3F)
ttynam, isatty: find name of a terminal	pointer.	lseek(2)
ttytype: data base of terminal types by	pop shell directory stack.	csh(1)
prmail: print out mail in the	popd: pop shell directory stack.	csh(1)
root. exp, log, log10,	popen, pclose: initiate I/O to/from a process.	popen(3)
exp, log, log10, pow, sqrt: exponential, logarithm,	port.	ttynam(3F)
name rule file..	port.	ttytype(5)
print:	post office.	prmail(1)
prec: C	pow, sqrt: exponential, logarithm, power, square	exp(3M)
monitor, monstartup, moncontrol:	power, square root.	exp(3M)
fpu: determine	PP dstrules: Daylight savings time and time zone	dstrules(5)
colert: filter nroff output for CRT	pr: print file.	pr(1)
types:	pr to the printer.	print(1)
cat: catenate and	prec: C precedence chart.	prec(7)
lpr: off line	precedence chart.	prec(7)
fortune:	prepare execution profile.	monitor(3)
date:	presence of the floating point coprocessor.	fpu(1)
cal:	previewing.	colcrt(1)
hashstat:	primitive system data types.	types(5)
jobs:	print.	cat(1)
whoami:	print.	lpr(1)
pr:	print a random, hopefully interesting, adage.	fortune(6)
tac: concatenate and	date:	date(1)
fpr:	cal:	cal(1)
history:	hashstat:	csh(1)
hostid: set or	jobs:	csh(1)
banner:	whoami:	whoami(1)
files. rlog:	pr:	pr(1)
nm:	tac: concatenate and	tac(1)
hostname: set or	fpr:	fpr(1)
prmail:	history:	csh(1)
tmail:	hostid: set or	hostid(1)
printenv:	banner:	banner(6)
man: find manual information by keywords;	files. rlog:	rlog(1)
pstat:	nm:	nm(1)
pagesize:	hostname: set or	hostname(1)
vers:	prmail:	prmail(1)
diction,explain:	tmail:	tmail(1)
explain, diction-	printenv:	printenv(1)
file. strings: find the	man: find manual information by keywords;	man(1)
banner: print large banner on	pstat:	print(1)
print: pr to the	pagesize:	pstat(8)
printcap:	vers:	pagesize(1)
lpc: line	diction,explain:	vers(1)
lpd: line	explain, diction-	diction(1)
cleanlpd: clean line	file. strings: find the	explain(1)
lprm: remove jobs from the line	banner: print large banner on	strings(1)
pac:	print: pr to the	printcap(5)
destinations for spooled output from SLPT	printcap:	printenv(1)
conversion.	lpc: line	banner(6)
setpriority: get/set program scheduling	lpd: line	print(1)
nice: set program	cleanlpd: clean line	printcap(5)
nice, nohup: run a command at low	lprm: remove jobs from the line	lpc(8)
renice: alter	pac:	lpd(8)
nice: run low	destinations for spooled output from SLPT	cleanlpd(8)
dosc: connect to	conversion.	lprm(1)
adduser:	setpriority: get/set program scheduling	pac(8)
reboot: UNDX bootstrapping	nice: set program	dosprinters(5)
nice: run low priority	nice, nohup: run a command at low	printf(3S)
	renice: alter	getpriority(2)
	nice: run low	nice(3C)
	dosc: connect to	nice(1)
	adduser:	renice(8)
	reboot: UNDX bootstrapping	csh(1)
	nice: run low priority	prmail(1)
		dosc(1)
		adduser(8)
		reboot(8)
		csh(1)

stop: halt a job or	process.	cs(1)
_exit: terminate a	process.	exit(2)
fork: create a new	process.	fork(2)
fork: create a copy of this	process.	fork(3F)
implogd: IMP logger	process.	implogd(8C)
kill: send signal to a	process.	kill(2)
kill: send a signal to a	process.	kill(3F)
popen, pclose: initiate I/O to/from a	process.	popen(3)
wait: await completion of	process.	wait(1)
exit: terminate a	process after flushing any pending output.	exit(3)
init:	process control initialization.	init(8)
getpgrp: get	process group.	getpgrp(2)
killpg: send signal to a	process group.	killpg(2)
setpgrp: set	process group.	setpgrp(2)
getpid: get	process id.	getpid(3F)
getpid, getppid: get	process identification.	getpid(2)
vfork: spawn new	process in a virtual memory efficient way.	vfork(2)
onintr:	process interrupts in command scripts.	cs(1)
ps:	process status.	ps(1)
times: get	process times.	times(3C)
wait, wait3: wait for	process to terminate.	wait(2)
wait: wait for a	process to terminate.	wait(3F)
ptrace:	process trace.	ptrace(2)
kill: terminate a	process with extreme prejudice.	kill(1)
exit: terminate	process with status.	exit(3F)
kill: kill jobs and	processes.	cs(1)
gcore: get core images of running	processes.	gcore(1)
renice: alter priority of running	processes.	renice(8)
wait: wait for background	processes to complete.	cs(1)
awk: pattern scanning and	processing language.	awk(1)
halt: stop the	processor.	halt(8)
m4: macro	processor.	m4(1)
reboot: reboot system or halt	processor.	reboot(2)
	prof: display profile data.	prof(1)
	profil: execution time profile.	profil(2)
monitor, monstartup, moncontrol: prepare execution	profile.	monitor(3)
profil: execution time	profile.	profil(2)
kgmon: generate a dump of the operating system's	profile buffers.	kgmon(8)
gprof: display call graph	profile data.	gprof(1)
prof: display	profile data.	prof(1)
pxp: Pascal execution	profiler.	pxp(1)
copy: standalone copy	program.	copy(8)
end, etext, edata: last locations in	program.	end(3)
finger: user information lookup	program.	finger(1)
ftp: file transfer	program.	ftp(1C)
liszt: compile a Franz Lisp	program.	liszt(1)
lpc: line printer control	program.	lpc(8)
lpq: spool queue examination	program.	lpq(1)
lxref: lisp cross reference	program.	lxref(1)
msgs: system messages and junk mail	program.	msgs(1)
mt: magnetic tape manipulating	program.	mt(1)
pxref: Pascal cross-reference	program.	pxref(1)
rdist: remote file distribution	program.	rdist(1)
tftp: trivial file transfer	program.	tftp(1C)
units: conversion	program.	units(1)
whereis: locate source, binary, and or manual for	program.	whereis(1)
cb: C	program beautifier.	cb(1)
only). which: locate a	program file including aliases and paths (cs)	which(1)
make: maintain	program groups.	make(1)
nice: set	program priority.	nice(3C)
getpriority, setpriority: get/set	program scheduling priority.	getpriority(2)
indent: indent and format C	program source.	indent(1)
MPS/DOS vdisks. dosdisk:	program to create and display information for	dosdisk(8)
binstl:	program to install bootloader on disk.	binstl(8)
bload:	program to load standalone programs.	bload(8)
mkfs:	program to make UNIX file systems.	mkfs(8)
park:	program to park the hard disk heads.	park(8)
assert:	program verification.	assert(3X)
lint: a C	program verifier.	lint(1)
fp: Functional	Programming language compiler/interpreter.	fp(1)
bload: program to load standalone	programs.	bload(8)
lex: generator of lexical analysis	programs.	lex(1)
struct: structure Fortran	programs.	struct(1)
vgrind: grind nice listings of	programs.	vgrind(1)
xstr: extract strings from C	programs to implement shared strings.	xstr(1)

telnet: user interface to the TELNET	protocol.	telnet(1C)
getprotobyname, setprotoent, endprotoent: get	protocol entry. getprotoent, getprotobyname,	getprotoent(3n)
rmt: remote magtape	protocol module.	rmt(8C)
protocols:	protocol name data base.	protocols(5)
ftpd: DARPA Internet File Transfer	Protocol server.	ftpd(8C)
telnetd: DARPA TELNET	protocol server.	telnetd(8C)
tftpd: DARPA Trivial File Transfer	Protocol server.	tftpd(8C)
trpt: transliterate	protocol trace.	trpt(8C)
	protocols: protocol name data base.	protocols(5)
mkproto: construct a	prototype file system.	mkproto(8)
arithmetic:	provide drill in number facts.	arithmetic(6)
false, true:	provide truth values.	false(1)
true, false:	provide truth values.	true(1)
	ps: process status.	ps(1)
	psignal, sys_siglist: system signal messages.	psignal(3)
	pstat: print system facts.	pstat(8)
doctor: interact with a	psychoanalyst.	doctor(6)
	pti: phototypesetter interpreter.	pti(1)
	ptrace: process trace.	ptrace(2)
	ptx: permuted index.	ptx(1)
ungetc:	push character back into input stream.	ungetc(3S)
pushd:	push shell directory stack.	pushd(1)
	pushd: push shell directory stack.	pushd(1)
	put a string on a stream.	puta(3S)
puts, fputs:	put character or word on a stream.	putc(3S)
putc, putchar, fputc, putw:	putc, fputc: write a character to a fortran logical	putc(3F)
unit.	putc, putchar, fputc, putw: put character or word	putc(3S)
on a stream.	putchar, fputc, putw: put character or word on a	putc(3S)
stream. putc,	puts, fputs: put a string on a stream.	puts(3S)
	putw: put character or word on a stream.	putc(3S)
putc, putchar, fputc,	pwd: working directory name.	pwd(1)
	px: Pascal interpreter.	px(1)
	pxp: Pascal execution profiler.	pxp(1)
	pxref: Pascal cross-reference program.	pxref(1)
	qsort: quick sort.	qsort(3F)
	qsort: quicker sort.	qsort(3)
	queue.	insque(3)
	queue.	lprm(1)
	queue.	msgget(2)
	queue examination program.	lpq(1)
	queue, semaphore set or shared memory id.	ipcrm(1)
	quick sort.	qsort(3F)
	quicker sort.	qsort(3)
	quiz: test your knowledge.	quiz(6)
	quot: summarize file system ownership.	quot(8)
quotacheck: file system	quota consistency checker.	quotacheck(8)
	quota: display disc usage and limits.	quota(1)
	quota: manipulate disk quotas.	quota(2)
	quotacheck: file system quota consistency checker.	quotacheck(8)
	quotaon: turn file system quotas on and off.	quotaon(8)
	quotaon, quotaoff: turn file system quotas on and	quotaon(8)
	quotas.	edquota(8)
	quotas.	quota(2)
	quotas for a file system.	repquota(8)
	quotas on a file system.	setquota(2)
	quotas on and off.	quotaon(8)
	rain: animated raindrops display.	rain(6)
	raindrops display.	rain(6)
	rand, drand, irand: return random values.	rand(3F)
	rand, srand: random number generator.	rand(3C)
	random, hopefully interesting, adage.	fortune(6)
	random libraries.	ranlib(1)
	random number generator.	rand(3C)
	random number generator; routines for changing/	random(3)
	random, srand, initstate, setstate: better random	random(3)
	random values.	rand(3F)
	ranlib: convert archives to random libraries.	ranlib(1)
	ratfor: rational Fortran dialect.	ratfor(1)
	ratfor: rational Fortran dialect.	ratfor(1)
	rc: command script for auto-reboot and daemons.	rc(8)
stream to a remote command.	rcmd, rresvport, ruserok: routines for returning a	rcmd(3X)
	rep: remote file copy.	rcp(1C)
	rcs: change RCS file attributes.	rcs(1)
	RCS commands.	rcsintro(1)
resintro: introduction to	RCS file.	rcsfile(5)
rcsfile: format of		

rcs: change	RCS file attributes.	rcs(1)
scstores: build	RCS file from SCCS file.	scstores(1)
scstores: build	RCS file from SCCS file.	scstores(8)
print log messages and other information about	RCS files. rlog:	rlog(1)
ci: check in	RCS revisions.	ci(1)
co: check out	RCS revisions.	co(1)
readiff: compare	RCS revisions.	readiff(1)
rcsmerge: merge	RCS revisions.	rcsmerge(1)
	readiff: compare RCS revisions.	readiff(1)
	rcsfile: format of RCS file.	rcsfile(5)
	rcsintro: introduction to RCS commands.	rcsintro(1)
	rcsmerge: merge RCS revisions.	rcsmerge(1)
	rdist: remote file distribution program.	rdist(1)
	rdump: file system dump across the network.	rdump(8C)
getpass:	read a password.	getpass(3)
source:	read commands from file.	cs(1)
read, readv:	read input.	read(2)
/continue, cd, eval, exec, exit, export, login,	read, readonly, set, shift, times, trap, umask,/	sh(1)
	read, readv: read input.	read(2)
readlink:	read value of a symbolic link.	readlink(2)
directory operations. opendir,	readdir, telldir, seekdir, rewinddir, closedir:	directory(3)
directory operations. directory: opendir,	readdir, telldir, seekdir, rewinddir, closedir:	directory(3X)
open: open a file for	reading or writing, or create a new file.	open(2)
	readlink: read value of a symbolic link.	readlink(2)
command/ /cd, eval, exec, exit, export, login, read,	readonly, set, shift, times, trap, umask, wait:	sh(1)
	readv: read input.	read(2)
	lseek: move	lseek(2)
	setregid: set	setregid(2)
	real and effective user ID's.	setreuid(2)
	malloc, free,	malloc(3)
	symorder:	symorder(1)
	reboot: reboot system or halt processor.	reboot(2)
	reboot: reboot system or halt processor.	reboot(2)
	reboot: UNIX bootstrapping procedures.	reboot(8)
fastboot, fasthalt:	reboot/halt the system without checking the disks.	fastboot(8)
newaliases:	rebuild the data base for the mail aliases file.	newaliases(1)
recv, recvfrom, recvmsg:	receive a message from a socket.	recv(2)
mail: send and	receive mail.	mail(1)
binmail: send or	receive mail among users.	binmail(1)
rmail: handle remote mail	received via uuap.	rmail(1)
tmail: print out mail messages, most	recent first.	tmail(1)
	re_comp, re_exec: regular expression handler.	regex(3)
	rehash: recompute command hash table.	cs(1)
	lockf: record locking on files.	lockf(3C)
utmp, wtmp: login	records.	utmp(5)
jove_recover -	recover JOVE buffers after a system/editor crash.	jove_recover(1)
eyacc: modified yacc allowing much improved error	recovery.	eyacc(1)
socket.	recv, recvfrom, recvmsg: receive a message from a	recv(2)
recv,	recvfrom, recvmsg: receive a message from a socket.	recv(2)
recv, recvfrom,	recvmsg: receive a message from a socket.	recv(2)
eval:	re-evaluate shell data.	cs(1)
re_comp,	re_exec: regular expression handler.	regex(3)
documents.	refer: find and insert literature references in	refer(1)
lxref: lisp cross	reference program.	lxref(1)
build inverted index for a bibliography, find	references in a bibliography. indxib, lookbib:	lookbib(1)
refer: find and insert literature	references in documents.	refer(1)
re_comp, re_exec:	regular expression handler.	regex(3)
	rehash: recompute command hash table.	cs(1)
comm: select or	reject lines common to two sorted files.	comm(1)
lorder: find ordering	relation for an object library.	lorder(1)
join:	relational database operator.	join(1)
sigpause: atomically	release blocked signals and wait for interrupt.	sigpause(2)
strip: remove symbols and	relocation bits.	strip(1)
leave:	remind you when you have to leave.	leave(1)
calendar:	reminder service.	calendar(1)
ruserok: routines for returning a stream to a	remote command. rcmd, rresvport,	rcmd(3X)
rexec: return stream to a	remote command.	rexec(3X)
rexecc:	remote execution server.	rexecc(8C)
rcp:	remote file copy.	rcp(1C)
rdist:	remote file distribution program.	rdist(1)
uuserok: send a file to a	remote host.	uuserok(1C)
remote:	remote host description file.	remote(5)
phones:	remote host phone number data base.	phones(5)
rlogin:	remote login.	rlogin(1C)
rlogind:	remote login server.	rlogind(8C)

rmt:	remote magtape protocol module.	rmt(8C)
rmail: handle	remote mail received via uucp.	rmail(1)
	remote: remote host description file.	remote(5)
rsh:	remote shell.	rsh(1C)
rshd:	remote shell server.	rshd(8C)
tip, cu: connect to a	remote system.	tip(1C)
talkd:	remote user communication server.	talkd(8C)
unlink:	remove a directory entry.	unlink(3F)
rmdir:	remove a directory file.	rmdir(2)
memory id. ipcrm:	remove a message queue, semaphore set or shared	ipcrm(1)
unalias:	remove aliases.	csh(1)
flock: apply or	remove an advisory lock on an open file.	flock(2)
colrm:	remove columns from a file.	colrm(1)
unlink:	remove directory entry.	unlink(2)
unsetenv:	remove environment variables.	csh(1)
mount, umount: mount or	remove file system.	mount(2)
unifdef:	remove ifdef'ed lines.	unifdef(1)
lprm:	remove jobs from the line printer spooling queue.	lprm(1)
deroff:	remove rroff, troff, tbl and eqn constructs.	deroff(1)
unlimit:	remove resource limitations.	csh(1)
strip:	remove symbols and relocation bits.	strip(1)
rmdir, rm:	remove (unlink) directories or files.	rmdir(1)
rm, rmdir:	remove (unlink) files or directories.	rm(1)
insque,	remque: insert/remove element from a queue.	insque(3)
rename:	rename a file.	rename(3F)
	rename: change the name of a file.	rename(2)
mv: move or	rename files.	mv(1)
	rename: rename a file.	rename(3F)
	renice: alter priority of running processes.	renice(8)
fscck: file system consistency check and interactive	repair.	fscck(8)
trpfp, fpcnt: trap and	repair floating point faults.	trpfp(3F)
trapov: trap and	repair floating point overflow.	trapov(3F)
while:	repeat commands conditionally.	csh(1)
	repeat: execute command repeatedly.	csh(1)
uniq: report	repeated lines in a file.	uniq(1)
repeat: execute command	repeatedly.	csh(1)
yes: be	repetitively affirmative.	yes(1)
status. ipc:	report inter-process communication facilities	ipc(1)
iostat:	report I/O statistics.	iostat(1)
uniq:	report repeated lines in a file.	uniq(1)
sendbug: mail a system bug	report to 4bsd-bugs.	sendbug(1)
vmstat:	report virtual memory statistics.	vmstat(1)
bugfiler: file bug	reports in folders automatically.	bugfiler(8)
fseek, ftell:	reposition a file on a logical unit.	fseek(3F)
fseek, ftell, rewind:	reposition a stream.	fseek(3S)
	repquota: summarize quotas for a file system.	repquota(8)
notify:	request immediate notification.	csh(1)
lock:	reserve a terminal.	lock(1)
	reset: reset the teletype bits to a sensible state.	reset(1)
reset:	reset the teletype bits to a sensible state.	reset(1)
getrlimit, setrlimit: control maximum system	resource consumption.	getrlimit(2)
vlimit: control maximum system	resource consumption.	vlimit(3C)
limit: alter per-process	resource limitations.	csh(1)
unlimit: remove	resource limitations.	csh(1)
getrusage: get information about	resource utilization.	getrusage(2)
vtimes: get information about	resource utilization.	vtimes(3C)
restore: incremental file system	restore.	restore(8)
rrestore:	restore a file system dump across the network.	rrestore(8C)
	restore: incremental file system restore.	restore(8)
suspend: suspend a shell,	resuming its superior.	csh(1)
getarg, iargc:	return command line arguments.	getarg(3F)
fdate:	return date and time in an ASCII string.	fdate(3F)
idate, itime:	return date or time in numerical form.	idate(3F)
etime, dtime:	return elapsed execution time.	etime(3F)
fmin, fmax, frac, dfmin, dfmax, inmax:	return extreme values.	fmin(3F)
rand, drand, irand:	return random values.	rand(3F)
rexec:	return stream to a remote command.	rexec(3X)
time, ctime, ltime, gmtime:	return system time.	time(3F)
loc:	return the address of an object.	loc(3F)
rcmd, rresvport, ruserok: routines for	returning a stream to a remote command.	rcmd(3X)
	rev: reverse lines of a file.	rev(1)
col: filter	reverse line feeds.	col(1)
rev:	reverse lines of a file.	rev(1)
lastcomm: show last commands executed in	reverse order.	lastcomm(1)
tac: concatenate and print files in	reverse order.	tac(1)

ci: check in RCS	revisions.	ci(1)
co: check out RCS	revisions.	co(1)
rcsdiff: compare RCS	revisions.	rcsdiff(1)
rcsmerge: merge RCS	revisions.	rcsmerge(1)
fseek, ftell,	rewind: reposition a stream.	fseek(3S)
opendir, readdir, telldir, seekdir,	rewinddir, closedir: directory operations.	directory(3)
directory: opendir, readdir, telldir, seekdir,	rewinddir, closedir: directory operations.	directory(3X)
	rexec: return stream to a remote command.	rexec(3X)
	rexecd: remote execution server.	rexecd(8C)
	rindex, lnblink, len: tell about character objects.	index(3F)
	rindex: string operations. strcat, strncat,	string(3)
strem, strcmp, strepy, strncpy, strlen, index,	rlog: print log messages and other information	rlog(1)
about RCS files.	rlogin: remote login.	rlogin(1C)
	rlogind: remote login server.	rlogind(8C)
	rm: remove (unlink) directories or files.	rmdir(1)
	rm, rmdir: remove (unlink) files or directories.	rm(1)
	rmail: handle remote mail received via uucp.	rmail(1)
	rmdir: remove a directory file.	rmdir(2)
	rm, rmdir: remove (unlink) files or directories.	rm(1)
	rmdir, rm: remove (unlink) directories or files.	rmdir(1)
	rmt: remote magtape protocol module.	rmt(8C)
robots: fight off villainous	robots.	robots(6)
	robots: fight off villainous robots.	robots(6)
	roffbib: run off bibliographic database.	roffbib(1)
	rogue: Exploring The Dungeons of Doom.	rogue(6)
pow, sqrt: exponential, logarithm, power, square	root. exp, log, log10.	exp(3M)
chroot: change	root directory.	chroot(2)
	route: manually manipulate the routing tables.	route(8C)
	routed: network routing daemon.	routed(8C)
	routines. /inet_ntoa, inet_makeaddr, inet_lnaof,	inet(3n)
inet_netof: Internet address manipulation	routines. tgetent, tgetnum, tgetflag, tgetstr,	termcap(3X)
tgoto, tputs: terminal independent operation	routines for changing generators. /initstate,	random(3)
setstate: better random number generator;	routines for returning a stream to a remote	rcmd(3X)
command. rcmd, rresvport, ruserok:	routing daemon.	routed(8C)
routed: network	routing tables.	route(8C)
route: manually manipulate the	rrestore: restore a file system dump across the	rrestore(8C)
network.	rresvport, ruserok: routines for returning a stream	rcmd(3X)
to a remote command. rcmd,	rsh: remote shell.	rsh(1C)
	rshd: remote shell server.	rshd(8C)
	rshift, lshift bitwise functions.	bit(3F)
bit: and, or, xor, not,	rule file.. .PP	dstrules(5)
dstrules: Daylight savings time and time zone name	run a command at low priority (sh only).	nice(1)
nice, nohup:	run command immune to hangups.	csh(1)
nohup:	run low priority process.	csh(1)
nice:	run off bibliographic database.	roffbib(1)
roffbib:	running processes.	gcore(1)
gcore: get core images of	running processes.	renice(8)
renice: alter priority of	ruptime: show host status of local machines.	ruptime(1C)
	ruserok: routines for returning a stream to a	rcmd(3X)
remote command. rcmd, rresvport,	rwho: who's logged in on local machines.	rwho(1C)
	rwhod: system status server.	rwhod(8C)
	sa, accton: system accounting.	sa(8)
	Sanyo/ICON machine operation mode..	standalone(8)
Standalone mode: definition of this	savings time and time zone name rule file..	dstrules(5)
.PP dstrules: Daylight	sbrk: change data segment size.	brk(2)
brk:	scan a directory.	scandir(3)
scandir:	scandir: scan a directory.	scandir(3)
	scanf, fscanf, sscanf: formatted input conversion.	scanf(3S)
	scanning and processing language.	awk(1)
awk: pattern	SCCS file.	scstorcs(1)
scstorcs: build RCS file from	SCCS file.	scstorcs(8)
scstorcs: build RCS file from	scstorcs: build RCS file from SCCS file.	scstorcs(1)
	scstorcs: build RCS file from SCCS file.	scstorcs(8)
	alarm: schedule signal after specified time.	alarm(3C)
alarm:	scheduling priority.	getpriority(2)
getpriority, setpriority: get/set program	screen.	clear(1)
clear: clear terminal	screen functions with "optimal" cursor motion.	curses(3X)
curses:	screen oriented (visual) display editor based on	vi(1)
ex. vi:	script for auto-reboot and daemons.	rc(8)
rc: command	script: make typescript of terminal session.	script(1)
	scripts.	csh(1)
onintr: process interrupts in command	search a file for a pattern.	grep(1)
grep, egrep, fgrep:	secret mail.	xsend(1)
xsend, xget, enroll:	sed: stream editor.	sed(1)
	seekdir, rewinddir, closedir: directory operations.	directory(3)
opendir, readdir, telldir,		

directory: opendir, readdir, telldir,	seekdir, rewinddir, closedir: directory operations.	directory(3X)
shmget: get shared memory	segment.	shmget(2)
brk, sbrk: change data	segment size.	brk(2)
comm:	select or reject lines common to two sorted files.	comm(1)
	select: synchronous i/o multiplexing.	select(2)
case:	selector in switch.	csh(1)
semctl:	semaphore control operations.	semctl(2)
semop:	semaphore operations.	semop(2)
ipcrm: remove a message queue,	semaphore set or shared memory id.	ipcrm(1)
semget: get set of	semaphores.	semget(2)
	semctl: semaphore control operations.	semctl(2)
	semget: get set of semaphores.	semget(2)
	semop: semaphore operations.	semop(2)
uucsd:	send a file to a remote host.	uucsd(1C)
send, sendto, sendmsg:	send a message from a socket.	send(2)
kill:	send a signal to a process.	kill(3F)
mail:	send and receive mail.	mail(1)
ping:	send ICMP ECHO_REQUEST packets to network hosts.	ping(8)
sendmail:	send mail over the internet.	sendmail(8)
binmail:	send or receive mail among users.	binmail(1)
socket.	send, sendto, sendmsg: send a message from a	send(2)
kill:	send signal to a process.	kill(2)
killpg:	send signal to a process group.	killpg(2)
	sendbug: mail a system bug report to 4bsd-bugs.	sendbug(1)
aliases: aliases file for	sendmail.	aliases(5)
	sendmail: send mail over the internet.	sendmail(8)
send, sendto,	sendmsg: send a message from a socket.	send(2)
send,	sendto, sendmsg: send a message from a socket.	send(2)
reset: reset the teletype bits to a	sensible state.	reset(1)
diction, explain: print wordy	sentences; thesaurus for diction.	diction(1)
explain, diction- print wordy	sentences; thesaurus for diction.	explain(1)
slattach: attach	serial lines as network interfaces.	slattach(8C)
comsat: biff	server.	comsat(8C)
ftpd: DARPA Internet File Transfer Protocol	server.	ftpd(8C)
rexecd: remote execution	server.	rexecd(8C)
rlogind: remote login	server.	rlogind(8C)
rshd: remote shell	server.	rshd(8C)
rwhod: system status	server.	rwhod(8C)
talkd: remote user communication	server.	talkd(8C)
telnetd: DARPA TELNET protocol	server.	telnetd(8C)
tftpd: DARPA Trivial File Transfer Protocol	server.	tftpd(8C)
	services: service name data base.	services(5)
logout: end	session.	csh(1)
script: make typescript of terminal	session.	script(1)
ascii: map of ASCII character	set.	ascii(7)
stty, gtty:	set and get terminal state (defunct).	stty(3C)
sigstack:	set and/or get signal stack context.	sigstack(2)
	set: change value of shell variable.	csh(1)
sigsetmask:	set current signal mask.	sigsetmask(2)
umask:	set file creation mode mask.	umask(2)
utime:	set file times.	utime(3C)
utimes:	set file times.	utimes(2)
setgroups:	set group access list.	setgroups(2)
apply: apply a command to a	set of arguments.	apply(1)
semget: get	set of semaphores.	semget(2)
getsockopt, setsockopt: get and	set options on sockets.	getsockopt(2)
hostid:	set or print identifier of current host system.	hostid(1)
hostname:	set or print name of current host system.	hostname(1)
ipcrm: remove a message queue, semaphore	set or shared memory id.	ipcrm(1)
setprgrp:	set process group.	setprgrp(2)
nice:	set program priority.	nice(3C)
setregid:	set real and effective group ID.	setregid(2)
setreuid:	set real and effective user ID's.	setreuid(2)
eval, exec, exit, export, login, read, readonly,	set, shift, times, trap, umask, wait: command/ /cd,	sh(1)
getty:	set terminal mode.	getty(8)
stty:	set terminal options.	stty(1)
tabs:	set terminal tabs.	tabs(1)
date: print and	set the date.	date(1)
sfddate:	set the time/date of a file.	sfddate(1)
setuid, seteuid, setruid, setgid, setegid, setrgid:	set user and group ID.	setuid(3)
setenv:	set variable in environment.	csh(1)
a stream.	setbuf, setbuffer, setlinebuf: assign buffering to	setbuf(3S)
stream. setbuf,	setbuffer, setlinebuf: assign buffering to a	setbuf(3S)
setuid, seteuid, setruid, setgid,	setgid, setrgid: set user and group ID.	setuid(3)
	setenv: set variable in environment.	csh(1)

user and group ID.	setuid, seteuid, setruid, setgid, setegid, setrgid: set	setuid(3)
entry. getfsent, getfsspec, getfsfile, getfstype,	setfsent, endfsent: get file system descriptor file	getfsent(3X)
setuid, seteuid, setruid,	setgid, setegid, setrgid: set user and group ID.	setuid(3)
getgrent, getgrgid, getgrnam,	setgrent, endgrent: get group file entry.	setgrent(3)
gethostent, gethostbyaddr, gethostbyname,	setgroups: set group access list.	setgroups(2)
host. gethostid,	sethostent, endhostent: get network host entry.	gethostent(3n)
gethostname,	sethostid: get/set unique identifier of current	gethostid(2)
getitimer,	sethostname: get/set name of current host.	gethostname(2)
	setitimer: get/set value of interval timer.	getitimer(2)
	setjmp, longjmp: non-local goto.	setjmp(3)
crypt,	setkey, encrypt: DES encryption.	crypt(3)
setbuf, setbuffer,	setlinebuf: assign buffering to a stream.	setbuf(3S)
getnetent, getnetbyaddr, getnetbyname,	setnetent, endnetent: get network entry.	getnetent(3n)
	setpgrp: set process group.	setpgrp(2)
	setpriority: get/set program scheduling priority.	setpriority(2)
getprotoent, getprotobyname, getprotobynumber,	setprotoent, endprotoent: get protocol entry.	getprotoent(3n)
getpwent, getpwuid, getpwnam,	setpwent, endpwent: get password file entry.	getpwent(3)
	setquota: enable/disable quotas on a file system.	setquota(2)
	setregid: set real and effective group ID.	setregid(2)
setuid, seteuid, setruid, setgid, setegid,	setreuid: set real and effective user ID's.	setreuid(2)
consumption. getrlimit,	setrgid: set user and group ID.	setuid(3)
group ID. setuid, seteuid,	setrlimit: control maximum system resource	getrlimit(2)
getservent, getservbyport, getservbyname,	setruid, setgid, setegid, setrgid: set user and	setuid(3)
getsockopt,	setservent, endservent: get service entry.	getservent(3n)
for changing/ random, srandom, initstate,	setsockopt: get and set options on sockets.	getsockopt(2)
gettimeofday,	setstate: better random number generator; routines	random(3)
set user and group ID.	settimeofday: get/set date and time.	gettimeofday(2)
	setuid, seteuid, setruid, setgid, setegid, setrgid:	setuid(3)
continue, cd, eval, exec, exit, export, login,/	sfdate: set the time/date of a file.	sfdate(1)
shmctl:	sh, for, case, if, while, t, ., ., break,	sh(1)
ipcrm: remove a message queue, semaphore set or	shared memory control operations.	shmctl(2)
shmop:	shared memory id.	ipcrm(1)
shmget: get	shared memory operations.	shmop(2)
xstr: extract strings from C programs to implement	shared memory segment.	shmget(2)
chsh: change default login	shared strings.	xstr(1)
exit: leave	shell.	chsh(1)
rsh: remote	shell.	csh(1)
system: issue a	shell command.	rsh(1C)
csh: a	shell (command interpreter) with C-like syntax.	system(3)
eval: re-evaluate	shell data.	csh(1)
popd: pop	shell directory stack.	csh(1)
pushd: push	shell directory stack.	csh(1)
alias:	shell macros.	csh(1)
suspend: suspend a	shell, resuming its superior.	csh(1)
rshd: remote	shell server.	rshd(8C)
set: change value of	shell variable.	csh(1)
@: arithmetic on	shell variables.	csh(1)
unset: discard	shell variables.	csh(1)
exec: overlay	shell with specified command.	csh(1)
/exec, exit, export, login, read, readonly, set,	shift: manipulate argument list.	csh(1)
	shift, times, trap, umask, wait: command language.	sh(1)
	shmctl: shared memory control operations.	shmctl(2)
	shmget: get shared memory segment.	shmget(2)
	shmop: shared memory operations.	shmop(2)
	short: integer object conversion.	long(3F)
long,	show group memberships.	groups(1)
groups:	show host status of local machines.	ruptime(1C)
ruptime:	show how long system has been up.	uptime(1)
uptime:	show last commands executed in reverse order.	lastcomm(1)
lastcomm:	show network status.	netstat(1)
netstat:	show snapshot of the UUCP system.	uusnap(8C)
uusnap:	show what versions of object modules were used to	what(1)
construct a file. what:	shut down part of a full-duplex connection.	shutdown(2)
shutdown:	shutdown: close down the system at a given time.	shutdown(8)
	shutdown: shut down part of a full-duplex	shutdown(2)
connection.	sigblock: block signals.	sigblock(2)
	sign on.	login(1)
login:	signal.	pause(3C)
pause: stop until	signal.	signal(3F)
signal: change the action for a	signal after specified time.	alarm(3C)
alarm: schedule	signal: change the action for a signal.	signal(3F)
	signal facilities.	signal(3C)
signal: simplified software	signal facilities.	sigvec(2)
sigvec: software	signal mask.	sigsetmask(2)
sigsetmask: set current		

psignal, sya_siglist: system	signal messages.	psignal(3)
sigstack: set and/or get	signal: simplified software signal facilities.	signal(3C)
kill: send	signal stack context.	sigstack(2)
kill: send a	signal to a process.	kill(2)
killpg: send	signal to a process.	kill(3F)
sigblock: block	signal to a process group.	killpg(2)
sigpause: atomically release blocked	signals.	sigblock(2)
wait for interrupt.	signals and wait for interrupt.	sigpause(2)
	sigpause: atomically release blocked signals and	sigpause(2)
	sigsetmask: set current signal mask.	sigsetmask(2)
	sigstack: set and/or get signal stack context.	sigstack(2)
	sigvec: software signal facilities.	sigvec(2)
signal:	simplified software signal facilities.	signal(3C)
tc: phototypesetter	simulator.	tc(1)
trigonometric functions.	sin, cos, tan, asin, acos, atan, atan2:	sin(3M)
	sinh, cosh, tanh: hyperbolic functions.	sinh(3M)
null: data	sink.	null(4)
brk, sbrk: change data segment	size.	brk(2)
getdtablesize: get descriptor table	size.	getdtablesize(2)
getpagesize: get system page	size.	getpagesize(2)
pagesize: print system page	size.	pagesize(1)
size:	size of an object file.	size(1)
	size: size of an object file.	size(1)
interfaces.	slattach: attach serial lines as network	slattach(8C)
	sleep: suspend execution for an interval.	sleep(1)
	sleep: suspend execution for an interval.	sleep(3F)
	sleep: suspend execution for interval.	sleep(3)
dosprinters: destinations for spooled output from	SLPT printers.	dosprinters(5)
spline: interpolate	smooth curve.	spline(1G)
	snake, snscore: display chase game.	snake(6)
usnap: show	snapshot of the UUCP system.	usnap(8C)
snake,	nscore: display chase game.	snake(6)
accept: accept a connection on a	socket.	accept(2)
bind: bind a name to a	socket.	bind(2)
connect: initiate a connection on a	socket.	connect(2)
listen: listen for connections on a	socket.	listen(2)
recv, recvfrom, recvmsg: receive a message from a	socket.	recv(2)
send, sendto, sendmsg: send a message from a	socket.	send(2)
	socket: create an endpoint for communication.	socket(2)
getsockname: get	socket name.	getsockname(2)
	socketpair: create a pair of connected sockets.	socketpair(2)
getsockopt, setsockopt: get and set options on	sockets.	getsockopt(2)
socketpair: create a pair of connected	sockets.	socketpair(2)
	soelim: eliminate .so's from nroff input.	soelim(1)
signal: simplified	software signal facilities.	signal(3C)
sigvec:	software signal facilities.	sigvec(2)
canfield, cfscores: the	solitaire card game canfield.	canfield(6)
qsort: quicker	sort.	qsort(3)
qsort: quick	sort.	qsort(3F)
tsort: topological	sort.	tsort(1)
sortbib:	sort bibliographic database.	sortbib(1)
sort:	sort or merge files.	sort(1)
	sort: sort or merge files.	sort(1)
	sortbib: sort bibliographic database.	sortbib(1)
comm: select or reject lines common to two	sorted files.	comm(1)
look: find lines in a	sorted list.	look(1)
soelim: eliminate	.so's from nroff input.	soelim(1)
soelim: eliminate .	so's from nroff input.	soelim(1)
indent: indent and format C program	source.	indent(1)
mkstr: create an error message file by massaging C	source.	mkstr(1)
whereis: locate	source, binary, and or manual for program.	whereis(1)
	source: read commands from file.	cs(1)
line, circle, arc, move, cont, point, linemod,	space, closepl: graphics interface. /erase, label,	plot(3X)
expand, unexpand: expand tabs to	spaces, and vice versa.	expand(1)
way. vfork:	spawn new process in a virtual memory efficient	vfork(2)
exec: overlay shell with	specified command.	cs(1)
truncate: truncate a file to a	specified length.	truncate(2)
alarm: schedule signal after	specified time.	alarm(3C)
alarm: execute a subroutine after a	specified time.	alarm(3F)
swapon:	specify a swap directory.	swapon(2)
swapon:	specify additional device for paging and swapping.	swapon(8)
	spell, spellin, spellout: find spelling errors.	spell(1)
spell,	spellin, spellout: find spelling errors.	spell(1)
spell, spellin, spellout: find	spelling errors.	spell(1)
spell, spellin,	spellout: find spelling errors.	spell(1)

	spline: interpolate smooth curve.	spline(1G)
	split: split a file into pieces.	split(1)
	files. fsplit: split a multi-routine Fortran file into individual	fsplit(1)
	frexp, ldexp, modf: split into mantissa and exponent.	frexp(3)
	split: split a file into pieces.	split(1)
	uuclean: uucp	uuclean(8C)
	lpq: spool queue examination program.	lpq(1)
	dosprinters: destinations for spooled output from SLPT printers.	dosprinters(5)
	dosprint: MPS/DOS spooler daemon.	dosprint(8)
	lprm: remove jobs from the line printer spooling queue.	lprm(1)
	printf, sprintf: formatted output conversion.	printf(3S)
	exp, log, log10, pow, sqrt: exponential, logarithm, power, square root.	exp(3M)
	log10, pow, sqrt: exponential, logarithm, power, square root.	exp(3M)
	rand: random number generator.	rand(3C)
	generator; routines for changing/ random, scanf, fscanf, srand, initstate, setstate: better random number	random(3)
	sscanf: formatted input conversion.	scanf(3S)
	stab: symbol table types.	stab(5)
	popd: pop shell directory	cs(1)
	pushd: push shell directory	cs(1)
	sigstack: set and/or get signal	sigstack(2)
	copy: standalone copy program.	copy(8)
	dkfmt: standalone disk formatter.	dkfmt(8)
	machine operation mode.. Standalone mode: definition of this Sanyo/ICON	standalone(8)
	blood: program to load standalone programs.	blood(8)
	stdio: standard buffered input/output package.	intro(3S)
	htable: convert NIC standard format host tables.	htable(8)
	stat, lstat, fstat: get file status.	stat(2)
	stat, lstat, fstat: get file status.	stat(3F)
	state.	reset(1)
	state (defunct).	stty(3C)
	state with that on disk.	fsync(2)
	statement.	cs(1)
	static information about the filesystems.	fstab(5)
	statistics.	cs(1)
	statistics.	iocstat(1)
	statistics.	vmstat(1)
	status.	exit(3F)
	status.	ipc(1)
	status.	netstat(1)
	status.	ps(1)
	status.	stat(2)
	status.	stat(3F)
	status inquiries.	ferror(3S)
	status line of a terminal.	sysline(1)
	status of local machines.	ruptime(1C)
	status on status line of a terminal.	sysline(1)
	status server.	rwhod(8C)
	stdio: standard buffered input/output package.	intro(3S)
	sticky: executable files with persistent text.	sticky(8)
	stop: halt a job or process.	cs(1)
	halt: stop the processor.	halt(8)
	pause: stop until signal.	pause(3C)
	icheck: file system storage consistency check.	icheck(8)
	subroutines. dbminit, fetch, store, delete, firstkey, nextkey: data base	dbm(3X)
	strlen, index, rindex: string operations.	string(3)
	rindex: string operations. strcat, strncat, strcmp, strncmp, strcpy, strncpy, strlen, index, rindex: string	string(3)
	operations. strcat, strncat, strcmp, strncmp, stream.	fclose(3S)
	fclose, fflush: close or flush a stream.	fopen(3S)
	fopen, freopen, fdopen: open a stream.	fseek(3S)
	fseek, ftell, rewind: reposition a stream.	getc(3S)
	getchar, fgetc, getw: get character or word from stream.	gets(3S)
	gets, fgets: get a string from a stream.	putc(3S)
	putchar, fputc, putw: put character or word on a stream.	puts(3S)
	puts, fputs: put a string on a stream.	setbuf(3S)
	setbuffer, setlinebuf: assign buffering to a stream.	ungetc(3S)
	ungetc: push character back into input stream editor.	sed(1)
	sed: stream status inquiries.	ferror(3S)
	ferror, feof, clearerr, fileno: stream to a remote command.	rcmd(3X)
	rcmd, rresvport, ruserok: routines for returning stream to a remote command.	rindex(3X)
	rexec: return string.	fdate(3F)
	fdate: return date and time in an ASCII string from a stream.	getc(3S)
	gets, fgets: get a string on a stream.	puts(3S)
	puts, fputs: put a string on a stream.	bstring(3)
	string operations. strcat, strncat, strcmp, string operations. strcat, strncat, strcmp,	string(3)

extract strings from C programs to implement shared other binary, file.	strings. xstr: extract strings: find the printable basename:	strings. xstr: xstr(1) strings: find the printable strings in a object, or strings from C programs to implement shared strings in a object, or other binary, file. strings(1) strip filename affixes. basename(1)
streat, strncat, stremp, strncmp, strepy, strncpy, index, rindex: string operations. strcat, string operations. strcat, strncat, stremp, strcat, strncat, stremp, strncmp, strepy,	strip: remove symbols and relocation bits. strlen, index, rindex: string operations. strncat, stremp, strncmp, strepy, strncpy, strlen, strncmp, strepy, strncpy, strlen, index, rindex: string operations. struct: structure Fortran programs. structure Fortran programs.	strip(1) string(3) string(3) string(3) string(3) struct(1) struct(1)
struct:	stty, gtty: set and get terminal state (defunct). stty: set terminal options.	stty(3C) stty(1)
document.	style: analyze surface characteristics of a su: substitute user id temporarily.	style(1) su(1)
alarm: execute a	subroutine after a specified time.	alarm(3F)
fetch, store, delete, firstkey, nextkey: data base	subroutines. dbminit,	dbm(3X)
lib2648:	subroutines for the HP 2648 graphics terminal.	lib2648(3X)
su:	substitute user id temporarily.	su(1)
sum:	sum and count blocks in a file.	sum(1)
du:	sum: sum and count blocks in a file.	sum(1)
quot:	summarize disk usage.	du(1)
repquota:	summarize file system ownership.	quot(8)
sync: update the	summarize quotas for a file system.	repquota(8)
update: periodically update the	super block.	sync(8)
sync: update	super block.	update(8)
suspend: suspend a shell, resuming its	super-block.	sync(2)
style: analyze	superior.	cs(1)
suspend:	surface characteristics of a document.	style(1)
sleep:	suspend a shell, resuming its superior.	cs(1)
sleep:	suspend execution for an interval.	sleep(1)
sleep:	suspend execution for an interval.	sleep(3F)
sleep:	suspend execution for interval.	sleep(3)
	suspend: suspend a shell, resuming its superior.	cs(1)
swab:	swab: swap bytes.	swab(3)
swapon: specify a	swap bytes.	swab(3)
swapping.	swap directory.	swapon(2)
swapon: specify additional device for paging and	swapon: specify a swap directory.	swapon(2)
breaksw: exit from	swapon: specify additional device for paging and	swapon(8)
case: selector in	swapping.	swapon(8)
default: catchall clause in	switch.	cs(1)
endsw: terminate	switch.	cs(1)
	switch.	cs(1)
	switch.	cs(1)
	switch: multi-way command branch.	cs(1)
stab:	symbol table types.	stab(5)
readlink: read value of a	symbolic link.	readlink(2)
symlink: make	symbolic link to a file.	symlink(2)
strip: remove	symbols and relocation bits.	strip(1)
	symlink: make symbolic link to a file.	symlink(2)
	symorder: rearrange name list.	symorder(1)
	sync: update super-block.	sync(2)
	sync: update the super block.	sync(8)
disk. fsync:	synchronize a file's in-core state with that on	fsync(2)
select:	synchronous i/o multiplexing.	select(2)
cs(1): a shell (command interpreter) with C-like	syntax.	cs(1)
	syscall: indirect system call.	syscall(2)
perror, sys_errlist, sys_nerr: system error messages.	sys_errlist, sys_nerr: system error messages.	perror(3)
terminal.	sysline: display system status on status line of a	sysline(1)
	syslog: log systems messages.	syslog(8)
	syslog, openlog, closelog: control system log.	syslog(3)
perror, sys_errlist, psignal,	sys_nerr: system error messages.	perror(3)
psignal,	sys_siglist: system signal messages.	psignal(3)
jove_recover - recover JOVE buffers after a	system/editor crash.	jove_recover(1)
mkfs: program to make UNIX file	systems.	mkfs(8)
syslog: log	systems messages.	syslog(8)
kgmon: generate a dump of the operating	system's profile buffers.	kgmon(8)
rehash: recompute command hash	table.	cs(1)
unhash: discard command hash	table.	cs(1)
mtab: mounted file system	table.	mtab(5)
getdtablesize: get descriptor	table size.	getdtablesize(2)
stab: symbol	table types.	stab(5)
htable: convert NIC standard format host	tables.	htable(8)
route: manually manipulate the routing	tables.	route(8C)

tbl: format	tables for nroff or troff.	tbl(1)
gettable: get NIC format host	tables from a host.	gettable(8C)
tabs: set terminal	tabs.	tabs(1)
expand, unexpand: expand	tabs: set terminal tabs.	tabs(1)
ctags: create a	tabs to spaces, and vice versa.	expand(1)
	tags file.	ctags(1)
	tail: deliver the last part of a file.	tail(1)
	talk: talk to another user.	talk(1)
talk:	talk to another user.	talk(1)
	talkd: remote user communication server.	talkd(8C)
functions. sin, cos,	tan, asin, acos, atan, atan2: trigonometric	sin(3M)
sinh, cosh,	tanh: hyperbolic functions.	sinh(3M)
tecopy: copy a mag	tape.	tecopy(1)
	tar: tape archive file format.	tar(5)
	tar: tape archiver.	tar(1)
tclose, tread, twrite, trewin, tskipf, tstata: f77	tape I/O. topen,	topen(3F)
mt: magnetic	tape manipulating program.	mt(1)
	tar: tape archive file format.	tar(5)
	tar: tape archiver.	tar(1)
deroff: remove nroff, troff,	tbl and eqn constructs.	deroff(1)
	tbl: format tables for nroff or troff.	tbl(1)
	tc: phototypesetter simulator.	tc(1)
	tclose, tread, twrite, trewin, tskipf, tstata: f77	topen(3F)
	tecopy: copy a mag tape.	tecopy(1)
	TEACHJOVE - learn how to use the JOVE editor.	teachjove(1)
	tee: pipe fitting.	tee(1)
tk: paginator for the	Tektronix 4014.	tk(1)
reset: reset the	teletype bits to a sensible state.	reset(1)
last: indicate last logins of users and	teletypes.	last(1)
index, rindex, lnblnk, len:	tell about character objects.	index(3F)
operations. opendir, readdir,	telldir, seekdir, rewinddir, closedir: directory	directory(3)
operations. directory: opendir, readdir,	telldir, seekdir, rewinddir, closedir: directory	directory(3X)
telnet: user interface to the	TELNET protocol.	telnet(1C)
telnetd: DARPA	TELNET protocol server.	telnetd(8C)
	telnet: user interface to the TELNET protocol.	telnet(1C)
	telnetd: DARPA TELNET protocol server.	telnetd(8C)
su: substitute user id	temporarily.	su(1)
	term: conventional names for terminals.	term(7)
lib2648: subroutines for the HP 2648 graphics	termcap: terminal capability data base.	termcap(5)
lock: reserve a	terminal.	lib2648(3X)
sysline: display system status on status line of a	terminal.	lock(1)
ttynam, isatty, ttyslot: find name of a	terminal.	sysline(1)
vhangup: virtually "hangup" the current control	terminal.	ttynam(3)
worms: animate worms on a display	terminal.	vhangup(2)
	terminal.	worms(6)
	termcap: terminal capability data base.	termcap(5)
	gettytab: terminal configuration data base.	gettytab(5)
	tset: terminal dependent initialization.	tset(1)
tgetent, tgetnum, tgetflag, tgetstr, tgoto, tputs:	terminal independent operation routines.	termcap(3X)
	ttys: terminal initialization data.	ttys(5)
	tty: general	tty(4)
	getty: set	getty(8)
	tty: get	tty(1)
	atty: set	stty(1)
ttynam, isatty: find name of a	terminal port.	ttynam(3F)
clear: clear	terminal screen.	clear(1)
script: make typescript of	terminal session.	script(1)
atty, gtty: set and get	terminal state (defunct).	stty(3C)
tabs: set	terminal tabs.	tabs(1)
ttytype: data base of	terminal types by port.	ttytype(5)
term: conventional names for	terminals.	term(7)
wait, wait3: wait for process to	terminate.	wait(2)
wait: wait for a process to	terminate.	wait(3F)
_exit:	terminate a process.	exit(2)
output. exit:	terminate a process after flushing any pending	exit(3)
kill:	terminate a process with extreme prejudice.	kill(1)
abort:	terminate abruptly with memory image.	abort(3F)
endif:	terminate conditional.	csh(1)
end:	terminate loop.	csh(1)
exit:	terminate process with status.	exit(3F)
endsw:	terminate switch.	csh(1)
	test: condition command.	test(1)
quiz:	test your knowledge.	quiz(6)
sticky: executable files with persistent	text.	sticky(8)
ed:	text editor.	ed(1)

ex, edit:	text editor.	ex(1)
jove: an interactive display-oriented	text editor.	jove(1)
fmt: simple	text formatter.	fmt(1)
nroff:	text formatting.	nroff(1)
troff, nroff:	text formatting and typesetting.	troff(1)
ms:	text formatting macros.	ms(7)
	tftp: trivial file transfer program.	tftp(1C)
	tftpd: DARPA Trivial File Transfer Protocol server.	tftpd(8C)
terminal independent operation routines.	tgetent, tgetnum, tgetflag, tgetstr, tgoto, tputs:	termcap(3X)
independent operation routines. tgetent, tgetnum,	tgetflag, tgetstr, tgoto, tputs: terminal	termcap(3X)
independent operation routines. tgetent,	tgetnum, tgetflag, tgetstr, tgoto, tputs: terminal	termcap(3X)
operation routines. tgetent, tgetnum, tgetflag,	tgetstr, tgoto, tputs: terminal independent	termcap(3X)
routines. tgetent, tgetnum, tgetflag, tgetstr,	tgoto, tputs: terminal independent operation	termcap(3X)
diction, explain: print wordy sentences;	thesaurus for diction.	diction(1)
explain, diction- print wordy sentences;	thesaurus for diction.	explain(1)
merge:	three-way file merge.	merge(1)
alarm: schedule signal after specified	time.	alarm(3C)
alarm: execute a subroutine after a specified	time.	alarm(3F)
at: execute commands at a later	time.	at(1)
etime, dtime: return elapsed execution	time.	etime(3F)
gettimeofday, settimeofday: get/set date and	time.	gettimeofday(2)
shutdown: close down the system at a given	time.	shutdown(8)
time, ftime: get date and	time.	time(3C)
time, ctime, ltime, gmtime: return system	time.	time(3F)
time:	time a command.	time(1)
.PP dstrules: Daylight savings	time and time zone name rule file..	dstrules(5)
time:	time command.	cs(1)
	time, ctime, ltime, gmtime: return system time.	time(3F)
	time, ftime: get date and time.	time(3C)
	time in an ASCII string.	fdate(3F)
fdate: return date and	time in numerical form.	idate(3F)
idate, itime: return date or	time profile.	profil(2)
profil: execution	time: time a command.	time(1)
	time: time command.	cs(1)
gmtime, asctime, timezone: convert date and	time to ASCII. ctime, localtime,	ctime(3)
.PP dstrules: Daylight savings time and	time zone name rule file..	dstrules(5)
afdate: set the	time/date of a file.	sfdate(1)
getitimer, setitimer: get/set value of interval	timer.	getitimer(2)
times: get process	times.	times(3C)
utime: set file	times.	utime(3C)
utimes: set file	times.	utimes(2)
	times: get process times.	times(3C)
exit, export, login, read, readonly, set, shift,	times, trap, umask, wait: command language. /exec,	sh(1)
ctime, localtime, gmtime, asctime,	timezone: convert date and time to ASCII.	ctime(3)
	tip, cu: connect to a remote system.	tip(1C)
	tk: paginator for the Tektronix 4014.	tk(1)
	tmail: print out mail messages, most recent first.	tmail(1)
popen, pclose: initiate I/O	to/from a process.	popen(3)
tstate: f77 tape I/O.	topen, tclose, tread, twrite, twrewin, tskipf,	topen(3F)
tsort:	topological sort.	tsort(1)
	touch: update date last modified of a file.	touch(1)
tgetent, tgetnum, tgetflag, tgetstr, tgoto,	tputs: terminal independent operation routines.	termcap(3X)
	tr: translate characters.	tr(1)
ptrace: process	trace.	ptrace(2)
trpt: transliterate protocol	trace.	trpt(8C)
goto: command	transfer.	cs(1)
kermit: kermit file	transfer.	kermit(1)
ftp: file	transfer program.	ftp(1C)
tftp: trivial file	transfer program.	tftp(1C)
ftpd: DARPA Internet File	Transfer Protocol server.	ftpd(8C)
tftpd: DARPA Trivial File	Transfer Protocol server.	tftpd(8C)
tr:	translate characters.	tr(1)
macros. trman:	translate version 6 manual macros to version 7	trman(1)
pi: Pascal interpreter code	translator.	pi(1)
trpt:	transliterate protocol trace.	trpt(8C)
uuencode, uudecode: encode/decode a binary file for	transmission via mail.	uuencode(1C)
trpspe, specnt:	trap and repair floating point faults.	trpspe(3F)
trapov:	trap and repair floating point overflow.	trapov(3F)
traper:	trap arithmetic errors.	traper(3F)
export, login, read, readonly, set, shift, times,	trap, umask, wait: command language. /exec, exit,	sh(1)
	traper: trap arithmetic errors.	traper(3F)
	trapov: trap and repair floating point overflow.	trapov(3F)
I/O. topen, tclose,	tread, twrite, twrewin, tskipf, tstate: f77 tape	topen(3F)
	trek: trekkie game.	trek(6)
trek:	trekkie game.	trek(6)

topen, tclose, tread, twrite,	trewin, tskipf, tstate: f77 tape I/O.	topen(3F)
sin, cos, tan, asin, acos, atan, atan2:	trigonometric functions.	sin(3M)
tftp:	trivial file transfer program.	tftp(1C)
tftpd: DARPA	Trivial File Transfer Protocol server.	tftpd(8C)
7 macros.	trman: translate version 6 manual macros to version	trman(1)
tbl: format tables for nroff or	troff.	tbl(1)
	troff, nroff: text formatting and typesetting.	troff(1)
deroff: remove nroff,	troff, tbl and eqn constructs.	deroff(1)
battlestar: a	tropical adventure game.	battlestar(6)
faults.	trpffe, fpecnt: trap and repair floating point	trpffe(3F)
	trpt: transliterate protocol trace.	trpt(8C)
	true, false: provide truth values.	true(1)
	true: provide truth values.	false(1)
false,	truncate a file to a specified length.	truncate(2)
truncate:	truncate: truncate a file to a specified length.	truncate(2)
	truth values.	false(1)
false, true: provide	truth values.	true(1)
true, false: provide	tset: terminal dependent initialization.	tset(1)
	tskipf, tstate: f77 tape I/O.	topen(3F)
topen, tclose, tread, twrite, trewin,	tsort: topological sort.	tsort(1)
topen, tclose, tread, twrite, trewin, tskipf,	tstate: f77 tape I/O.	topen(3F)
	tty: general terminal interface.	tty(4)
	tty: get terminal name.	tty(1)
	ttynam, isatty: find name of a terminal port.	ttynam(3F)
	ttynam, isatty, ttyslot: find name of a terminal.	ttynam(3)
	ttys: terminal initialization data.	ttys(5)
ttynam, isatty,	ttyslot: find name of a terminal.	ttynam(3)
	ttytype: data base of terminal types by port.	ttytype(5)
tunefs:	tune up an existing file system.	tunefs(8)
	tunefs: tune up an existing file system.	tunefs(8)
topen, tclose, tread,	twrite, trewin, tskipf, tstate: f77 tape I/O.	topen(3F)
file: determine file	type.	file(1)
stab: symbol table	types.	stab(5)
types: primitive system data	types.	types(5)
ttytype: data base of terminal	types by port.	ttytype(5)
	types: primitive system data types.	types(5)
	typescript of terminal session.	script(1)
script: make	typeset manual.	man(7)
man: macros to	typeset mathematics.	eqn(1)
eqn, neqn, checkeq:	typesetting.	troff(1)
troff, nroff: text formatting and	uid.	getpw(3C)
getpw: get name from	ul: do underlining.	ul(1)
	umask: change or display file creation mask.	csh(1)
	umask: set file creation mode mask.	umask(2)
login, read, readonly, set, shift, times, trap,	umask, wait: command language. /exec, exit, export,	sh(1)
mount,	umount: mount and dismount file system.	mount(8)
mount,	umount: mount or remove file system.	mount(2)
	unalias: remove aliases.	csh(1)
cat them. compact,	uncompact, ccat: compress and uncompress files, and	compact(1)
compact, uncompact, ccat: compress and	uncompress files, and cat them.	compact(1)
compress,	uncompress, zcat: compress and expand data.	compress(1)
ul: do	underlining.	ul(1)
expand,	unexpand: expand tabs to spaces, and vice versa.	expand(1)
	ungetc: push character back into input stream.	ungetc(3S)
	unhash: discard command hash table.	csh(1)
	unifdef: remove ifdef'ed lines.	unifdef(1)
	uniq: report repeated lines in a file.	uniq(1)
	unique file name.	mktemp(3)
mktemp: make a	unique identifier of current host.	gethostid(2)
gethostid, sethostid: get/set	unit.	flush(3F)
flush: flush output to a logical	unit.	fseek(3F)
fseek, ftell: reposition a file on a logical	unit.	getc(3F)
getc, fgetc: get a character from a logical	unit.	putc(3F)
putc, fputc: write a character to a fortran logical	units: conversion program.	units(1)
	UNIX.	learn(1)
learn: computer aided instruction about	UNIX bootstrapping procedures.	reboot(8)
reboot:	UNIX command.	system(3F)
system: execute a	unix command execution.	uux(1C)
uux: unix to	unix copy.	uucp(1C)
uucp, uulog: unix to	UNIX file systems.	mkfs(8)
mkfs: program to make	unix to unix command execution.	uux(1C)
uux:	unix to unix copy.	uucp(1C)
uucp, uulog:	unlimit: remove resource limitations.	csh(1)
	(unlink) directories or files.	rmdir(1)
rmdir, rm: remove	(unlink) files or directories.	rm(1)
rm, rmdir: remove		

uptime: show how long system has been	unlink: remove a directory entry.	unlink(3F)
tunefs: tune	unlink: remove directory entry.	unlink(2)
touch:	unset: discard shell variables.	csh(1)
sync:	unsetenv: remove environment variables.	csh(1)
sync: update the super block.	up.	uptime(1)
update: periodically	up an existing file system.	tunefs(8)
sync:	update date last modified of a file.	touch(1)
sync: update the super block.	update: periodically update the super block.	update(8)
update: periodically	update super-block.	sync(2)
du: summarize disk	update the super block.	sync(8)
quota: display disc	update the super block.	update(8)
TEACHJOVE - learn how to	uptime: show how long system has been up.	uptime(1)
what: show what versions of object modules were	usage.	du(1)
miscellaneous: miscellaneous	usage and limits.	quota(1)
login: login new	use the JOVE editor.	teachjove(1)
talk: talk to another	used to construct a file.	what(1)
write: write to another	useful information pages.	intro(7)
seteuid, setruid, setgid, setegid, setrgid: set	user.	csh(1)
talkd: remote	user.	talk(1)
environ:	user.	write(1)
whoami: print effective current	user and group ID. seteuid,	setuid(3)
su: substitute	user communication server.	talkd(8C)
getuid, geteuid: get	user environment.	environ(7)
setreuid: set real and effective	user id.	whoami(1)
finger:	user id temporarily.	su(1)
telnet:	user identity.	getuid(2)
getuid, getgid: get	user ID's.	setreuid(2)
edquota: edit	user information lookup program.	finger(1)
adduser: procedure for adding new	user interface to the TELNET protocol.	telnet(1C)
binmail: send or receive mail among	user or group ID of the caller.	getuid(3F)
wall: write to all	user quotas.	edquota(8)
whodos: display information about doc	users.	adduser(8)
last: indicate last logins of	users.	binmail(1)
getlog: get	users.	wall(1)
users: compact list of	users.	whodos(1)
getrusage: get information about resource	users and teletypes.	last(1)
vtimes: get information about resource	users: compact list of users who are on the system.	users(1)
	user's login name.	getlog(3F)
	users who are on the system.	users(1)
	utilization.	getrusage(2)
	utilization.	vtimes(3C)
	utime: set file times.	utime(3C)
	utimes: set file times.	utimes(2)
	utmp, wtmp: login records.	utmp(5)
	uuclean: uucp spool directory clean-up.	uuclean(8C)
rmail: handle remote mail received via	uucp.	rmail(1)
uuclean:	uucp spool directory clean-up.	uuclean(8C)
uusnap: show snapshot of the	UUCP system.	uusnap(8C)
transmission via mail. uuencode,	uucp, uulog: unix to unix copy.	uucp(1C)
uuencode: format of an encoded	uudecode: encode/decode a binary file for	uuencode(1C)
transmission via mail.	uuencode file.	uuencode(5)
uucp,	uuencode: format of an encoded uuencode file.	uuencode(5)
uuencode, uudecode: encode/decode a binary file for	uuencode, uudecode: encode/decode a binary file for	uuencode(1C)
uulog: unix to unix copy.	uulog: unix to unix copy.	uucp(1C)
uusend: send a file to a remote host.	uusend: send a file to a remote host.	uusend(1C)
uusnap: show snapshot of the UUCP system.	uusnap: show snapshot of the UUCP system.	uusnap(8C)
uux: unix to unix command execution.	uux: unix to unix command execution.	uux(1C)
uxrc: configuration file for kernel.	uxrc: configuration file for kernel.	uxrc(8)
valloc: aligned memory allocator.	valloc: aligned memory allocator.	valloc(3)
value.	value.	abs(3)
value, floor, ceiling functions.	value, floor, ceiling functions.	floor(3M)
value for environment name.	value for environment name.	getenv(3)
value of a symbolic link.	value of a symbolic link.	readlink(2)
value of environment variables.	value of environment variables.	getenv(3F)
value of interval timer.	value of interval timer.	getitimer(2)
value of shell variable.	value of shell variable.	csh(1)
values.	values.	false(1)
values. fmin, fmax, frac,	values. fmin, fmax, frac,	fmin(3F)
values.	values.	rand(3F)
values.	values.	true(1)
values between host and network byte order.	values between host and network byte order.	byteorder(3n)
varargs: variable argument list.	varargs: variable argument list.	varargs(3)
variable.	variable.	csh(1)
variable argument list.	variable argument list.	varargs(3)
variable in environment.	variable in environment.	csh(1)
set: change value of shell		
varargs:		
setenv: set		

@: arithmetic on shell	variables.	csh(1)
unset: discard shell	variables.	csh(1)
unsetenv: remove environment	variables.	csh(1)
getenv: get value of environment	variables.	getenv(3F)
to create and display information for MPS/DOS	vdisk. dosdisk: program	dosdisk(8)
assert: program	verification.	assert(3X)
lint: a C program	verifier.	lint(1)
expand, unexpand: expand tabs to spaces, and vice	vers: print version number of the kernel.	vers(1)
trman: translate	versa.	expand(1)
trman: translate version 6 manual macros to	version 6 manual macros to version 7 macros.	trman(1)
vers: print	version 7 macros.	trman(1)
hangman: Computer	version number of the kernel.	vers(1)
file. what: show what	version of the game hangman.	hangman(6)
efficient way.	versions of object modules were used to construct a	what(1)
	vfork: spawn new process in a virtual memory	vfork(2)
	vgrind: grind nice listings of programs.	vgrind(1)
	vgrindefs: vgrind's language definition data base.	vgrindefs(5)
	vgrind's language definition data base.	vgrindefs(5)
	vhangup: virtually "hangup" the current control	vhangup(2)
	vi: screen oriented (visual) display editor based	vi(1)
	via mail. uuencode,uudecode:	uuencode(1C)
	via uucp.	rmail(1)
encode/decode a binary file for transmission	vice versa.	expand(1)
rmail: handle remote mail received	viewing.	more(1)
expand, unexpand: expand tabs to spaces, and	villainous robots.	robots(6)
more, page: file perusal filter for crt	vipw: edit the password file.	vipw(8)
robots: fight off	virtual disks.	dosdisks(5)
	virtual memory efficient way.	vfork(2)
	virtual memory statistics.	vmstat(1)
	virtually "hangup" the current control terminal.	vhangup(2)
	(visual) display editor based on ex.	vi(1)
	vlimit: control maximum system resource	vlimit(3C)
	vmstat: report virtual memory statistics.	vmstat(1)
	volume.	fs(5)
	vtimes: get information about resource utilization.	vtimes(3C)
	w: who is on and what they are doing.	w(1)
	wait: await completion of process.	wait(1)
	wait: command language. /exec, exit, export, login,	sh(1)
	wait for a process to terminate.	wait(3F)
	wait for background processes to complete.	csh(1)
	wait for interrupt.	sigpause(2)
	wait for process to terminate.	wait(2)
	wait: wait for a process to terminate.	wait(3F)
	wait: wait for background processes to complete.	csh(1)
	wait, wait3: wait for process to terminate.	wait(2)
	wait3: wait for process to terminate.	wait(2)
	wall: write to all users.	wall(1)
	wc: word count.	wc(1)
	were used to construct a file.	what(1)
	what: show what versions of object modules	what(1)
	whatis: describe	whatis(1)
	crash: what happens when the system crashes.	crash(8V)
	used to construct a file.	what(1)
	w: who is on and	w(1)
	construct a file. what: show	what(1)
	whatis: describe what a command is.	whatis(1)
	crash: what happens	crash(8V)
	leave: remind you	leave(1)
	program.	whereis(1)
	paths (csh only).	which(1)
	exec, exit, export, login, / sh, for, case, if,	sh(1)
	while, s, ., break, continue, cd, eval,	csh(1)
	while: repeat commands conditionally.	csh(1)
	while/foreach loop.	csh(1)
	who are on the system.	users(1)
	break: exit	from(1)
	users: compact list of users	w(1)
	from: who is my mail from?	who(1)
	w: who is on and what they are doing.	biff(1)
	who: who is on the system.	who(1)
	who it is from.	whoami(1)
	who: who is on the system.	whodos(1)
	whoami: print effective current user id.	rwho(1C)
	whodos: display information about dosc users.	fold(1)
	who's logged in on local machines.	window(1)
	rwho: who's logged in on local machines.	window(1)
	fold: fold long lines for finite	fastboot(8)
	window: window environment.	
	without checking the disks.	
fastboot, fasthalt: reboot/halt the system		

wc:	word count.	wc(1)
getc, getchar, fgetc, getw:	get character or word from stream.	getc(3S)
putc, putchar, fputc, putw:	put character or word on a stream.	putc(3S)
diction, explain:	print wordy sentences; thesaurus for diction.	diction(1)
explain, diction-	print wordy sentences; thesaurus for diction.	explain(1)
cd:	change working directory.	cd(1)
chdir:	change current working directory.	chdir(2)
getcwd:	get pathname of current working directory.	getcwd(3F)
pwd:	working directory name.	pwd(1)
getwd:	get current working directory pathname.	getwd(3)
worm:	Play the growing worm game.	worm(6)
	worm: Play the growing worm game.	worm(6)
	worms: animate worms on a display terminal.	worms(6)
worms: animate	worms on a display terminal.	worms(6)
putc, fputc:	write a character to a fortran logical unit.	putc(3F)
write, writev:	write on a file.	write(2)
wall:	write to all users.	wall(1)
write:	write to another user.	write(1)
	write: write to another user.	write(1)
	write, writev: write on a file.	write(2)
	writev: write on a file.	write(2)
open: open a file for reading or writing, or create a new file.		open(2)
utmp, wtmp:	login records.	utmp(5)
	wump: the game of hunt-the-wumpus.	wump(6)
xsend, xget, enroll:	secret mail.	xsend(1)
bit: and, or, xor, not, rshift, lshift	bitwise functions.	bit(3F)
	xsend, xget, enroll: secret mail.	xsend(1)
shared strings.	xstr: extract strings from C programs to implement	xstr(1)
j0, j1, jn,	y0, y1, yn: bessell functions.	j0(3M)
j0, j1, jn, y0,	y1, yn: bessell functions.	j0(3M)
eyacc: modified	yacc allowing much improved error recovery.	eyacc(1)
	yacc: yet another compiler-compiler.	yacc(1)
	yes: be repetitively affirmative.	yes(1)
j0, j1, jn, y0, y1,	yn: bessell functions.	j0(3M)
compress, uncompress,	zcat: compress and expand data.	compress(1)
.PP dstrules: Daylight savings time and time	zone name rule file..	dstrules(5)



SECTION II

**ICON/UXB
OPERATING
SYSTEM
INTRODUCTION
TO VOLUME 1**



INTRODUCTION TO VOLUME 1

INTRODUCTION

This volume gives descriptions of the publicly available features of the ICON/UXB operating system, as extended to provide a virtual memory environment and other enhancements. It does not attempt to provide perspective or tutorial information upon the ICON/UXB operating system, its facilities, or its implementation. Various documents on those topics are contained in Volumes 2 and 3, *Supplementary Documentation*. In particular, for an overview see *The UNIX® Time-Sharing System* by Ritchie and Thompson; for a tutorial see *UNIX® for Beginners* by Kernighan.

Within the area it surveys, this volume attempts to be timely, complete and concise. Where the latter two objectives conflict, the obvious is often left unsaid in favor of brevity. It is intended that each program be described as it is, not as it should be. Inevitably, this means that various sections will soon be out of date. As changes are made, updates will be forthcoming.

This volume is divided into eight sections:

1. Commands
2. System Calls
3. Subroutines
4. Special Files
5. File Formats and Conventions
6. Games
7. Miscellaneous
8. Maintenance Commands and Procedures

Commands are programs intended to be invoked directly by the user, in contradistinction to subroutines, which are intended to be called by the user's programs. Commands generally reside in directory */bin* (for *binary* programs). Some programs also reside in */usr/bin*, or in */usr/ucb*, to save space in */bin*. These directories are searched automatically by the command interpreters.

System Calls are entries into the ICON/UXB supervisor. The system call interface is identical to a C language procedure call; the equivalent C procedures are described in Section 2.

An assortment of **Subroutines** is available; they are described in section 3. The primary libraries in which they are kept are described in the introduction to section 3, *intro(3)*. The functions are described in terms of C, but most will work with Fortran as well.

UNIX is a registered trademark of AT&T.

The **Special Files** section discusses the characteristics of each system 'file' that actually refers to an I/O device. The names in this section refer to the ICON MultiMicro/MainFrame Architecture device names for the hardware, instead of the names of the special files themselves.

The **File Formats and Conventions** section documents the structure of particular kinds of files; for example, the form of the output of the loader and assembler is given. Excluded are files used by only one command, for example the assembler's intermediate files.

Games have been relegated to section 6 to keep them from contaminating the more staid information of section 1.

Section 7 is a **Miscellaneous** collection of information necessary to writing in various specialized languages: character codes, macro packages for typesetting, etc.

The **Maintenance** section discusses commands and procedures not intended for use by the ordinary user. The commands and files described in Section 8 are almost all kept in the directory */etc*.

SECTION FORMAT

Each section in the ICON/UXB Reference Manual consists of a number of independent entries of one or more pages each. The name of the entry is in the upper corners of its pages, together with the section number, and sometimes a letter characteristic of a subcategory, e.g. graphics is 1G, and the math library is 3M. Entries within each section are alphabetized. The page numbers of each entry start at 1; it is infeasible to number consecutively the pages of a document like this that is republished in many variant forms and is constantly being changed, updated, or obsoleted.

All entries are based on a common format, not all of whose subsections will always appear.

The *NAME* subsection lists the exact names of the commands and subroutines covered under the entry and gives a very short description of their purpose.

The *SYNOPSIS* summarizes the use of the program being described. A few conventions are used, particularly in the Commands subsection:

- **Boldface** words are considered literals, and are typed just as they appear.
- Square brackets [] around an argument indicate that the argument is optional. When an argument is given as 'name', it always refers to a file name.
- Ellipses '...' are used to show that the previous argument-prototype may be repeated.
- A final convention is used by the commands themselves. An argument beginning with a minus sign (-) is often taken to mean some sort of option-specifying argument even if it appears in a position where a file name could appear. Therefore, it is unwise to have files whose names begin with '-'

The *DESCRIPTION* subsection discusses in detail the subject at hand.

The *FILES* subsection gives the names of files which are built into the program.

A *SEE ALSO* subsection gives pointers to related information.

A *DIAGNOSTICS* subsection discusses the diagnostic indications which may be produced. Messages which are intended to be self-explanatory are not listed.

The *BUGS* subsection gives known bugs and sometimes deficiencies. Occasionally also the suggested fix is described.

At the beginning of the volume is a table of contents, organized by section and alphabetically within each section. There is also a permuted index derived from the table of contents. Within each index entry, the title of the writeup to which it refers is followed by the appropriate section number in parentheses. This fact is important because there is considerable name duplication among the sections, arising principally from commands which exist only to exercise a particular system call.

HOW TO GET STARTED

This section sketches the basic information you need to get started on the ICON/UXB operating system; how to log in and log out, how to communicate through your terminal, and how to run a program. See *UNIX® for Beginners* in Volume 2 for a more complete introduction to the system.

Logging In

You must call ICON/UXB from an appropriate terminal. Almost any ASCII terminal capable of full duplex operation and generating the entire character set can be used. You must also have a valid user name, which may be obtained, together with necessary telephone numbers, from the system administration. After a data connection is established, the login procedure depends on what kind of terminal you are using and local system conventions. The following examples are typical.

300-baud terminals: Such terminals include the GE Terminet 300, and most display terminals run with popular modems. These terminals generally have a speed switch which should be set at '300' (or '30' for 30 characters per second) and a half/full duplex switch which should be set at full-duplex. (This switch will often have to be changed since many other systems require half-duplex). When a connection is established, the system types **login:** you type your user name, followed by the 'return' key. If you have a password, the system asks for it and turns off the printer on the terminal so the password will not appear. After you have logged in, the 'return', 'new line', or 'linefeed' keys will give exactly the same results.

1200- and 150-baud terminals: If there is a half/full duplex switch, set it at full-duplex. When you have established a data connection, the system types out a few garbage characters (the **login:** message at the wrong speed). Depress the 'break' (or 'interrupt') key; this is a speed-independent signal to the ICON/UXB operating system that a different speed terminal is in use. The system then will type **login:** this time at another speed. Continue depressing the break key until **login:** appears in clear, then respond with your user name. From the TTY 37 terminal, and any other which has the 'newline' function

(combined carriage return and linefeed), terminate each line you type with the 'new line' key, otherwise use the 'return' key.

Hard-wired terminals. Hard-wired terminals usually begin at the right speed, up to 9600 baud; otherwise the preceding instructions apply.

For all these terminals, it is important that you type your name in lower-case if possible; if you type upper-case letters, ICON/UXB will assume that your terminal cannot generate lower-case letters and will translate all subsequent upper-case letters to lower case.

The evidence that you have successfully logged in is that a shell program will type a prompt (either a '\$' or a '%' depending on the shell program activated) to you. (The shells are described below under *How to Run a Program — the Shells.*)

For more information, consult *tset(1)*, and *stty(1)*, which tell how to adjust terminal behavior, *getty(8)*, which discusses the login sequence in more detail, and *tty(4)*, which discusses terminal I/O.

Logging Out

Logging out is a process where the active shell is terminated and all files accessed by the *logged in* user are saved to the system hard disk. This process maintains the integrity and security of the ICON/UXB file system.

There are three ways to log out of the ICON/UXB operating system:

- By typing an end-of-text indicator (EOT character, **control-d**, shown as d) to the shell. The shell will terminate and the **login:** message will appear again.
- You can type **logout**, if you are using the **C** shell. The shell will terminate and return the **login:** prompt. (If you have several shells active, you must enter the **exit** command to terminate the active shell until the actual *login* shell is reached.)
- You can login directly as yourself or another user by entering the *login* command. (See *login (1)*.)

If you are communicating over a phone line, and worse comes to worse, you can simply hang up the phone; but beware — some machines may lack the necessary hardware to detect that the phone has been hung up. Ask your system administrator if this is a problem on your machine. (NOTE: Turning off your terminal, however, **does not** log you off the system. You must still use the **control-d** key sequence, or enter **logout**, or **login** to terminate the *login* shell.)

HOW TO COMMUNICATE THROUGH YOUR TERMINAL

When you type characters, a gnome deep in the system gathers your characters and saves them in a secret place. The characters will not be given to a program until you type a return (or newline), as described above in *Logging In*.

ICON/UXB terminal I/O is full-duplex. It has full read-ahead, which means that you can type at any time, even while a program is typing at you. Of course, if you type

during output, the printed output will have the input characters interspersed. However, whatever you type will be saved up and interpreted in 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 in trouble. When the read-ahead limit is exceeded, the system throws away all the saved characters (or beeps, if your prompt was a %).

The character '@' in typed input kills all the preceding characters in the line, so typing mistakes can be repaired on a single line. Also, the character '#' erases the last character typed. (Most users prefer to use a backspace rather than '#', and many prefer control-U instead of '@'; *tset(1)* or *stty(1)* can be used to arrange this.) Successive uses of '#' erase characters back to, but not beyond, the beginning of the line. '@' and '#' can be transmitted to a program by preceding them with '\'. (So, to erase '\', you need two '#'s).

The 'break' or 'interrupt' key causes an *interrupt signal*, as does the ASCII 'delete' (or 'rubout') character, which is not passed to programs. This signal generally causes whatever program you are running to terminate. It is typically 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). The editor, for example, catches interrupts and stops what it is doing, instead of terminating, so that an interrupt can be used to halt an editor printout without losing the file being edited. Many users change this interrupt character to be ^C (control-C) using *stty(1)*.

It is also possible to suspend output temporarily using ^S (control-s) and later resume output with ^Q. In a newer terminal driver, it is possible to cause output to be thrown away without interrupting the program by typing ^O; see *tty(4)*.

The *quit* signal is generated by typing the ASCII FS character. (FS appears many places on different terminals, most commonly as control-\ or control-!). It not only causes a running program to terminate but also generates a file with the core image of the terminated process. Quit is useful for debugging.

Besides adapting to the speed of the terminal, ICON/UXB tries to be intelligent about whether you have a terminal with the newline function or whether it must be simulated with carriage-return and line-feed. In the latter case, all input carriage returns are turned to newline characters (the standard line delimiter) and both a carriage return and a line feed are echoed to the terminal. If you get into the wrong mode, the *reset(1)* command will rescue you.

Tab characters are used freely in ICON/UXB source programs. If your terminal does not have the tab function, you can arrange to have them turned into spaces during output, and echoed as spaces during input. The system assumes that tabs are set every eight columns. Again, the *tset(1)* or *stty(1)* command will set or reset this mode. *Tset(1)* can be used to set the tab stops automatically when necessary.

HOW TO RUN A PROGRAM — THE SHELLS

When you have successfully logged in, a program called a *Shell* is listening to your terminal. The shell reads typed-in lines, splits them up into a command name and arguments, and executes the command. A command is simply an executable program. The shell looks in several system directories to find the command. You can also place commands in your own directory and have the shell find them there. There is nothing

special about system-provided commands except that they are kept in a directory where the shell can find them.

The command name is always the first word on an input line; it and its arguments are separated from one another by spaces.

When a program terminates, the shell will ordinarily regain control and type a prompt at you to indicate that it is ready for another command.

The shells have many other capabilities, which are described in detail in sections *sh(1)* and *cs(1)*. If the shell prompts you with '\$', then it is an instance of *sh(1)* the standard Bell-Labs provided shell. If it prompts with '%' then it is an instance of *cs(1)*, a shell written at U.C. Berkeley. The shells are different for all but the most simple terminal usage. Most users at Berkeley choose *cs(1)* because of the *history* mechanism and the *alias* feature, which greatly enhance its power when used interactively. The Berkeley *cs* also supports the job-control facilities. Refer to *cs(1)* or the *Introduction to the C Shell* in Volume 2 for details.

You can change from one shell to the other by using the *chsh(1)* command, which takes effect at your next login.

THE CURRENT DIRECTORY

ICON/UXB has a file system arranged in a hierarchy of directories. When the system administrator gave you a user name, he also created a directory for you (ordinarily with the same name as your user name). When you log in, any file name you type is by default in this directory. Since you are the owner of this directory, you have full permission to read, write, alter, or destroy its contents. Permissions to have your will with other directories and files will have been granted or denied to you by their owners. As a matter of observed fact, few ICON/UXB users protect their files from perusal by other users.

To change the current directory (but not the set of permissions you were endowed with at login) use *cd(1)*.

PATH NAMES

To refer to files not in the current directory, you must use a path name. Full path names begin with '/', the name of the root directory of the whole file system. After the slash comes the name of each directory containing the next sub-directory (followed by a '/') until finally the file name is reached. For example, */usr/lem/filex* refers to the file *filex* in the directory *lem*; *lem* is itself a subdirectory of *usr*; *usr* springs directly from the root directory */*.

If your current directory has subdirectories, the path names of files therein begin with the name of the subdirectory with no prefixed '/'.

A path name may be used anywhere a file name is required.

Important commands which modify the contents of files are *cp(1)*, *mv(1)*, and *rm(1)*, which respectively copy, move (i.e. rename) and remove files. To find out the status of files or directories, use *ls(1)*. See *mkdir(1)* for making directories and *rmdir* (in *rm(1)*) for destroying them.

For a fuller discussion of the file system, see *The UNIX® Time-Sharing System*, by Ken Thompson and Dennis Ritchie in Volume 2. It may also be useful to glance through Section 2 of this manual, which discusses system calls, even if you don't intend to deal with the system at that level.

WRITING AN ICON/UXB PROGRAM

To enter the text of a source program into an ICON/UXB file, use the editor *ex*(1) or its display editing alias *vi*(1). (The old standard editor *ed*(1) is also available.) The principal languages in ICON/UXB are provided by the C compiler *cc*(1), the Fortran compiler *f77*(1), the Pascal compiler *pc*(1), and interpreter *pi*(1) and *px*(1), and the Lisp system *lisp*(1). User contributed software in the latest release of the system supports APL, the Functional Programming language, and Icon. Refer to *apl*(1), *fp*(1), and *icon*(1), respectively for more information about each. After the program text has been entered through the editor and written on a file, you can give the file to the appropriate language processor as an argument. The output of the language processor will be left on a file in the current directory named 'a.out'. (If the output is precious, use *mv* to move it to a less exposed name soon.)

When you have finally gone through this entire process without provoking any diagnostics, the resulting program can be run by giving its name to the shell in response to the shell ('\$' or '%') prompt.

Your programs can receive arguments from the command line just as system programs do. Refer to the *execve*(2) command.

TEXT PROCESSING

Almost all text is entered through the editor *ex*(1) or its display-oriented counterpart *vi*(1). The commands most often used to write text on a terminal are: *cat*, *pr*, *more* and *nroff*, all described in section 1.

The *cat* command simply dumps ASCII text on the terminal, with no processing at all. The *pr* command paginates the text, supplies headings, and has a facility for multi-column output.

nroff is an elaborate text formatting program. Used naked, it requires careful forethought, but for ordinary documents it has been tamed; see *me*(7) and *ms*(7).

troff prepares documents for a Graphics Systems phototypesetter or a Versatec Plotter; it is very similar to *nroff*, and often works from exactly the same source text. It was used to produce this manual.

script(1) lets you keep a record of your session in a file, which can then be printed, mailed, etc. It provides the advantages of a hard-copy terminal even when using a display terminal.

more(1) is useful for preventing the output of a command from zipping off the top of your screen. It is also well suited to perusing files.

STATUS INQUIRIES

Various commands exist to provide you with useful information. *w(1)* prints a list of users presently logged in, and what they are doing. *date(1)* prints the current time and date. *ls(1)* will list the files in your directory or give summary information about particular files.

SURPRISES

Certain commands provide inter-user communication. Even if you do not plan to use them, it would be well to learn something about them, because someone else may aim them at you.

To communicate with another user currently logged in, *write(1)* is used; *mail(1)* will leave a message whose presence will be announced to another user when he next logs in. The write-ups in the manual also suggest how to respond to the two commands if you are a target.

If you use *cs(1)* the key *^Z* (control-Z) will cause jobs to *stop*. If this happens before you learn about it, you can simply continue by entering *fg* (for foreground) to bring the job back.

When you log in, a *message-of-the-day* may greet you before the first prompt. Be patient, especially if the message is *long*. Your shell prompt will eventually appear and you can go to work.

OTHER MANUALS AND BOOKS ON THE UNIX® OPERATING SYSTEM

To assist you in learning more about the UNIX operating system, the following is a recommended list of some excellent books and manuals for novice users of the UNIX® operating system. These are also invaluable resources and references for experienced UNIX® operating system users.

- *UNIX Primer Plus*, Waite, Martin, and Prata, Howard W. Sams, 1983
- *Understanding UNIX, A Conceptual Guide*, Groff and Weinberg, Que Corp., 1983
- *Exploring the UNIX System*, Kochan and Wood, Hayden Book, 1983
- *A User Guide to the UNIX System*, Thomas and Yates, Osborne/McGraw-Hill, 1985
- *UNIX Programmer's Manual*, Volumes 1 and 2, Bell Labs, AT&T, Holt, Rinehart and Winston, 1983
- *Introducing the UNIX System*, McGilton and Morgan, McGraw-Hill, 1983
- *Starting With UNIX*, Brown, Addison-Wesley, 1984
- *The UNIX Programming Environment*, Kernighan and Pike, Prentice-Hall, 1984

**ICON/UXB
OPERATING
SYSTEM
COMMANDS
& APPLICATION
PROGRAMS**



NAME

intro - introduction to commands

DESCRIPTION

This section describes publicly accessible commands in alphabetic order. Certain distinctions of purpose are made in the headings:

- (1) Commands of general utility.
- (1C) Commands for communication with other systems.
- (1G) Commands used primarily for graphics and computer-aided design.

N.B.: Commands related to system maintenance used to appear in section 1 manual pages and were distinguished by (1M) at the top of the page. These manual pages now appear in section 8.

SEE ALSO

Section (6) for computer games.

How to get started, in the Introduction.

DIAGNOSTICS

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, see *wait* and *exit(2)*. The former byte is 0 for normal termination, the latter is customarily 0 for successful execution, nonzero to indicate troubles such as erroneous parameters, bad or inaccessible data, or other inability to cope with the task at hand. It is called variously 'exit code', 'exit status' or 'return code', and is described only where special conventions are involved.

NAME

adb - debugger

SYNOPSIS

adb [-w] [-k] [-d] [-Idir] [objfil [corfil]]

DESCRIPTION

Adb is a general purpose debugging program. It may be used to examine files and to provide a controlled environment for the execution of UNIX programs.

Objfil is normally an executable program file, preferably containing a symbol table; if not then 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*; the default for *corfil* is **core**.

Requests to *adb* are read from the standard input and responses are to the standard output. If the **-w** flag is present then both *objfil* and *corfil* are created if necessary and opened for reading and writing so that files can be modified using *adb*.

The **-k** option allows *adb* to examine the running UNIX kernel. This option, on the ICON machines, selects the main-side division of UNIX, for disk-side operation, see the **-d** option. If this option is selected, *objfil* must be the directory (with an appended '/') where the running kernel resides, and *corfil* must be present and set to */dev/kmem*.

The **-d** option allows *adb* to examine the running disk cache processor's instruction/data space. If this option is selected, *objfil* must be the directory (with an appended '/') where the disk cache kernel resides, and *corfil* must be present and set to */dev/dmem*.

The **-I** option specifies a directory where files to be read with \$< or \$<< (see below) will be sought; the default is */usr/lib/adb*.

Adb ignores QUIT; INTERRUPT causes return to the next *adb* command.

In general requests to *adb* are of the form

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

If *address* is present then *dot* is set to *address*. Initially *dot* is set to 0. For most commands *count* specifies how many times the command will be executed. The default *count* is 1. *Address* and *count* are expressions.

The interpretation of an address depends on the context it is used in. If a subprocess is being debugged then addresses are interpreted in the usual way in the address space of the subprocess. If the operating system is being debugged, using the special file */dev/kmem* to interactive examine and/or modify memory, the maps are set to map the kernel addresses which start at 0x4000000. see ADDRESSES.

EXPRESSIONS

- . The value of *dot*.
- + The value of *dot* incremented by the current increment.
- ^ The value of *dot* decremented by the current increment.
- " The last *address* typed.

integer A number. The prefixes 0o and 0O ("zero oh") force interpretation in octal radix; the prefixes 0t and 0T force interpretation in decimal radix; the prefixes 0x and 0X force interpretation in hexadecimal radix. Thus 0o20 = 0t16 = 0x10 = sixteen. If no prefix appears, then the *default radix* is used; see the \$d command. The default radix

is initially hexadecimal. The hexadecimal digits are 0123456789abcdefABCDEF with the obvious values. Note that a hexadecimal number whose most significant digit would otherwise be an alphabetic character must have a 0x (or 0X) prefix (or a leading zero if the default radix is hexadecimal).

integer.fraction

A 32 bit floating point number.

'cccc' The ASCII value of up to 4 characters. \ may be used to escape a '.

< name

The value of *name*, which is either a variable name or a register name. *Adb* maintains a number of variables (see VARIABLES) named by single letters or digits. If *name* is a register name then the value of the register is obtained from the system header in *corfil*. The register names are those printed by the \$r command.

symbol A *symbol* is a sequence of upper or lower case letters, underscores or digits, not starting with a digit. The backslash character \ may be used to escape other characters. The value of the *symbol* is taken from the symbol table in *objfil*. An initial _ will be prepended to *symbol* if needed.

_ symbol

In C, the 'true name' of an external symbol begins with _. It may be necessary to utter this name to distinguish it from internal or hidden variables of a program.

(exp) The value of the expression *exp*.

Monadic operators

*exp The contents of the location addressed by *exp* in *corfil*.

@exp The contents of the location addressed by *exp* in *objfil*.

-exp Integer negation.

~ exp Bitwise complement.

#exp Logical negation.

Dyadic operators are left associative and are less binding than monadic operators.

e1+e2 Integer addition.

e1-e2 Integer subtraction.

e1*e2 Integer multiplication.

e1%e2 Integer division.

e1&e2 Bitwise conjunction.

e1|e2 Bitwise disjunction.

e1#e2 E1 rounded up to the next multiple of e2.

COMMANDS

Most commands consist of a verb followed by a modifier or list of modifiers. The following verbs are available. (The commands '?' and '/' may be followed by '*'; see ADDRESSES for further details.)

?f Locations starting at *address* in *objfil* are printed according to the format *f*. *dot* is incremented by the sum of the increments for each format letter (q.v.).

/f Locations starting at *address* in *corfil* are printed according to the format *f* and *dot* is incremented as for '?'.

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

A *format* consists of one or more characters that specify a style of printing. Each format character may be preceded by a decimal integer that is a repeat count for the format character. While stepping through a format *dot* is incremented by the amount given for each format letter. If no format is given then the last format is used. The format letters available are as follows.

- o** 2 Print 2 bytes in octal. All octal numbers output by *adb* are preceded by 0.
- O** 4 Print 4 bytes in octal.
- q** 2 Print in signed octal.
- Q** 4 Print long signed octal.
- d** 2 Print in decimal.
- D** 4 Print long decimal.
- x** 2 Print 2 bytes in hexadecimal.
- X** 4 Print 4 bytes in hexadecimal.
- u** 2 Print as an unsigned decimal number.
- U** 4 Print long unsigned decimal.
- f** 4 Print the 32 bit value as a floating point number.
- F** 8 Print double floating point.
- b** 1 Print the addressed byte in octal.
- c** 1 Print the addressed character.
- C** 1 Print the addressed character using the standard escape convention where control characters are printed as ^X and the delete character is printed as ^?.
- s** *n* Print the addressed characters until a zero character is reached.
- S** *n* Print a string using the ^X escape convention (see **C** above). *n* is the length of the string including its zero terminator.
- Y** 4 Print 4 bytes in date format (see *ctime*(3)).
- i** *n* Print as machine instructions. *n* is the number of bytes occupied by the instruction. This style of printing causes variables 1 and 2 to be set to the offset parts of the source and destination respectively.
- a** 0 Print the value of *dot* in symbolic form. Symbols are checked to ensure that they have an appropriate type as indicated below.
 - / local or global data symbol
 - ? local or global text symbol
 - = local or global absolute symbol
- p** 4 Print the addressed value in symbolic form using the same rules for symbol lookup as **a**.
- t** 0 When preceded by an integer tabs to the next appropriate tab stop. For example, **8t** moves to the next 8-space tab stop.
- r** 0 Print a space.
- n** 0 Print a newline.
- "..."** 0 Print the enclosed string.
- ^** *Dot* is decremented by the current increment. Nothing is printed.
- +** *Dot* is incremented by 1. Nothing is printed.
- *Dot* is decremented by 1. Nothing is printed.

newline

Repeat the previous command with a *count* of 1.

[?/]l value mask

Words starting at *dot* are masked with *mask* and compared with *value* until a match is found. If **L** is used then the match is for 4 bytes at a time instead of 2. If no match

is found then *dot* is unchanged; otherwise *dot* is set to the matched location. If *mask* is omitted then -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 *b1 e1 f1*[?/]

New values for (*b1*, *e1*, *f1*) are recorded. If less than three expressions are given then the remaining map parameters are left unchanged. If the '?' or '/' is followed by '*' then the second segment (*b2*, *e2*, *f2*) of the mapping is changed. If the list is terminated by '?' or '/' then the file (*objfil* or *corfil* respectively) is used for subsequent requests. (So that, for example, '/m?' will cause '/' to refer to *objfil*.) see ADDRESSES.

> *name*

Dot is assigned to the variable or register named.

! A shell (/bin/sh) is called to read the rest of the line following '!.

\$*modifier*

Miscellaneous commands. The available *modifiers* are:

- < *f* Read commands from the file *f*. If this command is executed in a file, further commands in the file are not seen. If *f* is omitted, the current input stream is terminated. If a *count* is given, and is zero, the command will be ignored. The value of the count will be 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 during the execution of this command, and restored when it completes. There is a (small) finite limit to the number of << files that can be open at once.
- > *f* Append output to the file *f*, which is created if it does not exist. If *f* is omitted, output is returned to the terminal.
- ? Print process id, the signal which caused stoppage or termination, as well as the registers as \$r. This is the default if *modifier* is omitted.
- r Print the general registers and the instruction addressed by **pc**. *Dot* is set to **pc**.
- F If there is an MC68881 floating point coprocessor, print the value in each register, in double format.
- b Print all breakpoints and their associated counts and commands.
- D Delete all breakpoints and their associated counts and commands.
- c C stack backtrace. If *address* is given then it is taken as the address of the current frame instead of the contents of the frame-pointer register. If **C** is used then the names and (32 bit) values of all automatic and static variables are printed for each active function. (broken on the VAX). If *count* is given then 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 "0t10\$d".
- e The names and values of external variables are printed.
- w Set the page width for output to *address* (default 80).
- s Set the limit for symbol matches to *address* (default 255).
- o All integers input are regarded as octal.
- q Exit from *adb*.
- v Print all non zero variables in octal.
- m Print the address map.

- p** (*Kernel debugging*) Change the current kernel memory mapping to map the designated **user structure** to the address given by the symbol *_u*. The *address* argument is the address of the user's user page table entries (on the VAX).

:modifier

Manage a subprocess. Available modifiers are:

- bc** Set breakpoint at *address*. The breakpoint is executed *count*-1 times before causing a stop. Each time the breakpoint is encountered the command *c* is executed. If this command is omitted or sets *dot* to zero then the breakpoint causes a stop.
- d** Delete breakpoint at *address*.
- r** Run *objfil* as a subprocess. If *address* is given explicitly then the program is entered at this point; otherwise the program is entered at its standard entry point. *count* specifies how many breakpoints are to be 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.
- cs** The subprocess is continued with signal *s*, see *sigvec(2)*. If *address* is given then the subprocess is continued at this address. If no signal is specified then the signal that caused the subprocess to stop is sent. Breakpoint skipping is the same as for **r**.
- ss** As for **c** except that the subprocess is single stepped *count* times. If there is no current subprocess then *objfil* 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.
- n** Stops at the next address after the present one. If the next instruction is any kind of *jsr*, a breakpoint is set at the instruction after it, and the process is continued, otherwise a single step is taken. *Jmp*'s and *bcc*'s are followed.
- u** Looks at the current call frame for the return address, sets a breakpoint at that address, then continues the process, thus popping you back up the call tree. Because it looks at the current call frame, if the *up* command is used before the link instruction is executed, you will be popped up two call levels.
- k** The current subprocess, if any, is terminated.

VARIABLES

Adb provides a number of 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 \$< or \$<< command.

On entry the following are set from the system header in the *corfil*. If *corfil* does not appear to be a **core** file then these values are set from *objfil*.

- b** The base address of the data segment.
- d** The data segment size.
- e** The entry point.
- m** The 'magic' number (0407, 0410 or 0413).

s The stack segment size.
t The text segment size.

ADDRESSES

The address in a file associated with a written address is determined by a mapping associated with that file. Each mapping is represented by two triples ($b1$, $e1$, $f1$) and ($b2$, $e2$, $f2$) and the file address corresponding to a written address is calculated as follows.

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

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

otherwise, the requested address is not legal. In some cases (e.g. for programs with separated I and D space) the two segments for a file may overlap. If a ? or / is followed by an * then only the second triple is used.

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

FILES

a.out
core

SEE ALSO

cc(1), dbx(1), ptrace(2), a.out(5), core(5)

DIAGNOSTICS

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

NOTES

The only commands that make any sense in kernel mode are those for displaying and changing data. Printing backtrace and registers, setting and displaying breakpoints, and process manipulations do not work. Data and instruction space are not really differentiated. On the main side, because addresses under 0x40000000 are in user space, and nothing can be guaranteed about their contents, from moment to moment, 0x40000000 is the lowest address readable in /dev/kmem. This is a property of /dev/kmem, not adb. Obviously only those for whom /dev/kmem has write permission turned on may change anything in the running kernel.

In most cases it is difficult to change floating point variables. If the integer equivalents (whatever the radix is) for the floating point number is known, they may be written with the /w command.

A variant syntax is used to simplify writing to the floating point registers on the floating point coprocessor. Rather than:

```
value>regn
```

as is used to write data and address registers, the following syntax is used (fvalue = floating point format):

```
>fpn fvalue
```

Setuid and setgid programs do not run correctly under *adb* because ptrace inhibits the setuid and setgid mechanism. This is to enhance system security.

BUGS

Since no shell is invoked to interpret the arguments of the *:r* command, the customary wildcard and variable expansions cannot occur.

Because of the manner in which the current C compiler pushes parameters on the stack, there is no way to differentiate between subroutines with one parameter, and those with none. Thus we always assume one parameter.

NAME

addbib - create or extend bibliographic database

SYNOPSIS

addbib [-p promptfile] [-a] database

DESCRIPTION

When this program starts up, answering "y" to the initial "Instructions?" prompt yields directions; typing "n" or RETURN skips them. *Addbib* then prompts for various bibliographic fields, reads responses from the terminal, and sends output records to a *database*. A null response (just RETURN) means to leave out that field. A minus sign (-) means to go back to the previous field. A trailing backslash allows a field to be continued on the next line. The repeating "Continue?" prompt allows the user either to resume by typing "y" or RETURN, to quit the current session by typing "n" or "q", or to edit the *database* with any system editor (*vi*, *ex*, *edit*, *ed*).

The *-a* option suppresses prompting for an abstract; asking for an abstract is the default. Abstracts are ended with a CTRL-d. The *-p* option causes *addbib* to use a new prompting skeleton, defined in *promptfile*. This file should contain prompt strings, a tab, and the key-letters to be written to the *database*.

The most common key-letters and their meanings are given below. *Addbib* insulates you from these key-letters, since it gives you prompts in English, but if you edit the bibliography file later on, you will need to know this information.

%A	Author's name
%B	Book containing article referenced
%C	City (place of publication)
%D	Date of publication
%E	Editor of book containing article referenced
%F	Footnote number or label (supplied by <i>refer</i>)
%G	Government order number
%H	Header commentary, printed before reference
%I	Issuer (publisher)
%J	Journal containing article
%K	Keywords to use in locating reference
%L	Label field used by <i>-k</i> option of <i>refer</i>
%M	Bell Labs Memorandum (undefined)
%N	Number within volume
%O	Other commentary, printed at end of reference
%P	Page number(s)
%Q	Corporate or Foreign Author (unreversed)
%R	Report, paper, or thesis (unpublished)
%S	Series title
%T	Title of article or book
%V	Volume number
%X	Abstract — used by <i>roffbib</i> , not by <i>refer</i>
%Y,Z	ignored by <i>refer</i>

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

An example is:

```
%A Bill Tuthill
%T Refer — A Bibliography System
%I Computing Services
%C Berkeley
%D 1982
%O UNX 4.3.5.
```

FILES

promptfile optional file to define prompting

SEE ALSO

refer(1), sortbib(1), roffbib(1), indxbib(1), lookbib(1)

AUTHORS

Al Stangenberger, Bill Tuthill

NAME

apply - apply a command to a set of arguments

SYNOPSIS

apply [**-ac**] [**-n**] *command* *args* ...

DESCRIPTION

Apply runs the named *command* on each argument *arg* in turn. Normally arguments are chosen singly; the optional number *n* specifies the number of arguments to be passed to *command*. If *n* is zero, *command* is run without arguments once for each *arg*. Character sequences of the form *%d* in *command*, where *d* is a digit from 1 to 9, are replaced by the *d*'th following unused *arg*. If any such sequences occur, *n* is ignored, and the number of arguments passed to *command* is the maximum value of *d* in *command*. The character '%' may be changed by the **-a** option.

Examples:

```
    apply echo *
is similar to ls(1);
    apply -2 cmp a1 b1 a2 b2 ...
compares the 'a' files to the 'b' files;
    apply -0 who 1 2 3 4 5
runs who(1) 5 times; and
    apply 'ln %1 /usr/joe' *
links all files in the current directory to the directory /usr/joe.
```

SEE ALSO

sh(1)

AUTHOR

Rob Pike

BUGS

Shell metacharacters in *command* may have bizarre effects; it is best to enclose complicated commands in single quotes ' '.

There is no way to pass a literal '%2' if '%' is the argument expansion character.

NAME

apropos - locate commands by keyword lookup

SYNOPSIS

apropos keyword ...

DESCRIPTION

Apropos shows which manual sections contain instances of any of the given keywords in their title. Each word is considered separately and case of letters is ignored. Words which are part of other words are considered thus looking for compile will hit all instances of 'compiler' also. Try

apropos password

and

apropos editor

If the line starts 'name(section) ...' you can do 'man section name' to get the documentation for it. Try 'apropos format' and then 'man 3s printf' to get the manual on the subroutine *printf*.

Apropos is actually just the **-k** option to the *man(1)* command.

FILES

/usr/lib/whatis

data base

SEE ALSO

man(1), whatis(1), catman(8)

AUTHOR

William Joy

NAME

ar - archive and library maintainer

SYNOPSIS

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

DESCRIPTION

Ar maintains groups of files combined into a single archive file. Its main use is to create and update library files as used by the loader. It can be used, though, for any similar purpose. **N.B:** This version of *ar* uses a ASCII-format archive which is portable among the various machines running UNIX. Programs for dealing with older formats are available: see *arcv*(8).

Key is one character from the set **drqtpmx**, optionally concatenated with one or more of **vuaibclo**. *Afile* is the archive file. The *names* are constituent files in the archive file. The meanings of the *key* characters are:

- d** Delete the named files from the archive file.
- r** Replace the named files in the archive file. If the optional character **u** is used with **r**, then only those files with 'last-modified' dates later than the archive files are replaced. If an optional positioning character from the set **abi** is used, then the *posname* argument must be present and specifies that new files are to be placed after (**a**) or before (**b** or **i**) *posname*. Otherwise new files are placed at the end.
- q** Quickly append the named files to the end of the archive file. Optional positioning characters are invalid. The command does not check whether the added members are already in the archive. Useful only to avoid quadratic behavior when creating a large archive piece-by-piece.
- t** Print a table of contents of the archive file. If no names are given, all files in the archive are tabled. If names are given, only those files are tabled.
- p** Print the named files in the archive.
- m** Move the named files to the end of the archive. If a positioning character is present, then the *posname* argument must be present and, as in **r**, specifies where the files are to be moved.
- x** Extract the named files. If no names are given, all files in the archive are extracted. In neither case does **x** alter the archive file. Normally the 'last-modified' date of each extracted file is the date when it is extracted. However, if **o** is used, the 'last-modified' date is reset to the date recorded in the archive.
- v** Verbose. Under the verbose option, *ar* gives a file-by-file description of the making of a new archive file from the old archive and the constituent files. When used with **t**, it gives a long listing of all information about the files. When used with **p**, it precedes each file with a name.
- c** Create. Normally *ar* will create *afile* when it needs to. The create option suppresses the normal message that is produced when *afile* is created.
- l** Local. Normally *ar* places its temporary files in the directory */tmp*. This option causes them to be placed in the local directory.

FILES

*/tmp/v** temporaries

SEE ALSO

lorder(1), ld(1), ranlib(1), ar(5), arcv(8)

BUGS

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

The 'last-modified' date of a file will not be altered by the `o` option if the user is not the owner of the extracted file, or the super-user.

NAME

as - M68020 assembler

SYNOPSIS

as [**-o** *objfile*] [*name ...*]

DESCRIPTION

As assembles the named files.

The output of the assembly is left on the file *objfile*; if that is omitted, *name.o* is used.

FILES

<i>/tmp/as*</i>	default temporary files
<i>a.out</i>	default resultant object file

SEE ALSO

ld(1), *nm(1)*, *adb(1)*, *dbx(1)*, *a.out(5)*
Auxiliary documentation *Assembler Reference Manual*.

NAME

at - execute commands at a later time

SYNOPSIS

at time [day] [file]

DESCRIPTION

At squirrels away a copy of the named *file* (standard input default) to be used as input to *sh*(1) (or *cs**h*(1) if you normally use it) at a specified later time. A *cd* command to the current directory is inserted at the beginning, followed by assignments to all environment variables (excepting the variable *TERM*, which is useless in this context.) When the script is run, it uses the user and group ID of the creator of the copy file.

The *time* is 1 to 4 digits, with an optional following 'A', 'P', 'N' or 'M' for AM, PM, noon or midnight. One and two digit numbers are taken to be hours, three and four digits to be hours and minutes. If no letters follow the digits, a 24 hour clock time is understood.

The optional *day* is either (1) a month name followed by a day number, or (2) a day of the week; if the word 'week' follows invocation is moved seven days further off. Names of months and days may be recognizably truncated. Examples of legitimate commands are

```
at 8am jan 24
at 1530 fr week
```

At programs are executed by periodic execution of the command */usr/lib/atrun* from *cron*(8). The granularity of *at* depends upon how often *atrun* is executed.

Standard output or error output is lost unless redirected.

FILES

<i>/usr/lib/atrun</i>	executor (run by <i>cron</i> (8)).
in <i>/usr/spool/at</i> :	
yy.ddd.hhhh.*	activity for year yy, day dd, hour hhhh.
lasttimedone	last hhhh
past	activities in progress

SEE ALSO

calendar(1), *pwd*(1), *sleep*(1), *cron*(8)

DIAGNOSTICS

Complains about various syntax errors and times out of range.

BUGS

Due to the granularity of the execution of */usr/lib/atrun*, there may be bugs in scheduling things almost exactly 24 hours into the future.

NAME

awk - pattern scanning and processing language

SYNOPSIS

awk [-Fc] [prog] [file] ...

DESCRIPTION

Awk scans each input *file* for lines that match any of a set of patterns specified in *prog*. With each pattern in *prog* there can be an associated action that will be performed when a line of a *file* matches the pattern. The set of patterns may appear literally as *prog*, or in a file specified as *-f file*.

Files are read in order; if there are no files, the standard input is read. The file name '-' means the standard input. Each line is matched against the pattern portion of every pattern-action statement; the associated action is performed for each matched pattern.

An input line is made up of fields separated by white space. (This default can be changed by using FS, *vide infra*.) The fields are denoted \$1, \$2, ... ; \$0 refers to the entire line.

A pattern-action statement has the form

```
pattern { action }
```

A missing { action } means print the line; a missing pattern always matches.

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

```
if ( conditional ) statement [ else statement ]
while ( conditional ) statement
for ( expression ; conditional ; expression ) statement
break
continue
{ [ statement ] ... }
variable = expression
print [ expression-list ] [ >expression ]
printf format [ , expression-list ] [ >expression ]
next # skip remaining patterns on this input line
exit # skip the rest of the input
```

Statements are terminated by semicolons, newlines or right braces. An empty expression-list stands for the whole line. Expressions take on string or numeric values as appropriate, and are built using the operators +, -, *, /, %, and concatenation (indicated by a blank). The C operators ++, --, +=, -=, *=, /=, and %= are also available in expressions. Variables may be scalars, array elements (denoted x[i]) or fields. Variables are initialized to the null string. Array subscripts may be any string, not necessarily numeric; this allows for a form of associative memory. String constants are quoted "...".

The *print* statement prints its arguments on the standard output (or on a file if *>file* is present), separated by the current output field separator, and terminated by the output record separator. The *printf* statement formats its expression list according to the format (see *printf(3S)*).

The built-in function *length* returns the length of its argument taken as a string, or of the whole line if no argument. There are also built-in functions *exp*, *log*, *sqrt*, and *int*. The last truncates its argument to an integer. *substr(s, m, n)* returns the *n*-character substring of *s* that begins at position *m*. The function *sprintf(fmt, expr, expr, ...)* formats the expressions according to the *printf(3S)* format given by *fmt* and returns the resulting string.

Patterns are arbitrary Boolean combinations (!, ||, &&, and parentheses) of regular expressions and relational expressions. Regular expressions must be surrounded by slashes and are as in *egrep*. Isolated regular expressions in a pattern apply to the entire line. Regular expressions may also occur in relational expressions.

A pattern may consist of two patterns separated by a comma; in this case, the action is performed for all lines between an occurrence of the first pattern and the next occurrence of the second.

A relational expression is one of the following:

```
expression matchop regular-expression
expression relop expression
```

where a relop is any of the six relational operators in C, and a matchop is either ~ (for contains) or !~ (for does not contain). A conditional is an arithmetic expression, a relational expression, or a Boolean combination of these.

The special patterns BEGIN and END may be used to capture control before the first input line is read and after the last. BEGIN must be the first pattern, END the last.

A single character *c* may be used to separate the fields by starting the program with

```
BEGIN { FS = "c" }
```

or by using the `-Fc` option.

Other variable names with special meanings include NF, the number of fields in the current record; NR, the ordinal number of the current record; FILENAME, the name of the current input file; OFS, the output field separator (default blank); ORS, the output record separator (default newline); and OFMT, the output format for numbers (default "%0.6g").

EXAMPLES

Print lines longer than 72 characters:

```
length > 72
```

Print first two fields in opposite order:

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

Add up first column, print sum and average:

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

Print fields in reverse order:

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

Print all lines between start/stop pairs:

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

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

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

SEE ALSO

`lex(1)`, `sed(1)`

A. V. Aho, B. W. Kernighan, P. J. Weinberger, *Awk - a pattern scanning and processing language*

BUGS

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

NAME

basename – strip filename affixes

SYNOPSIS

basename string [suffix]

DESCRIPTION

Basename deletes any prefix ending in '/' and the *suffix*, if present in *string*, from *string*, and prints the result on the standard output. It is normally used inside substitution marks `` in shell procedures.

This shell procedure invoked with the argument */usr/src/bin/cat.c* compiles the named file and moves the output to *cat* in the current directory:

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

SEE ALSO

sh(1)

NAME

bc - arbitrary-precision arithmetic language

SYNOPSIS

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

DESCRIPTION

Bc is an interactive processor for a language which resembles *C* but provides unlimited precision arithmetic. It takes input from any files given, then reads the standard input. The *-l* argument stands for the name of an arbitrary precision math library. The syntax for *bc* programs is as follows; L means letter a-z, E means expression, S means statement.

Comments

are enclosed in /* and */.

Names

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

Other operands

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)

Operators

+ - * / % ^ (% is remainder; ^ is power)
++ -- (prefix and postfix; apply to names)
== <= >= != < >
= += -= *= /= %= ^=

Statements

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

Function definitions

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

Functions in -l math library

s(x) sine
c(x) cosine
e(x) exponential
l(x) log

a(x) arctangent
j(n,x) Bessel function

All function arguments are passed by value.

The value of a statement that is an expression is printed unless the main operator is an assignment. Either semicolons or newlines may 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.

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

For example

```
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
    }
}
```

defines a function to compute an approximate value of the exponential function and

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

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

Bc is actually a preprocessor for *dc(1)*, which it invokes automatically, unless the *-c* (compile only) option is present. In this case the *dc* input is sent to the standard output instead.

FILES

/usr/lib/lib.b mathematical library
dc(1) desk calculator proper

SEE ALSO

dc(1)
L. L. Cherry and R. Morris, *BC - An arbitrary precision desk-calculator language*

BUGS

No *&&*, *||* or *!* operators.
For statement must have all three E's.
Quit is interpreted when read, not when executed.

NAME

biff - be notified if mail arrives and who it is from

SYNOPSIS

biff [yn]

DESCRIPTION

Biff informs the system whether you want to be notified when mail arrives during the current terminal session. The command

biff y

enables notification; the command

biff n

disables it. When mail notification is enabled, the header and first few lines of the message will be printed on your screen whenever mail arrives. A "biff y" command is often included in the file *.login* or *.profile* to be executed at each login.

Biff operates asynchronously. For synchronous notification use the MAIL variable of *sh*(1) or the *mail* variable of *cs*h(1).

SEE ALSO

*cs*h(1), *sh*(1), *mail*(1), *comsat*(8C)

NAME

binmail – send or receive mail among users

SYNOPSIS

```
/bin/mail [ + ] [ -i ] [ person ] ...
/bin/mail [ + ] [ -i ] -f file
```

DESCRIPTION

Note: This is the old version 7 UNIX system mail program. The default *mail* command is described in *Mail(1)*, and its binary is in the directory */usr/ucb*.

mail with no argument prints a user's mail, message-by-message, in last-in, first-out order; the optional argument *+* displays the mail messages in first-in, first-out order. For each message, it reads a line from the standard input to direct disposition of the message.

newline

Go on to next message.

d Delete message and go on to the next.

p Print message again.

- Go back to previous message.

s [*file*] ...

Save the message in the named *files* ('mbox' default).

w [*file*] ...

Save the message, without a header, in the named *files* ('mbox' default).

m [*person*] ...

Mail the message to the named *persons* (yourself is default).

EOT (control-D)

Put unexamined mail back in the mailbox and stop.

q Same as EOT.

!*command*

Escape to the Shell to do *command*.

* Print a command summary.

An interrupt normally terminates the *mail* command; the mail file is unchanged. The optional argument *-i* tells *mail* to continue after interrupts.

When *persons* are named, *mail* takes the standard input up to an end-of-file (or a line with just '.') and adds it to each *person's* 'mail' file. The message is preceded by the sender's name and a postmark. Lines that look like postmarks are prepended with '>'. A *person* is usually a user name recognized by *login(1)*. To denote a recipient on a remote system, prefix *person* by the system name and exclamation mark (see *uucp(1C)*).

The *-f* option causes the named file, for example, 'mbox', to be printed as if it were the mail file.

When a user logs in he is informed of the presence of mail.

FILES

/etc/passwd	to identify sender and locate persons
/usr/spool/mail/*	incoming mail for user *
mbox	saved mail
/tmp/ma*	temp file
/usr/spool/mail/*.lock	lock for mail directory
dead.letter	unmailable text

SEE ALSO

Mail(1), write(1), uucp(1C), uux(1C), xsend(1), sendmail(8)

BUGS

Race conditions sometimes result in a failure to remove a lock file.

Normally anybody can read your mail, unless it is sent by *xsend*(1). An installation can overcome this by making *mail* a set-user-id command that owns the mail directory.

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. *Year* can be between 1 and 9999. The *month* is a number between 1 and 12. The calendar produced is that for England and her colonies.

Try September 1752.

BUGS

The year is always considered to start in January even though this is historically naive. Beware that 'cal 78' refers to the early Christian era, not the 20th century.

NAME

calendar - reminder service

SYNOPSIS

calendar [-]

DESCRIPTION

Calendar consults the file 'calendar' in the current directory and prints out lines that contain today's or tomorrow's date anywhere in the line. Most reasonable month-day dates such as 'Dec. 7,' 'december 7,' '12/7,' etc., are recognized, but not '7 December' or '7/12'. If you give the month as "*" with a date, i.e. "* 1", that day in any month will do. On weekends 'tomorrow' extends through Monday.

When an argument is present, *calendar* does its job for every user who has a file 'calendar' in his login directory and sends him any positive results by *mail*(1). Normally this is done daily in the wee hours under control of *cron*(8).

The file 'calendar' is first run through the "C" preprocessor, */lib/cpp*, to include any other calendar files specified with the usual "#include" syntax. Included calendars will usually be shared by all users, maintained and documented by the local administration.

FILES

calendar
/usr/lib/calendar to figure out today's and tomorrow's dates
/etc/passwd
/tmp/cal*
/lib/cpp, egrep, sed, mail as subprocesses

SEE ALSO

at(1), cron(8), mail(1)

BUGS

Calendar's extended idea of 'tomorrow' doesn't account for holidays.

NAME

`cat` - catenate and print

SYNOPSIS

`cat [-u] [-n] [-s] [-v] file ...`

DESCRIPTION

Cat reads each *file* in sequence and displays it on the standard output. Thus

```
cat file
```

displays the file on the standard output, and

```
cat file1 file2 >file3
```

concatenates the first two files and places the result on the third.

If no input file is given, or if the argument '-' is encountered, *cat* reads from the standard input file. Output is buffered in 1024-byte blocks unless the standard output is a terminal, in which case it is line buffered. The `-u` option makes the output completely unbuffered.

The `-n` option displays the output lines preceded by lines numbers, numbered sequentially from 1. Specifying the `-b` option with the `-n` option omits the line numbers from blank lines.

The `-s` option crushes out multiple adjacent empty lines so that the output is displayed single spaced.

The `-v` option displays non-printing characters so that they are visible. Control characters print like `^X` for control-x; the delete character (octal 0177) prints as `^?`. Non-ascii characters (with the high bit set) are printed as `M-` (for meta) followed by the character of the low 7 bits. A `-e` option may be given with the `-v` option, which displays a '\$' character at the end of each line. Specifying the `-t` option with the `-v` option displays tab characters as `^I`.

SEE ALSO

`cp(1)`, `ex(1)`, `more(1)`, `pr(1)`, `tail(1)`

BUGS

Beware of `'cat a b >a'` and `'cat a b >b'`, which destroy the input files before reading them.

NAME

cb - C program beautifier

SYNOPSIS

cb

DESCRIPTION

Cb places a copy of the C program from the standard input on the standard output with spacing and indentation that displays the structure of the program.

NAME

`cc` - C compiler

SYNOPSIS

`cc` [option] ... file ...

DESCRIPTION

`Cc` is the UNIX C compiler. `Cc` accepts several types of arguments:

Arguments whose names end with `'c'` are taken to be C source programs; they are compiled, and each object program is left on the file whose name is that of the source with `'o'` substituted for `'c'`. The `'o'` file is normally deleted, however, if a single C program is compiled and loaded all at one go.

In the same way, arguments whose names end with `'s'` are taken to be assembly source programs and are assembled, producing a `'o'` file.

The following options are interpreted by `cc`. See `ld(1)` for load-time options.

- `-c` Suppress the loading phase of the compilation, and force an object file to be produced even if only one program is compiled.
- `-f` Compiles floating point operations to use the MC68881 floating point coprocessor. Also switches to versions of `libc.a` and `/usr/lib/libm.a` that use the floating point chip. Setting the environment variable `FP` to `m68881` has the same effect as specifying this flag. Code generated with this option will cause an "Illegal instruction" trap when executed on machines that do not have the floating point coprocessor chip installed.
- `-g` Have the compiler produce additional symbol table information for `dbx(1)`. Also pass the `-lg` flag to `ld(1)`.
- `-w` Suppress warning diagnostics.
- `-p` Arrange for the compiler to produce code which counts the number of times each routine is called. If loading takes place, replace the standard startup routine by one which automatically calls `monitor(3)` at the start and arranges to write out a `mon.out` file at normal termination of execution of the object program. An execution profile can then be generated by use of `prof(1)`.
- `-pg` Causes the compiler to produce counting code in the manner of `-p`, but invokes a run-time recording mechanism that keeps more extensive statistics and produces a `gmon.out` file at normal termination. Also, a profiling library is searched, in lieu of the standard C library. An execution profile can then be generated by use of `gprof(1)`.
- `-O` Invoke an object-code improver.
- `-R` Passed on to `as`, making initialized variables shared and read-only.
- `-S` Compile the named C programs, and leave the assembler-language output on corresponding files suffixed `'s'`.
- `-E` Run only the macro preprocessor on the named C programs, and send the result to the standard output.
- `-C` prevent the macro preprocessor from eliding comments.
- `-#` Debug flag. Prints the phases (including arguments) of the compiler that would be executed if this flag was not present.

-o *output*

Name the final output file *output*. If this option is used the file 'a.out' will be left undisturbed.

-D*name=def*

-D*name* Define the *name* to the preprocessor, as if by '#define'. If no definition is given, the name is defined as "1".

-U*name* Remove any initial definition of *name*.

-I*dir* '#include' files whose names do not begin with '/' are always sought first in the directory of the *file* argument, then in directories named in -I options, then in directories on a standard list.

Other arguments are taken to be either loader option arguments, or 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 loaded (in the order given) to produce an executable program with name **a.out**.

FILES

file.c	input file
file.o	object file
a.out	loaded output
/tmp/ctm?	temporary
/lib/cpp	preprocessor
/lib/c0	pass 1 of the compiler
/lib/c1	pass 2 of the compiler
/lib/c2	optional optimizer
/lib/crt0.o	runtime startoff
/lib/mcrt0.o	startoff for profiling
/usr/lib/gcrt0.o	startoff for gprof-profiling
/lib/libc.a	standard library, see <i>intro(3)</i>
/usr/lib/libc_p.a	profiling library, see <i>intro(3)</i>
/usr/include	standard directory for '#include' files
mon.out	file produced for analysis by <i>prof(1)</i>
gmon.out	file produced for analysis by <i>gprof(1)</i>

SEE ALSO

B. W. Kernighan and D. M. Ritchie, *The C Programming Language*, Prentice-Hall, 1978
 B. W. Kernighan, *Programming in C—a tutorial*
 D. M. Ritchie, *C Reference Manual*
 monitor(3), prof(1), gprof(1), adb(1), ld(1), dbx(1), as(1)

DIAGNOSTICS

The diagnostics produced by C itself are intended to be self-explanatory. Occasional messages may be produced by the assembler or loader.

NAME

cd - change working directory

SYNOPSIS

cd directory

DESCRIPTION

Directory becomes the new working directory. The process must have execute (search) permission in *directory*.

Because a new process is created to execute each command, *cd* would be ineffective if it were written as a normal command. It is therefore recognized and executed by the shells. In *cs(1)* you may specify a list of directories in which *directory* is to be sought as a subdirectory if it is not a subdirectory of the current directory; see the description of the *cdpath* variable in *cs(1)*.

SEE ALSO

cs(1), *sh(1)*, *pwd(1)*, *chdir(2)*

NAME

checknr - check nroff/troff files

SYNOPSIS

checknr [**-s**] [**-f**] [**-a.x1.y1.x2.y2.xn.yn**] [**-c.x1.x2.x3xn**] [*file ...*]

DESCRIPTION

checknr checks a list of *nroff*(1) or *troff*(1) input files for certain kinds of errors involving mismatched opening and closing delimiters and unknown commands. If no files are specified, *checknr* checks the standard input. Delimiters checked are:

- (1) Font changes using `\fx ... \fP`.
- (2) Size changes using `\sx ... \s0`.
- (3) Macros that come in open ... close forms, for example, the `.TS` and `.TE` macros which must always come in pairs.

checknr knows about the *ms*(7) and *me*(7) macro packages.

Additional pairs of macros can be added to the list using the `-a` option. This must be followed by groups of six characters, each group defining a pair of macros. The six characters are a period, the first macro name, another period, and the second macro name. For example, to define a pair `.BS` and `.ES`, use `-a.BS.ES`

The `-c` option defines commands which would otherwise be complained about as undefined.

The `-f` option requests *checknr* to ignore `\f` font changes.

The `-s` option requests *checknr* to ignore `\s` size changes.

checknr is intended to be used on documents that are prepared with *checknr* in mind, much the same as *lint*. It expects a certain document writing style for `\f` and `\s` commands, in that each `\fx` must be terminated with `\fP` and each `\sx` must be terminated with `\s0`. While it will work to directly go into the next font or explicitly specify the original font or point size, and many existing documents actually do this, such a practice will produce complaints from *checknr*. Since it is probably better to use the `\fP` and `\s0` forms anyway, you should think of this as a contribution to your document preparation style.

SEE ALSO

nroff(1), *troff*(1), *checkeq*(1), *ms*(7), *me*(7)

DIAGNOSTICS

Complaints about unmatched delimiters.

Complaints about unrecognized commands.

Various complaints about the syntax of commands.

AUTHOR

Mark Horton

BUGS

There is no way to define a 1 character macro name using `-a`.

Does not correctly recognize certain reasonable constructs, such as conditionals.

NAME

chfn - change finger entry

SYNOPSIS

chfn [loginname]

DESCRIPTION

Chfn is used to change information about users. This information is used by the finger program, among others. It consists of the user's "real life" name, office room number, office phone number, and home phone number. *Chfn* prompts the user for each field. Included in the prompt is a default value, which is enclosed between brackets. The default value is accepted simply by typing <return>. To enter a blank field, type the word 'none'. Below is a sample run:

```
Name [Biff Studsworth II]:
Room number (Exs: 597E or 197C) []: 521E
Office Phone (Ex: 1632) []: 1863
Home Phone (Ex: 987532) [5771546]: none
```

Chfn allows phone numbers to be entered with or without hyphens. Because *finger* only knows about UCB extensions, *chfn* will insist upon a four digit number (after the hyphens are removed) for office phone numbers. Also, room numbers must be in Evans or Cory; again, this is because of *finger*.

It is a good idea to run *finger* after running *chfn* to make sure everything is the way you want it.

The optional argument **loginname** is used to change another person's finger information. This can only be done by the super-user.

FILES

/etc/passwd, /etc/ptmp

SEE ALSO

finger(1), passwd(5)

BUGS

The encoding of the office and extension information is installation dependent.

For historical reasons, the user's name, etc are stored in the passwd file. This is a bad place to store the information. Rumors are that a data base is being developed to store this information, but don't hold your breath.

Because two users may try to write the passwd file at once, a synchronization method was developed. On rare occasions, a message that the password file is "busy" will be printed. In this case, *chfn* sleeps for a while and then tries to write to the passwd file again.

NAME

`chgrp` - change group

SYNOPSIS

`chgrp` [-f] group file ...

DESCRIPTION

Chgrp changes the group-ID of the *files* to *group*. The group may be either a decimal GID or a group name found in the group-ID file.

The user invoking *chgrp* must belong to the specified group and be the owner of the file, or be the super-user.

No errors are reported when the `-f` (force) option is given.

FILES

`/etc/group`

SEE ALSO

`chown(2)`, `passwd(5)`, `group(5)`

NAME

chmod - change mode

SYNOPSIS

chmod mode file ...

DESCRIPTION

The mode of each named file is changed according to *mode*, which may be absolute or symbolic. An absolute *mode* is an octal number constructed from the OR of the following modes:

4000	set user ID on execution
2000	set group ID on execution
1000	sticky bit, see <i>chmod(2)</i>
0400	read by owner
0200	write by owner
0100	execute (search in directory) by owner
0070	read, write, execute (search) by group
0007	read, write, execute (search) by others

A symbolic *mode* has the form:

[*who*] *op permission [op permission] ...*

The *who* part is a combination of the letters **u** (for user's permissions), **g** (group) and **o** (other). The letter **a** stands for all, or **ugo**. If *who* is omitted, the default is *a* but the setting of the file creation mask (see *umask(2)*) is taken into account.

Op can be **+** to add *permission* to the file's mode, **-** to take away *permission* and **=** to assign *permission* absolutely (all other bits will be reset).

Permission is any combination of the letters **r** (read), **w** (write), **x** (execute), **s** (set owner or group id) and **t** (save text - sticky). Letters **u**, **g** or **o** indicate that *permission* is to be taken from the current mode. Omitting *permission* is only useful with **=** to take away all permissions.

EXAMPLES

The first example denies write permission to others, the second makes a file executable:

```
chmod o-w file
chmod +x file
```

Multiple symbolic modes separated by commas may be given. Operations are performed in the order specified. The letter **s** is only useful with **u** or **g**.

Only the owner of a file (or the super-user) may change its mode.

NOTE

Currently all programs behave as if the sticky bit were set. But in the future this may have some functionality.

SEE ALSO

ls(1), chmod(2), stat(2), umask(2), chown(8)

NAME

chsh - change default login shell

SYNOPSIS

chsh name [shell]

DESCRIPTION

Chsh is a command similar to *passwd*(1) except that it is used to change the login shell field of the password file rather than the password entry. If no *shell* is specified then the shell reverts to the default login shell */bin/sh*. Otherwise only */bin/csh*, */bin/oldcsh*, or */usr/new/csh* can be specified as the shell unless you are the super-user.

An example use of this command would be

```
chsh bill /bin/csh
```

SEE ALSO

csh(1), passwd(1), passwd(5)

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 taken to be an RCS file, all others are assumed to be working files containing new revisions. *ci* deposits the contents of each working file into the corresponding RCS file.

Pairs of RCS files and working files may be specified in 3 ways (see also the example section of *co* (1)).

1) Both the RCS file and the working file are given. The RCS file name is of the form *path1/workfile,v* and the working file name is of the form *path2/workfile*, where *path1/* and *path2/* are (possibly different or empty) paths and *workfile* is a file name.

2) Only the RCS file is given. Then the working file is assumed to be in the current directory and its name is derived from the name of the RCS file by removing *path1/* and the suffix 'v'.

3) Only the working file is given. Then the name of the RCS file is derived from the name of the working file by removing *path2/* and appending the suffix 'v'.

If the RCS file is omitted or specified without a path, then *ci* looks for the RCS file first in the directory *./RCS* and then in the current directory.

For *ci* to work, the caller's login must be on the access list, except if the access list is empty or the caller is the superuser or the owner of the file. To append a new revision to an existing branch, the tip 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 may be broken with the *rcs* command.

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.

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 checked in, *ci* asks whether to reuse the previous log message. If the std. input is not a terminal, *ci* suppresses the prompt and uses the same log message for all files. See also *-m*.

The number of the deposited revision can be given by any of the options *-r*, *-f*, *-k*, *-l*, *-u*, or *-q* (see *-r*).

If the RCS file does not exist, *ci* creates it and deposits the contents of the working file as the initial revision (default number: 1.1). The access list is initialized to empty. Instead of the log message, *ci* requests descriptive text (see *-t* below).

-r[rev] assigns the revision number *rev* to the checked-in revision, releases the corresponding lock, and deletes the working file. This is also the default.

If *rev* is omitted, *ci* derives the new revision number from the caller's last lock. If the caller has locked the tip revision of a branch, the new revision is appended to that branch. The new revision number is obtained by incrementing the tip revision number. If the caller locked a non-tip revision, a new branch is started at

that revision by incrementing the highest branch number at that revision. The default initial branch and level numbers are 1. If the caller holds no lock, but he is the owner of the file and locking is not set to *strict*, then the revision is appended to 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 appended to that branch. The level number is obtained by incrementing the tip revision number of that branch. If *rev* indicates a non-existing branch, that branch is created with the initial revision numbered *rev.1*.

Exception: On the trunk, revisions can be appended to the end, but not inserted.

- f[*rev*]** forces a deposit; the new revision is deposited even 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 by 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 checkin.
- u[*rev*]** works like **-l**, except that the deposited revision is not locked. This is useful if one wants to process (e.g., compile) the revision immediately after checkin.
- 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.
- mmsg** uses the string *msg* as the log message for all revisions checked in.
- nname** assigns the symbolic name *name* to the number of the checked-in revision. *Ci* prints an error message if *name* is already assigned to another number.
- Nname** same as **-n**, except that it overrides a previous assignment of *name*.
- sstate** 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 supplied from the std. input, terminated with a line containing a single '.' or control-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 std. input is not a terminal.

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 otherwise.

FILE MODES

An RCS file created by *ci* inherits the read and execute permissions from the working file. If the RCS file exists already, *ci* preserves its read and execute permissions. *ci* always turns off all write permissions of RCS files.

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. This strategy makes links to RCS files useless.

IDENTIFICATION

Author: Walter F. Tichy, Purdue University, West Lafayette, IN, 47907.
Revision Number: 3.1 ; Release Date: 83/04/04 .
Copyright © 1982 by Walter F. Tichy.

SEE ALSO

co (1), *ident*(1), *rcs* (1), *rcsdiff* (1), *rcsintro* (1), *rcsmerge* (1), *rlog* (1), *rcsfile* (5), *scstores* (8).
Walter F. Tichy, "Design, Implementation, and Evaluation of a Revision Control System," in *Proceedings of the 6th International Conference on Software Engineering*, IEEE, Tokyo, Sept. 1982.

BUGS

NAME

clear - clear terminal screen

SYNOPSIS

clear

DESCRIPTION

Clear clears your screen if this is possible. It looks in the environment for the terminal type and then in */etc/termcap* to figure out how to clear the screen.

FILES

/etc/termcap terminal capability data base

NAME

cmp - compare two files

SYNOPSIS

cmp [-l] [-s] file1 file2

DESCRIPTION

The two files are compared. (If *file1* 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.

Options:

- l Print the byte number (decimal) and the differing bytes (octal) for each difference.
- s Print nothing for differing files; return codes only.

SEE ALSO

diff(1), comm(1)

DIAGNOSTICS

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

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 into the corresponding working file.

Pairs of RCS files and working files may be specified in 3 ways (see also the example section).

- 1) Both the RCS file and the working file are given. The RCS file name is of the form *path1/workfile,v* and the working file name is of the form *path2/workfile*, where *path1/* and *path2/* are (possibly different or empty) paths and *workfile* is a file name.
- 2) Only the RCS file is given. Then the working file is created in the current directory and its name is derived from the name of the RCS file by removing *path1/* and the suffix *'v'*.
- 3) Only the working file is given. Then the name of the RCS file is derived from the name of the working file by removing *path2/* and appending the suffix *'v'*.

If the RCS file is omitted or specified without a path, then *co* looks for the RCS file first in the directory *./RCS* and then in the current directory.

Revisions of an RCS file may 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 checkin must normally be locked. Locking a revision currently locked by another user fails. (A lock may be broken with the *rcs* (1) command.) *Co* with locking requires the caller to be on the access list of the RCS file, unless he is the owner of the file or the superuser, or the access list is empty. *Co* without locking is not subject to accesslist restrictions.

A revision is selected by number, checkin 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 may be attached to one of the options *-l*, *-p*, *-q*, or *-r*.

A *co* command applied to an RCS file with no revisions creates a zero-length file. *Co* always performs keyword substitution (see below).

- 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 std. 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.
- d**date* retrieves the latest revision on the selected branch whose checkin 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*:

22-April-1982, 17:20-CDT,

2:25 AM, Dec. 29, 1983,

Tue-PDT, 1981, 4pm Jul 21

(free format),

Fri, April 16 15:52:25 EST 1982 (output of *ctime*).

Most fields in the date and time may be defaulted. *Co* determines the defaults in the order year, month, day, hour, minute, and second (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. The date/time must be quoted if it contains spaces.

- 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 **-n** option of the commands *ci* and *rcs*.
- sstate** retrieves the latest revision on the selected branch whose state is set to *state*.
- w[*login*]** retrieves the latest revision on the selected branch which was checked in by the user with login name *login*. If the argument *login* is omitted, the caller's login is assumed.
- jjoinlist** generates a new revision which is the join 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 such 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 which is like *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 by the lines <<<<<<< *rev1*, =====, and >>>>>>> *rev3*.

For the initial pair, *rev2* may 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 option **-l** 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 checkout, *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 will be 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. 300. Class\$
\$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 <i>rcs -s</i> or <i>ci -s</i> .

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 otherwise.

EXAMPLES

Suppose the current directory contains a subdirectory 'RCS' with an RCS file 'io.c,v'. Then all of the following commands retrieve the latest revision from 'RCS/io.c,v' and store 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;
```

FILE MODES

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 checkout 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 checkout.

FILES

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 which contains the RCS file.

A number of temporary files are created. A semaphore file is created in the directory of the RCS file to prevent simultaneous update.

IDENTIFICATION

Author: Walter F. Tichy, Purdue University, West Lafayette, IN, 47907.
Revision Number: 3.1 ; Release Date: 83/04/04 .

Copyright © 1982 by Walter F. Tichy.

SEE ALSO

ci (1), ident(1), rcs (1), rcsdiff (1), rcsintro (1), rcsmerge (1), rlog (1), rcsfile (5), sccstorcs (8).
Walter F. Tichy, "Design, Implementation, and Evaluation of a Revision Control System," in *Proceedings of the 6th International Conference on Software Engineering*, IEEE, Tokyo, Sept. 1982.

LIMITATIONS

The option **-d** gets confused in some circumstances, and accepts no date before 1970. 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.

BUGS

The option **-j** does not work for files that contain lines with a single '.'.

NAME

col - filter reverse line feeds

SYNOPSIS

col [*-bfx*]

DESCRIPTION

Col reads the standard input and writes the standard output. It performs the line overlays implied by reverse line feeds (ESC-7 in ASCII) and by forward and reverse half line feeds (ESC-9 and ESC-8). *Col* is particularly useful for filtering multicolumn output made with the 'rt' command of *nroff* and output resulting from use of the *tbl(1)* preprocessor.

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.

If the *-b* option is given, *col* assumes that the output device in use is not capable of backspacing. In this case, if several characters are to appear in the same place, only the last one read will be taken.

The control characters SO (ASCII code 017), and SI (016) are assumed to start and end text in an alternate character set. The character set (primary or alternate) associated with each printing character read is remembered; on output, SO and SI characters are generated where necessary to maintain the correct treatment of each character.

Col normally converts white space to tabs to shorten printing time. If the *-x* option is given, this conversion is suppressed.

All control characters are removed from the input except space, backspace, tab, return, new-line, ESC (033) followed by one of 7, 8, 9, SI, SO, and VT (013). This last character is an alternate form of full reverse line feed, for compatibility with some other hardware conventions. All other non-printing characters are ignored.

SEE ALSO

troff(1), *tbl(1)*

BUGS

Can't back up more than 128 lines.

No more than 800 characters, including backspaces, on a line.

NAME

`colcrt` - filter nroff output for CRT previewing

SYNOPSIS

`colcrt` [-] [-2] [file ...]

DESCRIPTION

Colcrt provides virtual half-line and reverse line feed sequences for terminals without such capability, and on which overstriking is destructive. Half-line characters and underlining (changed to dashing '-') are placed on new lines in between the normal output lines.

The optional - suppresses all underlining. It is especially useful for previewing *allboxed* tables from *tbl(1)*.

The option -2 causes all half-lines to be printed, effectively double spacing the output. Normally, a minimal space output format is used which will suppress empty lines. The program never suppresses two consecutive empty lines, however. The -2 option is useful for sending output to the line printer when the output contains superscripts and subscripts which would otherwise be invisible.

A typical use of *colcrt* would be

```
tbl exum2.n | nroff -ms | colcrt - | more
```

SEE ALSO

nroff/troff(1), *col(1)*, *more(1)*, *ul(1)*

AUTHOR

William Joy

BUGS

Should fold underlines onto blanks even with the '-' option so that a true underline character would show; if we did this, however, *colcrt* wouldn't get rid of *cu'd* underlining completely.

Can't back up more than 102 lines.

General overstriking is lost; as a special case '†' overstruck with '-' or underline becomes '+'.
†

Lines are trimmed to 132 characters.

Some provision should be made for processing superscripts and subscripts in documents which are already double-spaced.

NAME

`colrm` - remove columns from a file

SYNOPSIS

`colrm` [`startcol` [`endcol`]]

DESCRIPTION

Colrm removes selected columns from a file. Input is taken from standard input. Output is sent to standard output.

If called with one parameter the columns of each line will be removed starting with the specified column. If called with two parameters the columns from the first column to the last column will be removed.

Column numbering starts with column 1.

SEE ALSO

`expand(1)`

AUTHOR

Jeff Schriebman

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 ASCII collating sequence, and produces a three column output: lines only in *file1*; lines only in *file2*; and lines in both files. The filename '-' means the standard input.

Flags 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` is a no-op.

SEE ALSO

`cmp(1)`, `diff(1)`, `uniq(1)`

NAME

`compact`, `uncompact`, `ccat` – compress and uncompress 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, the standard input is compacted 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 prepend a decoding tree to the compressed file since 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 and the resulting file is placed in *file.C*; *file* is unlinked. The first two bytes of the compacted file code the fact that the file is compacted. This code is used to prohibit recompaction.

The amount of compression to be expected depends on the type of file being compressed. Typical values of compression are: Text (38%), Pascal Source (43%), C Source (36%) and Binary (19%). These values are the percentages of file bytes reduced.

Uncompact restores the original file from a file compressed by *compact*. If no file names are given, the standard input is uncompact to the standard output.

Ccat cats the original file from a file compressed by *compact*, without uncompressing the file.

RESTRICTION

The last segment of the filename must contain fewer than thirteen characters to allow space for the appended '.C'.

FILES

*.C compacted file created by *compact*, removed by *uncompact*

SEE ALSO

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

Colin L. Mc Master

NAME

compress, uncompress, zcat - compress and expand data

SYNOPSIS

```
compress [-f] [-v] [-c] [-b bits] [ name ... ]
uncompress [-f] [-v] [-c] [ name ... ]
zcat [ name ... ]
```

DESCRIPTION

Compress reduces the size of the named files using adaptive Lempel-Ziv coding. Whenever possible, each file is replaced by one with the extension *.Z*, while keeping the same ownership modes, access and modification times. If no files are specified, the standard input is compressed to the standard output. Compressed files can be restored to their original form using *uncompress* or *zcat*.

The *-f* option will force compression of *name*, even if it does not actually shrink or the corresponding *name.Z* file already exists. Except when run in the background under */bin/sh*, if *-f* is not given the user is prompted as to whether an existing *name.Z* file should be overwritten.

The *-c* ("cat") option makes *compress/uncompress* write to the standard output; no files are changed. The nondestructive behavior of *zcat* is identical to that of *uncompress -c*.

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), pp. 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). *Bits* must be between 9 and 16. The default can be changed in the source to allow *compress* to be run on a smaller machine.

After the *bits* 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 decreases, *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 *bits* 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.

The amount of compression obtained depends on the size of the input, the number of *bits* per code, and the distribution of common substrings. Typically, text such as source code or English is reduced by 50-60%. Compression is generally much better than that achieved by Huffman coding (as used in *pack*), or adaptive Huffman coding (*compact*), and takes less time to compute.

The *-v* option causes the printing of the percentage reduction of each file.

If an error occurs, exit status is 1, else if the last file was not compressed because it became larger, the status is 2; else the status is 0.

DIAGNOSTICS

```
Usage: compress [-fvc] [-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 *xx* bits, can only handle *yy* bits

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

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 file cannot be compressed because its name is longer than 12 characters. Rename and try again. This message does not occur on BSD systems.

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

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

uncompress: corrupt input

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

Compression: *xx.xx*%

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

-- not a regular file: unchanged

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

-- has *xx* other links: unchanged

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

-- file unchanged

No savings is achieved by compression. The input remains virgin.

BUGS

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 (64KB or less, as exhibited by the DEC PDP series, the Intel 80286, etc.)

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

NAME

cp - copy

SYNOPSIS

```
cp [-i] [-r] [-t] file1 file2
cp [-i] [-r] [-t] file ... directory
```

DESCRIPTION

File1 is copied onto *file2*. The mode and owner of *file2* are preserved if it already existed; the mode of the source file is used otherwise.

In the second form, one or more *files* are copied into the *directory* with their original file-names.

Cp refuses to copy a file onto itself.

If the *-i* option is specified, *cp* will prompt the user with the name of the file whenever the copy will cause an old file to be overwritten. An answer of 'y' will cause *cp* to continue. Any other answer will prevent it from overwriting the file.

If the *-r* option is specified and any of the source files are directories, *cp* copies each subtree rooted at that name; in this case the destination must be a directory. Otherwise, if the *-r* flag is not specified, directories will not be copied.

If the *-t* option is specified, the time stamp on the destination file will be set to the time stamp of the source file.

SEE ALSO

cat(1), pr(1), mv(1)

NAME

cpio - copy file archives in and out

SYNOPSIS

```
cpio -o [ acBv ]
cpio -i [ BcdmrtuvfsSb6 ] [ patterns ]
cpio -p [ adlmruv ] directory
```

DESCRIPTION

Cpio -o (copy out) reads the standard input to get a list of path names and copies those files onto the standard output together with path name and status information.

Cpio -i (copy in) reads the standard input (which is assumed to be the product of a previous **Cpio -o** command), to get a list of files selected by zero or more *patterns* as defined in the name-generating notation of *sh*(1) or *cs*(1). In *patterns*, the meta-characters *?*, ***, and *[...]* match the slash (*/*) character. The default for *patterns* is *** (select all files).

Cpio -p (pass) copies out and in in a single operation. Destination pathnames are interpreted relative to the named *directory*.

OPTIONS

- a** Reset the access times of input files after they have been copied.
- B** Input/output is to be blocked at 5120 bytes to the record. This does not apply to the *pass* option. This option is only meaningful with data directed to or from */dev/rmt?*
- d** *Directories* should be created as needed.
- c** Write *header* information in ASCII character form for portability.
- r** Interactively *rename* files. If the user types a null line, the file is skipped.
- t** Print a *Table of contents* of the input. No files are created.
- u** Copy *unconditionally*. Normally, an older file will not replace a newer file with the same name.
- v** *Verbose* option. A list of filenames is displayed. When used with the **t** option, the table of contents looks like the output of an **ls -l** command (see *ls*(1)).
- l** Whenever possible, *link* files rather than copying them. Usable only with the **-p** option.
- m** Retain previous file modification time. This option is ineffective on directories that are being copied.
- f** Copy in all files except those in *patterns*.
- s** Swap bytes. User only with the **-i** option.
- S** Swap halfwords. Use only with the **-i** option.
- b** Swap both bytes and halfwords. Use only with the **-i** option.
- 6** Process an old (version 6 UNIX system) file. This is only useful with **-i** (copy in).

EXAMPLES

To copy the contents of a directory into an archive:

```
% ls | cpio -o > /dev/mt0
```

To duplicate the *olddir* directory hierarchy in the *newdir* directory:

```
% cd olddir  
% find . -print | cpio -pdl newdir
```

Some forms of *cpio* tapes from other sites have the bytes swapped in the file. The *s* option doesn't help since it only swaps the data bytes and not the header. To overcome this problem, use *dd* with the *conv=swab* option to swap *all* pairs of bytes (including the header), then pipe the output of *dd* through *cpio* with the *s* option to swap the data bytes back again:

```
% dd if=whatever the file is conv=swab | cpio -is
```

SEE ALSO

ar(1), find(1), cpio(5)

BUGS

Pathnames are restricted to 128 characters. If there are too many unique linked files, *cpio* runs out of memory to keep track of them and linking information is lost thereafter. Only the super-user can copy special files.

NAME

`crypt` - encode/decode

SYNOPSIS

`crypt` [*password*]

DESCRIPTION

Crypt reads from the standard input and writes on the standard output. The *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
```

will print the clear.

Files encrypted by *crypt* are compatible with those treated by the editor *ed* in encryption mode.

The security of encrypted files depends on three factors: the fundamental method must be hard to solve; direct search of the key space must be infeasible; 'sneak paths' by which keys or cleartext can become visible must be minimized.

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 known, but not widely; moreover the amount of work required is likely to be large.

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 (say) three lower-case 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 *ps*(1) or a derivative. To minimize this possibility, *crypt* takes care to destroy any record of the key immediately upon entry. No doubt the choice of keys and key security are the most vulnerable aspect of *crypt*.

FILES

/dev/tty for typed key

SEE ALSO

ed(1), *makekey*(8)

BUGS

There is no warranty of merchantability nor any warranty of fitness for a particular purpose nor any other warranty, either express or implied, as to the accuracy of the enclosed materials or as to their suitability for any particular purpose. Accordingly, Bell Telephone Laboratories assumes no responsibility for their use by the recipient. Further, Bell Laboratories assumes no obligation to furnish any assistance of any kind whatsoever, or to furnish any additional information or documentation.

NAME

csh - a shell (command interpreter) with C-like syntax

SYNOPSIS

csh [**-cefinstvVxX**] [arg ...]

DESCRIPTION

Csh is a first implementation of a command language interpreter incorporating a history mechanism (see **History Substitutions**) job control facilities (see **Jobs**) and a C-like syntax. So as to be able to use its job control facilities, users of *cs*h must (and automatically) use the new tty driver fully described in *tty*(4). This new tty driver allows generation of interrupt characters from the keyboard to tell jobs to stop. See *stty*(1) for details on setting options in the new tty driver.

An instance of *cs*h begins by executing commands from the file *.cshrc* in the *home* directory of the invoker. If this is a login shell then it also executes commands from the file *.login* there. It is typical for users on crt's to put the command "stty crt" in their *.login* file, and to also invoke *tset*(1) there.

In the normal case, the shell will then begin reading commands from the terminal, prompting with '% '. Processing of arguments and the use of the shell to process files containing command scripts will be described later.

The shell then repeatedly performs the following actions: a line of command input is read and broken into *words*. This sequence of words is placed on the command history list and then parsed. Finally each command in the current line is executed.

When a login shell terminates it executes commands from the file *.logout* in the users home directory.

Lexical Structure

The shell splits input lines into words at blanks and tabs with the following exceptions. The characters '&' '| ';' '<' '>' '(' ')' form separate words. If doubled in '&&', '| |', '<<' or '>>' these pairs form single words. These parser metacharacters may be made part of other words, or prevented their special meaning, by preceding them with '\'. A newline preceded by a '\' is equivalent to a blank.

In addition strings enclosed in matched pairs of quotations, '"', '' or ''', form parts of a word; metacharacters in these strings, including blanks and tabs, do not form separate words. These quotations have semantics to be described subsequently. Within pairs of '" or ''' characters a newline preceded by a '\' gives a true newline character.

When the shell's input is not a terminal, the character '#' introduces a comment which continues to the end of the input line. It is prevented this special meaning when preceded by '\' and in quotations using '"', ''', and ''''.

Commands

A simple command is a sequence of words, the first of which specifies the command to be executed. A simple command or a sequence of simple commands separated by '|' characters forms a pipeline. The output of each command in a pipeline is connected to the input of the next. Sequences of pipelines may be separated by ';', and are then executed sequentially. A sequence of pipelines may be executed without immediately waiting for it to terminate by following it with an '&'.

Any of the above may be placed in '(' ')' to form a simple command (which may be a component of a pipeline, etc.) It is also possible to separate pipelines with '| |' or '&&' indicating.

as in the C language, that the second is to be executed only if the first fails or succeeds respectively. (See *Expressions*.)

Jobs

The shell associates a *job* with each pipeline. It keeps a table of current jobs, printed by the *jobs* command, and assigns them small integer numbers. When a job is started asynchronously with '&', the shell prints a line which looks like:

```
[1] 1234
```

indicating that the jobs which was started asynchronously was job number 1 and had one (top-level) process, whose process id was 1234.

If you are running a job and wish to do something else you may hit the key **^Z** (control-Z) which sends a STOP signal to the current job. The shell will then normally indicate that the job has been 'Stopped', and print another prompt. You can then manipulate the state of this job, putting it in the background with the *bg* command, or run some other commands and then eventually bring the job back into the foreground with the foreground command *fg*. A **^Z** takes effect immediately and is like an interrupt in that pending output and unread input are discarded when it is typed. There is another special key **^Y** 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 wish to stop after it has read them.

A job being run in the background will stop if it tries to read from the terminal. Background jobs are normally allowed to produce output, but this can be disabled by giving the command "stty tostop". If you set this tty option, then background jobs will stop when they try to produce output like they do when they try to read input.

There are several ways to refer to jobs in the shell. 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 saying '%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' would normally restart a suspended *ex(1)* job, if there were only one suspended job whose name began 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.

The shell 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.

Status Reporting

This shell learns immediately whenever a process changes state. It normally informs you whenever a job becomes blocked so that no further progress is possible, but only just before it prints a prompt. This is done so that it does not otherwise disturb your work. If, however, you set the shell variable *notify*, the shell will notify you immediately of changes of status in background jobs. There is also a shell command *notify* which marks a single process so that its status changes will be immediately reported. By default *notify* marks the current process; simply say 'notify' after starting a background job to mark it.

When you try to leave the shell while jobs are stopped, you will be warned that 'You have stopped jobs.' You may use the *jobs* command to see what they are. If you do this or immediately try to exit again, the shell will not warn you a second time, and the suspended jobs will be terminated.

Substitutions

We now describe the various transformations the shell performs on the input in the order in which they occur.

History Substitutions

History substitutions place words from previous command input as portions of new commands, making it easy to repeat commands, repeat arguments of a previous command in the current command, or fix spelling mistakes in the previous command with little typing and a high degree of confidence. History substitutions begin with the character '!' and may begin **anywhere** in the input stream (with the proviso that they **do not** nest.) This '!' may be preceded by an '\' to prevent its special meaning; for convenience, a '!' is passed unchanged when it is followed by a blank, tab, newline, '=' or '?'. (History substitutions also occur when an input line begins with '^'. This special abbreviation will be described later.) Any input line which contains history substitution is echoed on the terminal before it is executed as it could have been typed without history substitution.

Commands input from the terminal which 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 size of which is controlled by the *history* variable; the previous command is always retained, regardless of its value. Commands are numbered sequentially from 1.

For definiteness, consider the following output from the *history* command:

```

 9 write michael
10 ex write.c
11 cat oldwrite.c
12 diff *write.c
```

The commands are shown with their event numbers. It is not usually necessary to use event numbers, but the current event number can be made part of the *prompt* by placing an '!' in the prompt string.

With the current event 13 we can refer to previous events by event number '!11', relatively as in '!-2' (referring to the same event), by a prefix of a command word as in '!d' for event 12 or '!wri' for event 9, or by a string contained in a word in the command as in '!?mic?' also referring to event 9. These forms, without further modification, simply reintroduce the words of the specified events, each separated by a single blank. As a special case '!!' refers to the previous command; thus '!!' alone is essentially a *redo*.

To select words from an event we can follow the event specification by a ':' and a designator for the desired words. The words of a input line are numbered from 0, the first (usually command) word being 0, the second word (first argument) being 1, etc. The basic word designators are:

```

0      first (command) word
n      n'th argument
↑      first argument, i.e. '1'
$      last argument
%      word matched by (immediately preceding)
       ?s? search
x-y    range of words
-y     abbreviates '0-y'
*      abbreviates '↑-$', or nothing if only 1 word
       in event
x*     abbreviates 'x-$'
x-     like 'x*' but omitting word '$'
```


The ':' separating the event specification from the word designator can be omitted if the argument selector begins with a '↑', '\$', '*' '-' or '%'. After the optional word designator can be placed a sequence of modifiers, each preceded by a ':'. The following modifiers are defined:

h	Remove a trailing pathname component, leaving the head.
r	Remove a trailing '.xxx' component, leaving the root name.
e	Remove all but the extension '.xxx' part.
s/l/r/	Substitute <i>l</i> for <i>r</i>
t	Remove all leading pathname components, leaving the tail.
&	Repeat the previous substitution.
g	Apply the change globally, prefixing the above, e.g. 'g&'.
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.

Unless preceded by a 'g' the modification is applied only to the first modifiable word. With substitutions, it is an error for no word to be applicable.

The left hand side of substitutions are not regular expressions in the sense of the editors, but rather strings. Any character may be used as the delimiter in place of '/'; a '\ ' quotes the delimiter into the *l* and *r* strings. The character '&' in the right hand side is replaced by the text from the left. A '\ ' quotes '&' also. A null *l* uses the previous string either from a *l* or from a contextual scan string *s* in '!s?'. The trailing delimiter in the substitution may be omitted if a newline follows immediately as may the trailing '?' in a contextual scan.

A history reference may be given without an event specification, e.g. '\$'. 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 '↑'. 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 may be surrounded with '{' and '}' if necessary to insulate it from the characters which follow. Thus, after 'ls -ld ~paul' we might do '{l}a' to do 'ls -ld ~paula', while '!la' would look for a command starting 'la'.

Quotations with ' and "

The quotation of strings by "" and "" can be used to prevent all or some of the remaining substitutions. Strings enclosed in "" are prevented any further interpretation. Strings enclosed in "" may be 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 "" quoted string yield parts of more than one word; "" quoted strings never do.

Alias Substitution

The shell maintains a list of aliases which 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, then the text which is the alias for that command is reread with the history mechanism available as though that command were the previous input line. The resulting words replace the command and argument list. If no reference is made to the history list, then the argument list is left unchanged.

Thus if the alias for 'ls' is 'ls -l' the command 'ls /usr' would map to 'ls -l /usr', the argument list here being undisturbed. Similarly if the alias for 'lookup' was 'grep !↑ /etc/passwd' then 'lookup bill' would map 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 reformed 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 we can 'alias print 'pr \!* | lpr'' to make a command which *pr*'s its arguments to the line printer.

Variable Substitution

The shell maintains a set of variables, each of which has as value a list of zero or more words. Some of these variables are set by the shell or referred to by it. For instance, the *argv* variable is an image of the shell's argument list, and words of this variable's value are referred to in special ways.

The values of variables may be displayed and changed by using the *set* and *unset* commands. Of the variables referred to by the shell a number are toggles; the shell does not care what their value is, only whether they are set or not. For instance, the *verbose* variable is a toggle which causes command input to be echoed. The setting of this variable results from the *-v* command line option.

Other operations treat variables numerically. The '@' command permits numeric calculations to be performed and the result assigned to a variable. Variable values are, however, always represented as (zero or more) strings. For the purposes of numeric operations, the null string is considered to be zero, and the second and subsequent words of multiword values are ignored.

After the input line is aliased and parsed, and before each command is executed, variable substitution is performed keyed by '\$' characters. This expansion can be prevented by preceding the '\$' with a '\' except within ""s where it **always** occurs, and within ' 's where it **never** occurs. Strings quoted by "" are interpreted later (see *Command substitution* below) so '\$' substitution does not occur there until later, if at all. A '\$' 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. It is thus possible for the first (command) word to this point to generate more than one word, the first of which becomes the command name, and the rest of which become arguments.

Unless enclosed in "" or given the ':q' modifier the results of variable substitution may eventually be command and filename substituted. Within "" a variable whose value consists of multiple words expands to a (portion of) a single word, with the words of the variables value separated by blanks. When the ':q' modifier is applied to a substitution the variable will expand to multiple words with each word separated by a blank and quoted to prevent later command or filename 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 which is not set.

`$name`

`${name}`

Are replaced by the words of the value of variable *name*, each separated by a blank. Braces insulate *name* from following characters which would otherwise be part of it. Shell variables have names consisting of up to 20 letters and digits starting with a letter. The underscore character is considered a letter.

If *name* is not a shell variable, but is set in the environment, then that value is returned (but `:` modifiers and the other forms given below are not available in this case).

`$name[selector]`

`${name[selector]}`

May be used to select only some of the words from the value of *name*. The selector is subjected to `'$'` substitution and may consist of a single number or two numbers separated by a `'-'`. The first word of a variables value is numbered `'1'`. If the first number of a range is omitted it defaults to `'1'`. If the last member of a range is omitted it defaults to `'$#name'`. The selector `'*'` selects all words. It is not an error for a range to be empty if the second argument is omitted or in range.

`$#name`

`${#name}`

Gives the number of words in the variable. This is useful for later use in a `'[selector]'`.

`$0`

Substitutes the name of the file from which command input is being read. An error occurs if the name is not known.

`$number`

`${number}`

Equivalent to `'$argv[number]'`.

`$*`

Equivalent to `'$argv[*]'`.

The modifiers `':h'`, `':t'`, `':r'`, `':q'` and `':x'` may be applied to the substitutions above as may `':gh'`, `':gt'` and `':gr'`. If braces `'{ '}'` appear in the command form then the modifiers must appear within the braces. **The current implementation allows only one `':'` modifier on each `'$'` expansion.**

The following substitutions may not be modified with `':'` modifiers.

`$?name`

`${?name}`

Substitutes the string `'1'` if *name* is set, `'0'` if it is not.

`$?0`

Substitutes `'1'` if the current input filename is known, `'0'` if it is not.

`$$`

Substitute 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.

Command and Filename Substitution

The remaining substitutions, command and filename substitution, are applied selectively to the arguments of builtin commands. This means that portions of expressions which 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 ```. 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 ```s, 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.

Filename Substitution

If a word contains any of the characters `*`, `?`, `[` or `{` or begins with the character `~`, then that word is a candidate for filename substitution, also known as 'globbing'. This word is then regarded as a pattern, and replaced with an alphabetically sorted list of file names which match the pattern. In a list of words specifying filename substitution it is an error for no pattern to match an existing file name, but it is not required for each pattern to match. Only the metacharacters `*`, `?` and `[` imply pattern matching, the characters `~` and `{` being more akin to abbreviations.

In matching filenames, the character `.` at the beginning of a filename or immediately following a `/`, as well as the character `/` must be matched explicitly. The character `*` matches any string of characters, including the null string. The character `?` matches any single character. The sequence `[...]` matches any one of the characters enclosed. Within `[...]`, a pair of characters separated by `-` matches any character lexically between the two.

The character `~` at the beginning of a filename is used to refer to home directories. Standing alone, i.e. `~` it expands to the invokers home directory as reflected in the value of the variable `home`. When followed by a name consisting of letters, digits and `-` characters the shell searches for a user with that name and substitutes their home directory; thus `~ken` might expand to `/usr/ken` and `~ken/chmach` to `/usr/ken/chmach`. If the character `~` is followed by a character other than a letter or `/` or appears not at the beginning of a word, it is left undisturbed.

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 `/usr/source/s1/oldls.c /usr/source/s1/ls.c` whether or not these files exist without any chance of error if the home directory for 'source' is `/usr/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 may be redirected with the following syntax:

`< name`

Open file *name* (which is first variable, command and filename expanded) as the standard input.

<< word

Read the shell input up to a line which is identical to *word*. *Word* is not subjected to variable, filename or command substitution, and each input line is compared to *word* before any substitutions are done on this input line. Unless a quoting '\', '"', or '' appears in *word* 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.

> name

>! name

>& name

>&! name

The file *name* is used as standard output. If the file does not exist then it is created; if the file exists, its is truncated, its previous contents being lost.

If the variable *noclobber* is set, then 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 '!' forms can be used and suppress this check.

The forms involving '&' route the diagnostic output into the specified file as well as the standard output. *Name* is expanded in the same way as '<' input filenames are.

>> name

>>& name

>>! name

>>&! name

Uses file *name* as standard output like '>' but places output at the end of the file. If the variable *noclobber* is set, then it is an error for the file not to exist unless one of the '!' forms is given. Otherwise 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 run from a file of shell commands 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 command scripts to function as components of pipelines and allows the shell to block read its input. Note that the default standard input for a command run detached is **not** modified to be the empty file '/dev/null'; rather the standard input remains as the original standard input of the shell. If this is a terminal and if the process attempts to read from the terminal, then the process will block and the user will be notified (see **Jobs** above.)

Diagnostic output may be directed through a pipe with the standard output. Simply use the form '|&' rather than just '|'.

Expressions

A number of the builtin commands (to be described subsequently) 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:

|| && | ↑ & == != =~ !~ <= >= < > << >> + - * / % ! ~
()

Here the precedence increases to the right, '==' '!=' ' =~ ' and '!~', '<=' '>=' '<' and '>', '<<' and '>>', '+' and '-', '*' '/' and '%' being, in groups, at the same level. The '==' '!=' ' =~ ' and '!~' operators compare their arguments as strings; all others operate on numbers. The operators ' =~ ' and '!~' are like '! =' and ' == ' except that the right hand side

is a *pattern* (containing, e.g. '*'s, '?'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 which begin with '0' are considered octal numbers. Null or missing arguments are considered '0'. The result of all expressions are strings, which represent decimal numbers. It is important to note that no two components of an expression can appear in the same word; except when adjacent to components of expressions which are syntactically significant to the parser ('&' '!' '<' '>' '(' ')') they should be surrounded by spaces.

Also available in expressions as primitive operands are command executions enclosed in '{' and '}' and file enquiries of the form '-l name' 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 name is command and filename expanded and then tested to see if it has the specified relationship to the real user. If the file does not exist or is inaccessible then all enquiries return false, i.e. '0'. Command executions succeed, returning true, i.e. '1', if the command exits with status 0, otherwise they fail, returning false, i.e. '0'. If more detailed status information is required then the command should be executed outside of an expression and the variable *status* examined.

Control Flow

The shell contains a number of commands which 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 in 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 up 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 goto's will succeed on non-seekable inputs.)

Built-In Commands

Builtin commands are executed within the shell. If a builtin command occurs as any component of a pipeline except the last then it is executed in a subshell.

alias

alias name

alias name wordlist

The first form prints all aliases. The second form prints the alias for name. The final form assigns the specified *wordlist* as the alias of *name*; *wordlist* is command and filename substituted. *Name* is not allowed to be *alias* or *unalias*.

alloc

Shows the amount of dynamic core in use, broken down into used and free core, and address of the last location in the heap. With an argument shows each used and free block on the internal dynamic memory chain indicating its address, size, and whether it

is used or free. This is a debugging command and may not work in production versions of the shell; it requires a modified version of the system memory allocator.

bg

bg %job ...

Puts the current 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 name

chdir

chdir name

Change the shells working directory to directory *name*. If no argument is given then change to the home directory of the user.

If *name* is not found as a subdirectory of the current directory (and does not begin with '/', './' or '../'), then each component of the variable *cdpath* is checked to see if it has a subdirectory *name*. Finally, if all else fails but *name* is a shell variable whose value begins with '/', then this is tried to see if it is a directory.

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 *case* labels.

dirs

Prints the directory stack; the top of the stack is at the left, the first directory in the stack being the current directory.

echo wordlist

echo -n wordlist

The specified words are written to the shells standard output, separated by spaces, and terminated with a newline unless the -n option is specified.

else**end****endif****endsw**

See the description of the *foreach*, *if*, *switch*, and *while* statements below.

eval arg ...

(As in *sh*(1).) The arguments are read as input to the shell and the resulting command(s) executed in the context of the current shell. This is usually used to execute commands generated as the result of command or variable substitution, since parsing occurs before these substitutions. See *tset*(1) for an example of using *eval*.

exec command

The specified command is executed in place of the current shell.

exit**exit**(*expr*)

The shell exits either with the value of the *status* variable (first form) or with the value of the specified *expr* (second form).

fg**fg** %*job* ...

Brings the current 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 builtin command *continue* may be used to continue the loop prematurely and the builtin command *break* to terminate it prematurely. When this command is read from the terminal, the loop is read up once prompting with '?' before any statements in the loop are executed. If you make a mistake typing in a loop at the terminal you can rub it out.

glob *wordlist*

Like *echo* but no '\ ' escapes are recognized and words are delimited by null characters in the output. Useful for programs which wish to use the shell to filename expand a list of words.

goto *word*

The specified *word* is filename 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 *exec*'s). An *exec* is attempted for each component of the *path* where the hash function indicates a possible hit, and in each component which does not begin with a '/ '.

history**history** *n***history** -r *n***history** -h *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 causes the history list to be printed without leading numbers. This is used to produce files suitable for sourcing using the -h option to *source*.

if (*expr*) *command*

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

if (expr) then

...
else if (expr2) then

...
else

...
endif

If the specified *expr* is true then the commands to the first *else* are executed; else if *expr2* is true then the commands to the second *else* are executed, etc. Any number of *else-if* pairs are possible; 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

jobs -l

Lists the active jobs; given the **-l** options lists process id's in addition to the normal 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 prefix "SIG"). The signal names are listed by "kill -l". There is no default, saying just 'kill' does not send a signal to the current job. If the signal being sent is TERM (terminate) or HUP (hangup), then the job or process will be sent a CONT (continue) signal as well.

limit

limit resource

limit resource maximum-use

Limits the consumption by the current process and each process it creates to not individually exceed *maximum-use* on the specified *resource*. If no *maximum-use* is given, then the current limit is printed; if no *resource* is given, then all limitations are given.

Resources controllable currently include *cpulimit* (the maximum number of cpu-seconds to be used by each process), *filesize* (the largest single file which can be created), *datasize* (the maximum growth of the data+stack region via *sbrk(2)* beyond the end of the program text), *stacksize* (the maximum size of the automatically-extended stack region), and *coredumpsize* (the size of the largest core dump that will be created).

The *maximum-use* may be given as a (floating point or integer) number followed by a scale factor. For all limits other than *cpulimit* the default scale is 'k' or 'kilobytes' (1024 bytes); a scale factor of 'm' or 'megabytes' may also be used. For *cpulimit* the default scaling is 'seconds', while 'm' for minutes or 'h' for hours, or a time of the form 'mm:ss' giving minutes and seconds may be used.

For both *resource* names and scale factors, unambiguous prefixes of the names suffice.

login

Terminate a login shell, replacing it with an instance of */bin/login*. This is one way to log off, included for compatibility with *sh(1)*.

logout

Terminate a login shell. Especially useful if *ignoreeof* is set.

nice**nice** +number**nice** command**nice** +number command

The first form sets the *nice* for this shell to 4. The second form sets the *nice* to the given number. The final two forms run command at priority 4 and *number* respectively. The super-user may specify negative niceness by using 'nice -number ...'. Command is always executed in a sub-shell, and the restrictions place on commands in simple *if* statements apply.

nohup**nohup** command

The first form can be used in shell scripts to cause hangups to be ignored for the remainder of the script. The second form causes the specified command to be run with hangups ignored. All processes detached with '&' are effectively *nohup*'ed.

notify**notify** %job ...

Causes the shell to notify the user asynchronously when the status of the current or specified jobs changes; normally notification is presented before a prompt. This is automatic if the shell variable *notify* is set.

onintr**onintr** -**onintr** label

Control the action of the shell on interrupts. The first form restores the default action of the shell on interrupts which is to terminate shell scripts or to return to the terminal command input level. The second form 'onintr -' causes all interrupts to be ignored. The final form causes the shell to execute a 'goto label' when an interrupt is received or a child process terminates because it was interrupted.

In any case, if the shell is running detached and interrupts are being ignored, all forms of *onintr* have no meaning and interrupts continue to be ignored by the shell and all invoked commands.

popd**popd** +n

Pops the directory stack, returning to the new top directory. With a argument '+n' discards the *n*th entry in the stack. The elements of the directory stack are numbered from 0 starting at the top.

pushd**pushd** name**pushd** +n

With no arguments, *pushd* exchanges the top two elements of the directory stack. Given a *name* argument, *pushd* changes to the new directory (ala *cd*) and pushes the old current working directory (as in *csd*) onto the directory stack. With a numeric argument, rotates the *n*th argument of the directory stack around to be the top element and changes to it. The members of the directory stack are numbered from the top starting at 0.

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 the command shows the value of all shell variables. Variables which have other than a single word as value 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 filename expanded.

These arguments may 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 variable USER, TERM, and PATH are automatically imported to and exported from the *cs*h variables *user*, *term*, and *path*; there is no need to use *setenv* for these.

shift**shift** variable

The members of *argv* are shifted to the left, discarding *argv*[1]. It is an error for *argv* not to be set or to have less than one word as value. The second form performs the same function on the specified variable.

source name**source** -h name

The shell reads commands from *name*. *Source* commands may be nested; if they are 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 causes the commands to be placed in the history list without being executed.

stop**stop** %job ...

Stops the current or specified job which is executing in the background.

suspend

Causes the shell to stop in its tracks, much as if it had been sent a stop signal with ^Z. This is most often used to stop shells started by *su*(1).

switch (string)**case** str1:

...

breaksw

...

default:

...

breaksw
endsw

Each case label is successively matched, against the specified *string* which is first command and filename expanded. The file metacharacters '*', '?' and '['...' may be used in the case labels, which are variable expanded. If none of the labels match before a 'default' label is found, then the execution begins after the default label. Each case label and the default label must appear at the beginning of a line. The command *breaksw* 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
time command

With no argument, a summary of time used by this shell and its children is printed. If arguments are given the specified 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
umask value

The file creation mask is displayed (first form) or set to the specified value (second form). The mask is given in octal. Common values for the mask are 002 giving all access to the group and read and execute access to others or 022 giving all access except no write access for users in the group or others.

unalias pattern

All aliases whose names match the specified pattern are discarded. Thus all aliases are removed by 'unalias *'. It is not an error for nothing to be *unaliased*.

unhash

Use of the internal hash table to speed location of executed programs is disabled.

unlimit resource

unlimit

Removes the limitation on *resource*. If no *resource* is specified, then all *resource* limitations are removed.

unset pattern

All variables whose names match the specified pattern are removed. Thus all variables are removed by 'unset *'; this has noticeably distasteful side-effects. It is not an error for nothing to be *unset*.

unsetenv pattern

Removes all variables whose name match the specified pattern from the environment. See also the *setenv* command above and *printenv*(1).

wait

All background jobs are waited for. If the shell is interactive, then an interrupt can disrupt the wait, at which time the shell prints names and job numbers of all jobs known to be outstanding.

while (expr)

...
end

While the specified expression evaluates non-zero, the commands between the *while* and the matching *end* are evaluated. *Break* and *continue* may be used to terminate or continue the loop prematurely. (The *while* and *end* must appear alone on their input lines.)

Prompting occurs here the first time through the loop as for the *foreach* statement if the input is a terminal.

%job

Brings the specified job into the foreground.

%job &

Continues the specified job in the background.

@

@ name = expr

@ name[index] = expr

The first form prints the values of all the shell variables. The second form sets the specified *name* to the value of *expr*. If the expression contains '<', '>', '&' or '|' then at least this part of the expression must be placed within '(' ')'. The third form assigns the value of *expr* to the *index*'th argument of *name*. Both *name* and its *index*'th component must already exist.

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

Special postfix '++' and '--' operators increment and decrement *name* respectively, i.e. '@ i++'.

Pre-Defined and Environment Variables

The following variables have special meaning to the shell. Of these, *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; these variables will not then be modified unless this is done explicitly by the user.

This shell copies the environment variable *USER* into the variable *user*, *TERM* into *term*, and *HOME* into *home*, and copies these back into the environment whenever the normal shell variables are reset. The environment variable *PATH* is likewise handled; it is not necessary to worry about its setting other than in the file *.cshrc* as inferior *csh* processes will import the definition of *path* from the environment, and re-export it if you then change it.

- argv** Set to the arguments to the shell, it is from this variable that positional parameters are substituted, i.e. '\$1' is replaced by '\$argv[1]', etc.
- cdpath** Gives a list of alternate directories searched to find subdirectories in *chdir* commands.
- cwd** The full pathname of the current directory.
- echo** Set when the *-x* command line option is given. Causes each command and its arguments to be echoed just before it is executed. For non-builtin commands all expansions occur before echoing. Builtin commands are echoed before command and filename substitution, since these substitutions are then done selectively.
- histchars** Can be given a string value to change the characters used in history substitution. The first character of its value is used as the history substitution character, replacing the default character '!'. The second character of its value replaces the character ↑ in quick substitutions.
- history** Can be given a numeric value to control the size of the history list. Any command which has been referenced in this many events will not be discarded. Too large values of *history* may run the shell out of memory. The last executed command is always saved on the history list.

- home** The home directory of the invoker, initialized from the environment. The filename expansion of '~' refers to this variable.
- ignoreeof** If set the shell ignores end-of-file from input devices which are terminals. This prevents shells from accidentally being killed by control-D's.
- mail** The files where the shell checks for mail. This is done after each command completion which will result in a prompt, if a specified interval has elapsed. The shell says 'You have new mail.' if the file exists with an access time not greater than its modify time.
- If the first word of the value of *mail* is numeric it specifies a different mail checking interval, in seconds, than the default, which is 10 minutes.
- If multiple mail files are specified, then the shell says 'New mail in *name*' when there is mail in the file *name*.
- noclobber** As described in the section on *Input/output*, restrictions are placed on output redirection to insure that files are not accidentally destroyed, and that '>>' redirections refer to existing files.
- noglob** If set, filename expansion is inhibited. This is most useful in shell scripts which are not dealing with filenames, or after a list of filenames has been obtained and further expansions are not desirable.
- nonomatch** If set, it is not an error for a filename expansion to not match any existing files; rather the primitive pattern is returned. It is still an error for the primitive pattern to be malformed, i.e. 'echo [' still gives an error.
- notify** If set, the shell notifies asynchronously of job completions. The default is to rather present 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 the current directory. If there is no *path* variable then only full path names will execute. The usual search path is '.', '/bin' and '/usr/bin', but this may vary from system to system. For the super-user the default search path is '/etc', '/bin' and '/usr/bin'. A shell which is given neither the *-c* nor the *-t* option will normally hash 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 may be necessary to give the *rehash* or the commands may not be found.
- prompt** The string which is printed before each command is read from an interactive terminal input. If a '!' appears in the string it will be replaced by the current event number unless a preceding '\ ' is given. Default is '% ', or '# ' for the super-user.
- savehist** is given a numeric value to control the number of entries of the history list that are saved in *~/.history* when the user logs out. Any command which has been referenced in this many events will be saved. During start up the shell sources *~/.history* into the history list enabling history to be saved across logins. Too large values of *savehist* will slow down the shell during start up.
- shell** The file in which the shell resides. This is used in forking shells to interpret files which have execute bits set, but which are not executable by the system. (See the description of *Non-builtin Command Execution* below.) Initialized to the (system-dependent) home of the shell.

- status** The status returned by the last command. If it terminated abnormally, then 0200 is added to the status. Builtin commands which fail return exit status '1', all other builtin commands set status '0'.
- time** Controls automatic timing of commands. If set, then any command which takes more than this many cpu seconds will cause a line giving user, system, and real times and a utilization percentage which is the ratio of user plus system times to real time to be printed when it terminates.
- verbose** Set by the `-v` command line option, causes the words of each command to be printed after history substitution.

Non-Built-In Command Execution

When a command to be executed is found to not be a builtin command the shell attempts to execute the command via `execve(2)`. Each word in the variable `path` names a directory from which the shell will attempt to execute the command. If it is given neither a `-c` nor a `-t` option, the shell will hash the names in these directories into an internal table so that it will only try an `exec` in a directory if there is a possibility that the command resides there. 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 the shell was given a `-c` or `-t` argument, and in any case for each directory component of `path` which does not begin with a '/', the shell concatenates with the given command name to form a path name of a file which it then attempts to execute.

Parenthesized commands are always executed in a subshell. Thus `(cd ; pwd) ; pwd` prints the `home` directory; leaving you where you were (printing this after the `home` directory), while `cd ; pwd` leaves you in the `home` directory. Parenthesized commands are most often used to prevent `chdir` from affecting the current shell.

If the file has execute permissions but is not an executable binary to the system, then it is assumed to be a file containing shell commands and a new shell is spawned to read it.

If there is an `alias` for `shell` then the words of the alias will be prepended to the argument list to form the shell command. The first word of the `alias` should be the full path name of the shell (e.g. `$shell`). Note that this is a special, late occurring, case of `alias` substitution, and only allows words to be prepended to the argument list without modification.

Argument List Processing

If argument 0 to the shell is '-' then this is a login shell. The flag arguments are interpreted as follows:

- `-c` Commands are read from the (single) following argument which must be present. Any remaining arguments are placed in `argv`.
- `-e` The shell exits if any invoked command terminates abnormally or yields a non-zero exit status.
- `-f` The shell will start faster, because it will neither search for nor execute commands from the file `.cshrc` in the invokers home directory.
- `-i` The shell is interactive and prompts for its top-level input, even if it appears to not be a terminal. Shells are interactive without this option if their inputs and outputs are terminals.
- `-n` Commands are parsed, but not executed. This aids in syntactic checking of shell scripts.
- `-s` Command input is taken from the standard input.
- `-t` A single line of input is read and executed. A '\ ' may be used to escape the newline at the end of this line and continue onto another line.

- v Causes the *verbose* variable to be set, with the effect that command input is echoed after history substitution.
- x Causes the *echo* variable to be set, so that commands are echoed immediately before execution.
- V Causes the *verbose* variable to be set even before '.cshrc' is executed.
- X Is to -x as -V is to -v.

After processing of flag arguments if arguments remain but none of the -c, -i, -s, or -t options was given the first argument is taken as the name of a file of commands to be executed. The shell opens this file, and saves its name for possible resubstitution by '\$0'. Since many systems use either the standard version 6 or version 7 shells whose shell scripts are not compatible with this shell, the shell will execute such a 'standard' shell if the first character of a script is not a '#', i.e. if the script does not start with a comment. Remaining arguments initialize the variable *argv*.

Signal Handling

The shell normally ignores *quit* signals. Jobs running detached (either by '&' or the *bg* or *%...* & commands) are immune to signals generated from the keyboard, including hangups. Other signals have the values which the shell inherited from its parent. The shells 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'.

AUTHOR

William Joy. Job control and directory stack features first implemented by J.E. Kulp of I.I.A.S.A, Laxenburg, Austria, with different syntax than that used now.

FILES

~/ .cshrc Read at beginning of execution by each shell.
~/ .login Read by login shell, after '.cshrc' at login.
~/ .logout Read by login shell, at logout.
/bin/sh Standard shell, for shell scripts not starting
 with a '#'.
/tmp/sh* Temporary file for '<<'.
/etc/passwd Source of home directories for '~name'.

LIMITATIONS

Words can be no longer than 1024 characters. The system limits argument lists to 10240 characters. The number of arguments to a command which involves filename expansion is limited to 1/6'th 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.

SEE ALSO

sh(1), access(2), execve(2), fork(2), killpg(2), pipe(2), sigvec(2), umask(2), setrlimit(2), wait(2), tty(4), a.out(5), environ(7), 'An introduction to the C shell'

BUGS

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

Shell builtin functions are not stoppable/restartable. Command sequences of the form 'a ; b ; c' are also not handled gracefully when stopping is attempted. If you suspend 'b', the shell will then immediately execute 'c'. This is especially noticeable if this expansion results from an *alias*. It suffices to place the sequence of commands in ()'s to force it to a subshell, i.e. '(a ; b ; c)'.

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.

Symbolic links fool the shell. In particular, *dirs* and 'cd ..' don't work properly once you've crossed through a symbolic link.

NAME

`ctags` - create a tags file

SYNOPSIS

`ctags` [`-BFatuwvx`] `name` ...

DESCRIPTION

`Ctags` makes a tags file for `ex(1)` from the specified C, Pascal and Fortran sources. A tags file gives the locations of specified objects (in this case functions and typedefs) in a group of files. Each line of the tags file contains the object name, the file in which it is defined, and an address specification for the object definition. Functions are searched with a pattern, typedefs with a line number. Specifiers are given in separate fields on the line, separated by blanks or tabs. Using the `tags` file, `ex` can quickly find these objects definitions.

If the `-x` flag is given, `ctags` produces a list of object names, the line number and file name on which each is defined, as well as the text of that line and prints this on the standard output. This is a simple index which can be printed out as an off-line readable function index.

If the `-v` flag is given, an index of the form expected by `vgrind(1)` is produced on the standard output. This listing contains the function name, file name, and page number (assuming 64 line pages). Since the output will be sorted into lexicographic order, it may be desired to run the output through `sort -f`. Sample use:

```
ctags -v files | sort -f > index
vgrind -x index
```

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` append to tags file.
- `-t` create tags for typedefs.
- `-w` suppressing warning diagnostics.
- `-u` causing the specified files to be *updated* in tags, that is, all references to them are deleted, and the new values are appended to the file. (Beware: this option is implemented in a way which is rather slow; it is usually faster to simply rebuild the `tags` file.)

The tag `main` is treated specially in C programs. The tag formed is created by prepending `M` to the name of the file, with a trailing `.c` removed, if any, and leading pathname components also removed. This makes use of `ctags` practical in directories with more than one program.

FILES

`tags` output tags file

SEE ALSO

`ex(1)`, `vi(1)`

AUTHOR

Ken Arnold; FORTRAN added by Jim Kleckner; Bill Joy added Pascal and `-x`, replacing `cxref`; C typedefs added by Ed Pelegri-Llopart.

BUGS

Recognition of **functions**, **subroutines** and **procedures** for FORTRAN and Pascal is done in a very simpleminded way. No attempt is made to deal with block structure; if you have two Pascal procedures in different blocks with the same name you lose.

The method of deciding whether to look for C or Pascal and FORTRAN functions is a hack.

Does not know about `#ifdefs`.

Should know about Pascal types. Relies on the input being well formed to detect typedefs. Use of `-tx` shows only the last line of typedefs.

NAME

date - print and set the date

SYNOPSIS

date [-z zone] [-d dst] [-u] [yymmddhhmm [.ss]]

DESCRIPTION

If no arguments are given, the current date and time are printed. If a date is specified, the current date is set. The *-z* flag is used to set your local timezone. Timezone is specified in minutes west of Greenwich, England. Thus eastern timezone would be specified *-z 300* and pacific time would be specified *-z 480*. The *-d* flag is used to set the type of daylight savings correction to be applied if any. No correction would be specified *-d 0* and standard USA style correction would be specified *-d 1*. The *-u* flag is used to display the date in GMT (universal) time. This flag may also be used to set GMT time. *yy* is the last two digits of the year; the first *mm* is the month number; *dd* is the day number in the month; *hh* is the hour number (24 hour system); the second *mm* is the minute number; *.ss* is optional and is the seconds. For example:

date 10080045

sets the date to Oct 8, 12:45 AM. The year, month and day may be omitted, the current values being the defaults. The system operates in GMT. *Date* takes care of the conversion to and from local standard and daylight time.

FILES

/usr/adm/wtmp to record time-setting
/etc/rc.local to set default timezone and dst flag at boot time

SEE ALSO

utmp(5)

DIAGNOSTICS

'Failed to set date: Not owner' if you try to change the date but are not the super-user.

NAME

dbx - debugger

SYNOPSISdbx [-r] [-i] [-I *dir*] [*objfile* [*coredump*]]**DESCRIPTION**

Dbx is a tool for source level debugging and execution of programs under UNIX. The *objfile* is an object file produced by a compiler with the appropriate flag (usually “-g”) specified to produce symbol information in the object file. Currently, *cc*(1) and *f77*(1) produce the appropriate source information and it is expected that in the future the Pascal compiler will also be able to generate source level information. The machine level facilities of *dbx* can be used on any program.

If no *objfile* is specified, *dbx* looks for a file named “a.out” in the current directory. The object file contains a symbol table which includes the name of the all the source files translated by the compiler to create it. These files are available for perusal while using the debugger.

If a file named “core” exists in the current directory or a *coredump* file is specified, *dbx* can be used to examine the state of the program when it faulted.

If the file “.dbxinit” exists in the current directory then the debugger commands in it are executed. *Dbx* also checks for a “.dbxinit” in the user’s home directory if there isn’t one in the current directory.

The command line options and their meanings are:

- r Execute *objfile* immediately. If it terminates successfully *dbx* exits. Otherwise the reason for termination will be reported and the user offered the option of entering the debugger or letting the program fault. *Dbx* will read from “/dev/tty” when -r is specified and standard input is not a terminal.
- i Force *dbx* to act as though standard input is a terminal.
- I *dir* Add *dir* to the list of directories that are searched when looking for a source file. Normally *dbx* looks for source files in the current directory and in the directory where *objfile* is located. The directory search path can also be set with the **use** command.

Unless -r is specified, *dbx* just prompts and waits for a command.

Execution and Tracing Commands

run [*args*] [< *filename*] [> *filename*]

Start executing *objfile*, passing *args* as command line arguments; < or > can be used to redirect input or output in the usual manner. If *objfile* has been written since the last time the symbolic information was read in, *dbx* will read in the new information.

trace [*in procedure/function*] [*if condition*]

trace *source-line-number* [*if condition*]

trace *procedure/function* [*in procedure/function*] [*if condition*]

trace *expression at source-line-number* [*if condition*]

trace *variable* [**in** *procedure/function*] [**if** *condition*]

Have tracing information printed when the program is executed. A number is associated with the command that is used to turn the tracing off (see the **delete** command).

The first argument describes what is to be traced. If it is a *source-line-number*, then the line is printed immediately prior to being executed. Source line numbers in a file other than the current one must be preceded by the name of the file in quotes and a colon, e.g. "mumble.p":17.

If the argument is a procedure or function name then every time it is called, information is printed telling what routine called it, from what source line it was called, and what parameters were passed to it. In addition, its return is noted, and if it's a function then the value it is returning is also printed.

If the argument is an *expression* with an **at** clause then the value of the expression is printed whenever the identified source line is reached.

If the argument is a variable then the name and value of the variable is printed whenever it changes. Execution is substantially slower during this form of tracing.

If no argument is specified then all source lines are printed before they are executed. Execution is substantially slower during this form of tracing.

The clause "**in** *procedure/function*" restricts tracing information to be printed only while executing inside the given procedure or function.

Condition is a boolean expression and is evaluated prior to printing the tracing information; if it is false then the information is not printed.

stop *if condition*

stop **at** *source-line-number* [**if** *condition*]

stop **in** *procedure/function* [**if** *condition*]

stop *variable* [**if** *condition*]

Stop execution when the given line is reached, procedure or function called, variable changed, or condition true.

status [**>** *filename*]

Print out the currently active **trace** and **stop** commands.

delete *command-number*

The trace or stop corresponding to the given number is removed. The numbers associated with traces and stops are printed by the **status** command.

catch *number*

ignore *number*

Start or stop trapping signal *number* before it is sent to the program. This is useful when a program being debugged handles signals such as interrupts. Initially all signals are trapped except SIGCONT, SIGCHILD, SIGALRM and SIGKILL.

cont Continue execution from where it stopped. Execution cannot be continued if the process has "finished", that is, called the standard procedure "exit". *Dbx* does not allow the process to exit, thereby letting the user to examine the program state.

step Execute one source line.

next Execute up to the next source line. The difference between this and **step** is that if the line contains a call to a procedure or function the **step** command will stop at the beginning of that block, while the **next** command will not.

Displaying and Naming Data

print *expression* [, *expression* ...]

Print out the values of the expressions. Array expressions are always subscripted by brackets ("["]). Variables having the same identifier as one in the current block may be referenced as "*block-name . variable*". The field reference operator (".") can be used with pointers as well as records, making the C operator "->" unnecessary (although it is supported). The construct *typename(expression)* can be used to print the *expression* out in the format of the named *type*.

whatis *name*

Print the declaration of the given name, which may be qualified with block names as above.

which *identifier*

Print the full qualification of the given identifier, i.e. the outer blocks that the identifier is associated with.

whereis *identifier*

Print the full qualification of all the symbols whose name matches the given identifier. The order in which the symbols are printed is not meaningful.

assign *variable* = *expression*

set *variable* = *expression*

Assign the value of the expression to the variable.

call *procedure(parameters)*

Execute the object code associated with the named procedure or function. Currently, calls to a procedure with a variable number of arguments are not possible. Also, string parameters are not passed properly for C.

where Print out a list of the active procedures and function.

dump [> *filename*]

Print the names and values of all active variables.

Accessing Source Files

edit [*filename*]

edit *procedure/function-name*

Invoke an editor on *filename* or the current source file if none is specified. If a *procedure* or *function* name is specified, the editor is invoked on the file that contains it. Which editor is invoked by default depends on the installation. The default can be overridden by setting the environment variable EDITOR to the name of the desired editor.

file [*filename*]

Change the current source file name to *filename*. If none is specified then the current source file name is printed.

func [*procedure/function*]

Change the current function. If none is specified then print the current function. Changing the current function implicitly changes the current source file to the one that contains the function; it also changes the current scope used for name resolution.

list [*source-line-number* [, *source-line-number*]]**list** *procedure/function*

List the lines in the current source file from the first line number to the second inclusive. If no lines are specified, the next 10 lines are listed. If the name of a procedure or function is given lines *n-k* to *n+k* are listed where *n* is the first statement in the procedure or function and *k* is small.

use *directory-list*

Set the list of directories to be searched when looking for source files.

Machine Level Commands**tracei** [*address*] [*if cond*]**tracei** [*variable*] [*at address*] [*if cond*]**stopi** [*address*] [*if cond*]**stopi** [*at*] [*address*] [*if cond*]

Turn on tracing or set a stop using a machine instruction address.

stepi

nexti Single step as in **step** or **next**, but do a single instruction rather than source line.

address, *address/* [*mode*][*address*] / [*count*] [*mode*]

Print the contents of memory starting at the first *address* and continuing up to the second *address* or until *count* items are printed. If no address is specified, the address following the one printed most recently is used. The *mode* specifies how memory is to be printed; if it is omitted the previous mode specified is used. The initial mode is "X". The following modes are supported:

- i** print the machine instruction
- d** print a short word in decimal
- D** print a long word in decimal
- o** print a short word in octal
- O** print a long word in octal
- x** print a short word in hexadecimal
- X** print a long word in hexadecimal
- b** print a byte in octal
- c** print a byte as a character
- s** print a string of characters terminated by a null byte
- f** print a single precision real number
- g** print a double precision real number

Symbolic addresses are specified by preceding the name with an "&". Registers are denoted by, "\$dN" where N is the number of a data register, and "\$aN" where N is the number of an address register. Addresses may be expressions made up of other addresses and the operators "+", "-", and indirection (unary "*").

Miscellaneous Commands

sh *command-line*

Pass the command line to the shell for execution. The SHELL environment variable determines which shell is used.

alias *new-command-name old-command-name*

Respond to *new-command-name* as though it were *old-command-name*.

help Print out a synopsis of *dbx* commands.

gripe Invoke a mail program to send a message to the person in charge of *dbx*.

source *filename*

Read *dbx* commands from the given *filename*. Especially useful when the *filename* has been created by redirecting a **status** command from an earlier debugging session.

quit Exit *dbx*.

FILES

a.out	object file
.dbxinit	initial commands

SEE ALSO

cc(1), f77(1), pc(1)

COMMENTS

Non-local gotos can cause some trace/stops to be missed. Most of the command names are too long. The alias facility helps, but is really quite weak. A *cs*h-like history capability would improve the situation. But then, who wants to duplicate the c-shell in a debugger?

Dbx suffers from the same "multiple include" malady as does *sdb*. If you have a program consisting of a number of object files and each is built from source files that include header files, the symbolic information for the header files is replicated in each object file. Since about one debugger start-up is done for each link, having the linker (*ld*) re-organize the symbol information won't save much time, though it would reduce some of the disk space used. The problem is an artifact of the unrestricted semantics of *#include*'s in C; for example an include file can contain static declarations that are separate entities for each file in which they are included.

NAME

dc - desk calculator

SYNOPSIS

dc [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. 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. The following constructions are recognized:

number

The value of the number is pushed on the stack. A number is an unbroken string of the digits 0-9. It may be preceded by an underscore `_` to input a negative number. Numbers may contain decimal points.

+ - / * % ^

The top two values on the stack are added (+), subtracted (-), multiplied (*), divided (/), remaindered (%), or exponentiated (^). The two entries are popped off the stack; the result is pushed on the stack in their place. Any fractional part of an exponent is ignored.

sx The top of the stack is popped and stored into a register named *x*, where *x* may be any character. If the **s** is capitalized, *x* is treated as a stack and the value is pushed on it.

lx The value in register *x* is pushed on the stack. The register *x* is not altered. All registers start with zero value. If the **l** is capitalized, register *x* 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 and in registers 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.

< x > x = x

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

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

! interprets the rest of the line as a UNIX command.

c All values on the stack are popped.

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

- o The top value on the stack is popped and used as the number radix for further output.
- O pushes the output base on the top of the stack.
- k the top of the stack is popped, and that value is used as a non-negative scale factor: the appropriate number of places are printed on output, and maintained during multiplication, division, and exponentiation. The interaction of scale factor, input base, and output base will be reasonable if all are changed together.
- z The stack level is pushed onto the stack.
- Z replaces the number on the top of the stack with its length.
- ? A line of input is taken from the input source (usually the terminal) and executed.
- ;: are used by *bc* for array operations.

An example which prints the first ten values of $n!$ is

```
[la1+dsa*pla10>y]sy
0sa1
lyx
```

SEE ALSO

bc(1), which is a preprocessor for *dc* providing infix notation and a C-like syntax which implements functions and reasonable control structures for programs.

DIAGNOSTICS

- 'x is unimplemented' where x is an octal number.
- 'stack empty' for not enough elements on the stack to do what was asked.
- 'Out of space' when the free list is exhausted (too many digits).
- 'Out of headers' for too many numbers being kept around.
- 'Out of pushdown' for too many items on the stack.
- 'Nesting Depth' for too many levels of nested execution.

NAME

`dd` - convert and copy a file

SYNOPSIS

`dd [option=value] ...`

DESCRIPTION

Dd copies the specified input file to the specified output with possible conversions. The standard input and output are used by default. The input and output block size may be specified to take advantage of raw physical I/O.

<i>option</i>	<i>values</i>
<code>if=</code>	input file name; standard input is default
<code>of=</code>	output file name; standard output is default
<code>ibs=<i>n</i></code>	input block size <i>n</i> bytes (default 512)
<code>obs=<i>n</i></code>	output block size (default 512)
<code>bs=<i>n</i></code>	set both input and output block size, superseding <i>ibs</i> and <i>obs</i> ; also, if no conversion is specified, it is particularly efficient since no copy need be done
<code>cbs=<i>n</i></code>	conversion buffer size
<code>skip=<i>n</i></code>	skip <i>n</i> input records before starting copy
<code>files=<i>n</i></code>	copy <i>n</i> input files before terminating (makes sense only where input is a magtape or similar device).
<code>seek=<i>n</i></code>	seek <i>n</i> records from beginning of output file before copying
<code>count=<i>n</i></code>	copy only <i>n</i> input records
<code>conv=ascii</code>	convert EBCDIC to ASCII
<code>ebcdic</code>	convert ASCII to EBCDIC
<code>ibm</code>	slightly different map of ASCII to EBCDIC
<code>block</code>	convert variable length records to fixed length
<code>unblock</code>	convert fixed length records to variable length
<code>lcase</code>	map alphabets to lower case
<code>ucase</code>	map alphabets to upper case
<code>swab</code>	swap every pair of bytes
<code>noerror</code>	do not stop processing on an error
<code>sync</code>	pad every input record to <i>ibs</i>
<code>... , ...</code>	several comma-separated conversions

Where sizes are specified, a number of bytes is expected. A number may end with **k**, **b** or **w** to specify multiplication by 1024, 512, or 2 respectively; a pair of numbers may be separated by **x** to indicate a product.

Cbs is used only if *ascii*, *unblock*, *ebcdic*, *ibm*, or *block* conversion is specified. In the first two cases, *cbs* characters are placed into the conversion buffer, any specified character mapping is done, trailing blanks trimmed and new-line added before sending the line to the output. In the latter three cases, characters are read into the conversion buffer, and blanks added to make up an output record of size *cbs*.

After completion, *dd* reports the number of whole and partial input and output blocks.

For example, to read an EBCDIC tape blocked ten 80-byte EBCDIC card images per record into the ASCII file *x*:

```
dd if=/dev/rmt0 of=x ibs=800 cbs=80 \
    conv=ascii,lcase
```

Note the use of raw magtape. *Dd* is especially suited to I/O on the raw physical devices because it allows reading and writing in arbitrary record sizes.

SEE ALSO

cp(1), tr(1)

DIAGNOSTICS

f+p records in(out): numbers of full and partial records read(written)

BUGS

The ASCII/EBCDIC conversion tables are taken from the 256 character standard in the CACM Nov, 1968. The 'ibm' conversion, while less blessed as a standard, corresponds better to certain IBM print train conventions. There is no universal solution.

One must specify "conv=noerror, sync" when copying raw disks with bad sectors to insure *dd* stays synchronized.

NAME

deroff - remove nroff, troff, tbl and eqn constructs

SYNOPSIS

deroff [-w] file ...

DESCRIPTION

Deroff reads each file in sequence and removes all *nroff* and *troff* command lines, backslash constructions, macro definitions, *eqn* constructs (between '.EQ' and '.EN' lines or between delimiters), and table descriptions and writes the remainder on the standard output. *Deroff* follows chains of included files ('.so' and '.nx' commands); if a file has already been included, a '.so' is ignored and a '.nx' terminates execution. If no input file is given, *deroff* reads from the standard input file.

If the *-w* flag is given, the output is a word list, one 'word' (string of letters, digits, and apostrophes, beginning with a letter; apostrophes are removed) per line, and all other characters ignored. Otherwise, the output follows the original, with the deletions mentioned above.

SEE ALSO

troff(1), eqn(1), tbl(1)

BUGS

Deroff is not a complete *troff* interpreter, so it can be confused by subtle constructs. Most errors result in too much rather than too little output.

NAME

df - disk free

SYNOPSIS

df [-i] [filesystem ...] [file ...]

DESCRIPTION

Df prints out the amount of free disk space available on the specified *filesystem*, e.g. “/dev/sc0a”, or on the filesystem in which the specified *file*, e.g. “\$HOME”, is contained. If no file system is specified, the free space on all of the normally mounted file systems is printed. The reported numbers are in kilobytes.

Other options are:

-i Report also the number of inodes which are used.

FILES

/etc/fstab list of normally mounted filesystems

SEE ALSO

fstab(5), quot(8)

NAME

diction,explain - print wordy sentences; thesaurus for diction

SYNOPSIS

diction [**-ml**] [**-mm**] [**-n**] [**-f pfile**] file ...
explain

DESCRIPTION

Diction finds all sentences in a document that contain phrases from a data base of bad or wordy diction. Each phrase is bracketed with []. Because *diction* runs *deroff* before looking at the text, formatting header files should be included as part of the input. The default macro package **-ms** may be overridden with the flag **-mm**. The flag **-ml** which causes **deroff** to skip lists, should be used if the document contains many lists of non-sentences. The user may supply her/his own pattern file to be used in addition to the default file with **-f pfile**. If the flag **-n** is also supplied the default file will be suppressed.

Explain is an interactive thesaurus for the phrases found by *diction*.

SEE ALSO

deroff(1)

BUGS

Use of non-standard formatting macros may cause incorrect sentence breaks. In particular, *diction* doesn't grok **-me**.

NAME

`diff` - differential file and directory comparator

SYNOPSIS

```
diff [-l] [-r] [-s] [-cefh] [-b] dir1 dir2
diff [-cefh] [-b] file1 file2
diff [-Dstring] [-b] file1 file2
```

DESCRIPTION

If both arguments are directories, `diff` sorts the contents of the directories by name, and then runs the regular file `diff` algorithm (described below) on text files which are different. Binary files which differ, common subdirectories, and files which appear in only one directory are listed. Options when comparing directories are:

- `-l` long output format; each text file `diff` is piped through `pr(1)` to paginate it, other differences are remembered and summarized after all text file differences are reported.
- `-r` causes application of `diff` recursively to common subdirectories encountered.
- `-s` causes `diff` to report files which are the same, which are otherwise not mentioned.

-Sname

starts a directory `diff` in the middle beginning with file `name`.

When run on regular files, and when comparing text files which differ during directory comparison, `diff` tells what lines must be changed in the files to bring them into agreement. Except in rare circumstances, `diff` finds a smallest sufficient set of file differences. If neither `file1` nor `file2` is a directory, then either may be given as '-', in which case the standard input is used. If `file1` is a directory, then a file in that directory whose file-name is the same as the file-name of `file2` is used (and vice versa).

There are several options for output format; the default output format contains lines of these forms:

```
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 backward 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`, which may be given with any of the others, the following options are mutually exclusive:

- `-e` producing a script of `a`, `c` and `d` commands for the editor `ed`, which will recreate `file2` from `file1`. In connection with `-e`, the following shell program may help maintain multiple versions of a file. Only an ancestral file (`$1`) and a chain of version-to-version `ed` scripts (`$2, $3, ...`) made by `diff` need be on hand. A 'latest version' appears on the standard output.

```
(shift; cat $*; echo `1,$p` ) | ed - $1
```

Extra commands are added to the output when comparing directories with `-e`, so that the result is a `sh(1)` script for converting text files which are common to the two directories from their state in `dir1` to their state in `dir2`.

- f** produces a script similar to that of **-e**, not useful with *ed*, and in the opposite order.
- c** produces a diff with lines of context. The default is to present 3 lines of context and may be changed, e.g to 10, by **-c10**. With **-c** the output format is modified slightly: the output beginning with identification of the files involved and their creation dates and then each change is separated by a line with a dozen *'s. The lines removed from *file1* are marked with '-'; those added to *file2* are marked '+'. Lines which are changed from one file to the other are marked in both files with '!'.
-Dstring causes *diff* to 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 defining *string* will yield *file2*.
- b** causes trailing blanks (spaces and tabs) to be ignored, and other strings of blanks to compare equal.

FILES

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

SEE ALSO

cmp(1), cc(1), comm(1), ed(1), diff3(1)

DIAGNOSTICS

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

BUGS

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

When comparing directories with the **-b** option specified, *diff* first compares the files ala *cmp*, and then decides to run the *diff* algorithm if they are not equal. This may cause a small amount of spurious output if the files then turn out to be identical because the only differences are insignificant blank string differences.

NAME

diff3 - 3-way differential file comparison

SYNOPSIS

diff3 [-ex3] file1 file2 file3

DESCRIPTION

Diff3 compares three versions of a file, and publishes 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 suffered in converting a given range of a given file to some other is indicated in one of these ways:

```
f: n1 a      Text is to be appended after line number n1 in file f, where f = 1, 2, or 3.
f: n1 , n2 c Text is to be changed in the range line n1 to line n2. If n1 = n2, the range
              may 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.

Under the **-e** option, *diff3* publishes a script for the editor *ed* that will incorporate into *file1* all changes between *file2* and *file3*, *i.e.* the changes that normally would be flagged ===== and =====3. Option **-x** (**-3**) produces a script to incorporate only changes flagged ===== (=3). The following command will apply the resulting script to 'file1'.

```
(cat script; echo '1,$p') | ed - file1
```

FILES

```
/tmp/d3?????
/usr/lib/diff3
```

SEE ALSO

diff(1)

BUGS

Text lines that consist of a single '.' will defeat **-e**.

NAME

dis - an mc68020 disassembler

SYNOPSIS

dis file ...

DESCRIPTION

The *dis* command produces an assembly language listing of each of its object *file* arguments. The listing includes assembly statements and the hexadecimal objects that produced those statements.

SEE ALSO

as(1), cc(1), ld(1).

NAME

`dosc` - connect to proc/286 system

SYNOPSIS

`dosc` [*partition*]

DESCRIPTION

Dosc is used to connect to one of the Multi-Link partitions. The optional *partition* parameter may be specified to access a specific partition. If no partition is specified, the command will attempt to find an available partition and connect to it. If the terminal type is *dt1200* or *pcshad*, the terminal will be switched to make-break mode. Multi-Link must be configured properly to correspond to the terminal currently in use.

The total number of active partitions is set in the file `/etc/mttys`. If `/etc/mttys` does not exist *dosc* will attempt to access up to 8 partitions. If *dosc* is unable to access a partition, a message is printed and an exit status of 10 (decimal) is returned. A shell script could be implemented to wait for a partition to become available.

There are two commands that are responded to by the *dosc* program: `exit` and `suspend`. For non-PC-compatible terminals, `exit` is signaled with `^\`, and `suspend` with `^]`. For PC-compatible terminals, `exit` is `CONTROL-ALT-\`, and `suspend` is `CONTROL-ALT-]`. `Exit` disconnects and relinquishes the partition. Subsequent uses of the *dosc* command will re-use the partition and the Multi-Link session will be as it was when the `exit` was done. `Suspend`, on the other hand, disconnects but does not relinquish the partition. After other activities, the Multi-Link session can be resumed by using the `fg` command of *cs(1)*. No other *dosc* users can access the partition until an `exit` command is given.

Make sure when entering the multiple key sequences that the `CONTROL` and `ALT` keys are fully down before the `\` or `]` is pressed. Failing to do this can result in spurious characters being sent to the Multi-link session.

If the screen gets overwritten with system messages it can be re-painted by issuing `ALT-r` (from PC-compatible terminals) or `ESC-b` (from normal terminals).

FILES

`/usr/spool/uucp/LCK..mtty`, `/dev/mtty`, `/etc/mttys`

SEE ALSO

`cs(1)` (for `fg` command), `whodos(1)`

NOTE

If a *dosc* process is terminated other than by `exit`, it will leave a lock file in `/usr/spool/uucp`. This file must be deleted before the partition can be re-used.

NAME

`du` - summarize disk usage

SYNOPSIS

`du [-s] [-a] [name ...]`

DESCRIPTION

Du gives the number of kilobytes contained in all files and, recursively, directories within each specified directory or file *name*. If *name* is missing, '.' is used.

The argument `-s` causes only the grand total to be given. The argument `-a` causes an entry to be generated for each file. Absence of either causes an entry to be generated for each directory only.

A file which has two links to it is only counted once.

SEE ALSO

`df(1)`, `quot(8)`

BUGS

Non-directories given as arguments (not under `-a` option) are not listed.

If there are too many distinct linked files, *du* counts the excess files multiply.

NAME

echo - echo arguments

SYNOPSIS

echo [-n] [arg] ...

DESCRIPTION

Echo writes its arguments separated by blanks and terminated by a newline on the standard output. If the flag **-n** is used, no newline is added to the output.

Echo is useful for producing diagnostics in shell programs and for writing constant data on pipes. To send diagnostics to the standard error file, do 'echo ... 1>&2'.

NAME

ed - text editor

SYNOPSIS

ed [-] [-x] [name]

DESCRIPTION

Ed is the standard text editor.

If a *name* argument is given, *ed* simulates an *e* command (see below) on the named file; that is to say, the file is read into *ed*'s buffer so that it can be edited. If *-x* is present, an *x* command is simulated first to handle an encrypted file. The optional *-* suppresses the printing of explanatory output and should be used when the standard input is an editor script.

Ed operates on a copy of any file it is editing; changes made in the copy have no effect on the file until a *w* (write) command is given. The copy of the text being edited resides in a temporary file called the *buffer*.

Commands to *ed* have a simple and regular structure: zero or more *addresses* followed by a single character *command*, possibly followed by parameters to the command. These addresses specify one or more lines in the buffer. Missing addresses are supplied by default.

In general, only one command may appear on a line. Certain commands allow the addition of text to the buffer. While *ed* is accepting text, it is said to be in *input mode*. In this mode, no commands are recognized; all input is merely collected. Input mode is left by typing a period '.' alone at the beginning of a line.

Ed supports a limited form of *regular expression* notation. A regular expression specifies a set of strings of characters. A member of this set of strings is said to be *matched* by the regular expression. In the following specification for regular expressions the word 'character' means any character but newline.

1. Any character except a special character matches itself. Special characters are the regular expression delimiter plus \[. and sometimes ^*\$.
2. A . matches any character.
3. A \ followed by any character except a digit or () matches that character.
4. A nonempty string *s* bracketed [*s*] (or [^*s*]) matches any character in (or not in) *s*. In *s*, \ has no special meaning, and] may only appear as the first letter. A substring *a-b*, with *a* and *b* in ascending ASCII order, stands for the inclusive range of ASCII characters.
5. A regular expression of form 1-4 followed by * matches a sequence of 0 or more matches of the regular expression.
6. A regular expression, *x*, of form 1-8, bracketed \(*x*\) matches what *x* matches.
7. A \ followed by a digit *n* matches a copy of the string that the bracketed regular expression beginning with the *n*th \(\) matched.
8. A regular expression of form 1-8, *x*, followed by a regular expression of form 1-7, *y* matches a match for *x* followed by a match for *y*, with the *x* match being as long as possible while still permitting a *y* match.
9. A regular expression of form 1-8 preceded by ^ (or followed by \$), is constrained to matches that begin at the left (or end at the right) end of a line.
10. A regular expression of form 1-9 picks out the longest among the leftmost matches in a line.

11. An empty regular expression stands for a copy of the last regular expression encountered.

Regular expressions are used in addresses to specify lines and in one command (see *s* below) to specify a portion of a line which is to be replaced. If it is desired to use one of the regular expression metacharacters as an ordinary character, that character may be preceded by '\'. This also applies to the character bounding the regular expression (often '/') and to '\' itself.

To understand addressing in *ed* it is necessary to know that at any time there is a *current line*. Generally speaking, the current line is the last line affected by a command; however, the exact effect on the current line is discussed under the description of the command. Addresses are constructed as follows.

1. The character '.' addresses the current line.
2. The character '\$' addresses the last line of the buffer.
3. A decimal number *n* addresses the *n*-th line of the buffer.
4. '*x*' addresses the line marked with the name *x*, which must be a lower-case letter. Lines are marked with the *k* command described below.
5. A regular expression enclosed in slashes '/' addresses the line found by searching forward from the current line and stopping at the first line containing a string that matches the regular expression. If necessary the search wraps around to the beginning of the buffer.
6. A regular expression enclosed in queries '?' addresses the line found by searching backward from the current line and stopping at the first line containing a string that matches the regular expression. If necessary the search wraps around to the end of the buffer.
7. An address followed by a plus sign '+' or a minus sign '-' followed by a decimal number specifies that address plus (resp. minus) the indicated number of lines. The plus sign may be omitted.
8. If an address begins with '+' or '-' the addition or subtraction is taken with respect to the current line; e.g. '-5' is understood to mean '.-5'.
9. If an address ends with '+' or '-', then 1 is added (resp. subtracted). As a consequence of this rule and rule 8, the address '-' refers to the line before the current line. Moreover, trailing '+' and '-' characters have cumulative effect, so '--' refers to the current line less 2.
10. To maintain compatibility with earlier versions of the editor, the character '^' in addresses is equivalent to '-'.

Commands may require zero, one, or two addresses. Commands which require no addresses regard the presence of an address as an error. Commands which accept one or two addresses assume default addresses when insufficient are given. If more addresses are given than such a command requires, the last one or two (depending on what is accepted) are used.

Addresses are separated from each other typically by a comma ','. They may also be separated by a semicolon ';'. In this case the current line '.' is set to the previous address before the next address is interpreted. This feature can be used to determine the starting line for forward and backward searches ('/', '?'). The second address of any two-address sequence must correspond to a line following the line corresponding to the first address. The special form '%' is an abbreviation for the address pair '1,\$'.

In the following list of *ed* commands, the default addresses are shown in parentheses. The parentheses are not part of the address, but are used to show that the given addresses are the default.

As mentioned, it is generally illegal for more than one command to appear on a line. However, most commands may be suffixed by 'p' or by 'l', in which case the current line is either printed or listed respectively in the way discussed below. Commands may also be suffixed by 'n', meaning the output of the command is to be line numbered. These suffixes may be combined in any order.

(.)a
<text>

The append command reads the given text and appends it after the addressed line. '.' is left on the last line input, if there were any, otherwise at the addressed line. Address '0' is legal for this command; text is placed at the beginning of the buffer.

(.,.)c
<text>

The change command deletes the addressed lines, then accepts input text which replaces these lines. '.' is left at the last line input; if there were none, it is left at the line preceding the deleted lines.

(.,.)d

The delete command deletes the addressed lines from the buffer. The line originally after the last line deleted becomes the current line; if the lines deleted were originally at the end, the new last line becomes the current line.

e filename

The edit command causes the entire contents of the buffer to be deleted, and then the named file to be read in. '.' is set to the last line of the buffer. The number of characters read is typed. 'filename' is remembered for possible use as a default file name in a subsequent *r* or *w* command. If 'filename' is missing, the remembered name is used.

E filename

This command is the same as *e*, except that no diagnostic results when no *w* has been given since the last buffer alteration.

f filename

The filename command prints the currently remembered file name. If 'filename' is given, the currently remembered file name is changed to 'filename'.

(1,\$)g/regular expression/command list

In the global command, the first step is to mark every line which matches the given regular expression. Then for every such line, the given command list is executed with '.' initially set to that line. A single command or the first of multiple commands appears on the same line with the global command. All lines of a multi-line list except the last line must be ended with '\'. *A*, *i*, and *c* commands and associated input are permitted; the '.' terminating input mode may be omitted if it would be on the last line of the command list. The commands *g* and *v* are not permitted in the command list.

(.)i
<text>

This command inserts the given text before the addressed line. '.' is left at the last line input, or, if there were none, at the line before the addressed line. This command differs from the *a* command only in the placement of the text.

(.,.+1)j

This command joins the addressed lines into a single line; intermediate newlines simply disappear. '.' is left at the resulting line.

- (.)k*x*
The mark command marks the addressed line with name *x*, which must be a lower-case letter. The address form "*x*" then addresses this line.
- (.,.)l
The list command prints the addressed lines in an unambiguous way: non-graphic characters are printed in two-digit octal, and long lines are folded. The *l* command may be placed on the same line after any non-i/o command.
- (.,.)m*a*
The move command repositions the addressed lines after the line addressed by *a*. The last of the moved lines becomes the current line.
- (.,.)n
The number command prints the addressed lines with line numbers and a tab at the left.
- (.,.)p
The print command prints the addressed lines. '.' is left at the last line printed. The *p* command may be placed on the same line after any non-i/o command.
- (.,.)P
This command is a synonym for *p*.
- q The quit command causes *ed* to exit. No automatic write of a file is done.
- Q This command is the same as *q*, except that no diagnostic results when no *w* has been given since the last buffer alteration.
- (\$)*r* filename
The read command reads in the given file after the addressed line. If no file name is given, the remembered file name, if any, is used (see *e* and *f* commands). The file name is remembered if there was no remembered file name already. Address '0' is legal for *r* and causes the file to be read at the beginning of the buffer. If the read is successful, the number of characters read is typed. '.' is left at the last line read in from the file.
- (.,.)s/regular expression/replacement/ or,
(.,.)s/regular expression/replacement/g
The substitute command searches each addressed line for an occurrence of the specified regular expression. On each line in which a match is found, all matched strings are replaced by the replacement specified, 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. It is an error for the substitution to fail on all addressed lines. Any punctuation character may be used instead of '/' to delimit the regular expression and the replacement. '.' is left at the last line substituted.
- An ampersand '&' appearing in the replacement is replaced by the string matching the regular expression. The special meaning of '&' in this context may be suppressed by preceding it by '\'. The characters '\n' where *n* is a digit, are replaced by the text matched by the *n*-th regular subexpression enclosed between '\(' and '\)'. When nested, parenthesized subexpressions are present, *n* is determined by counting occurrences of '\(' starting from the left.
- Lines may be split by substituting new-line characters into them. The new-line in the replacement string must be escaped by preceding it by '\n'.
- One or two trailing delimiters may be omitted, implying the 'p' suffix. The special form 's' followed by *no* delimiters repeats the most recent substitute command on the addressed lines. The 's' may be followed by the letters *r* (use the most recent regular expression for the left hand side, instead of the most recent left hand side of a substitute

command), **p** (complement the setting of the *p* suffix from the previous substitution), or **g** (complement the setting of the *g* suffix). These letters may be combined in any order.

(.,.)**t** *a*

This command acts just like the *m* command, except that a copy of the addressed lines is placed after address *a* (which may be 0). '.' is left on the last line of the copy.

(.,.)**u**

The undo command restores the buffer to its state before the most recent buffer modifying command. The current line is also restored. Buffer modifying commands are *a*, *c*, *d*, *g*, *i*, *k*, and *v*. For purposes of undo, *g* and *v* are considered to be a single buffer modifying command. Undo is its own inverse.

When *ed* runs out of memory (at about 8000 lines on any 16 bit mini-computer such as the PDP-11) This full undo is not possible, and *u* can only undo the effect of the most recent substitute on the current line. This restricted undo also applies to editor scripts when *ed* is invoked with the - option.

(1, \$)**v**/regular expression/command list

This command is the same as the global command *g* except that the command list is executed *g* with '.' initially set to every line *except* those matching the regular expression.

(1, \$)**w** filename

The write command writes the addressed lines onto the given file. If the file does not exist, it is created. The file name is remembered if there was no remembered file name already. If no file name is given, the remembered file name, if any, is used (see *e* and *f* commands). '.' is unchanged. If the command is successful, the number of characters written is printed.

(1, \$)**W** filename

This command is the same as *w*, except that the addressed lines are appended to the file.

(1, \$)**wq** filename

This command is the same as *w* except that afterwards a *q* command is done, exiting the editor after the file is written.

x A key string is demanded from the standard input. Later *r*, *e* and *w* commands will encrypt and decrypt the text with this key by the algorithm of *crypt*(1). An explicitly empty key turns off encryption. (.,+1)**zor**,

(.,+1)**zn**

This command scrolls through the buffer starting at the addressed line. 22 (or *n*, if given) lines are printed. The last line printed becomes the current line. The value *n* is sticky, in that it becomes the default for future *z* commands.

(\$)=

The line number of the addressed line is typed. '.' is unchanged by this command.

!<shell command>

The remainder of the line after the '!' is sent to *sh*(1) to be interpreted as a command. '.' is unchanged.

(.,+1,.,+1) <newline>

An address alone on a line causes the addressed line to be printed. A blank line alone is equivalent to '.+1p'; it is useful for stepping through text. If two addresses are present with no intervening semicolon, *ed* prints the range of lines. If they are separated by a semicolon, the second line is printed.

If an interrupt signal (ASCII DEL) is sent, *ed* prints '?interrupted' and returns to its command level.

Some size limitations: 512 characters per line, 256 characters per global command list, 64 characters per file name, and, on mini computers, 128K characters in the temporary file. The limit on the number of lines depends on the amount of core: each line takes 2 words.

When reading a file, *ed* discards ASCII NUL characters and all characters after the last newline. It refuses to read files containing non-ASCII characters.

FILES

/tmp/e*

edhup: work is saved here if terminal hangs up

SEE ALSO

B. W. Kernighan, *A Tutorial Introduction to the ED Text Editor*

B. W. Kernighan, *Advanced editing on UNIX*

ex(1), sed(1), crypt(1)

DIAGNOSTICS

'?name' for inaccessible file; '?self-explanatory message' for other errors.

To protect against throwing away valuable work, a *q* or *e* command is considered to be in error, unless a *w* has occurred since the last buffer change. A second *q* or *e* will be obeyed regardless.

BUGS

The *l* command mishandles DEL.

The *undo* command causes marks to be lost on affected lines.

The *x* command, *-x* option, and special treatment of hangups only work on UNIX.

NAME

efl - Extended Fortran Language

SYNOPSIS

efl [option ...] [filename ...]

DESCRIPTION

Efl compiles a program written in the EFL language into clean Fortran. *Efl* provides the same control flow constructs as does *ratfor(1)*, which are essentially identical to those in C:

statement grouping with braces;

decision-making with if, if-else, and switch-case; while, for, Fortran do, repeat, and repeat...until loops; multi-level break and next. In addition, EFL has C-like data structures, and more uniform and convenient input/output syntax, generic functions. EFL also provides some syntactic sugar to make programs easier to read and write:

free form input:

multiple statements/line; automatic continuation statement label names (not just numbers),

comments:

this is a comment

translation of relationals:

>, >=, etc., become .GT., .GE., etc.

return (expression)

returns expression to caller from function

define: define name replacement

include:

include filename

The *Efl* command option **-w** suppresses warning messages. The option **-C** causes comments to be copied through to the Fortran output (default); **-#** prevents comments from being copied through. If a command argument contains an embedded equal sign, that argument is treated as if it had appeared in an **option** statement at the beginning of the program. *Efl* is best used with *f77(1)*.

SEE ALSO

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

S. I. Feldman, *The Programming Language EFL*, Bell Labs Computing Science Technical Report #78.

NAME

eqn, *neqn*, *checkeq* – typeset mathematics

SYNOPSIS

eqn [*-dxy*] [*-pn*] [*-sn*] [*-fn*] [*file*] ...
checkeq [*file*] ...

DESCRIPTION

Eqn is a troff(1) preprocessor for typesetting mathematics on a Graphic Systems phototypesetter, *neqn* on terminals. Usage is almost always

```
eqn file ... | troff
neqn file ... | nroff
```

If no files are specified, these programs reads from the standard input. A line beginning with '.EQ' marks the start of an equation; the end of an equation is marked by a line beginning with '.EN'. Neither of these lines is altered, so they may be defined in macro packages to get centering, numbering, etc. It is also possible to set two characters as 'delimiters'; subsequent text between delimiters is also treated as *eqn* input. Delimiters may be set to characters *x* and *y* with the command-line argument *-dxy* or (more commonly) with 'delim *xy*' between .EQ and .EN. The left and right delimiters may be identical. Delimiters are turned off by 'delim off'. All text that is neither between delimiters nor between .EQ and .EN is passed through untouched.

The program *checkeq* reports missing or unbalanced delimiters and .EQ/.EN pairs.

Tokens within *eqn* are separated by spaces, tabs, newlines, braces, double quotes, tildes or circumflexes. Braces {} are used for grouping; generally speaking, anywhere a single character like *x* could appear, a complicated construction enclosed in braces may be used instead. Tilde ~ represents a full space in the output, circumflex ^ half as much.

Subscripts and superscripts are produced with the keywords **sub** and **sup**. Thus *x sub i* makes x_i , *a sub i sup 2* produces a_i^2 , and *e sup {x sup 2 + y sup 2}* gives $e^{x^2+y^2}$.

Fractions are made with **over**: *a over b* yields $\frac{a}{b}$.

sqrt makes square roots: *1 over sqrt {ax sup 2 +bx+c}* results in $\frac{1}{\sqrt{ax^2+bx+c}}$.

The keywords **from** and **to** introduce lower and upper limits on arbitrary things: $\lim_{n \rightarrow \infty} \sum_0^n x_i$ is made with *lim from {n-> inf} sum from 0 to n x sub i*.

Left and right brackets, braces, etc., of the right height are made with **left** and **right**: *left / x sup 2 + y sup 2 over alpha right / ~ = ~ 1* produces $\left[x^2 + \frac{y^2}{\alpha} \right] = 1$. The **right** clause is optional.

Legal characters after **left** and **right** are braces, brackets, bars, **c** and **f** for ceiling and floor, and "" for nothing at all (useful for a right-side-only bracket).

Vertical piles of things are made with **pile**, **lpile**, **cpile**, and **rpile**: *pile {a above b above c}* produces $\begin{matrix} a \\ b \\ c \end{matrix}$. There can be an arbitrary number of elements in a pile. **lpile** left-justifies, **pile** and **cpile** center, with different vertical spacing, and **rpile** right justifies.

Matrices are made with **matrix**: $matrix \{ lcol \{ x \text{ sub } i \text{ above } y \text{ sub } 2 \} ccol \{ 1 \text{ above } 2 \} \}$ produces $\begin{matrix} x_i & 1 \\ y_2 & 2 \end{matrix}$. In addition, there is **rcol** for a right-justified column.

Diacritical marks are made with **dot**, **dotdot**, **hat**, **tilde**, **bar**, **vec**, **dyad**, and **under**: $x \text{ dot} = \dot{x}$, $f(t) \text{ bar}$ is $\overline{f(t)}$, $y \text{ dotdot}$ is \ddot{y} , $n \text{ under}$ is \underline{n} , and $x \text{ vec}$ is \vec{x} , $y \text{ dyad}$ is \overleftrightarrow{y} .

Sizes and font can be changed with **size** n or **size** $\pm n$, **roman**, **italic**, **bold**, and **font** n . Size and fonts can be changed globally in a document by **gsize** n and **gfont** n , or by the command-line arguments **-sn** and **-fn**.

Normally subscripts and superscripts are reduced by 3 point sizes from the previous size; this may be changed by the command-line argument **-pn**.

Successive display arguments can be lined up. Place **mark** before the desired lineup point in the first equation; place **lineup** at the place that is to line up vertically in subsequent equations.

Shorthands may be defined or existing keywords redefined with **define**: $define \text{ thing } \% \text{ replacement}$ defines a new token called *thing* which will be replaced by *replacement* whenever it appears thereafter. The $\%$ may be any character that does not occur in *replacement*.

Keywords like *sum* (Σ) *int* (\int) *inf* (∞) and shorthands like \geq (\geq) \rightarrow (\rightarrow), and \neq (\neq) are recognized. Greek letters are spelled out in the desired case, as in *alpha* or *GAMMA*. Mathematical words like *sin*, *cos*, *log* are made Roman automatically. *Troff*(1) four-character escapes like $\backslash(bs$ (\emptyset) can be used anywhere. Strings enclosed in double quotes "..." are passed through untouched; this permits keywords to be entered as text, and can be used to communicate with *troff* when all else fails.

SEE ALSO

troff(1), *tbl*(1), *ms*(7), *eqnchar*(7)
 B. W. Kernighan and L. L. Cherry, *Typesetting Mathematics—User's Guide*
 J. F. Ossanna, *NROFF/TROFF User's Manual*

BUGS

To embolden digits, parens, etc., it is necessary to quote them, as in 'bold "12.3"'.

NAME

error - analyze and disperse compiler error messages

SYNOPSIS

error [**-n**] [**-s**] [**-q**] [**-v**] [**-t** suffixlist] [**-I** ignorefile] [name]

DESCRIPTION

Error analyzes and optionally disperses the diagnostic error messages produced by a number of compilers and language processors to the source file and line where the errors occurred. It can replace the painful, traditional methods of scribbling abbreviations of errors on paper, and permits error messages and source code to be viewed simultaneously without machinations of multiple windows in a screen editor.

Error looks at the error messages, either from the specified file *name* or from the standard input, and attempts to determine which language processor produced each error message, determines the source file and line number to which the error message refers, determines if the error message is to be ignored or not, and inserts the (possibly slightly modified) error message into the source file as a comment on the line preceding to which the line the error message refers. Error messages which can't be categorized by language processor or content are not inserted into any file, but are sent to the standard output. *Error* touches source files only after all input has been read. By specifying the **-q** query option, the user is asked to confirm any potentially dangerous (such as touching a file) or verbose action. Otherwise *error* proceeds on its merry business. If the **-t** touch option and associated suffix list is given, *error* will restrict itself to touch only those files with suffices in the suffix list. Error also can be asked (by specifying **-v**) to invoke *vi*(1) on the files in which error messages were inserted; this obviates the need to remember the names of the files with errors.

Error is intended to be run with its standard input connected via a pipe to the error message source. Some language processors put error messages on their standard error file; others put their messages on the standard output. Hence, both error sources should be piped together into *error*. For example, when using the *cs*h syntax,

```
make -s lint |& error -q -v
```

will analyze all the error messages produced by whatever programs *make* runs when making *lint*.

Error knows about the error messages produced by: *make*, *cc*, *cpp*, *ccom*, *as*, *ld*, *lint*, *pi*, *pc* and *f77*. *Error* knows a standard format for error messages produced by the language processors, so is sensitive to changes in these formats. For all languages except *Pascal*, error messages are restricted to be on one line. Some error messages refer to more than one line in more than one files; *error* will duplicate the error message and insert it at all of the places referenced.

Error will do one of six things with error messages.

synchronize

Some language processors produce short errors describing which file it is processing. *Error* uses these to determine the file name for languages that don't include the file name in each error message. These synchronization messages are consumed entirely by *error*.

discard

Error messages from *lint* that refer to one of the two *lint* libraries, */usr/lib/llib-lc* and */usr/lib/llib-port* are discarded, to prevent accidentally touching these libraries. Again, these error messages are consumed entirely by *error*.

nullify Error messages from *lint* can be nullified if they refer to a specific function, which is known to generate diagnostics which are not interesting. Nullified error messages are not inserted into the source file, but are written to the standard output. The names of functions to ignore are taken from either the file named *.errorrc* in the users's home directory, or from the file named by the *-I* option. If the file does not exist, no error messages are nullified. If the file does exist, there must be one function name per line.

not file specific

Error messages that can't be intuited are grouped together, and written to the standard output before any files are touched. They will not be inserted into any source file.

file specific

Error message that refer to a specific file, but to no specific line, are written to the standard output when that file is touched.

true errors Error messages that can be intuited are candidates for insertion into the file to which they refer.

Only true error messages are candidates for inserting into the file they refer to. Other error messages are consumed entirely by *error* or are written to the standard output. *Error* inserts the error messages into the source file on the line preceding the line the language processor found in error. Each error message is turned into a one line comment for the language, and is internally flagged with the string "###" at the beginning of the error, and "%%%" at the end of the error. This makes pattern searching for errors easier with an editor, and allows the messages to be easily removed. In addition, each error message contains the source line number for the line the message refers to. A reasonably formatted source program can be recompiled with the error messages still in it, without having the error messages themselves cause future errors. For poorly formatted source programs in free format languages, such as C or Pascal, it is possible to insert a comment into another comment, which can wreak havoc with a future compilation. To avoid this, programs with comments and source on the same line should be formatted so that language statements appear before comments.

Options available with *error* are:

- n Do not touch any files; all error messages are sent to the standard output.
- q The user is *queried* whether s/he wants to touch the file. A "y" or "n" to the question is necessary to continue. Absence of the *-q* option implies that all referenced files (except those referring to discarded error messages) are to be touched.
- v After all files have been touched, overlay the visual editor *vi* with it set up to edit all files touched, and positioned in the first touched file at the first error. If *vi* can't be found, try *ex* or *ed* from standard places.
- t Take the following argument as a suffix list. Files whose suffixes do not appear in the suffix list are not touched. The suffix list is dot separated, and "*" wildcards work. Thus the suffix list:
 ".c.y.foo*.h"
 allows *error* to touch files ending with ".c", ".y", ".foo*" and ".y".
- s Print out *statistics* regarding the error categorization. Not too useful.

Error catches interrupt and terminate signals, and if in the insertion phase, will orderly terminate what it is doing.

AUTHOR

Robert Henry

FILES~/errorrc
/dev/ttyfunction names to ignore for *lint* error messages
user's teletype**BUGS**

Opens the teletype directly to do user querying.

Source files with links make a new copy of the file with only one link to it.

Changing a language processor's format of error messages may cause *error* to not understand the error message.

Error, since it is purely mechanical, will not filter out subsequent errors caused by 'floodgating' initiated by one syntactically trivial error. Humans are still much better at discarding these related errors.

Pascal error messages belong after the lines affected (*error* puts them before). The alignment of the '|' marking the point of error is also disturbed by *error*.

Error was designed for work on CRT's at reasonably high speed. It is less pleasant on slow speed terminals, and has never been used on hardcopy terminals.

NAME

ex, edit - text editor

SYNOPSIS

```
ex [-] [-v] [-t tag] [-r] [+command] [-l] name ...
edit [ex options]
```

DESCRIPTION

Ex is the root of a family of editors: *edit*, *ex* and *vi*. *Ex* is a superset of *ed*, with the most notable extension being a display editing facility. Display based editing is the focus of *vi*.

If you have not used *ed*, or are a casual user, you will find that the editor *edit* is convenient for you. It avoids some of the complexities of *ex* used mostly by systems programmers and persons very familiar with *ed*.

If you have a CRT terminal, you may wish to use a display based editor; in this case see *vi*(1), which is a command which focuses on the display editing portion of *ex*.

DOCUMENTATION

The document *Edit: A tutorial* provides a comprehensive introduction to *edit* assuming no previous knowledge of computers or the UNIX system.

The *Ex Reference Manual - Version 3.5* is a comprehensive and complete manual for the command mode features of *ex*, but you cannot learn to use the editor by reading it. For an introduction to more advanced forms of editing using the command mode of *ex* see the editing documents written by Brian Kernighan for the editor *ed*; the material in the introductory and advanced documents works also with *ex*.

An Introduction to Display Editing with Vi introduces the display editor *vi* and provides reference material on *vi*. All of these documents can be found in volume 2c of the Programmer's Manual. In addition, the *Vi Quick Reference* card summarizes the commands of *vi* in a useful, functional way, and is useful with the *Introduction*.

FILES

/usr/lib/ex?.?strings	error messages
/usr/lib/ex?.?recover	recover command
/usr/lib/ex?.?preserve	preserve command
/etc/termcap	describes capabilities of terminals
~/.exrc	editor startup file
/tmp/Exnnnnn	editor temporary
/tmp/Rxnnnnn	named buffer temporary
/usr/preserve	preservation directory

SEE ALSO

awk(1), ed(1), grep(1), sed(1), grep(1), vi(1), termcap(5), environ(7)

AUTHOR

Originally written by William Joy

Mark Horton has maintained the editor since version 2.7, adding macros, support for many unusual terminals, and other features such as word abbreviation mode.

BUGS

The *undo* command causes all marks to be lost on lines changed and then restored if the marked lines were changed.

Undo never clears the buffer modified condition.

The *z* command prints a number of logical rather than physical lines. More than a screen full of output may result if long lines are present.

File input/output errors don't print a name if the command line '-' option is used.

There is no easy way to do a single scan ignoring case.

The editor does not warn if text is placed in named buffers and not used before exiting the editor.

Null characters are discarded in input files, and cannot appear in resultant files.

NAME

`expand`, `unexpand` - expand tabs to spaces, and vice versa

SYNOPSIS

```
expand [ -tabstop ] [ -tab1,tab2,...,tabn ] [ file ... ]  
unexpand [ -a ] [ file ... ]
```

DESCRIPTION

Expand processes the named files or the standard input writing the standard output with tabs changed into blanks. Backspace characters are preserved into the output and decrement the column count for tab calculations. *Expand* is useful for pre-processing character files (before sorting, looking at specific columns, etc.) that contain tabs.

If a single *tabstop* argument is given then tabs are set *tabstop* spaces apart instead of the default 8. If multiple *tabstops* are given then the tabs are set at those specific columns.

Unexpand puts tabs back into the data from the standard input or the named files and writes the result on the standard output. By default only leading blanks and tabs are reconverted to maximal strings of tabs. If the `-a` option is given, then tabs are inserted whenever they would compress the resultant file by replacing two or more characters.

NAME

explain, diction- print wordy sentences; thesaurus for diction

SYNOPSIS

diction [**-ml**] [**-mm**] [**-n**] [**-f pfile**] file ...

explain

DESCRIPTION

Diction finds all sentences in a document that contain phrases from a data base of bad or wordy diction. Each phrase is bracketed with []. Because *diction* runs *deroff* before looking at the text, formatting header files should be included as part of the input. The default macro package **-ms** may be overridden with the flag **-mm**. The flag **-ml** which causes *deroff* to skip lists, should be used if the document contains many lists of non-sentences. The user may supply her/his own pattern file to be used in addition to the default file with **-f pfile**. If the flag **-n** is also supplied the default file will be suppressed.

Explain is an interactive thesaurus for the phrases found by *diction*.

SEE ALSO

deroff(1)

BUGS

Use of non-standard formatting macros may cause incorrect sentence breaks. In particular, *diction* doesn't grok **-me**.

NAME

expr - evaluate arguments as an expression

SYNOPSIS

expr arg ...

DESCRIPTION

The arguments are taken as an expression. After evaluation, the result is written on the standard output. Each token of the expression is a separate argument.

The operators and keywords are listed below. The list is in order of increasing precedence, with equal precedence operators grouped.

expr | *expr*

yields the first *expr* if it is neither null nor '0', otherwise yields the second *expr*.

expr & *expr*

yields the first *expr* if neither *expr* is null or '0', otherwise yields '0'.

expr *relop* *expr*

where *relop* is one of < <= = != >= >, yields '1' if the indicated comparison is true, '0' if false. The comparison is numeric if both *expr* are integers, otherwise lexicographic.

expr + *expr*

expr - *expr*

addition or subtraction of the arguments.

expr * *expr*

expr / *expr*

expr % *expr*

multiplication, division, or remainder of the arguments.

expr : *expr*

The matching operator compares the string first argument with the regular expression second argument; regular expression syntax is the same as that of *ed*(1). The `\(...\)` pattern symbols can be used to select a portion of the first argument. Otherwise, the matching operator yields the number of characters matched ('0' on failure).

(*expr*)

parentheses for grouping.

Examples:

To add 1 to the Shell variable *a*:

```
a=`expr $a + 1`
```

To find the filename part (least significant part) of the pathname stored in variable *a*, which may or may not contain '/':

```
expr $a : '.*\/(.*)' '|' $a
```

Note the quoted Shell metacharacters.

SEE ALSO

sh(1), *test*(1)

DIAGNOSTICS

Expr returns the following exit codes:

- 0 if the expression is neither null nor '0',
- 1 if the expression is null or '0',
- 2 for invalid expressions.

NAME

eyacc - modified *yacc* allowing much improved error recovery

SYNOPSIS

eyacc [-v] [grammar]

DESCRIPTION

Eyacc is an old version of *yacc*(1), which produces tables used by the Pascal system and its error recovery routines. *Eyacc* fully enumerates test actions in its parser when an error token is in the look-ahead set. This prevents the parser from making undesirable reductions when an error occurs before the error is detected. The table format is different in *eyacc* than it was in the old *yacc*, as minor changes had been made for efficiency reasons.

SEE ALSO

yacc(1)

"Practical LR Error Recovery" by Susan L. Graham, Charles B. Haley and W. N. Joy; SIG-PLAN Conference on Compiler Construction, August 1979.

AUTHOR

S. C. Johnson

Eyacc modifications by Charles Haley and William Joy.

BUGS

Pc and its error recovery routines should be made into a library of routines for the new *yacc*.

NAME

f77 – Fortran 77 compiler

SYNOPSIS

f77 [option] ... file ...

DESCRIPTION

F77 is the UNIX Fortran 77 compiler. It accepts several types of arguments:

Arguments whose names end with ‘.f’ are taken to be Fortran 77 source programs; they are compiled, and each object program is left on the file in the current directory whose name is that of the source with ‘.o’ substituted for ‘.f’.

Arguments whose names end with ‘.F’ are also taken to be Fortran 77 source programs; these are first processed by the C preprocessor before being compiled by *f77*.

Arguments whose names end with ‘.r’ or ‘.e’ are taken to be Ratfor or EFL source programs respectively; these are first transformed by the appropriate preprocessor, then compiled by *f77*.

Arguments whose names end with ‘.c’ or ‘.s’ are taken to be C or assembly source programs and are compiled or assembled, producing a ‘.o’ file.

The following options have the same meaning as in *cc*(1). See *ld*(1) for load-time options.

- c Suppress loading and produce ‘.o’ files for each source file.
- g Have the compiler produce additional symbol table information for *dbx*(1). Also pass the –lg flag to *ld*(1).
- o output
Name the final output file *output* instead of ‘a.out’.
- p Prepare object files for profiling, see *prof*(1).
- pg Causes the compiler to produce counting code in the manner of –p, but invokes a run-time recording mechanism that keeps more extensive statistics and produces a *gmon.out* file at normal termination. An execution profile can then be generated by use of *gprof*(1).
- f Compiles floating point operations to use the MC68881 floating point coprocessor. Also switches to versions of *libc.a*, *libm.a*, *libF77.a*, and *libI77.a* that use the floating point chip. Setting the environment variable **FP** to **m68881** has the same effect as specifying this flag. Code generated with this option will cause an “Illegal instruction” trap on machines without the floating point coprocessor.
- w Suppress all warning messages. If the option is ‘–w66’, only Fortran 66 compatibility warnings are suppressed.
- Dname=def
- Dname
Define the *name* to the C preprocessor, as if by ‘#define’. If no definition is given, the name is defined as “1”. (‘.F’ suffix files only).
- Idir ‘#include’ files whose names do not begin with ‘/’ are always sought first in the directory of the *file* argument, then in directories named in –I options, then in directories on a standard list. (‘.F’ suffix files only).
- O Invoke an object-code optimizer.
- S Compile the named programs, and leave the assembler-language output on corresponding files suffixed ‘.s’. (No ‘.o’ is created.)

The following options are peculiar to *f77*.

- I2 On machines which support short integers, make the default integer constants and variables short. (-I4 is the standard value of this option). All logical quantities will be short.
- m Apply the M4 preprocessor to each 'r' file before transforming it with the Ratfor or EFL preprocessor.
- onetrip
Compile DO loops that are performed at least once if reached. (Fortran 77 DO loops are not performed at all if the upper limit is smaller than the lower limit.)
- u Make the default type of a variable 'undefined' rather than using the default Fortran rules.
- v Print the version number of the compiler, and the name of each pass as it executes.
- C Compile code to check that subscripts are within declared array bounds.
- F Apply the C, EFL, or Ratfor preprocessors to relevant files, put the result in the file with the suffix changed to 'f', but do not compile.
- Ex Use the string *x* as an EFL option in processing 'e' files.
- Rx Use the string *x* as a Ratfor option in processing 'r' files.
- N[qxscn]nnn
Make static tables in the compiler bigger. The compiler will complain if it overflows its tables and suggest you apply one or more of these flags. These flags have the following meanings:
 - q Maximum number of equivalenced variables. Default is 150.
 - x Maximum number of external names (common block names, subroutine and function names). Default is 200.
 - s Maximum number of statement numbers. Default is 401.
 - c Maximum depth of nesting for control statements (e.g. DO loops). Default is 20.
 - n Maximum number of identifiers. Default is 1009.
- U Do not convert upper case letters to lower case. The default is to convert Fortran programs to lower case except within character string constants.

Other arguments are taken to be either loader option arguments, or F77-compatible object programs, typically produced by an earlier run, or perhaps libraries of F77-compatible routines. These programs, together with the results of any compilations specified, are loaded (in the order given) to produce an executable program with name 'a.out'.

FILES

file.[fF]resc]	input file
file.o	object file
a.out	loaded output
/usr/lib/f77pass1	compiler
/lib/f1	pass 2
/lib/c2	optional optimizer
/lib/cpp	C preprocessor
/usr/lib/libF77.a	intrinsic function library
/usr/lib/libI77.a	Fortran I/O library
/usr/lib/libU77.a	UNIX interface library

/usr/lib/libF77_p.a profiling intrinsic function library
/usr/lib/libI77_p.a profiling Fortran I/O library
/usr/lib/libU77_p.a profiling UNIX interface library
/lib/libc.a C library, see section 3
mon.out file produced for analysis by prof(1).
gmon.out file produced for analysis by gprof(1).

SEE ALSO

S. I. Feldman, P. J. Weinberger, *A Portable Fortran 77 Compiler*
D. L. Wasley, *Introduction to the f77 I/O Library*
prof(1), gprof(1), cc(1), ld(1), eff(1), ratfor(1)

DIAGNOSTICS

The diagnostics produced by *f77* itself are intended to be self-explanatory. Occasional messages may be produced by the loader.

BUGS

This compiler is still somewhat experimental. The optimizer occasionally makes mistakes: it should be avoided when debugging if apparently incorrect results are obtained. Because of an assembler error, complaints about long branches may occur with very large source files; such errors can be avoided by splitting the sources into smaller sections. If necessary, the old version of *f77* can be resurrected from /usr/src/old.

NAME

false, true – provide truth values

SYNOPSIS

true

false

DESCRIPTION

True and *false* are usually used in a Bourne shell script. They test for the appropriate status "true" or "false" before running (or failing to run) a list of commands.

EXAMPLE

```
while false
do
    command list
done
```

SEE ALSO

csh(1), sh(1), true(1)

DIAGNOSTICS

False has exit status nonzero.

NAME

file - determine file type

SYNOPSIS

file file ...

DESCRIPTION

File performs a series of tests on each argument in an attempt to classify it. If an argument appears to be ascii, *file* examines the first 512 bytes and tries to guess its language.

BUGS

It often makes mistakes. In particular it often suggests that command files are C programs. Does not recognize Pascal or LISP.

NAME

find - find files

SYNOPSIS

find pathname-list expression

DESCRIPTION

Find recursively descends the directory hierarchy for each pathname in the *pathname-list* (i.e., one or more pathnames) seeking files that match a boolean *expression* written in the primaries given below. In the descriptions, the argument *n* is used as a decimal integer where *+n* means more than *n*, *-n* means less than *n* and *n* means exactly *n*.

-name filename

True if the *filename* argument matches the current file name. Normal Shell argument syntax may be used if escaped (watch out for '[', '?' and '*').

-perm onum

True if the file permission flags exactly match the octal number *onum* (see *chmod(1)*). If *onum* is prefixed by a minus sign, more flag bits (017777, see *stat(2)*) become significant and the flags are compared: $(flags \& onum) == onum$.

-type c True if the type of the file is *c*, where *c* is **b**, **c**, **d**, **f** or **l** for block special file, character special file, directory, plain file, or symbolic link.

-links n True if the file has *n* links.

-user uname

True if the file belongs to the user *uname* (login name or numeric user ID).

-group gname

True if the file belongs to group *gname* (group name or numeric group ID).

-size n True if the file is *n* blocks long (512 bytes per block).

-inum n True if the file has inode number *n*.

-atime n True if the file has been accessed in *n* days.

-mtime n True if the file has been modified in *n* days.

-exec command

True if the executed command returns a zero value as exit status. The end of the command must be punctuated by an escaped semicolon. A command argument '{' is replaced by the current pathname.

-ok command

Like **-exec** except that the generated command is written on the standard output, then the standard input is read and the command executed only upon response *y*.

-print Always true; causes the current pathname to be printed.

-newer file

True if the current file has been modified more recently than the argument *file*.

The primaries may be combined using the following operators (in order of decreasing precedence):

- 1) A parenthesized group of primaries and operators (parentheses are special to the Shell and must be escaped).
- 2) The negation of a primary ('!' is the unary *not* operator).

- 3) Concatenation of primaries (the *and* operation is implied by the juxtaposition of two primaries).
- 4) Alternation of primaries ('*o*' is the *or* operator).

EXAMPLE

To 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 {} \;
```

FILES

/etc/passwd
/etc/group

SEE ALSO

sh(1), test(1), fs(5)

BUGS

The syntax is painful.

NAME

finger - user information lookup program

SYNOPSIS

finger [options] name ...

DESCRIPTION

By default *finger* lists the login name, full name, terminal name and write status (as a '*' before the terminal name if write permission is denied), idle time, login time, and office location and phone number (if they are known) for each current UNIX user. (Idle time is minutes if it is a single integer, hours and minutes if a ':' is present, or days and hours if a 'd' is present.)

A longer format also exists and is used by *finger* whenever a list of peoples names is given. (Account names as well as first and last names of users are accepted.) This format is multi-line, and includes all the information described above as well as the user's home directory and login shell, any plan which the person has placed in the file *.plan* in their home directory, and the project on which they are working from the file *.project* also in the home directory.

Finger options include:

- m Match arguments only on user name.
- l Force long output format.
- p Suppress printing of the *.plan* files
- s Force short output format.

FILES

/etc/utmp	who file
/etc/passwd	for users names, offices, ...
/usr/adm/lastlog	last login times
~/plan	plans
~/project	projects

SEE ALSO

w(1), who(1)

AUTHOR

Earl T. Cohen

BUGS

Only the first line of the *.project* file is printed.

The encoding of the *gcos* field is UCB dependant - it knows that an office '197MC' is '197M Cory Hall', and that '529BE' is '529B Evans Hall'.

A user information data base is in the works and will radically alter the way the information that *finger* uses is stored. *Finger* will require extensive modification when this is implemented.

NAME

fold - fold long lines for finite width output device

SYNOPSIS

fold [-width] [file ...]

DESCRIPTION

Fold is a filter which will fold the contents of the specified files, or the standard input if no files are specified, breaking the lines to have maximum width *width*. The default for *width* is 80. *Width* should be a multiple of 8 if tabs are present, or the tabs should be expanded using *expand(1)* before coming to *fold*.

SEE ALSO

expand(1)

BUGS

If underlining is present it may be messed up by folding.

NAME

fp - Functional Programming language compiler/interpreter

SYNOPSIS

fp

DESCRIPTION

Fp is an interpreter/compiler that implements the applicative language proposed by John Backus. It is written in FRANZ LISP.

In a functional programming language intent is expressed in a mathematical style devoid of assignment statements and variables. Functions compute by value only; there are no side-effects since the result of a computation depends solely on the inputs.

Fp "programs" consist of *functional expressions* - primitive and user-defined *fp* functions combined by *functional forms*. These forms take functional arguments and return functional results. For example, the composition operator '@' takes two functional arguments and returns a function which represents their composition.

There exists a single operation in *fp* - *application*. This operation causes the system to evaluate the indicated function using the single argument as input (all functions are monadic).

GETTING STARTED

Fp invokes the system. *Fp* compiles functions into *lisp(1)* source code; *lisp(1)* interprets this code (the user may compile this code using the *liszt(1)* compiler to gain a factor of 10 in performance). *Control D* exits back to the shell. *Break* terminates any computation in progress and resets any open file units. *)/help* provides a short summary of all user commands.

FILES

/usr/ucb/lisp the FRANZ LISP interpreter
/usr/ucb/liszt the liszt compiler
/usr/doc/fp the User's Guide

SEE ALSO

lisp(1), *liszt(1)*.

The Berkeley FP user's manual, available on-line. The language is described in the August 1978 issue of *CACM* (Turing award lecture by John Backus).

BUGS

If a non-terminating function is applied as the result of loading a file, then control is returned to the user immediately, everything after that position in the file is ignored.

FP incorrectly marks the location of a syntax error on large, multi-line function definitions or applications.

AUTHOR

Scott B. Baden

NAME

fpr - print Fortran file

SYNOPSIS

fpr

DESCRIPTION

Fpr is a filter that transforms files formatted according to Fortran's carriage control conventions into files formatted according to UNIX line printer conventions.

Fpr copies its input onto its output, replacing the carriage control characters with characters that will produce the intended effects when printed using *lpr*(1). The first character of each line determines the vertical spacing as follows:

Character	Vertical Space Before Printing
Blank	One line
0	Two lines
1	To first line of next page
+	No advance

A blank line is treated as if its first character is a blank. A blank that appears as a carriage control character is deleted. A zero is changed to a newline. A one is changed to a form feed. The effects of a "+" are simulated using backspaces.

EXAMPLES

a.out | *fpr* | *lpr*

fpr < *f77.output* | *lpr*

AUTHOR

Robert P. Corbett

BUGS

Results are undefined for input lines longer than 170 characters.

NAME

`fpu` - determine presence of the floating point coprocessor

SYNOPSIS

`fpu [-s]`

DESCRIPTION

Fpu prints whether or not the MC68881 floating point coprocessor (unit) is installed. The `-s` (silent) flag suppresses printing (except for error messages). It is used for checking error status (in shell scripts for example).

EXAMPLE

The the following is a shell script that demonstrates the use of the `-s` flag:

```
#!/bin/csh -f
fpu -s
if ($status) then
    echo No FPU installed.
else
    echo FPU installed.
endif
```

NAME

from - who is my mail from?

SYNOPSIS

from [**-s** sender] [user]

DESCRIPTION

From prints out the mail header lines in your mailbox file to show you who your mail is from. If *user* is specified, then *user's* mailbox is examined instead of your own. If the *-s* option is given, then only headers for mail sent by *sender* are printed.

FILES

/usr/spool/mail/*

SEE ALSO

biff(1), mail(1), prmail(1)

NAME

`fsplit` - split a multi-routine Fortran file into individual files

SYNOPSIS

`fsplit` [`-e efile`] ... [file]

DESCRIPTION

`Fsplit` takes as input either a file or standard input containing Fortran source code. It attempts to split the input into separate routine files of the form *name.f*, where *name* is the name of the program unit (e.g. function, subroutine, block data or program). The name for unnamed block data subprograms has the form *blkdaNNN.f* where NNN is three digits and a file of this name does not already exist. For unnamed main programs the name has the form *mainNNN.f*. If there is an error in classifying a program unit, or if *name.f* already exists, the program unit will be put in a file of the form *zzzNNN.f* where *zzzNNN.f* does not already exist.

Normally each subprogram unit is split into a separate file. When the `-e` option is used, only the specified subprogram units are split into separate files. E.g.:

```
fsplit -e readit -e doit prog.f
```

will split `readit` and `doit` into separate files.

DIAGNOSTICS

If names specified via the `-e` option are not found, a diagnostic is written to *standard error*.

AUTHOR

Asa Romberger and Jerry Berkman

BUGS

`Fsplit` assumes the subprogram name is on the first noncomment line of the subprogram unit. Nonstandard source formats may confuse `fsplit`.

It is hard to use `-e` for unnamed main programs and block data subprograms since you must predict the created file name.

NAME

ftp - file transfer program

SYNOPSIS

ftp [-v] [-d] [-i] [-n] [-g] [host]

DESCRIPTION

Ftp is the user interface to the ARPANET standard File Transfer Protocol. The program allows a user to transfer files to and from a remote network site.

The client host with which *ftp* is to communicate may be specified on the command line. If this is done, *ftp* will immediately attempt to establish a connection to an FTP server on that host; otherwise, *ftp* will enter its command interpreter and await instructions from the user. When *ftp* is awaiting commands from the user the prompt "ftp>" is provided the user. The following commands are recognized by *ftp*:

! Invoke a shell on the local machine.

append *local-file* [*remote-file*]

Append a local file to a file on the remote machine. If *remote-file* is left unspecified, the local file name is used in naming the remote file. File transfer uses the current settings for *type*, *format*, *mode*, and *structure*.

ascii Set the file transfer *type* to network ASCII. This is the default type.

bell Arrange that a bell be sounded after each file transfer command is completed.

binary

Set the file transfer *type* to support binary image transfer.

bye Terminate the FTP session with the remote server and exit *ftp*.

cd *remote-directory*

Change the working directory on the remote machine to *remote-directory*.

close Terminate the FTP session with the remote server, and return to the command interpreter.

delete *remote-file*

Delete the file *remote-file* on the remote machine.

debug [*debug-value*]

Toggle debugging mode. If an optional *debug-value* is specified it is used to set the debugging level. When debugging is on, *ftp* prints each command sent to the remote machine, preceded by the string "-->".

dir [*remote-directory*] [*local-file*]

Print a listing of the directory contents in the directory, *remote-directory*, and, optionally, placing the output in *local-file*. If no directory is specified, the current working directory on the remote machine is used. If no local file is specified, output comes to the terminal.

form *format*

Set the file transfer *form* to *format*. The default format is "file".

get *remote-file* [*local-file*]

Retrieve the *remote-file* and store it on the local machine. If the local file name is not specified, it is given the same name it has on the remote machine. The current settings for *type*, *form*, *mode*, and *structure* are used while transferring the file.

- hash** Toggle hash-sign (“#”) printing for each data block transferred. The size of a data block is 1024 bytes.
- glob** Toggle file name globbing. With file name globbing enabled, each local file or path-name is processed for *cs*(1) metacharacters. These characters include “*?[]~{}”. Remote files specified in multiple item commands, e.g. *mput*, are globbed by the remote server. With globbing disabled all files and pathnames are treated literally.
- help** [*command*]
Print an informative message about the meaning of *command*. If no argument is given, *ftp* prints a list of the known commands.
- lcd** [*directory*]
Change the working directory on the local machine. If no *directory* is specified, the user’s home directory is used.
- ls** [*remote-directory*] [*local-file*]
Print an abbreviated listing of the contents of a directory on the remote machine. If *remote-directory* is left unspecified, the current working directory is used. If no local file is specified, the output is sent to the terminal.
- mdelete** *remote-files*
Delete the specified files on the remote machine. If globbing is enabled, the specification of remote files will first be expanded using *ls*.
- mdir** *remote-files local-file*
Obtain a directory listing of multiple files on the remote machine and place the result in *local-file*.
- mget** *remote-files*
Retrieve the specified files from the remote machine and place them in the current local directory. If globbing is enabled, the specification of remote files will first be expanding using *ls*.
- mkdir** *directory-name*
Make a directory on the remote machine.
- mls** *remote-files local-file*
Obtain an abbreviated listing of multiple files on the remote machine and place the result in *local-file*.
- mode** [*mode-name*]
Set the file transfer *mode* to *mode-name*. The default mode is “stream” mode.
- mput** *local-files*
Transfer multiple local files from the current local directory to the current working directory on the remote machine.
- open** *host* [*port*]
Establish a connection to the specified *host* FTP server. An optional port number may be supplied, in which case, *ftp* will attempt to contact an FTP server at that port. If the *auto-login* option is on (default), *ftp* will also attempt to automatically log the user in to the FTP server (see below).
- prompt**
Toggle interactive prompting. Interactive prompting occurs during multiple file transfers to allow the user to selectively retrieve or store files. If prompting is turned off (default), any *mget* or *mput* will transfer all files.

put *local-file* [*remote-file*]

Store a local file on the remote machine. If *remote-file* is left unspecified, the local file name is used in naming the remote file. File transfer uses the current settings for *type*, *format*, *mode*, and *structure*.

pwd Print the name of the current working directory on the remote machine.

quit A synonym for *bye*.

quote *arg1 arg2 ...*

The arguments specified are sent, verbatim, to the remote FTP server. A single FTP reply code is expected in return.

recv *remote-file* [*local-file*]

A synonym for *get*.

remotehelp [*command-name*]

Request help from the remote FTP server. If a *command-name* is specified it is supplied to the server as well.

rename [*from*] [*to*]

Rename the file *from* on the remote machine, to the file *to*.

rmdir *directory-name*

Delete a directory on the remote machine.

send *local-file* [*remote-file*]

A synonym for *put*.

sendport

Toggle the use of PORT commands. By default, *ftp* will attempt to use a PORT command when establishing a connection for each data transfer. If the PORT command fails, *ftp* will use the default data port. When the use of PORT commands is disabled, no attempt will be made to use PORT commands for each data transfer. This is useful for certain FTP implementations which do ignore PORT commands but, incorrectly, indicate they've been accepted.

status Show the current status of *ftp*.

struct [*struct-name*]

Set the file transfer *structure* to *struct-name*. By default "stream" structure is used.

tenex Set the file transfer type to that needed to talk to TENEX machines.

trace Toggle packet tracing.

type [*type-name*]

Set the file transfer *type* to *type-name*. If no type is specified, the current type is printed. The default type is network ASCII.

user *user-name* [*password*] [*account*]

Identify yourself to the remote FTP server. If the password is not specified and the server requires it, *ftp* will prompt the user for it (after disabling local echo). If an account field is not specified, and the FTP server requires it, the user will be prompted for it. Unless *ftp* is invoked with "auto-login" disabled, this process is done automatically on initial connection to the FTP server.

verbose

Toggle verbose mode. In verbose mode, all responses from the FTP server are displayed to the user. In addition, if verbose is on, when a file transfer completes, statistics regarding the efficiency of the transfer are reported. By default, verbose is on.

? [*command*]

A synonym for help.

Command arguments which have embedded spaces may be quoted with quote (") marks.

FILE NAMING CONVENTIONS

Files specified as arguments to *ftp* commands are processed according to the following rules.

- 1) If the file name "-" is specified, the *stdin* (for reading) or *stdout* (for writing) is used.
- 2) If the first character of the file name is "|", the remainder of the argument is interpreted as a shell command. *Ftp* then forks a shell, using *popen(3)* with the argument supplied, and reads (writes) from the *stdout* (*stdin*). If the shell command includes spaces, the argument must be quoted; e.g. "'| ls -lt'". A particularly useful example of this mechanism is: "dir |more".
- 3) Failing the above checks, if "globbing" is enabled, local file names are expanded according to the rules used in the *cs(1)*; c.f. the *glob* command.

FILE TRANSFER PARAMETERS

The FTP specification specifies many parameters which may affect a file transfer. The *type* may be one of "ascii", "image" (binary), "ebcdic", and "local byte size" (for PDP-10's and PDP-20's mostly). *Ftp* supports the ascii and image types of file transfer.

Ftp supports only the default values for the remaining file transfer parameters: *mode*, *form*, and *struct*.

OPTIONS

Options may be specified at the command line, or to the command interpreter.

The *-v* (verbose on) option forces *ftp* to show all responses from the remote server, as well as report on data transfer statistics.

The *-n* option restrains *ftp* from attempting "auto-login" upon initial connection. If auto-login is enabled, *ftp* will check the *.netrc* file in the user's home directory for an entry describing an account on the remote machine. If no entry exists, *ftp* will use the login name on the local machine as the user identity on the remote machine, and prompt for a password and, optionally, an account with which to login.

The *-i* option turns off interactive prompting during multiple file transfers.

The *-d* option enables debugging.

The *-g* option disables file name globbing.

BUGS

Many FTP server implementations do not support the experimental operations such as print working directory. Aborting a file transfer does not work right; if one attempts this the local *ftp* will likely have to be killed by hand.

NAME

gcore - get core images of running processes

SYNOPSIS

gcore process-id ...

DESCRIPTION

Gcore creates a core image of each specified process, suitable for use with *adb*(1) or *dbx*(1).

FILES

core.<process-id> core images

BUGS

Paging activity that occurs while *gcore* is running may cause the program to become confused. For best results, the desired processes should be stopped.

NAME

gprof - display call graph profile data

SYNOPSIS

gprof [options] [a.out [gmon.out ...]]

DESCRIPTION

gprof produces an execution profile of C, Pascal, or Fortran77 programs. The effect of called routines is incorporated in the profile of each caller. The profile data is taken from the call graph profile file (*gmon.out* default) which is created by programs which are compiled with the **-pg** option of *cc*, *pc*, and *f77*. That option also links in versions of the library routines which 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, the *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(1)*. This listing gives the total execution times and call counts for each of the functions in the program, sorted by decreasing time.

Next, these times are propagated along the edges of the call graph. Cycles are discovered, and calls into a cycle are made to share the time of the cycle. A second listing shows the functions sorted according to the time they represent including the time of their call graph descendents. Below each function entry is shown its (direct) call graph children, and how their times are propagated to this function. A similar display above the function shows how this function's time and the time of its descendents is 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 and their contributions to the time and call counts of the cycle.

The following options are available:

- a** suppresses the printing of statically declared functions. If this option is given, all relevant information about the static function (*e.g.*, time samples, calls to other functions, calls from other functions) belongs to the function loaded just before the static function in the *a.out* file.
- b** suppresses the printing of a description of each field in the profile.
- c** the static call graph of the program is discovered by a heuristic which examines the text space of the object file. Static-only parents or children are indicated with call counts of 0.
- e name** suppresses the printing of the graph profile entry for routine *name* and all its descendents (unless they have other ancestors that aren't suppressed). More than one **-e** option may be given. Only one *name* may be given with each **-e** option.
- E name** suppresses the printing of the graph profile entry for routine *name* (and its descendents) as **-e**, above, and also excludes the time spent in *name* (and its descendents) from the total and percentage time computations. (For example, **-E mcount -E mcleanup** is the default.)
- f name** prints the graph profile entry of only the specified routine *name* and its descendents. More than one **-f** option may be given. Only one *name* may be given with each **-f** option.

- F** *name*
prints the graph profile entry of only the routine *name* and its descendants (as **-f**, above) and also uses only the times of the printed routines in total time and percentage computations. More than one **-F** option may be given. Only one *name* may be given with each **-F** option. The **-F** option overrides the **-E** option.
- s**
a profile file *gmon.sum* is produced which represents the sum of the profile information in all the specified profile files. This summary profile file may be given to subsequent executions of *gprof* (probably also with a **-s**) to accumulate profile data across several runs of an *a.out* file.
- z**
displays routines which have zero usage (as indicated by call counts and accumulated time). This is useful in conjunction with the **-c** option for discovering which routines were never called.

FILES

<i>a.out</i>	the namelist and text space.
<i>gmon.out</i>	dynamic call graph and profile.
<i>gmon.sum</i>	summarized dynamic call graph and profile.

SEE ALSO

monitor(3), profil(2), cc(1), prof(1)

"gprof: A Call Graph Execution Profiler", by 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.

BUGS

Beware of quantization errors. The granularity of the sampling is shown, but remains statistical at best. We assume 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 which are not themselves profiled will have the time of their profiled children propagated to them, but they will appear to be spontaneously invoked in the call graph listing, and will not have their time propagated further. Similarly, signal catchers, even though profiled, will appear to be spontaneous (although for more obscure reasons). Any profiled children of signal catchers should have their times propagated properly, unless the signal catcher was invoked during the execution of the profiling routine, in which case all is lost.

The profiled program must call *exit(2)* or return normally for the profiling information to be saved in the *gmon.out* file.

NAME

graph - draw a graph

SYNOPSIS

graph [option] ...

DESCRIPTION

Graph with no options takes pairs of numbers from the standard input as abscissas and ordinates of a graph. Successive points are connected by straight lines. The graph is encoded on the standard output for display by the *plot(1G)* filters.

If the coordinates of a point are followed by a nonnumeric string, that string is printed as a label beginning on the point. Labels may be surrounded with quotes "...", in which case they may be empty or contain blanks and numbers; labels never contain newlines.

The following options are recognized, each as a separate argument.

- a Supply abscissas automatically (they are missing from the input); spacing is given by the next argument (default 1). A second optional argument is the starting point for automatic abscissas (default 0 or lower limit given by **-x**).
- b Break (disconnect) the graph after each label in the input.
- c Character string given by next argument is default label for each point.
- g Next argument is grid style, 0 no grid, 1 frame with ticks, 2 full grid (default).
- l Next argument is label for graph.
- m Next argument is mode (style) of connecting lines: 0 disconnected, 1 connected (default). Some devices give distinguishable line styles for other small integers.
- s Save screen, don't erase before plotting.
- x [1]
If 1 is present, x axis is logarithmic. Next 1 (or 2) arguments are lower (and upper) x limits. Third argument, if present, is grid spacing on x axis. Normally these quantities are determined automatically.
- y [1]
Similarly for y.
- h Next argument is fraction of space for height.
- w Similarly for width.
- r Next argument is fraction of space to move right before plotting.
- u Similarly to move up before plotting.
- t Transpose horizontal and vertical axes. (Option **-x** now applies to the vertical axis.)

A legend indicating grid range is produced with a grid unless the **-s** option is present.

If a specified lower limit exceeds the upper limit, the axis is reversed.

SEE ALSO

spline(1G), plot(1G)

BUGS

Graph stores all points internally and drops those for which there isn't room.

Segments that run out of bounds are dropped, not windowed.

Logarithmic axes may not be reversed.

NAME

grep, **egrep**, **fgrep** – search a file for a pattern

SYNOPSIS

grep [option] ... expression [file] ...
egrep [option] ... [expression] [file] ...
fgrep [option] ... [strings] [file]

DESCRIPTION

Commands of the *grep* family search the input *files* (standard input default) for lines matching a pattern. Normally, each line found is copied to the standard output. *Grep* patterns are limited regular expressions in the style of *ex(1)*; it uses a compact nondeterministic algorithm. *Egrep* patterns are full regular expressions; it uses a fast deterministic algorithm that sometimes needs exponential space. *Fgrep* patterns are fixed strings; it is fast and compact. The following options are recognized.

- v** All lines but those matching are printed.
- x** (Exact) only lines matched in their entirety are printed (*fgrep* only).
- c** Only a count of matching lines is printed.
- l** The names of files with matching lines are listed (once) separated by newlines.
- n** Each line is preceded by its relative line number in the file.
- b** Each line is preceded by the block number on which it was found. This is sometimes useful in locating disk block numbers by context.
- i** The case of letters is ignored in making comparisons — that is, upper and lower case are considered identical. This applies to *grep* and *fgrep* only.
- s** Silent mode. Nothing is printed (except error messages). This is useful for checking the error status.
- w** The expression is searched for as a word (as if surrounded by ‘\<’ and ‘\>’, see *ex(1)*.) (*grep* only)
- e expression**
Same as a simple *expression* argument, but useful when the *expression* begins with a **-**.
- f file** The regular expression (*egrep*) or string list (*fgrep*) is taken from the *file*.

In all cases the file name is shown if there is more than one input file. Care should be taken when using the characters \$ * [^ | () and \ in the *expression* as they are also meaningful to the Shell. It is safest to enclose the entire *expression* argument in single quotes ‘ ’.

Fgrep searches for lines that contain one of the (newline-separated) *strings*.

Egrep accepts extended regular expressions. In the following description ‘character’ excludes newline:

A \ followed by a single character other than newline matches that character.

The character ^ matches the beginning of a line.

The character \$ matches the end of a line.

A . (period) matches any character.

A single character not otherwise endowed with special meaning matches that character.

A string enclosed in brackets `[]` matches any single character from the string. Ranges of ASCII character codes may be abbreviated as in `'a-z0-9'`. A `]` may occur only as the first character of the string. A literal `-` must be placed where it can't be mistaken as a range indicator.

A regular expression followed by an `*` (asterisk) matches a sequence of 0 or more matches of the regular expression. A regular expression followed by a `+` (plus) matches a sequence of 1 or more matches of the regular expression. A regular expression followed by a `?` (question mark) matches a sequence of 0 or 1 matches of the regular expression.

Two regular expressions concatenated match a match of the first followed by a match of the second.

Two regular expressions separated by `|` or newline match either a match for the first or a match for the second.

A regular expression enclosed in parentheses matches a match for the regular expression.

The order of precedence of operators at the same parenthesis level is `[]` then `*+?` then concatenation then `|` and newline.

Ideally there should be only one *grep*, but we don't know a single algorithm that spans a wide enough range of space-time tradeoffs.

SEE ALSO

`ex(1)`, `sed(1)`, `sh(1)`

DIAGNOSTICS

Exit status is 0 if any matches are found, 1 if none, 2 for syntax errors or inaccessible files.

BUGS

Lines are limited to 256 characters; longer lines are truncated.

NAME

`groups` - show group memberships

SYNOPSIS

`groups [user]`

DESCRIPTION

The `groups` command shows the groups to which you or the optionally specified user belong. Each user belongs to a group specified in the password file `/etc/passwd` and possibly to other groups as specified in the file `/etc/group`. If you do not own a file but belong to the group which it is owned by then you are granted group access to the file.

When a new file is created it is given the group of the containing directory.

SEE ALSO

`setgroups(2)`

FILES

`/etc/passwd`, `/etc/group`

BUGS

More groups should be allowed.

NAME

head - give first few lines

SYNOPSIS

head [-count] [file ...]

DESCRIPTION

This filter gives the first *count* lines of each of the specified files, or of the standard input. If *count* is omitted it defaults to 10.

SEE ALSO

tail(1)

NAME

`hostid` - set or print identifier of current host system

SYNOPSIS

`hostid` [identifier]

DESCRIPTION

The *hostid* command prints the identifier of the current host in hexadecimal. This numeric value is expected to be unique across all hosts and is normally set to the host's Internet address. The super-user can set the *hostid* by giving a hexadecimal argument; this is usually done in the startup script `/etc/rc.local`.

SEE ALSO

`gethostid(2)`, `sethostid(2)`

NAME

hostname - set or print name of current host system

SYNOPSIS

hostname [nameofhost]

DESCRIPTION

The *hostname* command prints the name of the current host, as given before the "login" prompt. The super-user can set the hostname by giving an argument; this is usually done in the startup script */etc/rc.local*.

SEE ALSO

gethostname(2), sethostname(2)

NAME

ident — identify files

SYNOPSIS

ident file ...

DESCRIPTION

Ident searches the named files for all occurrences of the pattern *\$keyword:...\$*, where *keyword* is one of

Author
Date
Header
Locker
Log
Revision
Source
State

These patterns are normally inserted automatically by the RCS command *co (1)*, but can also be inserted manually.

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*, then the command

```
ident f.c f.o
```

will print

```
f.c:      $Header: Header information $
f.o:      $Header: Header information $
```

IDENTIFICATION

Author: Walter F. Tichy, Purdue University, West Lafayette, IN, 47907.
Revision Number: 3.0 ; Release Date: 82/12/04 .
Copyright © 1982 by Walter F. Tichy.

SEE ALSO

ci (1), *co (1)*, *rcs (1)*, *rcsdiff(1)*, *rcsintro (1)*, *rcsmerge (1)*, *rlog (1)*, *rcsfile (5)*.
Walter F. Tichy, "Design, Implementation, and Evaluation of a Revision Control System," in *Proceedings of the 6th International Conference on Software Engineering*, IEEE, Tokyo, Sept. 1982.

IDENT (1)

USER COMMANDS

IDENT (1)

BUGS

NAME

indent - indent and format C program source

SYNOPSIS

indent input [output] [flags]

DESCRIPTION

Indent is intended primarily as a C program formatter. Specifically, *indent* will:

- indent code lines
- align comments
- insert spaces around operators where necessary
- break up declaration lists as in "int a,b,c;"

Indent will not break up long statements to make them fit within the maximum line length, but it will flag lines that are too long. Lines will be broken so that each statement starts a new line, and braces will appear alone on a line. (See the *-br* option to inhibit this.) Also, an attempt is made to line up identifiers in declarations.

The *flags* which can be specified follow. They may appear before or after the file names. If the *output* file is omitted, the formatted file will be written back into *input* and a "backup" copy of *input* will be written in the current directory. If *input* is named "/blah/blah/file", the backup file will be named ".Bfile". If *output* is specified, *indent* checks to make sure it is different from *input*.

The following flags may be used to control the formatting style imposed by *indent*.

- lnnn* Maximum length of an output line. The default is 75.
- cnnn* The column in which comments will start. The default is 33.
- cdnnn* The column in which comments on declarations will start. The default is for these comments to start in the same column as other comments.
- innn* The number of spaces for one indentation level. The default is 4.
- dj,-ndj* *-dj* will cause declarations to be left justified. *-ndj* will cause them to be indented the same as code. The default is *-ndj*.
- v,-nv* *-v* turns on "verbose" mode, *-nv* turns it off. When in verbose mode, *indent* will report when it splits one line of input into two or more lines of output, and it will give some size statistics at completion. The default is *-nv*.
- bc,-nbc* If *-bc* is specified, then a newline will be forced after each comma in a declaration. *-nbc* will turn off this option. The default is *-bc*.
- dnnn* This option controls the placement of comments which are not to the right of code. Specifying *-d2* means that such comments will be placed two indentation levels to the left of code. The default *-d0* lines up these comments with the code. See the section on comment indentation below.
- br,-bl* Specifying *-bl* will cause complex statements to be lined up like this:


```

if (...)
{
    code
}
      
```

Specifying `-br` (the default) will make them look like this:

```
if (...) {
    code
}
```

You may set up your own "profile" of defaults to *indent* by creating the file ".indent.pro" in your login directory and including whatever switches you like. If *indent* is run and a profile file exists, then it is read to set up the program's defaults. Switches on the command line, though, will always override profile switches. The profile file must be a single line of not more than 127 characters. The switches should be separated on the line by spaces or tabs.

Multi-line expressions

Indent will not break up complicated expressions that extend over multiple lines, but it will usually correctly indent such expressions which have already been broken up. Such an expression might end up looking like this:

```
x =
    (
      (Arbitrary parenthesized expression)
    +
      (
        (Parenthesized expression)
      *
        (Parenthesized expression)
      )
    );
```

Comments

Indent recognizes four kinds of comments. They are: straight text, "box" comments, UNIX-style comments, and comments that should be passed through unchanged. The action taken with these various types are as follows:

"Box" comments. *Indent* assumes that any comment with a dash immediately after the start of comment (i.e. `/*-`) is a comment surrounded by a box of stars. Each line of such a comment will be left unchanged, except that the first non-blank character of each successive line will be lined up with the beginning slash of the first line. Box comments will be indented (see below).

"Unix-style" comments. This is the type of section header which is used extensively in the UNIX system source. If the start of comment (`/*`) appears on a line by itself, *indent* assumes that it is a UNIX-style comment. These will be treated similarly to box comments, except the first non-blank character on each line will be lined up with the `*` of the `/*`.

Unchanged comments. Any comment which starts in column 1 will be left completely unchanged. This is intended primarily for documentation header pages. The check for unchanged comments is made before the check for UNIX-style comments.

Straight text. All other comments are treated as straight text. *Indent* will fit as many words (separated by blanks, tabs, or newlines) on a line as possible. Straight text comments will be indented.

Comment indentation

Box, UNIX-style, and straight text comments may be indented. If a comment is on a line with code it will be started in the "comment column", which is set by the `-cnnn` command line parameter. Otherwise, the comment will be started at *nnn* indentation levels less than where code is currently being placed, where *nnn* is specified by the `-dnnn` command line

parameter. (Indented comments will never be placed in column 1.) If the code on a line extends past the comment column, the comment will be moved to the next line.

DIAGNOSTICS

Diagnostic error messages, mostly to tell that a text line has been broken or is too long for the output line.

FILES

.indent.pro profile file

BUGS

Does not know how to format "long" declarations.

NAME

install - install binaries

SYNOPSIS

install [**-c**] [**-m mode**] [**-o owner**] [**-g group**] [**-s**] *binary destination*

DESCRIPTION

Binary is moved (or copied if **-c** is specified) to *destination*. If *destination* already exists, it is removed before *binary* is moved. If the *destination* is a directory then *binary* is moved into the *destination* directory with its original file-name.

The mode for *Destination* is set to 755; the **-m mode** option may be used to specify a different mode.

Destination is changed to owner root; the **-o owner** option may be used to specify a different owner.

Destination is changed to group staff; the **-g group** option may be used to specify a different group.

If the **-s** option is specified the binary is stripped after being installed.

Install refuses to move a file onto itself.

SEE ALSO

chgrp(1), chmod(1), cp(1), mv(1), strip(1), chown(8)

NAME

iostat - report I/O statistics

SYNOPSIS

iostat [interval [count]]

DESCRIPTION

Iostat iteratively reports the number of characters read and written to terminals, and, for each disk, the number of seeks transfers per second, kilobytes transfered per second, and the milliseconds per average seek. It also gives the percentage of time the system has spent in user mode, in user mode running low priority (niced) processes, in system mode, and idling.

To compute this information, for each disk, seeks and data transfer completions and number of words transferred are counted; for terminals collectively, the number of input and output characters are counted. Also, each sixtieth of a second, the state of each disk is examined and a tally is made if the disk is active. From these numbers and given the transfer rates of the devices it is possible to determine average seek times for each device.

The optional *interval* argument causes *iostat* to report once each *interval* seconds. The first report is for all time since a reboot and each subsequent report is for the last interval only.

The optional *count* argument restricts the number of reports.

FILES

/dev/kmem
/vmunix

SEE ALSO

vmstat(1)

NAME

`ipcrm` - remove a message queue, semaphore set or shared memory id

SYNOPSIS

`ipcrm` [*options*]

DESCRIPTION

Ipcrm will remove one or more specified messages, semaphore or shared memory identifiers. The identifiers are specified by the following *options*:

- q** *msgid* removes the message queue identifier *msgid* from the system and destroys the message queue and data structure associated with it.
- m** *shmid* removes the shared memory identifier *shmid* from the system. The shared memory segment and data structure associated with it are destroyed after the last detach.
- s** *semid* removes the semaphore identifier *semid* from the system and destroys the set of semaphores and data structure associated with it.
- Q** *msgkey* removes the message queue identifier, created with key *msgkey*, from the system and destroys the message queue and data structure associated with it.
- M** *shmkey* removes 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.
- S** *semkey* removes the semaphore identifier, created with key *semkey*, from the system and destroys the set of semaphores and data structure associated with it.

The details of the removes are described in *msgctl(2)*, *shmctl(2)*, and *semctl(2)*. The identifiers and keys may be found by using *ipcs(1)*.

SEE ALSO

ipcs(1), *msgget(2)*, *msgop(2)*, *semctl(2)*, *semget(2)*, *semop(2)*, *shmctl(2)*, *shmget(2)*, *shmop(2)*

NAME

`ipcs` – report inter-process communication facilities status

SYNOPSIS

`ipcs` [options]

DESCRIPTION

Ipcs prints certain information about active inter-process communication facilities. Without *options*, information is printed in short format for message queues, shared memory, and semaphores that are currently active in the system. Otherwise, the information that is displayed is controlled by the following *options*:

- q** Print information about active message queues.
- m** Print information about active shared memory segments.
- s** Print information about active semaphores.

If any of the options **-q**, **-m**, or **-s** are specified, information about only those indicated will be printed. If none of these three are specified, information about all three will be printed.

- b** Print biggest allowable size information. (Maximum number of bytes in messages on queue for message queues, size of segments for shared memory, and number of semaphores in each set for semaphores.) See below for meaning of columns in a listing.
- c** Print creator's login name and group name. See below.
- o** Print information on outstanding usage. (Number of messages on queue and total number of bytes in messages on queue for message queues and number of processes attached to shared memory segments.)
- p** Print process number information. (Process ID of last process to send a message and process ID of last process to receive a message on message queues and process ID of creating process and process ID of last process to attach or detach on shared memory segments) See below.
- t** Print time information. (Time of the last control operation that changed the access permissions for all facilities. Time of last *msgsnd* and last *msgrcv* on message queues, last *shmat* and last *shmdt* on shared memory, last *semop(2)* on semaphores.) See below.
- a** Use all print *options*. (This is a shorthand notation for **-b**, **-c**, **-o**, **-p**, and **-t**.)
- C** *corefile*
Use the file *corefile* in place of `/dev/kmem`.
- N** *namelist*
The argument will be taken as the name of an alternate *namelist* (`/unix` is the default).

The column headings and the meaning of the columns in an *ipcs* listing are given below; the letters in parentheses indicate the *options* that cause the corresponding heading to appear; **all** means that the heading always appears. Note that these *options* only determine what infor-

mation is provided for each facility; they do *not* determine which facilities will be listed.

- | | | |
|-----------|-------|--|
| T | (all) | Type of the facility: |
| | | q message queue; |
| | | m shared memory segment; |
| | | s semaphore. |
| ID | (all) | The identifier for the facility entry. |

KEY	(all)	The key used as an argument to <i>msgget</i> , <i>semget</i> , or <i>shmget</i> to create the facility entry. (Note: The key of a shared memory segment is changed to IPC_PRIVATE when the segment has been removed until all processes attached to the segment detach it.)
MODE	(all)	The facility access modes and flags: The mode consists of 11 characters that are interpreted as follows: The first two characters are: <ul style="list-style-type: none"> R if a process is waiting on a <i>msgrcv</i>; S if a process is waiting on a <i>msgsnd</i>; D if the associated shared memory segment has been removed. It will disappear when the last process attached to the segment detaches it; C if the associated shared memory segment is to be cleared when the first attach is executed; - if the corresponding special flag is not set. The next 9 characters are interpreted as three sets of three bits each. The first set refers to the owner's permissions; the next to permissions of others in the user-group of the 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. The permissions are indicated as follows: <ul style="list-style-type: none"> r if read permission is granted; w if write permission is granted; a if alter permission is granted; - if the indicated permission is <i>not</i> granted.
OWNER	(all)	The login name of the owner of the facility entry.
GROUP	(all)	The group name of the group of the owner of the facility entry.
CREATOR	(a,c)	The login name of the creator of the facility entry.
CGROUP	(a,c)	The group name of the group of the creator of the facility entry.
CBYTES	(a,o)	The number of bytes in messages currently outstanding on the associated message queue.
QNUM	(a,o)	The number of messages currently outstanding on the associated message queue.
QBYTES	(a,b)	The maximum number of bytes allowed in messages outstanding on the associated message queue.
LSPID	(a,p)	The process ID of the last process to send a message to the associated queue.
LRPID	(a,p)	The process ID of the last process to receive a message from the associated queue.
STIME	(a,t)	The time the last message was sent to the associated queue.
RTIME	(a,t)	The time the last message was received from the associated queue.
CTIME	(a,t)	The time when the associated entry was created or changed.
NATTCH	(a,o)	The number of processes attached to the associated shared memory segment.
SEGSZ	(a,b)	The size of the associated shared memory segment.
CPID	(a,p)	The process ID of the creator of the shared memory entry.
LPID	(a,p)	The process ID of the last process to attach or detach the shared memory segment.
ATIME	(a,t)	The time the last attach was completed to the associated shared memory segment.
DTIME	(a,t)	The time the last detach was completed on the associated shared memory segment.
NSEMS	(a,b)	The number of semaphores in the set associated with the semaphore entry.

OTIME (a,t) The time the last semaphore operation was completed on the set associated with the semaphore entry.

FILES

/vmunix	system namelist
/dev/kmem	memory
/etc/passwd	user names
/etc/group	group names

SEE ALSO

msgop(2), semop(2), shmop(2)

BUGS

Things can change while *ipcs* is running; the picture it gives is only a close approximation to reality.

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* is '-', the standard input is used.

File1 and *file2* must be sorted in increasing ASCII collating sequence on the fields on which they are to be joined, normally the first in each line.

There is one line in the output for each pair of lines in *file1* and *file2* that have identical join fields. The output line normally consists of the common field, then the rest of the line from *file1*, then the rest of the line from *file2*.

Fields are normally separated by blank, tab or newline. In this case, multiple separators count as one, and leading separators are discarded.

These options are recognized:

- 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 *n m* Join on the *m*th field of file *n*. If *n* is missing, use the *m*th field in each file.
- 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.
- t *c* Use character *c* as a separator (tab character). Every appearance of *c* in a line is significant.

SEE ALSO

sort(1), comm(1), awk(1)

BUGS

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*, *look* and *awk(1)* are wildly incongruous.

NAME

jove - an interactive display-oriented text editor

SYNOPSIS

```
jove [-d directory] [-w] [-t tag] [+n file] [-p file] [files]
jove -r
```

DESCRIPTION

JOVE is Jonathan's Own Version of Emacs. It is based on the original EMACS editor written at MIT by Richard Stallman. Although JOVE is meant to be compatible with EMACS, there are some major differences between the two editors and you shouldn't rely on their behaving identically.

JOVE works on any reasonable display terminal that is described in the *termcap* file (see TERMCAP(5) for more details). When you start up JOVE, it checks to see whether you have your *TERM* environment variable set. On most systems that will automatically be set up for you, but if it's not JOVE will ask you what kind of terminal you are using. To avoid having to type this every time you run JOVE you can set your *TERM* environment variable yourself. How you do this depends on which shell you are running. If you are running the C Shell, as most of you are, you type

```
% setenv TERM type
```

and with the Bourne Shell, you type

```
$ TERM= type ; export TERM
```

where *type* is the name of the kind of terminal you are using (e.g., vt100). If neither of these works get somebody to help you.

INVOKING JOVE

If you run JOVE with no arguments you will be placed in an empty buffer, called *Main*. Otherwise, any arguments you supply are considered file names and each is "given" its own buffer. Only the first file is actually read in--reading other files is deferred until you actually try to use the buffers they are attached to. This is for efficiency's sake: most of the time, when you run JOVE on a big list of files, you end up editing only a few of them.

The names of all of the files specified on the command line are saved in a buffer, called *Minibuf*. The mini-buffer is a special JOVE buffer that is used when JOVE is prompting for some input to many commands (for example, when JOVE is prompting for a file name). When you are being prompted for a file name, you can type C-N (that's Control-N) and C-P to cycle through the list of files that were specified on the command line. The file name will be inserted where you are typing and then you can edit it as if you typed it in yourself.

JOVE recognizes the following switches:

- d The following argument is taken to be the name of the current directory. This is for systems that don't have a version of C shell that automatically maintains the *CWD* environment variable. If *-d* is not specified on a system without a modified C shell, JOVE will have to figure out the current directory itself, and that can be VERY slow.
- +n Reads the file, designated by the following argument, and positions point at the *n*'th line instead of the (default) 1'st line. This can be specified more than once but it doesn't make sense to use it twice on the same file; in that case the second one wins.

- p Parses the error messages in the file designated by the following argument. The error messages are assumed to be in a format similar to the C compiler, LINT, or GREP output.
- t Runs the *find-tag* command on the following argument (see *ctags(1)*).
- w Divides the window in two. When this happens, either the same file is displayed in both windows, or the second file in the list is read in and displayed in its window.

As a special case, invoking JOVE with the -r option runs JOVE_RECOVER. Use this when the system crashes, or JOVE crashes, or you accidentally get logged out while in JOVE. If there are any buffers to be recovered, this will find them. Read the documentation for JOVE_RECOVER.

GETTING HELP

Once in JOVE, there are several commands available to get help. To execute any JOVE command, you type "<ESC> X command-name" followed by <Return>. To get a list of all the JOVE commands you type "<ESC> X" followed by "?". The *describe-bindings* command can be used to get a list containing each key, and its associated command (that is, the command that gets executed when you type that key). If you want to save the list of bindings, you can set the jove variable *send-typeout-to-buffer* to ON (using the *set* command), and then execute the *describe-bindings* command. This will create a buffer and put in it the bindings list it normally would have printed on the screen. Then you can save that buffer to a file and print it to use as a quick reference card. (See VARIABLES below.)

Once you know the name of a command, you can find out what it does with the *describe-command* command, which you can invoke quickly by typing "ESC ?". The *apropos* command will give you a list of all the command with a specific string in their names. For example, if you want to know the names of all the commands that are concerned with windows, you can run "apropos" with the keyword *window*.

If you're not familiar with the EMACS command set, it would be worth your while to use run TEACHJOVE. Do do that, just type "teachjove" to your shell and you will be placed in JOVE in a file which contains directions. I highly recommend this for beginners; you may save yourself a lot of time and headaches.

KEY BINDINGS and VARIABLES

You can alter the key bindings in JOVE to fit your personal tastes. That is, you can change what a key does every time you strike it. For example, by default the C-N key is bound to the command *next-line* and so when you type it you move down a line. If you want to change a binding or add a new one, you use the *bind-to-key* command. The syntax is "bind-to-key <command> key".

You can also change the way JOVE behaves in little ways by changing the value of some variables with the *set* command. The syntax is "set <variable> value", where value is a number or a string, or "on" or "off", depending on the context. For example, if you want JOVE to make backup files, you set the "make-backup-files" variable to "on". To see the value of a variable, use the "print <variable>" command.

INITIALIZATION

JOVE automatically reads commands from an initialization file in your HOME directory, called ".joverc". In this file you can place commands that you would normally type in JOVE. If you like to rearrange the key bindings and set some variables every time you get into JOVE, you should put them in your initialization file. Here are a few lines from mine:

```
set match-regular-expressions on
```

```

auto-execute-command auto-fill /tmp/Re\|.*drft
bind-to-key i-search-forward ^\
bind-to-key i-search-reverse ^R
bind-to-key find-tag-at-point ^[T
bind-to-key scroll-down ^C
bind-to-key grow-window ^Xg
bind-to-key shrink-window ^Xs

```

(Note that the Control Characters can be either two character sequences (e.g. ^ and C together as ^C) or the actual control character. If you want to use an ^ by itself you must BackSlash it (e.g., bind-to-key grow-window ^X\^ binds grow-window to ""X^").

SOME MINOR DETAILS

You should type C-\ instead of C-S in many instances. For example, the way to search for a string is documented as being "C-S" but in reality you should type "C-\". This is because C-S is the XOFF character (what gets sent when you type the NO SCROLL key), and clearly that won't work. The XON character is "C-Q" (what gets sent when you type NO SCROLL again) which is documented as the way to do a quoted-insert. The alternate key for this is "C-'" (typed as "C-'" on vt100's and its look-alikes). If you want to enable C-S and C-Q and you know what you are doing, you can put the line:

```
set allow-^S-and-^Q on
```

in your initialization file.

If your terminal has a metakey, JOVE will use it if you turn on the "meta-key" variable. JOVE will automatically turn on "meta-key" if the METAKEY environment variable exists. This is useful for if you have different terminals (e.g., one at home and one at work) and one has a metakey and the other doesn't.

FILES

```

/usr/new/lib/jove/.joverc - system wide initialization file
~/joverc - personal initialization file
/tmp - where temporary files are stored
/usr/new/lib/jove/teach-jove - the interactive tutorial
/usr/new/lib/jove/portsrv - for running shells in windows (pdp11 only)

```

SEE ALSO

```

jove_recover(1) - to recover buffers after a
    system/editor crash
ed(1) - for a description of regular expressions
teachjove(1) - for an interactive JOVE tutorial.

```

DIAGNOSTICS

JOVE diagnostics are meant to be self-explanatory, but you are advised to seek help whenever you are confused. You can easily lose a lot of work if you don't know EXACTLY what you are doing.

BUGS

Lines can't be more than 1024 characters long.

Searches can't cross line boundaries.

AUTHOR

Jonathan Payne

NAME

jove_recover - recover JOVE buffers after a system/editor crash

SYNOPSIS

jove_recover [-syscrash] [-d directory] jove -r

DESCRIPTION

JOVE_RECOVER lets you recover your work in the JOVE editor in the event of a system or JOVE crash. It is designed to put invoked through JOVE with the "-r" switch. JOVE_RECOVER looks for JOVE buffers that are left around and are owned by you. (You cannot recover other peoples' buffers, obviously.) When the system is rebooted after a crash, you type "jove -r" after you've logged in. If there were no buffers that were modified at the time of the crash or there were but JOVE_RECOVER can't get its hands on them, you will be informed with the message, There is nothing here for you. Otherwise, JOVE_RECOVER prints the date and time of the version of the buffers it has, and then waits for you type a command.

To get a list of the buffers JOVE_RECOVER knows about, use the *list* command. This will list all the buffers and the files and the number of lines associated with them. Next to each buffer is a number. When you want to recover a buffer, use the *get* command. The syntax is *get buffer filename* where *buffer* is either the buffer's name or the number at the beginning of the line. If you don't type the buffer name or the filename, JOVE_RECOVER will prompt you for them.

If there are a lot of buffers and you want to recover all of them, use the *recover* command. This will recover each buffer to the name of the buffer with ".#" prepended to the name (so that the original isn't over-written). It asks for each file and if you want to restore that buffer to that name you type "yes". If you want to recover the file but to a different name, just type that name in. If you type "no" JOVE_RECOVER will skip that file and go on to the next one.

If you want to look at a buffer before deciding to recover it, use the *print* command. The syntax for this is *print buffer* where *buffer* again is either its name or the number. You can type ^C if you want to abort printing the file to the terminal, and JOVE_RECOVER will respond with an appropriate message.

When you're done and have all the buffers you want, type the *quit* command to leave. You will then be asked whether it's okay to delete the tmp files. Most of the time that's okay and you should type "yes". When you say that, JOVE removes all traces of those buffers and you won't be able to look at them again. (If you recovered some buffers they will still be around, so don't worry.) So, if you're not sure whether you've gotten all the buffers, you should answer "no" so that you'll be able to run JOVE_RECOVER again at a later time (presumably after you've figured out which ones you want to save).

If you type ^C at any time other than when you're printing a file to the terminal, JOVE_RECOVER will exit without a word. If you do this but wish you hadn't, just type "jove -r" to the shell again, and you will be put back with no loss.

A SAMPLE SESSION

```
% jove -r
Found 2 buffers (last updated: Sun Apr 14 14:13:38 1985).
(Type '?' for options): list
1) buffer recover.1 "/u/staff/jpay/doc/recover.1" (120 lines)
2) buffer recover.c "/u/staff/jpay/jove/recover.c" (635 lines)
```

(Type '?' for options): get recover.1 recover.save
"recover.save" 53 lines, 1821 characters.
(Type '?' for options): quit
Should I delete the tmp files? yes
%

Here I "got" the buffer *recover.1* and restored it to the temporary file *recover.save*. It's generally a good idea to recover buffers to temporary files and then compare them to the original or at least look them over before putting them back in their real name. This is just in case you were restoring what you thought you were restoring.

FILES

/tmp - where temporary files are stored.

SEE ALSO

JOVE(1) - for this to make any sense to you.

DIAGNOSTICS

JOVE_RECOVER diagnostics are meant to be self-explanatory.

BUGS

It works well enough...

AUTHOR

Jonathan Payne

NAME

kermit - kermit file transfer

SYNOPSIS

kermit [option ...] [file ...]

DESCRIPTION

Kermit is a public domain file transfer program that allows files to be moved between machines of many different operating systems and architectures. This man page describes version 4C of the program.

Arguments are optional. If *Kermit* is executed without arguments, it will enter command mode. Otherwise, *kermit* will read the arguments off the command line and interpret them.

The following notation is used in command descriptions:

fn A Unix file specification, possibly containing either of the "wildcard" characters '*' or '?' ('*' matches all character strings, '?' matches any single character).

fn1 A Unix file specification which may not contain '*' or '?'.

rfn A remote file specification in the remote system's own syntax, which may denote a single file or a group of files.

rfn1 A remote file specification which should denote only a single file.

n A decimal number between 0 and 94.

c A decimal number between 0 and 127 representing the value of an ASCII character.

cc A decimal number between 0 and 31, or else exactly 127, representing the value of an ASCII control character.

[] Any field in square braces is optional.

{*x,y,z*} Alternatives are listed in curly braces.

Kermit command line options may specify either actions or settings. If *Kermit* is invoked with a command line that specifies no actions, then it will issue a prompt and begin interactive dialog. Action options specify either protocol transactions or terminal connection.

COMMAND LINE OPTIONS

-s *fn* Send the specified file or files. If *fn* contains wildcard (meta) characters, the Unix shell expands it into a list. If *fn* is '-' then *Kermit* sends from standard input, which must come from a file:

```
kermit -s - < foo.bar
```

or a parallel process:

```
ls -l | kermit -s -
```

You cannot use this mechanism to send terminal typein. If you want to send a file whose name is "-" you can precede it with a path name, as in

```
kermit -s ./-
```

-r Receive a file or files. Wait passively for files to arrive.

- k** Receive (passively) a file or files, sending them to standard output. This option can be used in several ways:

```
kermit -k
```

Displays the incoming files on your screen; to be used only in "local mode" (see below).

```
kermit -k > fn1
```

Sends the incoming file or files to the named file, *fn1*. If more than one file arrives, all are concatenated together into the single file *fn1*.

```
kermit -k | command
```

Pipes the incoming data (single or multiple files) to the indicated command, as in

```
kermit -k | sort > sorted.stuff
```

- a *fn1*** If you have specified a file transfer option, you may specify an alternate name for a single file with the **-a** option. For example,

```
kermit -s foo -a bar
```

sends the file *foo* telling the receiver that its name is *bar*. If more than one file arrives or is sent, only the first file is affected by the **-a** option:

```
kermit -ra baz
```

stores the first incoming file under the name *baz*.

- x** Begin server operation. May be used in either local or remote mode.

Before proceeding, a few words about remote and local operation are necessary. *Kermit* is "local" if it is running on a PC or workstation that you are using directly, or if it is running on a multiuser system and transferring files over an external communication line — not your job's controlling terminal or console. *Kermit* is remote if it is running on a multiuser system and transferring files over its own controlling terminal's communication line, connected to your PC or workstation.

If you are running *Kermit* on a PC, it is in local mode by default, with the "back port" designated for file transfer and terminal connection. If you are running *Kermit* on a multiuser (timesharing) system, it is in remote mode unless you explicitly point it at an external line for file transfer or terminal connection. The following command sets *Kermit*'s "mode":

- l *dev*** Line — Specify a terminal line to use for file transfer and terminal connection, as in

```
kermit -l /dev/ttyi5
```

When an external line is being used, you might also need some additional options for successful communication with the remote system:

- b *n*** Baud — Specify the baud rate for the line given in the **-l** option, as in

```
kermit -l /dev/ttyi5 -b 9600
```

This option should always be included with the **-l** option, since the speed of an external line is not necessarily what you expect.

- p *x*** Parity — **e**, **o**, **m**, **s**, **n** (even, odd, mark, space, or none). If parity is other than none, then the 8th-bit prefixing mechanism will be used for transferring 8-bit binary data, provided the opposite *Kermit* agrees. The default parity is none.

- t** Specifies half duplex, line turnaround with XON as the handshake character.

The following commands may be used only with a *Kermit* which is local — either by default or else because the `-l` option has been specified.

`-g rfn` Actively request a remote server to send the named file or files; *rfn* is a file specification in the remote host's own syntax. If *fn* happens to contain any special shell characters, like '*', these must be quoted, as in

```
kermit -g x\*.\?
```

`-f` Send a 'finish' command to a remote server.

`-c` Establish a terminal connection over the specified or default communication line, before any protocol transaction takes place. Get back to the local system by typing the escape character (normally Control-Backslash) followed by the letter 'c'.

`-n` Like `-c`, but after a protocol transaction takes place; `-c` and `-n` may both be used in the same command. The use of `-n` and `-c` is illustrated below.

On a timesharing system, the `-l` and `-b` options will also have to be included with the `-r`, `-k`, or `-s` options if the other *Kermit* is on a remote system.

If *kermit* is in local mode, the screen (stdout) is continuously updated to show the progress of the file transfer. A dot is printed for every four data packets, other packets are shown by type (e.g. 'S' for Send-Init), 'T' is printed when there's a timeout, and '%' for each retransmission. In addition, you may type (to stdin) certain "interrupt" commands during file transfer:

Control-F: Interrupt the current File, and go on to the next (if any).

Control-B: Interrupt the entire Batch of files, terminate the transaction.

Control-R: Resend the current packet

Control-A: Display a status report for the current transaction.

These interrupt characters differ from the ones used in other *Kermit* implementations to avoid conflict with Unix shell interrupt characters. With System III and System V implementations of Unix, interrupt commands must be preceded by the escape character (e.g. control-\).

Several other command-line options are provided:

`-i` Specifies that files should be sent or received exactly "as is" with no conversions. This option is necessary for transmitting binary files. It may also be used to slightly boost efficiency in Unix-to-Unix transfers of text files by eliminating CRLF/newline conversion.

`-w` Write-Protect — Avoid filename collisions for incoming files.

`-q` Quiet — Suppress screen update during file transfer, for instance to allow a file transfer to proceed in the background.

`-d` Debug — Record debugging information in the file debug.log in the current directory. Use this option if you believe the program is misbehaving, and show the resulting log to your local *Kermit* maintainer.

`-h` Help — Display a brief synopsis of the command line options.

The command line may contain no more than one protocol action option.

INTERACTIVE OPERATION

Kermit's interactive command prompt is "C-Kermit>". In response to this prompt, you may type any valid command. *Kermit* executes the command and then prompts you for another command. The process continues until you instruct the program to terminate.

Commands begin with a keyword, normally an English verb, such as "send". You may omit trailing characters from any keyword, so long as you specify sufficient characters to distinguish it from any other keyword valid in that field. Certain commonly-used keywords (such as "send", "receive", "connect") have special non-unique abbreviations ("s" for "send", "r" for "receive", "c" for "connect").

Certain characters have special functions in interactive commands:

- ?** Question mark, typed at any point in a command, will produce a message explaining what is possible or expected at that point. Depending on the context, the message may be a brief phrase, a menu of keywords, or a list of files.
- ESC** (The Escape or Altmode key) — Request completion of the current keyword or filename, or insertion of a default value. The result will be a beep if the requested operation fails.
- DEL** (The Delete or Rubout key) — Delete the previous character from the command. You may also use BS (Backspace, Control-H) for this function.
- ^W** (Control-W) — Erase the rightmost word from the command line.
- ^U** (Control-U) — Erase the entire command.
- ^R** (Control-R) — Redisplay the current command.
- SP** (Space) — Delimits fields (keywords, filenames, numbers) within a command. HT (Horizontal Tab) may also be used for this purpose.
- CR** (Carriage Return) — Enters the command for execution. LF (Linefeed) or FF (formfeed) may also be used for this purpose.
- ** (Backslash) — Enter any of the above characters into the command, literally. To enter a backslash, type two backslashes in a row (\\). A single backslash immediately preceding a carriage return allows you to continue the command on the next line.

You may type the editing characters (DEL, ^W, etc) repeatedly, to delete all the way back to the prompt. No action will be performed until the command is entered by typing carriage return, linefeed, or formfeed. If you make any mistakes, you will receive an informative error message and a new prompt — make liberal use of "?" and ESC to feel your way through the commands. One important command is "help" — you should use it the first time you run *Kermit*.

Interactive *Kermit* accepts commands from files as well as from the keyboard. When you enter interactive mode, *Kermit* looks for the file .kermrc in your home or current directory (first it looks in the home directory, then in the current one) and executes any commands it finds there. These commands must be in interactive format, not Unix command-line format. A "take" command is also provided for use at any time during an interactive session. Command files may be nested to any reasonable depth.

Here is a brief list of *Kermit* interactive commands:

- !** Execute a Unix shell command.
- bye** Terminate and log out a remote *Kermit* server.
- close** Close a log file.
- connect** Establish a terminal connection to a remote system.
- cwd** Change Working Directory.

dial	Dial a telephone number.
directory	Display a directory listing.
echo	Display arguments literally.
exit	Exit from the program, closing any open logs.
finish	Instruct a remote <i>Kermit</i> server to exit, but not log out.
get	Get files from a remote <i>Kermit</i> server.
help	Display a help message for a given command.
log	Open a log file — debugging, packet, session, transaction.
quit	Same as 'exit'.
receive	Passively wait for files to arrive.
remote	Issue file management commands to a remote <i>Kermit</i> server.
script	Execute a login script with a remote system.
send	Send files.
server	Begin server operation.
set	Set various parameters.
show	Display values of 'set' parameters.
space	Display current disk space usage.
statistics	Display statistics about most recent transaction.
take	Execute commands from a file.

The 'set' parameters are:

block-check	Level of packet error detection.
delay	How long to wait before sending first packet.
duplex	Specify which side echoes during 'connect'.
escape-character	Character to prefix "escape commands" during 'connect'.
file	Set various file parameters.
flow-control	Communication line full-duplex flow control.
handshake	Communication line half-duplex turnaround character.
line	Communication line device name.
modem-dialer	Type of modem-dialer on communication line.
parity	Communication line character parity.
prompt	Change the <i>Kermit</i> program's prompt.
receive	Set various parameters for inbound packets.
send	Set various parameters for outbound packets.

speed Communication line speed.

The 'remote' commands are:

cwd Change remote working directory.
delete Delete remote files.
directory Display a listing of remote file names.
help Request help from a remote server.
host Issue a command to the remote host in its own command language.
space Display current disk space usage on remote system.
type Display a remote file on your screen.
who Display who's logged in, or get information about a user.

FILES

`$HOME/.kermrc` *Kermit* initialization commands
`./kermrc` more *Kermit* initialization commands

SEE ALSO

`cu(1C)`, `uucp(1C)`
Frank da Cruz and Bill Catchings, *Kermit User's Guide*, Columbia University, 6th Edition

DIAGNOSTICS

The diagnostics produced by *Kermit* itself are intended to be self-explanatory.

BUGS

See recent issues of the Info-Kermit digest (on ARPANET or Usenet), or the file `ckuker.bwr`, for a list of bugs.

COPYRIGHT

Copyright (C) 1985, Trustess of Columbia University in the City of New York. Permission is granted to any individual or institution to use, copy, or redistribute this (*Kermit*) software so long as it is not sold for profit, provided this copyright notice is retained.

NAME

kill - terminate a process with extreme prejudice

SYNOPSIS

kill [-sig] processid ...

kill -l

DESCRIPTION

Kill sends the TERM (terminate, 15) signal to the specified processes. If a signal name or number preceded by '-' is given as first argument, that signal is sent instead of terminate (see *sigvec(2)*). The signal names are listed by 'kill -l', and are as given in */usr/include/signal.h*, stripped of the common SIG prefix.

The terminate signal will kill processes that do not catch the signal; 'kill -9 ...' is a sure kill, as the KILL (9) signal cannot be caught. By convention, if process number 0 is specified, all members in the process group (i.e. processes resulting from the current login) are signaled (but beware: this works only if you use *sh(1)*; not if you use *cs(1)*.) The killed processes must belong to the current user unless he is the super-user.

The process number of an asynchronous process started with '&' is reported by the shell. Process numbers can also be found by using *Kill* is a built-in to *cs(1)*; it allows job specifiers "%..." so process id's are not as often used as *kill* arguments. See *cs(1)* for details.

SEE ALSO

cs(1), *ps(1)*, *kill(2)*, *sigvec(2)*

BUGS

An option to kill process groups ala *killpg(2)* should be provided; a replacement for "kill 0" for *cs(1)* users should be provided.

NAME

last - indicate last logins of users and teletypes

SYNOPSIS

last [-N] [name ...] [tty ...]

DESCRIPTION

Last will look back in the *wtmp* file which records all logins and logouts for information about a user, a teletype or any group of users and teletypes. Arguments specify names of users or teletypes of interest. Names of teletypes may be given fully or abbreviated. For example 'last 0' is the same as 'last tty0'. If multiple arguments are given, the information which applies to any of the arguments is printed. For example 'last root console' would list all of "root's" sessions as well as all sessions on the console terminal. *Last* will print the sessions of the specified users and teletypes, most recent first, indicating the times at which the session began, the duration of the session, and the teletype which the session took place on. If the session is still continuing or was cut short by a reboot, *last* so indicates.

The pseudo-user **reboot** logs in at reboots of the system, thus

last reboot

will give an indication of mean time between reboot.

Last with no arguments prints a record of all logins and logouts, in reverse order. The **-N** option limits the report to N lines.

If *last* is interrupted, it indicates how far the search has progressed in *wtmp*. If interrupted with a quit signal (generated by a control-\) *last* indicates how far the search has progressed so far, and the search continues.

FILES

/usr/adm/wtmp	login data base
/usr/adm/shutdownlog	which records shutdowns and reasons for same

SEE ALSO

wtmp(5), ac(8), lastcomm(1)

AUTHOR

Howard Katseff

NAME

`lastcomm` - show last commands executed in reverse order

SYNOPSIS

`lastcomm` [command name] ... [user name] ... [terminal name] ...

DESCRIPTION

Lastcomm gives information on previously executed commands. With no arguments, *lastcomm* prints information about all the commands recorded 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. So, for example,

`lastcomm a.out root ttyd0`

would produce a listing of all the executions of commands named *a.out* by user *root* on the terminal *ttyd0*.

For each process entry, the following are printed.

The name of the user who ran the process.

Flags, as accumulated by the accounting facilities in the system.

The command name under which the process was called.

The amount of cpu time used by the process (in seconds).

The time the process exited.

The flags are encoded as follows: "S" indicates the command was executed by the super-user, "F" indicates the command ran after a fork, but without a following *exec*, "D" indicates the command terminated with the generation of a *core* file, and "X" indicates the command was terminated with the signal SIGTERM.

SEE ALSO

`last(1)`, `sigvec(2)`, `acct(5)`, `core(5)`

NAME

ld - link editor

SYNOPSIS

ld [option] ... file ...

DESCRIPTION

Ld combines several object programs into one, resolves external references, and searches libraries. In the simplest case several object *files* are given, and *ld* combines them, producing an object module which can be either executed or become the input for a further *ld* run. (In the latter case, the *-r* option must be given to preserve the relocation bits.) The output of *ld* is left on **a.out**. This file is made executable only if no errors occurred during the load.

The argument routines are concatenated in the order specified. The entry point of the output is the beginning of the first routine (unless the *-e* option is specified).

If any argument is a library, it is searched exactly once at the point it is encountered in the argument list. Only those routines defining an unresolved external reference are loaded. If a routine from a library references another routine in the library, and the library has not been processed by *ranlib*(1), the referenced routine must appear after the referencing routine in the library. Thus the order of programs within libraries may be important. The first member of a library should be a file named **'__SYMDEF'**, which is understood to be a dictionary for the library as produced by *ranlib*(1); the dictionary is searched iteratively to satisfy as many references as possible.

The symbols **'_etext'**, **'_edata'** and **'_end'** (**'etext'**, **'edata'** and **'end'** in C) are reserved, and if referred to, are set to the first location above the program, the first location above initialized data, and the first location above all data respectively. It is erroneous to define these symbols.

Ld understands several options. Except for *-l*, they should appear before the file names.

- A** This option specifies incremental loading, i.e. linking is to be done in a manner so that the resulting object may be read into an already executing program. The next argument is the name of a file whose symbol table will be taken as a basis on which to define additional symbols. Only newly linked material will be entered into the text and data portions of **a.out**, but the new symbol table will reflect every symbol defined before and after the incremental load. This argument must appear before any other object file in the argument list. The *-T* option may be used as well, and will be taken to mean that the newly linked segment will commence at the corresponding address (which must be a multiple of 1024). The default value is the old value of **_end**.
- D** Take the next argument as a hexadecimal number and pad the data segment with zero bytes to the indicated length.
- d** Force definition of common storage even if the *-r* flag is present.
- e** The following argument is taken to be the name of the entry point of the loaded program; location 0 is the default.
- lx** This option is an abbreviation for the library name **'/lib/libx.a'**, where *x* is a string. If that does not exist, *ld* tries **'/usr/lib/libx.a'** A library is searched when its name is encountered, so the placement of a *-l* is significant.
- M** produce a primitive load map, listing the names of the files which will be loaded.
- N** Do not make the text portion read only or sharable. (Use "magic number" 0407.)

- n** Arrange (by giving the output file a 0410 "magic number") that when the output file is executed, the text portion will be read-only and shared among all users executing the file. This involves moving the data areas up to the first possible 1024 byte boundary following the end of the text.
- o** The *name* argument after **-o** is used as the name of the *ld* output file, instead of **a.out**.
- r** Generate relocation bits in the output file so that it can be the subject of another *ld* run. This flag also prevents final definitions from being given to common symbols, and suppresses the 'undefined symbol' diagnostics.
- S** 'Strip' the output by removing all symbols except locals and globals.
- s** 'Strip' the output, that is, remove the symbol table and relocation bits to save space (but impair the usefulness of the debuggers). This information can also be removed by *strip(1)*.
- T** The next argument is a hexadecimal number which sets the text segment origin. The default origin is 0.
- t** ("trace") Print the name of each file as it is processed.
- u** Take the following argument as a symbol and enter it as undefined in the symbol table. This is useful for loading wholly from a library, since initially the symbol table is empty and an unresolved reference is needed to force the loading of the first routine.
- U** Suppress "Undefined symbol" messages.
- X** Save local symbols except for those whose names begin with 'L'. This option is used by *cc(1)* to discard internally-generated labels while retaining symbols local to routines.
- x** Do not preserve local (non-globl) symbols in the output symbol table; only enter external symbols. This option saves some space in the output file.
- ysym** Indicate each file in which *sym* appears, its type and whether the file defines or references it. Many such options may be given to trace many symbols. (It is usually necessary to begin *sym* with an '_', as external C, FORTRAN and Pascal variables begin with underscores.)
- z** Arrange for the process to be loaded on demand from the resulting executable file (413 format) rather than preloaded. This is the default. Results in a 1024 byte header on the output file followed by a text and data segment each of which have size a multiple of 1024 bytes (being padded out with nulls in the file if necessary). With this format the first few BSS segment symbols may actually appear (from the output of *size(1)*) to live in the data segment; this to avoid wasting the space resulting from data segment size roundup.

FILES

/lib/lib*.a	libraries
/usr/lib/lib*.a	more libraries
/usr/local/lib/lib*.a	still more libraries
a.out	output file

SEE ALSO

as(1), ar(1), cc(1), ranlib(1)



BUGS

There is no way to force data to be page aligned. *Ld* pads images which are to be demand loaded from the file system to the next page boundary to avoid a bug in the system.

NAME

`learn` - computer aided instruction about UNIX

SYNOPSIS

`learn` [`-directory`] [`subject` [`lesson`]]

DESCRIPTION

Learn gives Computer Aided Instruction courses and practice in the use of UNIX, the C Shell, and the Berkeley text editors. To get started simply type `learn`. The program will ask questions to find out what you want to do. Some questions may be bypassed by naming a *subject*, and more yet by naming a *lesson*. You may enter the *lesson* as a number that *learn* gave you in a previous session. If you do not know the lesson number, you may enter the *lesson* as a word, and *learn* will look for the first lesson containing it. If the *lesson* is '-', *learn* prompts for each lesson; this is useful for debugging.

The *subject*'s presently handled are

- files
- editor
- vi
- morefiles
- macros
- eqn
- C

There are a few special commands. The command 'bye' terminates a *learn* session and 'where' tells you of your progress, with 'where m' telling you more. The command 'again' re-displays the text of the lesson and 'again *lesson*' lets you review *lesson*.

The `-directory` option allows one to exercise a script in a nonstandard place.

FILES

- `/usr/lib/learn` subtree for all dependent directories and files
- `/usr/tmp/pl*` playpen directories

SEE ALSO

`csh(1)`, `ex(1)`

BUGS

The main strength of *learn*, that it asks the student to use the real UNIX, also makes possible baffling mistakes. It is helpful, especially for nonprogrammers, to have a UNIX initiate near at hand during the first sessions.

Occasionally lessons are incorrect, sometimes because the local version of a command operates in a non-standard way. Such lessons may be skipped with the 'skip' command, but it takes some sophistication to recognize the situation.

To find a *lesson* given as a word, *learn* does a simple *fgrep(1)* through the lessons. It is unclear whether this sort of subject indexing is better than none.

Spawning a new shell is required for each of many user and internal functions.

NAME

leave - remind you when you have to leave

SYNOPSIS

leave [hhmm]

DESCRIPTION

Leave waits until the specified time, then reminds you that you have 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 just before it would have printed the next message.

The time of day is in the form hhmm where hh is a time in hours (on a 12 or 24 hour clock). All times are converted to a 12 hour clock, and assumed to be in the next 12 hours.

If no argument is given, *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*.

Leave ignores interrupts, quits, and terminates. To get rid of it you should either log off or use "kill -9" giving its process id.

SEE ALSO

calendar(1)

AUTHOR

Mark Horton

NAME

lex - generator of lexical analysis programs

SYNOPSIS

lex [-tvfn] [file] ...

DESCRIPTION

Lex generates programs to be used in simple lexical analysis of text. The input *files* (standard input default) contain regular expressions to be searched for, and actions written in C to be executed when expressions are found.

A C source program, 'lex.yy.c' is generated, to be compiled thus:

```
cc lex.yy.c -ll
```

This program, when run, copies unrecognized portions of the input to the output, and executes the associated C action for each regular expression that is recognized.

The options have the following meanings.

- t Place the result on the standard output instead of in file "lex.yy.c".
- v Print a one-line summary of statistics of the generated analyzer.
- n Opposite of -v; -n is default.
- f "Faster" compilation: don't bother to pack the resulting tables; limited to small programs.

EXAMPLE

```
lex lexcommands
```

would draw *lex* instructions from the file *lexcommands*, and place the output in *lex.yy.c*

```
%%
[A-Z]putchar(yytext[0]+'a'-'A');
[ ]+$
[ ]+putchar(' ');
```

is an example of a *lex* program that would be put into a *lex* command file. This program converts upper case to lower, removes blanks at the end of lines, and replaces multiple blanks by single blanks.

SEE ALSO

yacc(1), sed(1)

M. E. Lesk and E. Schmidt, *LEX - Lexical Analyzer Generator*

NAME

`lint` - a C program verifier

SYNOPSIS

`lint` [`-abchnpuvx`] file ...

DESCRIPTION

Lint attempts to detect features of the C program *files* which are likely to be bugs, or non-portable, or wasteful. It also checks the type usage of the program more strictly than the compilers. Among the things which are currently found are unreachable statements, loops not entered at the top, automatic variables declared and not used, and logical expressions whose value is constant. Moreover, the usage of functions is checked to find functions which return values in some places and not in others, functions called with varying numbers of arguments, and functions whose values are not used.

By default, it is assumed that all the *files* are to be loaded together; they are checked for mutual compatibility. Function definitions for certain libraries are available to *lint*; these libraries are referred to by a conventional name, such as `'-lm'`, in the style of *ld(1)*. Arguments ending in *.ln* are also treated as library files. To create lint libraries, use the `-C` option:

```
lint -Cfoo files . . .
```

where *files* are the C sources of library *foo*. The result is a file *llib-lfoo.ln* in the correct library format suitable for linting programs using *foo*.

Any number of the options in the following list may be used. The `-D`, `-U`, and `-I` options of *cc(1)* are also recognized as separate arguments.

- p** Attempt to check portability to the *IBM* and *GCOS* dialects of C.
- h** Apply a number of heuristic tests to attempt to intuit bugs, improve style, and reduce waste.
- b** Report *break* statements that cannot be reached. (This is not the default because, unfortunately, most *lex* and many *yacc* outputs produce dozens of such comments.)
- v** Suppress complaints about unused arguments in functions.
- x** Report variables referred to by extern declarations, but never used.
- a** Report assignments of long values to int variables.
- c** Complain about casts which have questionable portability.
- u** Do not complain about functions and variables used and not defined, or defined and not used (this is suitable for running *lint* on a subset of files out of a larger program).
- n** Do not check compatibility against the standard library.
- z** Do not complain about structures that are never defined (e.g. using a structure pointer without knowing its contents.).

Exit(2) and other functions which do not return are not understood; this causes various lies.

Certain conventional comments in the C source will change the behavior of *lint*:

```
/*NOTREACHED*/
```

at appropriate points stops comments about unreachable code.

/*VARARGSn*/

suppresses the usual checking for variable numbers of arguments in the following function declaration. The data types of the first *n* arguments are checked; a missing *n* is taken to be 0.

/*NOSTRICT*/

shuts off strict type checking in the next expression.

/*ARGSUSED*/

turns on the **-v** option for the next function.

/*LINTLIBRARY*/

at the beginning of a file shuts off complaints about unused functions in this file.

AUTHOR

S.C. Johnson. Lint library construction implemented by Edward Wang.

FILES

/usr/lib/lint/lint[12]	programs
/usr/lib/lint/l-lib-lc.ln	declarations for standard functions
/usr/lib/lint/l-lib-lc	human readable version of above
/usr/lib/lint/l-lib-port.ln	declarations for portable functions
/usr/lib/lint/l-lib-port	human readable . . .
lib-l*.ln	library created with -C

SEE ALSO

cc(1)

S. C. Johnson, *Lint, a C Program Checker*

BUGS

There are some things you just **can't** get lint to shut up about.

NAME

`lisp` - lisp interpreter

SYNOPSIS

`lisp`

DESCRIPTION

Lisp is a lisp interpreter for a dialect which closely resembles MIT's MACLISP. This lisp, known as FRANZ LISP, features an I/O facility which allows the user to change the input and output syntax, add macro characters, and maintain compatibility with upper-case only lisp systems; infinite precision integer arithmetic, and an error facility which allows the user to trap system errors in many different ways. Interpreted functions may be mixed with code compiled by *liszt*(1) and both may be debugged using the "Joseph Lister" trace package. A *lisp* containing compiled and interpreted code may be dumped into a file for later use.

There are too many functions to list here; one should refer to the manuals listed below.

AUTHORS

An early version was written by Jeff Levinsky, Mike Curry, and John Breedlove. Keith Sklower wrote and is maintaining the current version, with the assistance of John Foderaro. The garbage collector was implemented by Bill Rowan.

FILES

<code>/usr/lib/lisp/trace.l</code>	Joseph Lister trace package
<code>/usr/lib/lisp/toplevel.l</code>	top level read-eval-print loop

SEE ALSO

`liszt`(1), `lxref`(1)
'FRANZ LISP Manual, Version 1' by John K. Foderaro
MACLISP Manual

BUGS

The error system is in a state of flux and not all error messages are as informative as they could be.

NAME

`liszt` - compile a Franz Lisp program

SYNOPSIS

`liszt` [`-mpqruxwCQST`] [`-e form`] [`-o objfile`] [`name`]

DESCRIPTION

Liszt takes a file whose name ends in '.l' and compiles the FRANZ LISP code there leaving an object program on the file whose name is that of the source with '.o' substituted for '.l'.

The following options are interpreted by *liszt*.

- `-e` Evaluate the given form before compilation begins.
- `-m` Compile a MACLISP file, by changing the readtable to conform to MACLISP syntax and including a macro-defined compatibility package.
- `-o` Put the object code in the specified file, rather than the default '.o' file.
- `-p` places profiling code at the beginning of each non-local function. If the lisp system is also created with profiling in it, this allows function calling frequency to be determined (see *prof*(1).)
- `-q` Only print warning and error messages. Compilation statistics and notes on correct but unusual constructs will not be printed.
- `-r` place bootstrap code at the beginning of the object file, which when the object file is executed will cause a lisp system to be invoked and the object file fasl'ed in.
- `-u` Compile a UCI-lispfile, by changing the readtable to conform to UCI-Lisp syntax and including a macro-defined compatibility package.
- `-w` Suppress warning diagnostics.
- `-x` Create a lisp cross reference file with the same name as the source file but with '.x' appended. The program *lxref*(1) reads this file and creates a human readable cross reference listing.
- `-C` put comments in the assembler output of the compiler. Useful for debugging the compiler.
- `-Q` Print compilation statistics and warn of strange constructs. This is the default.
- `-S` Compile the named program and leave the assembler-language output on the corresponding file suffixed '.s'. This will also prevent the assembler language file from being assembled.
- `-T` send the assembler output to standard output.

If no source file is specified, then the compiler will run interactively. You will find yourself talking to the *lisp*(1) top-level command interpreter. You can compile a file by using the function *liszt* (an *nlambda*) with the same arguments as you use on the command line. For example to compile 'foo', a MACLISP file, you would use:

```
(liszt -m foo)
```

Note that *liszt* supplies the ".l" extension for you.

FILES

<code>/usr/lib/lisp/machacks.l</code>	MACLISP compatibility package
<code>/usr/lib/lisp/syscall.l</code>	macro definitions of Unix system calls
<code>/usr/lib/lisp/ucifnc.l</code>	UCI Lisp compatibility package

AUTHOR

John Foderaro

SEE ALSO

lisp(1), lxref(1)

NAME

ln - make links

SYNOPSIS

ln [-s] *name1* [*name2*]
ln *name* ... *directory*

DESCRIPTION

A link is a directory entry referring to a file; the same file (together with its size, all its protection information, etc.) may have several links to it. There are two kinds of links: hard links and symbolic links.

By default *ln* makes hard links. A hard link to a file is indistinguishable from the original directory entry; any changes to a file are effective independent of the name used to reference the file. Hard links may not span file systems and may not refer to directories.

The *-s* option causes *ln* to create symbolic links. A symbolic link contains the name of the file to which it is linked. The referenced file is used when an *open*(2) operation is performed on the link. A *stat*(2) on a symbolic link will return the linked-to file; an *lstat*(2) must be done to obtain information about the link. The *readlink*(2) call may be used to read the contents of a symbolic link. Symbolic links may span file systems and may refer to directories.

Given one or two arguments, *ln* creates a link to an existing file *name1*. If *name2* is given, the link has that name; *name2* may also be a directory in which to place the link; otherwise it is placed in the current directory. If only the directory is specified, the link will be made to the last component of *name1*.

Given more than two arguments, *ln* makes links to all the named files in the named directory. The links made will have the same name as the files being linked to.

SEE ALSO

rm(1), *cp*(1), *mv*(1), *link*(2), *readlink*(2), *stat*(2), *symlink*(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, he has no other recourse but to login elsewhere and kill the lock process.

AUTHOR

Kurt Shoens

BUGS

Should timeout after 15 minutes.

NAME

login - sign on

SYNOPSIS

login [username]

DESCRIPTION

The *login* command is used when a user initially signs on, or it may be used at any time to change from one user to another. The latter case is the one summarized above and described here. See "How to Get Started" for how to dial up initially.

If *login* is invoked without an argument, it asks for a user name, and, if appropriate, a password. Echoing is turned off (if possible) during the typing of the password, so it will not appear on the written record of the session.

After a successful login, accounting files are updated and the user is informed of the existence of mail, and the message of the day is printed, as is the time he last logged in (unless he has a ".hushlogin" file in his home directory - this is mostly used to make life easier for non-human users, such as *uucp*).

Login initializes the user and group IDs and the working directory, then executes a command interpreter (usually *sh*(1)) according to specifications found in a password file. Argument 0 of the command interpreter is "-sh", or more generally the name of the command interpreter with a leading dash ("-") prepended.

Login also initializes the environment *environ*(7) with information specifying home directory, command interpreter, terminal type (if available) and user name.

If the file */etc/nologin* exists *login* prints its contents on the user's terminal and exits. This is used by *shutdown*(8) to stop users logging in when the system is about to go down.

Login is recognized by *sh*(1) and *csh*(1) and executed directly (without forking).

FILES

<i>/etc/utmp</i>	accounting
<i>/usr/adm/wtmp</i>	accounting
<i>/usr/spool/mail/*</i>	mail
<i>/etc/motd</i>	message-of-the-day
<i>/etc/passwd</i>	password file
<i>/etc/nologin</i>	stops logins
<i>.hushlogin</i>	makes login quieter
<i>/etc/securetty</i>	lists ttys that root may log in on

SEE ALSO

init(8), *getty*(8), *mail*(1), *passwd*(1), *passwd*(5), *environ*(7), *shutdown*(8)

DIAGNOSTICS

"Login incorrect," if the name or the password is bad.

"No Shell", "cannot open password file", "no directory": consult a programming counselor.

BUGS

An undocumented option, *-r* is used by the remote login server, *rlogind*(8C) to force *login* to enter into an initial connection protocol.

NAME

look - find lines in a sorted list

SYNOPSIS

look [**-df**] string [file]

DESCRIPTION

Look consults a sorted *file* and prints all lines that begin with *string*. It uses binary search.

The options **d** and **f** affect comparisons as in *sort(1)*:

d 'Dictionary' order: only letters, digits, tabs and blanks participate in comparisons.

f Fold. Upper case letters compare equal to lower case.

If no *file* is specified, */usr/dict/words* is assumed with collating sequence **-df**.

FILES

/usr/dict/words

SEE ALSO

sort(1), *grep(1)*

NAME

indxib, lookbib – build inverted index for a bibliography, find references in a bibliography

SYNOPSIS

indxib database ...
lookbib database

DESCRIPTION

Indxib makes an inverted index to the named *databases* (or files) for use by *lookbib(1)* and *refer(1)*. These files contain bibliographic references (or other kinds of information) separated by blank lines.

A bibliographic reference is a set of lines, constituting fields of bibliographic information. Each field starts on a line beginning with a “%”, followed by a key-letter, then a blank, and finally the contents of the field, which may continue until the next line starting with “%”.

Indxib is a shell script that calls */usr/lib/refer/mkey* and */usr/lib/refer/inv*. The first program, *mkey*, truncates words to 6 characters, and maps upper case to lower case. It also discards words shorter than 3 characters, words among the 100 most common English words, and numbers (dates) < 1900 or > 2000. These parameters can be changed; see page 4 of the *Refer* document by Mike Lesk. The second program, *inv*, creates an entry file (.ia), a posting file (.ib), and a tag file (.ic), all in the working directory.

Lookbib uses an inverted index made by *indxib* to find sets of bibliographic references. It reads keywords typed after the “>” prompt on the terminal, and retrieves records containing all these keywords. If nothing matches, nothing is returned except another “>” prompt.

It is possible to search multiple databases, as long as they have a common index made by *indxib*. In that case, only the first argument given to *indxib* is specified to *lookbib*.

If *lookbib* does not find the index files (the .i[abc] files), it looks for a reference file with the same name as the argument, without the suffixes. It creates a file with a '.ig' suffix, suitable for use with *fgrep*. It then uses this *fgrep* file to find references. This method is simpler to use, but the .ig file is slower to use than the .i[abc] files, and does not allow the use of multiple reference files.

FILES

x.ia, *x.ib*, *x.ic*, where *x* is the first argument, or if these are not present, then *x.ig*, *x*

SEE ALSO

refer(1), *addbib(1)*, *sortbib(1)*, *roffb(1)*, *lookbib(1)*

BUGS

Probably all dates should be indexed, since many disciplines refer to literature written in the 1800s or earlier.

NAME

lorder - find ordering relation for an object library

SYNOPSIS

lorder file ...

DESCRIPTION

The input is one or more object or library archive (see *ar(1)*) files. 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. The output may be processed by *tsort(1)* to find an ordering of a library suitable for one-pass access by *ld(1)*.

This brash one-liner intends to build a new library from existing '.o' files.

```
ar cr library `lorder *.o | tsort`
```

The need for lorder may be vitiated by use of *ranlib(1)*, which converts an ordered archive into a randomly accessed library.

FILES

*symref, *symdef
nm(1), sed(1), sort(1), join(1)

SEE ALSO

tsort(1), ld(1), ar(1), ranlib(1)

BUGS

The names of object files, in and out of libraries, must end with '.o'; nonsense results otherwise.

NAME

lpq – spool queue examination program

SYNOPSIS

lpq [+[*n*]] [-l] [-Pprinter] [job # ...] [user ...]

DESCRIPTION

lpq examines the spooling area used by *lpd*(8) for printing files on the line printer, and reports the status of the specified jobs or all jobs associated with a user. *lpq* invoked without any arguments reports on any jobs currently in the queue. A **-P** flag may be used to specify a particular printer, otherwise the default line printer is used (or the value of the **PRINTER** variable in the environment). If a **+** argument is supplied, *lpq* displays the spool queue until it empties. Supplying a number immediately after the **+** sign indicates that *lpq* should sleep *n* seconds in between scans of the queue. All other arguments supplied are interpreted as user names or job numbers to filter out only those jobs of interest.

For each job submitted (i.e. invocation of *lpr*(1)) *lpq* reports the user's name, current rank in the queue, the names of files comprising the job, the job identifier (a number which may be supplied to *lprm*(1) for removing a specific job), and the total size in bytes. The **-l** option causes information about each of the files comprising the job to be printed. Normally, only as much information as will fit on one line is displayed. Job ordering is dependent on the algorithm used to scan the spooling directory and is supposed to be FIFO (First in First Out). File names comprising a job may be unavailable (when *lpr*(1) is used as a sink in a pipeline) in which case the file is indicated as "(standard input)".

If *lpq* warns that there is no daemon present (i.e. due to some malfunction), the *lpc*(8) command can be used to restart the printer daemon.

FILES

<i>/etc/termcap</i>	for manipulating the screen for repeated display
<i>/etc/printcap</i>	to determine printer characteristics
<i>/usr/spool/*</i>	the spooling directory, as determined from <i>printcap</i>
<i>/usr/spool/*/cf*</i>	control files specifying jobs
<i>/usr/spool/*/lock</i>	the lock file to obtain the currently active job

SEE ALSO

lpr(1), *lprm*(1), *lpc*(8), *lpd*(8)

BUGS

Due to the dynamic nature of the information in the spooling directory *lpq* may report unreliably. Output formatting is sensitive to the line length of the terminal; this can result in widely spaced columns.

DIAGNOSTICS

Unable to open various files. The lock file being malformed. Garbage files when there is no daemon active, but files in the spooling directory.

NAME

`lpr` - off line print

SYNOPSIS

```
lpr [ -Pprinter ] [ -#num ] [ -C class ] [ -J job ]
    [ -T title ] [ -i [ numcols ] ] [ -1234 font ] [ -wnum ]
    [ -pltndgvcfrms ] [ name ... ]
```

DESCRIPTION

`Lpr` uses a spooling daemon to print the named files when facilities become available. If no names appear, the standard input is assumed. The `-P` option may be used to force output to a specific printer. Normally, the default printer is used (site dependent), or the value of the environment variable `PRINTER` is used.

The following single letter options are used to notify the line printer spooler that the files are not standard text files. The spooling daemon will use the appropriate filters to print the data accordingly.

- `-p` Use `pr(1)` to format the files (equivalent to `print`).
- `-l` Use a filter which allows control characters to be printed and suppresses page breaks.
- `-t` The files are assumed to contain data from `troff(1)` (cat phototypesetter commands).
- `-n` The files are assumed to contain data from `ditroff` (device independent troff).
- `-d` The files are assumed to contain data from `tex(1)` (DVI format from Stanford).
- `-g` The files are assumed to contain standard plot data as produced by the `plot(3X)` routines (see also `plot(1G)` for the filters used by the printer spooler).
- `-v` The files are assumed to contain a raster image for devices like the Benson Varian.
- `-c` The files are assumed to contain data produced by `cifplot(1)`.
- `-f` Use a filter which interprets the first character of each line as a standard FORTRAN carriage control character.

The remaining single letter options have the following meaning.

- `-r` Remove the file upon completion of spooling or upon completion of printing (with the `-s` option).
- `-m` Send mail upon completion.
- `-h` Suppress the printing of the burst page.
- `-s` Use symbolic links. Usually files are copied to the spool directory.

The `-C` option takes the following argument as a job classification for use on the burst page. For example,

```
lpr -C EECS foo.c
```

causes the system name (the name returned by `hostname(1)`) to be replaced on the burst page by `EECS`, and the file `foo.c` to be printed.

The `-J` option takes the following argument as the job name to print on the burst page. Normally, the first file's name is used.

The `-T` option uses the next argument as the title used by `pr(1)` instead of the file name.

To get multiple copies of output, use the `-#num` option, where *num* is the number of copies desired of each file named. For example,

```
lpr -#3 foo.c bar.c more.c
```

would result in 3 copies of the file `foo.c`, followed by 3 copies of the file `bar.c`, etc. On the other hand,

```
cat foo.c bar.c more.c | lpr -#3
```

will give three copies of the concatenation of the files.

The `-i` option causes the output to be indented. If the next argument is numeric, it is used as the number of blanks to be printed before each line; otherwise, 8 characters are printed.

The `-w` option takes the immediately following number to be the page width for *pr*.

The `-s` option will use *symlink(2)* to link data files rather than trying to copy them so large files can be printed. This means the files should not be modified or removed until they have been printed.

The option `-1234` Specifies a font to be mounted on font position *i*. The daemon will construct a *.railmag* file referencing `/usr/lib/vfont/name.size`.

FILES

<code>/etc/passwd</code>	personal identification
<code>/etc/printcap</code>	printer capabilities data base
<code>/usr/lib/lpd*</code>	line printer daemons
<code>/usr/spool/*</code>	directories used for spooling
<code>/usr/spool/*/cf*</code>	daemon control files
<code>/usr/spool/*/df*</code>	data files specified in "cf" files
<code>/usr/spool/*/tf*</code>	temporary copies of "cf" files

SEE ALSO

`lpq(1)`, `lprm(1)`, `pr(1)`, `symlink(2)`, `printcap(5)`, `lpc(8)`, `lpd(8)`, `cleanlpd(8)`

DIAGNOSTICS

If you try to spool too large a file, it will be truncated. *Lpr* will object to printing binary files. If a user other than root prints a file and spooling is disabled, *lpr* will print a message saying so and will not put jobs in the queue. If a connection to *lpd* on the local machine cannot be made, *lpr* will say that the daemon cannot be started. Diagnostics may be printed in the daemon's log file regarding missing spool files by *lpd*. If *lpd* is not configured properly or has become out of sync with its lock files, `cleanlpd(8)` may be used to restart *lpd* in a new environment.

BUGS

Fonts for *troff* and *tex* reside on the host with the printer. It is currently not possible to use local font libraries.

NAME

lprm - remove jobs from the line printer spooling queue

SYNOPSIS

lprm [**-P***printer*] [**-**] [*job # ...*] [*user ...*]

DESCRIPTION

Lprm will remove a job, or jobs, from a printer's spool queue. Since the spooling directory is protected from users, using *lprm* is normally the only method by which a user may remove a job.

Lprm without any arguments will delete the currently active job if it is owned by the user who invoked *lprm*.

If the **-** flag is specified, *lprm* will remove all jobs which a user owns. If the super-user employs this flag, the spool queue will be emptied entirely. The owner is determined by the user's login name and host name on the machine where the *lpr* command was invoked.

Specifying a user's name, or list of user names, will cause *lprm* to attempt to remove any jobs queued belonging to that user (or users). This form of invoking *lprm* is useful only to the super-user.

A user may dequeue an individual job by specifying its job number. This number may be obtained from the *lpq*(1) program, e.g.

```
% lpq -l
```

```
1st: ken                [job #013ucbarpa]
      (standard input)  100 bytes
% lprm 13
```

Lprm will announce the names of any files it removes and is silent if there are no jobs in the queue which match the request list.

Lprm will kill off an active daemon, if necessary, before removing any spooling files. If a daemon is killed, a new one is automatically restarted upon completion of file removals.

The **-P** option may be used to specify the queue associated with a specific printer (otherwise the default printer, or the value of the **PRINTER** variable in the environment is used).

FILES

/etc/printcap	printer characteristics file
/usr/spool/*	spooling directories
/usr/spool/*/lock	lock file used to obtain the pid of the current daemon and the job number of the currently active job

SEE ALSO

lpr(1), *lpq*(1), *lpd*(8)

DIAGNOSTICS

"Permission denied" if the user tries to remove files other than his own.

BUGS

Since there are race conditions possible in the update of the lock file, the currently active job may be incorrectly identified.

NAME

ls – list contents of directory

SYNOPSIS

ls [**-acdfgilqrstu1ACLFR**] name ...

DESCRIPTION

For each directory argument, *ls* lists the contents of the directory; for each file argument, *ls* repeats its name and any other information requested. By default, the output is sorted alphabetically. When no argument is given, the current directory is listed. When several arguments are given, the arguments are first sorted appropriately, but file arguments are processed before directories and their contents.

There are a large number of options:

- l** List in long format, giving mode, number of links, owner, size in bytes, and time of last modification for each file. (See below.) If the file is a special file the size field will instead contain the major and minor device numbers. If the file is a symbolic link the pathname of the linked-to file is printed preceded by “->”.
- g** Include the group ownership of the file in a long output.
- t** Sort by time modified (latest first) instead of by name.
- a** List all entries; in the absence of this option, entries whose names begin with a period (.) are *not* listed.
- s** Give size in kilobytes of each file.
- d** If argument is a directory, list only its name; often used with **-l** to get the status of a directory.
- L** If argument is a symbolic link, list the file or directory the link references rather than the link itself.
- r** Reverse the order of sort to get reverse alphabetic or oldest first as appropriate.
- u** Use time of last access instead of last modification for sorting (with the **-t** option) and/or printing (with the **-l** option).
- c** Use time of file creation for sorting or printing.
- i** For each file, print the i-number in the first column of the report.
- f** Force each argument to be interpreted as a directory and list the name found in each slot. This option turns off **-l**, **-t**, **-s**, and **-r**, and turns on **-a**; the order is the order in which entries appear in the directory.
- F** cause directories to be marked with a trailing ‘/’, sockets with a trailing ‘=’, symbolic links with a trailing ‘@’, and executable files with a trailing ‘*’.
- R** recursively list subdirectories encountered.
- l** force one entry per line output format; this is the default when output is not to a terminal.
- C** force multi-column output; this is the default when output is to a terminal.
- q** force printing of non-graphic characters in file names as the character ‘?’; this is the default when output is to a terminal.

The mode printed under the `-l` option contains 11 characters which are interpreted as follows: the first character is

- d** if the entry is a directory;
- b** if the entry is a block-type special file;
- c** if the entry is a character-type special file;
- l** if the entry is a symbolic link;
- s** if the entry is a socket, or
- if the entry is a plain file.

The next 9 characters are interpreted as three sets of three bits each. The first set refers to owner permissions; the next to permissions to others in the same user-group; and the last to all others. Within each set the three characters indicate permission respectively to read, to write, or to execute the file as a program. For a directory, 'execute' permission is interpreted to mean permission to search the directory. The permissions are indicated as follows:

- r** if the file is readable;
- w** if the file is writable;
- x** if the file is executable;
- if the indicated permission is not granted.

The group-execute permission character is given as **s** if the file has the set-group-id bit set; likewise the user-execute permission character is given as **S** if the file has the set-user-id bit set.

The last character of the mode (normally 'x' or '-') is **t** if the 1000 bit of the mode is on. See `chmod(1)` for the meaning of this mode.

When the sizes of the files in a directory are listed, a total count of blocks, including indirect blocks is printed.

FILES

`/etc/passwd` to get user id's for '`ls -l`'.
`/etc/group` to get group id's for '`ls -g`'.

BUGS

Newline and tab are considered printing characters in file names.

The output device is assumed to be 80 columns wide.

The option setting based on whether the output is a teletype is undesirable as "`ls -s`" is much different than "`ls -s |lpr`". On the other hand, not doing this setting would make old shell scripts which used `ls` almost certain losers.

NAME

`lxref` - lisp cross reference program

SYNOPSIS

`lxref` [`-N`] `xref-file` ... [`-a` `source-file` ...]

DESCRIPTION

Lxref reads cross reference file(s) written by the lisp compiler *liszt* and prints a cross reference listing on the standard output. *Liszt* will create a cross reference file during compilation when it is given the `-x` switch. Cross reference files usually end in `.x` and consequently *lxref* will append a `.x` to the file names given if necessary. The first option to *lxref* is a decimal integer, `N`, which sets the *ignorelevel*. If a function is called more than *ignorelevel* times, the cross reference listing will just print the number of calls instead of listing each one of them. The default for *ignorelevel* is 50.

The `-a` option causes *lxref* to put limited cross reference information in the sources named. *lxref* will scan the source and when it comes across a definition of a function (that is a line beginning with `(def`) it will precede that line with a list of the functions which call this function, written as a comment preceded by `;;`. All existing lines beginning with `;;` will be removed from the file. If the source file contains a line beginning `;;-` then this will disable this annotation process from this point on until a `;;+` is seen (however, lines beginning with `;;` will continue to be deleted). After the annotation is done, the original file `'foo.l'` is renamed to `"#.foo.l"` and the new file with annotation is named `'foo.l'`

AUTHOR

John Foderaro

SEE ALSO

`lisp(1)`, `liszt(1)`

NAME

m4 – macro processor

SYNOPSIS

m4 [files]

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 arguments, or if an argument is '-', the standard input is read. The processed text is written on the standard output.

Macro calls have the form

```
name(arg1,arg2, . . . , argn)
```

The '(' must immediately follow the name of the macro. If a defined macro name is not followed by a '(', it is deemed to have no arguments. Leading unquoted blanks, tabs, and newlines are ignored while collecting arguments. Potential macro names consist of alphabetic letters, digits, and underscore '_', where the first character is not a digit.

Left and right single quotes (`) 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. 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.

M4 makes available the following built-in macros. They may be redefined, but once this is done the original meaning is lost. Their values are null unless otherwise stated.

define The second argument is installed as the value of the macro whose name is the first argument. Each occurrence of \$*n* in the replacement text, where *n* is a digit, is replaced by the *n*-th argument. Argument 0 is the name of the macro; missing arguments are replaced by the null string.

undefine removes the definition of the macro named in its argument.

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 UNIX versions of *m4*.

changequote

Change quote characters to the first and second arguments. *Changequote* without arguments restores the original values (i.e., ` `).

divert *M4* maintains 10 output streams, numbered 0-9. The final output is the concatenation of the streams in numerical order; initially stream 0 is the current stream. The *divert* macro changes the current output stream to its (digit-string) argument. Output diverted to a stream other than 0 through 9 is discarded.

undivert causes immediate output of text from diversions named as arguments, or all diversions if no argument. Text may be undiverted into another diversion. Undiverting discards the diverted text.

divnum returns the value of the current output stream.

dnl reads and discards characters up to and including the next newline.

- 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.
- 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.
- eval** evaluates its argument as an arithmetic expression, using 32-bit arithmetic. Operators include +, -, *, /, %, ^ (exponentiation); relationals; parentheses.
- len** returns the number of characters in its argument.
- index** returns the position in its first argument where the second argument begins (zero origin), or -1 if the second argument does not occur.
- substr** returns a substring of its first argument. The second argument is a zero origin number selecting the first character; the third argument indicates the length of the substring. A missing third argument is taken to be large enough to extend to the end of the first string.
- translit** transliterates the characters in its first argument from the set given by the second argument to the set given by the third. No abbreviations are permitted.
- include** returns the contents of the file named in the argument.
- sinclude** is identical to *include*, except that it says nothing if the file is inaccessible.
- syscmd** executes the UNIX command given in the first argument. No value is returned.
- maketemp** fills in a string of XXXXX in its argument with the current process id.
- errprint** prints its argument on the diagnostic output file.
- dumpdef** prints current names and definitions, for the named items, or for all if no arguments are given.

SEE ALSO

B. W. Kernighan and D. M. Ritchie, *The M4 Macro Processor*

NAME

mail – send and receive mail

SYNOPSIS

```
mail [-v] [-i] [-n] [-s subject] [ user ... ]
mail [-v] [-i] [-n] -f [ name ]
mail [-v] [-i] [-n] -u user
```

INTRODUCTION

Mail is a intelligent mail processing system, which has a command syntax reminiscent of *ed* with lines replaced by messages.

The **-v** flag puts mail into verbose mode; the details of delivery are displayed on the users terminal. The **-i** flag causes tty interrupt signals to be ignored. This is particularly useful when using *mail* on noisy phone lines. The **-n** flag inhibits the reading of `/usr/lib/Mail.rc`.

Sending mail. To send a message to one or more other people, *mail* can be invoked with arguments which are the names of people to send to. You are then expected to type in your message, followed by an EOT (control-D) at the beginning of a line. A subject may be specified on the command line by using the **-s** flag. (Only the first argument after the **-s** flag is used as a subject; be careful to quote subjects containing spaces.) The section below, labeled *Replying to or originating mail*, describes some features of *mail* available to help you compose your letter.

Reading mail. In normal usage *mail* is given no arguments and checks your mail out of the post office, then prints out a one line header of each message there. The current message is initially the first message (numbered 1) and can be printed using the **print** command (which can be abbreviated **p**). You can move among the messages much as you move between lines in *ed*, with the commands '+' and '-' moving backwards and forwards, and simple numbers.

Disposing of mail. After examining a message you can **delete** (**d**) the message or **reply** (**r**) to it. Deletion causes the *mail* program to forget about the message. This is not irreversible; the message can be **undeleted** (**u**) by giving its number, or the *mail* session can be aborted by giving the **exit** (**x**) command. Deleted messages will, however, usually disappear never to be seen again.

Specifying messages. Commands such as **print** and **delete** can be given a list of message numbers as arguments to apply to a number of messages at once. Thus "delete 1 2" deletes messages 1 and 2, while "delete 1-5" deletes messages 1 through 5. The special name "*" addresses all messages, and "\$" addresses the last message; thus the command **top** which prints the first few lines of a message could be used in "top *" to print the first few lines of all messages.

Replying to or originating mail. You can use the **reply** command to set up a response to a message, sending it back to the person who it was from. Text you then type in, up to an end-of-file, defines the contents of the message. While you are composing a message, *mail* treats lines beginning with the character '~' specially. For instance, typing "~m" (alone on a line) will place a copy of the current message into the response right shifting it by a tabstop. Other escapes will set up subject fields, add and delete recipients to the message and allow you to escape to an editor to revise the message or to a shell to run some commands. (These options are given in the summary below.)

Ending a mail processing session. You can end a *mail* session with the **quit** (**q**) command. Messages which have been examined go to your *mbox* file unless they have been deleted in which case they are discarded. Unexamined messages go back to the post office. The **-f** option causes *mail* to read in the contents of your *mbox* (or the specified file) for processing; when you **quit**, *mail* writes undeleted messages back to this file. The **-u** flag is a short way

of doing "mail -f /usr/spool/mail/user".

Personal and systemwide distribution lists. It is also possible to create a personal distribution lists so that, for instance, you can send mail to "cohorts" and have it go to a group of people. Such lists can be defined by placing a line like

```
alias cohorts bill ozalp jkf mark kridle@ucbcory
```

in the file `.mailrc` in your home directory. The current list of such aliases can be displayed with the **alias** (**a**) command in *mail*. System wide distribution lists can be created by editing `/usr/lib/aliases`, see *aliases*(5) and *sendmail*(8); these are kept in a different syntax. In mail you send, personal aliases will be expanded in mail sent to others so that they will be able to **reply** to the recipients. System wide *aliases* are not expanded when the mail is sent, but any reply returned to the machine will have the system wide alias expanded as all mail goes through *sendmail*.

Network mail (ARPA, UUCP, Berknet) See *mailaddr*(7) for a description of network addresses.

Mail has a number of options which can be set in the `.mailrc` file to alter its behavior; thus "set askcc" enables the "askcc" feature. (These options are summarized below.)

SUMMARY

(Adapted from the 'Mail Reference Manual')

Each command is typed on a line by itself, and may take arguments following the command word. The command need not be typed in its entirety - the first command which matches the typed prefix is used. For commands which take message lists as arguments, if no message list is given, then the next message forward which satisfies the command's requirements is used. If there are no messages forward of the current message, the search proceeds backwards, and if there are no good messages at all, *mail* types "No applicable messages" and aborts the command.

- Goes to the previous message and prints it out. If given a numeric argument *n*, goes to the *n*-th previous message and prints it.
- ? Prints a brief summary of commands.
- ! Executes the UNIX shell command which follows.
- Print** (P) Like **print** but also prints out ignored header fields. See also **print** and **ignore**.
- Reply** (R) Reply to originator. Does not reply to other recipients of the original message.
- Type** (T) Identical to the **Print** command.
- alias** (a) With no arguments, prints out all currently-defined aliases. With one argument, prints out that alias. With more than one argument, creates a new or changes an old alias.
- alternates** (alt) The **alternates** command is useful if you have accounts on several machines. It can be used to inform *mail* that the listed addresses are really you. When you **reply** to messages, *mail* will not send a copy of the message to any of the addresses listed on the *alternates* list. If the **alternates** command is given with no argument, the current set of alternate names is displayed.
- chdir** (c) Changes the user's working directory to that specified, if given. If no directory is given, then changes to the user's login directory.

- copy** (co) The **copy** command does the same thing that **save** does, except that it does not mark the messages it is used on for deletion when you quit.
- delete** (d) Takes a list of messages as argument and marks them all as deleted. Deleted messages will not be saved in *mbox*, nor will they be available for most other commands.
- dp** (also dt) Deletes the current message and prints the next message. If there is no next message, *mail* says "at EOF."
- edit** (e) Takes a list of messages and points the text editor at each one in turn. On return from the editor, the message is read back in.
- exit** (ex or x) Effects an immediate return to the Shell without modifying the user's system mailbox, his *mbox* file, or his edit file in *-f*.
- file** (fi) The same as **folder**.
- folders** List the names of the folders in your folder directory.
- folder** (fo) The **folder** command switches to a new mail file or folder. With no arguments, it tells you which file you are currently reading. If you give it an argument, it will write out changes (such as deletions) you have made in the current file and read in the new file. Some special conventions are recognized for the name. # means the previous file, % means your system mailbox, %user means user's system mailbox, & means your ~/mbox file, and +folder means a file in your folder directory.
- from** (f) Takes a list of messages and prints their message headers.
- headers** (h) Lists the current range of headers, which is an 18 message group. If a "+" argument is given, then the next 18 message group is printed, and if a "-" argument is given, the previous 18 message group is printed.
- help** A synonym for ?
- hold** (ho, also **preserve**) Takes a message list and marks each message therein to be saved in the user's system mailbox instead of in *mbox*. Does not override the **delete** command.
- ignore** Add the list of header fields named to the *ignored list*. Header fields in the ignore list are not printed on your terminal when you print a message. This command is very handy for suppression of certain machine-generated header fields. The **Type** and **Print** commands can be used to print a message in its entirety, including ignored fields. If **ignore** is executed with no arguments, it lists the current set of ignored fields.
- mail** (m) Takes as argument login names and distribution group names and sends mail to those people.
- mbox** Indicate that a list of messages be sent to *mbox* in your home directory when you quit. This is the default action for messages if you do *not* have the *hold* option set.
- next** (n like + or CR) Goes to the next message in sequence and types it. With an argument list, types the next matching message.
- preserve** (pre) A synonym for **hold**.
- print** (p) Takes a message list and types out each message on the user's terminal.
- quit** (q) Terminates the session, saving all undeleted, unsaved messages in the user's *mbox* file in his login directory, preserving all messages marked with **hold** or **preserve** or never referenced in his system mailbox, and removing all other

messages from his system mailbox. If new mail has arrived during the session, the message "You have new mail" is given. If given while editing a mailbox file with the `-f` flag, then the edit file is rewritten. A return to the Shell is effected, unless the rewrite of edit file fails, in which case the user can escape with the `exit` command.

- reply** (r) Takes a message list and sends mail to the sender and all recipients of the specified message. The default message must not be deleted.
- respond** A synonym for `reply`.
- save** (s) Takes a message list and a filename and appends each message in turn to the end of the file. The filename in quotes, followed by the line count and character count is echoed on the user's terminal.
- set** (se) With no arguments, prints all variable values. Otherwise, sets option. Arguments are of the form "option=value" or "option."
- shell** (sh) Invokes an interactive version of the shell.
- size** Takes a message list and prints out the size in characters of each message.
- source** (so) The `source` command reads *mail* commands from a file.
- top** Takes a message list and prints the top few lines of each. The number of lines printed is controlled by the variable `toplines` and defaults to five.
- type** (t) A synonym for `print`.
- unalias** Takes a list of names defined by `alias` commands and discards the remembered groups of users. The group names no longer have any significance.
- undelete** (u) Takes a message list and marks each one as *not* being deleted.
- unset** Takes a list of option names and discards their remembered values; the inverse of `set`.
- visual** (v) Takes a message list and invokes the display editor on each message.
- write** (w) A synonym for `save`.
- xit** (x) A synonym for `exit`.
- z** *Mail* presents message headers in windowfuls as described under the `headers` command. You can move *mail*'s attention forward to the next window with the `z` command. Also, you can move to the previous window by using `z-`.

Here is a summary of the tilde escapes, which are used when composing messages to perform special functions. Tilde escapes are only recognized at the beginning of lines. The name "tilde escape" is somewhat of a misnomer since the actual escape character can be set by the option `escape`.

- `~!command` Execute the indicated shell command, then return to the message.
- `~c name ...` Add the given names to the list of carbon copy recipients.
- `~d` Read the file "dead.letter" from your home directory into the message.
- `~e` Invoke the text editor on the message collected so far. After the editing session is finished, you may continue appending text to the message.
- `~f messages` Read the named messages into the message being sent. If no messages are specified, read in the current message.
- `~h` Edit the message header fields by typing each one in turn and allowing the user to append text to the end or modify the field by using the current terminal erase and kill characters.

- `~m messages` Read the named messages into the message being sent, shifted right one tab. If no messages are specified, read the current message.
- `~p` Print out the message collected so far, prefaced by the message header fields.
- `~q` Abort the message being sent, copying the message to "dead.letter" in your home directory if `save` is set.
- `~r filename` Read the named file into the message.
- `~s string` Cause the named string to become the current subject field.
- `~t name ...` Add the given names to the direct recipient list.
- `~v` Invoke an alternate editor (defined by the `VISUAL` option) on the message collected so far. Usually, the alternate editor will be a screen editor. After you quit the editor, you may resume appending text to the end of your message.
- `~w filename` Write the message onto the named file.
- `~|command` Pipe the message through the command as a filter. If the command gives no output or terminates abnormally, retain the original text of the message. The command `fmt(1)` is often used as `command` to rejustify the message.
- `~string` Insert the string of text in the message prefaced by a single `~`. If you have changed the escape character, then you should double that character in order to send it.

Options are controlled via the `set` and `unset` commands. Options may be either binary, in which case it is only significant to see whether they are set or not, or string, in which case the actual value is of interest. The binary options include the following:

- `append` Causes messages saved in `mbox` to be appended to the end rather than prepended. (This is set in `/usr/lib/Mail.rc` on version 7 systems.)
- `ask` Causes `mail` to prompt you for the subject of each message you send. If you respond with simply a newline, no subject field will be sent.
- `askcc` Causes you to be prompted for additional carbon copy recipients at the end of each message. Responding with a newline indicates your satisfaction with the current list.
- `autoprint` Causes the `delete` command to behave like `dp` – thus, after deleting a message, the next one will be typed automatically.
- `debug` Setting the binary option `debug` is the same as specifying `-d` on the command line and causes `mail` to output all sorts of information useful for debugging `mail`.
- `dot` The binary option `dot` causes `mail` to interpret a period alone on a line as the terminator of a message you are sending.
- `hold` This option is used to hold messages in the system mailbox by default.
- `ignore` Causes interrupt signals from your terminal to be ignored and echoed as `@`'s.
- `ignoreeof` An option related to `dot` is `ignoreeof` which makes `mail` refuse to accept a control-d as the end of a message. `Ignoreeof` also applies to `mail` command mode.
- `metoo` Usually, when a group is expanded that contains the sender, the sender is removed from the expansion. Setting this option causes the sender to be included in the group.

- nosave** Normally, when you abort a message with two RUBOUT, *mail* copies the partial letter to the file "dead.letter" in your home directory. Setting the binary option *nosave* prevents this.
- quiet** Suppresses the printing of the version when first invoked.
- verbose** Setting the option *verbose* is the same as using the *-v* flag on the command line. When *mail* runs in verbose mode, the actual delivery of messages is displayed on the users terminal.

The following options have string values:

- EDITOR** Pathname of the text editor to use in the **edit** command and *~e* escape. If not defined, then a default editor is used.
- SHELL** Pathname of the shell to use in the **!** command and the *~!* escape. A default shell is used if this option is not defined.
- VISUAL** Pathname of the text editor to use in the **visual** command and *~v* escape.
- crt** The valued option *crt* is used as a threshold to determine how long a message must be before *more* is used to read it.
- escape** If defined, the first character of this option gives the character to use in the place of *~* to denote escapes.
- folder** The name of the directory to use for storing folders of messages. If this name begins with a '/', *mail* considers it to be an absolute pathname; otherwise, the folder directory is found relative to your home directory.
- record** If defined, gives the pathname of the file used to record all outgoing mail. If not defined, then outgoing mail is not so saved.
- toplines** If defined, gives the number of lines of a message to be printed out with the **top** command; normally, the first five lines are printed.

FILES

/usr/spool/mail/*	post office
~/mbox	your old mail
~/.mailrc	file giving initial mail commands
/tmp/R#	temporary for editor escape
/usr/lib/Mail.help*	help files
/usr/lib/Mail.rc	system initialization file
Message*	temporary for editing messages

SEE ALSO

binmail(1), fmt(1), newaliases(1), aliases(5),
mailaddr(7), sendmail(8)
'The Mail Reference Manual'

BUGS

There are many flags that are not documented here. Most are not useful to the general user. Usually, *mail* is just a link to *Mail*, which can be confusing.

AUTHOR

Kurt Shoens

NAME

make - maintain program groups

SYNOPSIS

make [**-f** *makefile*] [*option*] ... *file* ...

DESCRIPTION

Make executes commands in *makefile* to update one or more target *names*. *Name* is typically a program. If no **-f** option is present, 'makefile' and 'Makefile' are tried in order. If *makefile* is '-', the standard input is taken. More than one **-f** option may appear.

Make updates a target if it depends on prerequisite files that have been modified since the target was last modified, or if the target does not exist.

Makefile contains a sequence of entries that specify dependencies. The first line of an entry is a blank-separated list of targets, then a colon, then a list of prerequisite files. Text following a semicolon, and all following lines that begin with a tab, are shell commands to be executed to update the target. If a name appears on the left of more than one 'colon' line, then it depends on all of the names on the right of the colon on those lines, but only one command sequence may be specified for it. If a name appears on a line with a double colon :: then the command sequence following that line is performed only if the name is out of date with respect to the names to the right of the double colon, and is not affected by other double colon lines on which that name may appear.

Two special forms of a name are recognized. A name like *a(b)* means the file named *b* stored in the archive named *a*. A name like *a((b))* means the file stored in archive *a* containing the entry point *b*.

Sharp and newline surround comments.

The following makefile says that 'pgm' depends on two files 'a.o' and 'b.o', and that they in turn depend on '.c' files and a common file 'incl'.

```
pgm: a.o b.o
    cc a.o b.o -lm -o pgm
a.o: incl a.c
    cc -c a.c
b.o: incl b.c
    cc -c b.c
```

Makefile entries of the form

```
string1 = string2
```

are macro definitions. Subsequent appearances of $\$(string1)$ or $\${string1}$ are replaced by *string2*. If *string1* is a single character, the parentheses or braces are optional.

Make infers prerequisites for files for which *makefile* gives no construction commands. For example, a '.c' file may be inferred as prerequisite for a '.o' file and be compiled to produce the '.o' file. Thus the preceding example can be done more briefly:

```
pgm: a.o b.o
    cc a.o b.o -lm -o pgm
a.o b.o: incl
```

Prerequisites are inferred according to selected suffixes listed as the 'prerequisites' for the special name '.SUFFIXES'; multiple lists accumulate; an empty list clears what came before.

Order is significant; the first possible name for which both a file and a rule as described in the next paragraph exist is inferred. The default list is

`.SUFFIXES: .out .o .c .e .r .f .y .l .s .p`

The rule to create a file with suffix *s2* that depends on a similarly named file with suffix *s1* is specified as an entry for the 'target' *s1s2*. In such an entry, the special macro `$(*)` stands for the target name with suffix deleted, `$(@)` for the full target name, `$(<)` for the complete list of prerequisites, and `$(?)` for the list of prerequisites that are out of date. For example, a rule for making optimized '.o' files from '.c' files is

`.c.o: ; cc -c -O -o $(@) $.c`

Certain macros are used by the default inference rules to communicate optional arguments to any resulting compilations. In particular, 'CFLAGS' is used for *cc*(1) options, 'FFLAGS' for *f77*(1) options, 'PFLAGS' for *pc*(1) options, and 'LFLAGS' and 'YFLAGS' for *lex* and *yacc*(1) options. In addition, the macro 'MFLAGS' is filled in with the initial command line options supplied to *make*. This simplifies maintaining a hierarchy of makefiles as one may then invoke *make* on makefiles in subdirectories and pass along useful options such as `-k`.

Command lines are executed one at a time, each by its own shell. A line is printed when it is executed unless the special target '.SILENT' is in *makefile*, or the first character of the command is '@'.

Commands returning nonzero status (see *intro*(1)) cause *make* to terminate unless the special target '.IGNORE' is in *makefile* or the command begins with `<tab><hyphen>`.

Interrupt and quit cause the target to be deleted unless the target is a directory or depends on the special name '.PRECIOUS'.

Other options:

- `-i` Equivalent to the special entry '.IGNORE:'.
- `-k` When a command returns nonzero status, abandon work on the current entry, but continue on branches that do not depend on the current entry.
- `-n` Trace and print, but do not execute the commands needed to update the targets.
- `-t` Touch, i.e. update the modified date of targets, without executing any commands.
- `-r` Equivalent to an initial special entry '.SUFFIXES:' with no list.
- `-s` Equivalent to the special entry '.SILENT:'.

FILES

makefile, Makefile

SEE ALSO

sh(1), *touch*(1), *f77*(1), *pc*(1)

S. I. Feldman *Make - A Program for Maintaining Computer Programs*

BUGS

Some commands return nonzero status inappropriately. Use `-i` to overcome the difficulty.

Commands that are directly executed by the shell, notably *cd*(1), are ineffectual across new-lines in *make*.

NAME

man - find manual information by keywords; print out the manual

SYNOPSIS

man **-k** keyword ...
man **-f** file ...
man [**-**] [**-t**] [section] title ...

DESCRIPTION

Man is a program which gives information from the programmers manual. It can be asked for one line descriptions of commands specified by name, or for all commands whose description contains any of a set of keywords. It can also provide on-line access to the sections of the printed manual.

When given the option **-k** and a set of keywords, *man* prints out a one line synopsis of each manual sections whose listing in the table of contents contains that keyword.

When given the option **-f** and a list of file names, *man* attempts to locate manual sections related to those files, printing out the table of contents lines for those sections.

When neither **-k** nor **-f** is specified, *man* formats a specified set of manual pages. If a section specifier is given *man* looks in the that section of the manual for the given *titles*. *Section* is an Arabic section number (3 for instance). The number may followed by a single letter classifier (1g for instance) indicating a graphics program in section 1. If *section* is omitted, *man* searches all sections of the manual, giving preference to commands over subroutines in system libraries, and printing the first section it finds, if any.

If the standard output is a teletype, or if the flag **-** is given, *man* pipes its output through *cat*(1) with the option **-s** to crush out useless blank lines, *ul*(1) to create proper underlines for different terminals, and through *more*(1) to stop after each page on the screen. Hit a space to continue, a control-D to scroll 11 more lines when the output stops.

The **-t** flag causes *man* to arrange for the specified section to be *troff*'ed to a suitable raster output device; see *vtroff*(1).

FILES

/usr/man/man?/*
/usr/man/cat?/*

SEE ALSO

more(1), *ul*(1), *whereis*(1), *catman*(8)

BUGS

The manual is supposed to be reproducible either on the phototypesetter or on a typewriter. However, on a typewriter some information is necessarily lost.

NAME

merge - three-way file merge

SYNOPSIS

merge [-p] file1 file2 file3

DESCRIPTION

Merge incorporates all changes that lead from *file2* to *file3* into *file1*. The result goes to std. output if **-p** is present, into *file1* otherwise. *Merge* is useful for combining separate changes to an original. Suppose *file2* is the original, and both *file1* and *file3* are modifications of *file2*. Then *merge* combines both changes.

An overlap occurs if both *file1* and *file3* have changes in a common segment of lines. *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, the user should edit the result and delete one of the alternatives.

IDENTIFICATION

Author: Walter F. Tichy, Purdue University, West Lafayette, IN, 47907.

Revision Number: 3.0 ; Release Date: 82/11/25 .

Copyright © 1982 by Walter F. Tichy.

SEE ALSO

diff3 (1), diff (1), rcsmerge (1), co (1).

NAME

mesg - permit or deny messages

SYNOPSIS

mesg [*n*] [*y*]

DESCRIPTION

Mesg with argument *n* forbids messages via *write* and *talk*(1) by revoking non-user write permission on the user's terminal. *Mesg* with argument *y* reinstates permission. All by itself, *mesg* reports the current state without changing it.

FILES

/dev/tty*

SEE ALSO

write(1), *talk*(1)

DIAGNOSTICS

Exit status is 0 if messages are receivable, 1 if not, 2 on error.

NAME

`mkdir` - make a directory

SYNOPSIS

`mkdir` dirname ...

DESCRIPTION

Mkdir creates specified directories in mode 777. Standard entries, '.', for the directory itself, and '..' for its parent, are made automatically.

Mkdir requires write permission in the parent directory.

SEE ALSO

`rmdir`(1)

NAME

mkstr - create an error message file by massaging C source

SYNOPSIS

mkstr [-] messagefile prefix file ...

DESCRIPTION

Mkstr is used to create files of error messages. Its use can make programs with large numbers of error diagnostics much smaller, and reduce system overhead in running the program as the error messages do not have to be constantly swapped in and out.

Mkstr will process each of the specified *files*, placing a massaged version of the input file in a file whose name consists of the specified *prefix* and the original name. A typical usage of *mkstr* would be

```
mkstr pistrings 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 *pistrings* and processed copies of the source for these files to be placed in files whose names are prefixed with *xx*.

To process the error messages in the source to the message file *mkstr* keys on the string 'error(' in the input stream. Each time it occurs, the C string starting at the '(' is placed in the message file followed by a null character and a new-line character; the null character terminates the message so it can be easily used when retrieved, the new-line character makes it possible to sensibly *cat* the error message file to see its contents. The massaged copy of the input file then contains a *lseek* pointer into the file which can be used to retrieve the message, i.e.:

```
char  efilename[] = "/usr/lib/pi_strings";
int   efil = -1;

error(a1, a2, a3, a4)
{
    char buf[256];

    if (efil < 0) {
        efil = open(efilename, 0);
        if (efil < 0) {
oops:
            perror(efilename);
            exit(1);
        }
    }
    if (lseek(efil, (long) a1, 0) || read(efil, buf, 256) <= 0)
        goto oops;
    printf(buf, a2, a3, a4);
}
```

The optional - causes the error messages to be placed at the end of the specified message file for recompiling part of a large *mkstr* ed program.

SEE ALSO

lseek(2), xstr(1)

AUTHORS

William Joy and Charles Haley

NAME

more, *page* - file perusal filter for crt viewing

SYNOPSIS

more [**-cdfisu**] [**-n**] [**+linenumber**] [**+/pattern**] [name ...]

page *more options*

DESCRIPTION

More is a filter which allows examination of a continuous text one screenful at a time on a soft-copy terminal. It normally pauses after each screenful, printing --More-- at the bottom of the screen. If the user then types a carriage return, one more line is displayed. If the user hits a space, another screenful is displayed. Other possibilities are enumerated later.

The command line options are:

- n** An integer which is the size (in lines) of the window which *more* will use instead of the default.
- c** *More* will draw each page by beginning at the top of the screen and erasing each line just before it draws on it. This avoids scrolling the screen, making it easier to read while *more* is writing. This option will be ignored if the terminal does not have the ability to clear to the end of a line.
- d** *More* will prompt the user with the message "Hit space to continue, Rubout to abort" at the end of each screenful. This is useful if *more* is being used as a filter in some setting, such as a class, where many users may be unsophisticated.
- f** This causes *more* to count logical, 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 may generate escape sequences. These escape sequences contain characters which would ordinarily occupy screen positions, but which do not print when they are sent to the terminal as part of an escape sequence. Thus *more* may think that lines are longer than they actually are, and fold lines erroneously.
- l** Do not treat $\text{\^}L$ (form feed) specially. If this option is not given, *more* will pause after any line that contains a $\text{\^}L$, as if the end of a screenful had been reached. Also, if a file begins with a form feed, the screen will be cleared before the file is printed.
- 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* will handle underlining such as produced by *nroff* in a manner appropriate to the particular terminal: if the terminal can perform underlining or has a stand-out mode, *more* will output appropriate escape sequences to enable underlining or stand-out mode for underlined information in the source file. The **-u** option suppresses this processing.

+linenumber

Start up at *linenumber*.

+/pattern

Start up two lines before the line containing the regular expression *pattern*.

If the program is invoked as *page*, then the screen is cleared before each screenful is printed (but only if a full screenful is being printed), and $k - 1$ rather than $k - 2$ lines are printed in each screenful, where k is the number of lines the terminal can display.

More looks in the file */etc/termcap* to determine terminal characteristics, and to determine the default window size. On a terminal capable of displaying 24 lines, the default window size is 22 lines.

More looks in the environment variable *MORE* to pre-set any flags desired. For example, if you prefer to view files using the *-c* mode of operation, the *cs*h command *setenv MORE -c* or the *sh* command sequence *MORE='-c'*; *export MORE* would cause all invocations of *more*, including invocations by programs such as *man* and *msgs*, to use this mode. Normally, the user will place the command sequence which sets up the *MORE* environment variable in the *.cshrc* or *.profile* file.

If *more* is reading from a file, rather than a pipe, then a percentage is displayed along with the *--More--* prompt. This gives the fraction of the file (in characters, not lines) that has been read so far.

Other sequences which may be typed when *more* pauses, and their effects, are as follows (*i* is an optional integer argument, defaulting to 1) :

- i* <space>
display *i* more lines, (or another screenful if no argument is given)
- ^D* display 11 more lines (a "scroll"). If *i* is given, then the scroll size is set to *i*.
- d* same as *^D* (control-D)
- iz* same as typing a space except that *i*, if present, becomes the new window size.
- is* skip *i* lines and print a screenful of lines
- if* skip *i* screenfuls and print a screenful of lines
- q* or *Q* Exit from *more*.
- =* Display the current line number.
- v* Start up the editor *vi* at the current line.
- h* Help command; give a description of all the *more* commands.
- i/expr* search for the *i*-th occurrence of the regular expression *expr*. If there are less than *i* occurrences of *expr*, and the input is a file (rather than a pipe), then the position in the file remains unchanged. Otherwise, a screenful is displayed, starting two lines before the place where the expression was found. The user's erase and kill characters may be used to edit the regular expression. Erasing back past the first column cancels the search command.
- in* search for the *i*-th occurrence of the last regular expression entered.
- '* (single quote) Go to the point from which the last search started. If no search has been performed in the current file, this command goes back to the beginning of the file.
- !command*
invoke a shell with *command*. The characters *'%*' and *'!*' in "command" 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.
- i:n* skip to the *i*-th next file given in the command line (skips to last file if *n* doesn't make sense)

- i:p** skip to the *i*-th previous file given in the command line. If this command is given in the middle of printing out a file, then *more* goes back to the beginning of the file. If *i* doesn't make sense, *more* skips back to the first file. If *more* is not reading from a file, the bell is rung and nothing else happens.
- :f** display the current file name and line number.
- :q** or **:Q** exit from *more* (same as **q** or **Q**).
- .** (dot) repeat the previous command.

The commands take effect immediately, i.e., it is not necessary to type a carriage return. Up to the time when the command character itself is given, the user may hit the line kill character to cancel the numerical argument being formed. In addition, the user may hit the erase character to redisplay the --More--(xx%) message.

At any time when output is being sent to the terminal, the user can hit the quit key (normally control-****). *More* will stop sending output, and will display the usual --More-- prompt. The user may then enter one of the above commands in the normal manner. Unfortunately, some output is lost when this is done, due to the fact that any characters waiting in the terminal's output queue are flushed when the quit signal occurs.

The terminal is set to *noecho* mode by this program so that the output can be continuous. What you type will thus not show on your terminal, except for the **/** and **!** commands.

If the standard output is not a teletype, then *more* acts just like *cat*, except that a header is printed before each file (if there is more than one).

A sample usage of *more* in previewing *nroff* output would be

```
nroff -ms +2 doc.n | more -s
```

AUTHOR

Eric Shienbrood, minor revisions by John Foderaro and Geoffrey Peck

FILES

/etc/termcap	Terminal data base
/usr/lib/more.help	Help file

SEE ALSO

csh(1), man(1), msgs(1), script(1), sh(1), environ(7)

NAME

msgs - system messages and junk mail program

SYNOPSIS

msgs [**-fhlpq**] [**number**] [**-number**]

DESCRIPTION

Msgs is used to read system messages. These messages are sent by mailing to the login 'msgs' and should be short pieces of information which are suitable to be read once by most users of the system.

Msgs is normally invoked each time you login, by placing it in the file *.login* (*.profile* if you use */bin/sh*). It will then prompt you with the source and subject of each new message. If there is no subject line, the first few non-blank lines of the message will be displayed. If there is more to the message, you will be told how long it is and asked whether you wish to see the rest of the message. The possible responses are:

y type the rest of the message

RETURN

synonym for **y**.

n skip this message and go on to the next message.

- redisplay the last message.

q drops you out of *msgs*; the next time you run the program it will pick up where you left off.

s append the current message to the file "Messages" in the current directory; 's-' will save the previously displayed message. A 's' or 's-' may be followed by a space and a filename to receive the message replacing the default "Messages".

m or 'm-' causes a copy of the specified message to be placed in a temporary mailbox and *mail(1)* to be invoked on that mailbox. Both 'm' and 's' accept a numeric argument in place of the '-'.

Msgs keeps track of the next message you will see by a number in the file *.msgsrc* in your home directory. In the directory */usr/msgs* it keeps a set of files whose names are the (sequential) numbers of the messages they represent. The file */usr/msgs/bounds* shows the low and high number of the messages in the directory so that *msgs* can quickly determine if there are no messages for you. If the contents of *bounds* is incorrect it can be fixed by removing it; *msgs* will make a new *bounds* file the next time it is run.

Options to *msgs* include:

-f which causes it not to say "No new messages.". This is useful in your *.login* file since this is often the case here.

-q Queries whether there are messages, printing "There are new messages." if there are. The command "*msgs -q*" is often used in login scripts.

-h causes *msgs* to print the first part of messages only.

-l option causes only locally originated messages to be reported.

num A message number can be given on the command line, causing *msgs* to start at the specified message rather than at the next message indicated by your *.msgsrc* file.

Thus

`msgs -h 1`

prints the first part of all messages.

-number

will cause *msgs* to start *number* messages back from the one indicated by your *.msgsrc* file, useful for reviews of recent messages.

-p causes long messages to be piped through *more(1)*.

Within *msgs* you can also go to any specific message by typing its number when *msgs* requests input as to what to do.

FILES

`/usr/msgs/*`
`~/.msgsrc`

database
number of next message to be presented

AUTHORS

William Joy
David Wasley

SEE ALSO

`mail(1)`, `more(1)`

NAME

mt - magnetic tape manipulating program

SYNOPSIS

mt [**-f** *tapename*] *command* [*count*]

DESCRIPTION

Mt is used to give commands to the cassette tape drive. If a tape name is not specified, the environment variable **TAPE** is used; if **TAPE** does not exist, *mt* uses the device */dev/rct0*. By default *mt* performs the requested operation once. Operations may be performed multiple times by specifying *count*.

The available commands are listed below. Only as many characters as are required to uniquely identify a command need be specified.

eof, weof

Write *count* end-of-file marks at the current position on the tape. Note that is currently necessary to write end-of-file marks to separate multiple files on the same tape.

fsf Forward space *count* files.

fsr Forward space *count* records.

fseof Forward space to end of tape. End of tape is defined as the last position written on tape. Note that it is necessary to issue this command (instead of **fsf**) to add data to a tape.

rewind

Rewind the tape (*Count* is ignored.)

status Print status information about the tape unit.

Mt returns a 0 exit status when the operation(s) were successful, 1 if the command was unrecognized, and 2 if an operation failed.

FILES

*/dev/rct** Cassette tape interface

SEE ALSO

mtio(4), *dd*(1), *ioctl*(2), *environ*(7)

NAME

mv - move or rename files

SYNOPSIS

mv [-i] [-f] [-] file1 file2

mv [-i] [-f] [-] file ... directory

DESCRIPTION

Mv moves (changes the name of) *file1* to *file2*.

If *file2* already exists, it is removed before *file1* is moved. If *file2* has a mode which forbids writing, *mv* prints the mode (see *chmod(2)*) and reads the standard input to obtain a line; if the line begins with *y*, the move takes place; if not, *mv* exits.

In the second form, one or more *files* (plain files or directories) are moved to the *directory* with their original file-names.

Mv refuses to move a file onto itself.

Options:

- i stands for interactive mode. Whenever a move is to supercede an existing file, the user is prompted by the name of the file followed by a question mark. If he answers with a line starting with 'y', the move continues. Any other reply prevents the move from occurring.
- f stands for force. This option overrides any mode restrictions or the -i switch.
- means interpret all the following arguments to *mv* as file names. This allows file names starting with minus.

SEE ALSO

cp(1), *ln(1)*

BUGS

If *file1* and *file2* lie on different file systems, *mv* must copy the file and delete the original. In this case the owner name becomes that of the copying process and any linking relationship with other files is lost.

NAME

netstat - show network status

SYNOPSIS

netstat [**-A**himnrs] [**-p** *protocol*] [**-a**] [*interval*] [*system*] [*core*]

DESCRIPTION

The *netstat* command symbolically displays the contents of various network-related data structures. The options have the following meaning:

- A** show the address of any associated protocol control blocks; used for debugging
- a** show the state of all sockets; normally sockets used by server processes are not shown
- h** show the state of the IMP host table
- i** show the state of interfaces which have been auto-configured (interfaces statically configured into a system, but not located at boot time are not shown)
- m** show statistics recorded by the memory management routines (the network manages a "private share" of memory)
- n** show network addresses as numbers (normally *netstat* interprets addresses and attempts to display them symbolically)
- p** *proto*
show the state of sockets utilizing protocol *proto*; the protocol is specified symbolically, and may be any protocol listed in the file */etc/protocols*.
- s** show per-protocol statistics
- r** show the routing tables

The arguments, *system* and *core* allow substitutes for the defaults *"/vmunix"* and *"/dev/kmem"*.

If an *interval* is specified, *netstat* will continuously display the information regarding packet traffic on the configured network interfaces, pausing *interval* seconds before refreshing the screen.

There are a number of display formats, depending on the information presented. The default display, for active sockets, shows the local and remote addresses, send and receive queue sizes (in bytes), protocol, and, optionally, the internal state of the protocol.

Address formats are of the form "host.port" or "network.port" if a socket's address specifies a network but no specific host address. When known the host and network addresses are displayed symbolically according to the data bases */etc/hosts* and */etc/networks*, respectively. If a symbolic name for an address is unknown, or if the **-n** option is specified, the address is printed in the Internet "dot format"; refer to *inet(3N)* for more information regarding this format. Unspecified, or "wildcard", addresses and ports appear as "*".

The interface display provides a table of cumulative statistics regarding packets transferred, errors, and collisions. The network address (currently Internet specific) of the interface and the maximum transmission unit ("mtu") are also displayed.

The routing table display indicates the available routes and their status. Each route consists of a destination host or network and a gateway to use in forwarding packets. The flags field shows the state of the route ("U" if "up"), and whether the route is to a gateway ("G"). Direct routes are created for each interface attached to the local host. The refcnt field gives the current number of active uses of the route. Connection oriented protocols normally hold on to a single route for the duration of a connection while connectionless protocols obtain a

route then discard it. The use field provides a count of the number of packets sent using that route. The interface entry indicates the network interface utilized for the route.

When *netstat* is invoked with an *interval* argument, it displays a running count of statistics related to network interfaces. This display consists of a column summarizing information for all interfaces, and a column for the interface with the most traffic since the system was last rebooted. The first line of each screen of information contains a summary since the system was last rebooted. Subsequent lines of output show values accumulated over the preceding interval.

SEE ALSO

ioostat(1), vmstat(1), hosts(5), networks(5), protocols(5), services(5), trpt(8C)

BUGS

The notion of errors is ill-defined. Collisions mean something else for the IMP.

NAME

newaliases - rebuild the data base for the mail aliases file

SYNOPSIS

newaliases

DESCRIPTION

Newaliases rebuilds the random access data base for the mail aliases file */usr/lib/aliases*. It must be run each time */usr/lib/aliases* is changed in order for the change to take effect.

SEE ALSO

aliases(5), *sendmail*(8)

BUGS

NAME

nice, *nohup* – run a command at low priority (*sh* only)

SYNOPSIS

nice [*-number*] *command* [*arguments*]

nohup *command* [*arguments*]

DESCRIPTION

Nice executes *command* with low scheduling priority. If the *number* argument is present, the priority is incremented (higher numbers mean lower priorities) by that amount up to a limit of 20. The default *number* is 10.

The super-user may run commands with priority higher than normal by using a negative priority, e.g. ‘-10’.

Nohup executes *command* immune to hangup and terminate signals from the controlling terminal. The priority is incremented by 5. *Nohup* should be invoked from the shell with ‘&’ in order to prevent it from responding to interrupts by or stealing the input from the next person who logs in on the same terminal. The syntax of *nice* is also different.

FILES

nohup.out standard output and standard error file under *nohup*

SEE ALSO

csh(1), *setpriority*(2), *renice*(8)

DIAGNOSTICS

Nice returns the exit status of the subject command.

BUGS

Nice and *nohup* are particular to *sh*(1). If you use *csh*(1), then commands executed with “&” are automatically immune to hangup signals while in the background. There is a builtin command *nohup* which provides immunity from terminate, but it does not redirect output to *nohup.out*.

Nice is built into *csh*(1) with a slightly different syntax than described here. The form “*nice +10*” nices to positive *nice*, and “*nice -10*” can be used by the super-user to give a process more of the processor.

NAME

nm - print name list

SYNOPSIS

nm [-gnopru] [file ...]

DESCRIPTION

Nm prints the name list (symbol table) of each object *file* in the argument list. If an argument is an archive, a listing for each object file in the archive will be produced. If no *file* is given, the symbols in "a.out" are listed.

Each symbol name is preceded by its value (blanks if undefined) and one of the letters U (undefined), A (absolute), T (text segment symbol), D (data segment symbol), B (bss segment symbol), C (common symbol), f file name, or - for sdb symbol table entries (see -a below). If the symbol is local (non-external) the type letter is in lower case. The output is sorted alphabetically.

Options are:

- g Print only global (external) symbols.
- n Sort numerically rather than alphabetically.
- o Prepend file or archive element name to each output line rather than only once.
- p Don't sort; print in symbol-table order.
- r Sort in reverse order.
- u Print only undefined symbols.

SEE ALSO

ar(1), ar(5), a.out(5), stab(5)

NAME

`fmt` - simple text formatter

SYNOPSIS

`fmt` [`-width`] [`name ...`]

DESCRIPTION

Fmt is a simple text formatter which reads the concatenation of input files (or standard input if none are given) and produces on standard output a version of its input with lines as close to 72 characters long as possible. The spacing at the beginning of the input lines is preserved in the output, as are blank lines and interword spacing. A width can be specified to change the default of 72 characters.

Fmt is meant to format mail messages prior to sending, but may also be useful for other simple tasks. For instance, within visual mode of the *ex* editor (e.g. *vi*) the command

```
!}fmt
```

will reformat a paragraph, evening the lines.

SEE ALSO

`nroff(1)`, `mail(1)`

AUTHOR

Kurt Shoens

BUGS

The program was designed to be simple and fast - for more complex operations, the standard text processors are likely to be more appropriate.



NAME

nroff - text formatting

SYNOPSIS

nroff [option] ... [file] ...

DESCRIPTION

Nroff formats text in the named *files* for typewriter-like devices. See also *troff(1)*. The full capabilities of *nroff* are described in the *Nroff/Troff User's Manual*.

If no *file* argument is present, the standard input is read. An argument consisting of a single minus (-) is taken to be a file name corresponding to the standard input.

The options, which may appear in any order so long as they appear *before* the files, are:

- olist Print only pages whose page numbers appear in the comma-separated *list* of numbers and ranges. A range *N-M* means pages *N* through *M*; an initial *-N* means from the beginning to page *N*; and a final *N-* means from *N* to the end.
- n*N* Number first generated page *N*.
- s*N* Stop every *N* pages. *Nroff* will halt prior to every *N* pages (default *N=1*) to allow paper loading or changing, and will resume upon receipt of a newline.
- m*name* Prepend the macro file */usr/lib/tmac/tmac.name* to the input *files*.
- ra*N* Set register *a* (one-character) to *N*.
- i Read standard input after the input files are exhausted.
- q Invoke the simultaneous input-output mode of the **rd** request.
- T*name* Prepare output for specified terminal. Known *names* are **37** for the (default) Teletype Corporation Model 37 terminal, **tn300** for the GE TermiNet 300 (or any terminal without half-line capability), **300S** for the DASI-300S, **300** for the DASI-300, and **450** for the DASI-450 (Diablo Hyterm).
- e Produce equally-spaced words in adjusted lines, using full terminal resolution.
- h Use output tabs during horizontal spacing to speed output and reduce output character count. Tab settings are assumed to be every 8 nominal character widths.

FILES

<i>/tmp/ta*</i>	temporary file
<i>/usr/lib/tmac/tmac.*</i>	standard macro files
<i>/usr/lib/term/*</i>	terminal driving tables for <i>nroff</i>

SEE ALSO

J. F. Ossanna, *Nroff/Troff user's manual*
 B. W. Kernighan, *A TROFF Tutorial*
troff(1), *eqn(1)*, *tbl(1)*, *ms(7)*, *me(7)*, *man(7)*, *col(1)*

NAME

od - octal, decimal, hex, ascii dump

SYNOPSIS

od [-format] [file] [[+]offset[.][b] [label]]

DESCRIPTION

Od displays *file*, or it's standard input, in one or more dump formats as selected by the first argument. If the first argument is missing, **-o** is the default. Dumping continues until end-of-file.

The meanings of the format argument characters are:

- a** Interpret bytes as characters and display them with their ACSII names. If the **p** character is given also, then bytes with even parity are underlined. The **P** character causes bytes with odd parity to be underlined. Otherwise the parity bit is ignored.
- b** Interpret bytes as unsigned octal.
- c** Interpret bytes as ASCII characters. Certain non-graphic characters appear as C escapes: null=**\0**, backspace=**\b**, formfeed=**\f**, newline=**\n**, return=**\r**, tab=**\t**; others appear as 3-digit octal numbers. Bytes with the parity bit set are displayed in octal.
- d** Interpret (short) words as unsigned decimal.
- f** Interpret long words as floating point.
- h** Interpret (short) words as unsigned hexadecimal.
- i** Interpret (short) words as signed decimal.
- l** Interpret long words as signed decimal.
- o** Interpret (short) words as unsigned octal.
- s[n]** Look for strings of ascii graphic characters, terminated with a null byte. *N* specifies the minimum length string to be recognized. By default, the minimum length is 3 characters.
- v** Show all data. By default, display lines that are identical to the last line shown are not output, but are indicated with an "*" in column 1.
- w[n]** Specifies the number of input bytes to be interpreted and displayed on each output line. If **w** is not specified, 16 bytes are read for each display line. If *n* is not specified, it defaults to 32.
- x** Interpret (short) words as hexadecimal.

An upper case format character implies the long or double precision form of the object.

The *offset* argument specifies the byte offset into the file where dumping is to commence. By default this argument is interpreted in octal. A different radix can be specified; If "." is appended to the argument, then *offset* is interpreted in decimal. If *offset* begins with "x" or "0x", it is interpreted in hexadecimal. If "b" ("B") is appended, the offset is interpreted as a block count, where a block is 512 (1024) bytes. If the *file* argument is omitted, an *offset* argument must be preceded by "+".

The radix of the displayed address will be the same as the radix of the *offset*, if specified; otherwise it will be octal.

Label will be interpreted as a pseudo-address for the first byte displayed. It will be shown in "(" following the file offset. It is intended to be used with core images to indicate the real memory address. The syntax for *label* is identical to that for *offset*.

SEE ALSO

adb(1)

BUGS

A file name argument can't start with "+". A hexadecimal offset can't be a block count. Only one file name argument can be given.

It is an historical botch to require specification of object, radix, and sign representation in a single character argument.

NAME

pagesize – print system page size

SYNOPSIS

pagesize

DESCRIPTION

Pagesize prints the size of a page of memory in bytes, as returned by *getpagesize(2)*. This program is useful in constructing portable shell scripts.

SEE ALSO

getpagesize(2)

NAME

passwd - change login password

SYNOPSIS

passwd [name]

DESCRIPTION

This command changes (or installs) a password associated with the user *name* (your own name by default).

The program prompts for the old password and then for the new one. The caller must supply both. The new password must be typed twice, to forestall mistakes.

New passwords must be at least four characters long if they use a sufficiently rich alphabet and at least six characters long if monospace. These rules are relaxed if you are insistent enough.

Only the owner of the name or the super-user may change a password; the owner must prove he knows the old password.

FILES

/etc/passwd

SEE ALSO

login(1), passwd(5), crypt(3)

Robert Morris and Ken Thompson, *UNIX password security*

BUGS

The password file information should be kept in a different data structure allowing indexed access; *dbm(3X)* would probably be suitable.

NAME

pc - Pascal compiler

SYNOPSIS

pc [option] [*-i* name ...] name ...

DESCRIPTION

Pc is a Pascal compiler. If given an argument file ending with *.p*, it will compile the file and load it into an executable file called, by default, *a.out*.

A program may be separated into more than one *.p* file. *Pc* will compile a number of argument *.p* files into object files (with the extension *.o* in place of *.p*). Object files may then be loaded into an executable *a.out* file. Exactly one object file must supply a **program** statement to successfully create an executable *a.out* file. The rest of the files must consist only of declarations which logically nest within the program. References to objects shared between separately compiled files are allowed if the objects are declared in **included** header files, whose names must end with *.h*. Header files may only be included at the outermost level, and thus declare only globally available objects. To allow **functions** and **procedures** to be declared, an **external** directive has been added, whose use is similar to the **forward** directive but restricted to appear only in *.h* files. **Function** and **procedure** bodies may not appear in *.h* files. A binding phase of the compiler checks that declarations are used consistently, to enforce the type checking rules of Pascal.

Object files created by other language processors may be loaded together with object files created by *pc*. The **functions** and **procedures** they define must have been declared in *.h* files included by all the *.p* files which call those routines. Calling conventions are as in C, with **var** parameters passed by address.

See the Berkeley Pascal User's Manual for details.

The following options have the same meaning as in *cc(1)* and *f77(1)*. See *ld(1)* for load-time options.

- c** Suppress loading and produce '*.o*' file(s) from source file(s).
- g** Have the compiler produce additional symbol table information for *dbx(1)*.
- w** Suppress warning messages.
- p** Prepare object files for profiling, see *prof(1)*.
- O** Invoke an object-code improver.
- S** Compile the named program, and leave the assembler-language output on the corresponding file suffixed '*.s*'. (No '*.o*' is created.)
- o output**
Name the final output file *output* instead of *a.out*.

The following options are peculiar to *pc*.

- C** Compile code to perform runtime checks, verify **assert** calls, and initialize all variables to zero as in *pi*.
- b** Block buffer the file *output*.
- i** Produce a listing for the specified procedures, functions and **include** files.
- l** Make a program listing during translation.
- s** Accept standard Pascal only; non-standard constructs cause warning diagnostics.

-t directory

Use the given *directory* for compiler temporary files.

-z Allow execution profiling with *pxp* by generating statement counters, and arranging for the creation of the profile data file *pmon.out* when the resulting object is executed.

Other arguments are taken to be loader option arguments, perhaps libraries of *pc* compatible routines. Certain flags can also be controlled in comments within the program as described in the *Berkeley Pascal User's Manual*.

FILES

file.p	pascal source files
/usr/lib/pc0	compiler
/lib/f1	code generator
/usr/lib/pc2	runtime integrator (inline expander)
/lib/c2	peephole optimizer
/usr/lib/pc3	separate compilation consistency checker
/usr/lib/pc2.*strings	text of the error messages
/usr/lib/how_pc	basic usage explanation
/usr/lib/libpc.a	intrinsic functions and I/O library
/usr/lib/libm.a	math library
/lib/libc.a	standard library, see <i>intro(3)</i>

SEE ALSO

Berkeley Pascal User's Manual
pi(1), *pxp(1)*, *pxref(1)*, *sdb(1)*

DIAGNOSTICS

For a basic explanation do

pc

See *pi(1)*. for an explanation of the error message format. Internal errors cause messages containing the word SNARK.

AUTHORS

Charles B. Haley, William N. Joy, and Ken Thompson
 Retargetted to the second pass of the portable *C* compiler by Peter Kessler
 Runtime library and inline optimizer by M. Kirk McKusick
 Separate compilation consistency checking by Louise Madrid

BUGS

The keyword **packed** is recognized but has no effect.

The binder is not as strict as described here, with regard to the rules about external declarations only in '.h' files and including '.h' files only at the outermost level. It will be made to perform these checks in its next incarnation, so users are warned not to be sloppy.

The **-z** flag doesn't work for separately compiled files.

Because the **-s** option is usurped by the compiler, it is not possible to pass the *strip* option to the loader. Thus programs which are to be stripped, must be run through *strip(1)* after they are compiled.

NAME

pdx - pascal debugger

SYNOPSIS

pdx [-r] [*objfile*]

DESCRIPTION

Pdx is a tool for source level debugging and execution of Pascal programs. The *objfile* is an object file produced by the Pascal translator *pi*(1). If no *objfile* is specified, *pdx* looks for a file named "obj" in the current directory. The object file contains a symbol table which includes the name of the all the source files translated by *pi* to create it. These files are available for perusal while using the debugger.

If the file ".pdxinit" exists in the current directory, then the debugger commands in it are executed.

The *-r* option causes the *objfile* to be executed immediately; if it terminates successfully *pdx* exits. Otherwise it reports the reason for termination and offers the user the option of entering the debugger or simply letting *px* continue with a traceback. If *-r* is not specified, *pdx* just prompts and waits for a command.

The commands are:

run [*args*] [< *filename*] [> *filename*]

Start executing *objfile*, passing *args* as command line arguments; < or > can be used to redirect input or output in the usual manner.

trace [*in procedure/function*] [*if condition*]

trace *source-line-number* [*if condition*]

trace *procedure/function* [*in procedure/function*] [*if condition*]

trace *expression at source-line-number* [*if condition*]

trace *variable* [*in procedure/function*] [*if condition*]

Have tracing information printed when the program is executed. A number is associated with the command that is used to turn the tracing off (see the **delete** command).

The first argument describes what is to be traced. If it is a *source-line-number*, then the line is printed immediately prior to being executed. Source line numbers in a file other than the current one must be preceded by the name of the file and a colon, e.g. "mumble.p:17".

If the argument is a procedure or function name then every time it is called, information is printed telling what routine called it, from what source line it was called, and what parameters were passed to it. In addition, its return is noted, and if it's a function then the value it is returning is also printed.

If the argument is an *expression* with an **at** clause then the value of the expression is printed whenever the identified source line is reached.

If the argument is a variable then the name and value of the variable is printed whenever it changes. Execution is substantially slower during this form of tracing.

If no argument is specified then all source lines are printed before they are executed. Execution is substantially slower during this form of tracing.

The clause "*in procedure/function*" restricts tracing information to be printed only while executing inside the given procedure or function.

Condition is a Pascal boolean expression and is evaluated prior to printing the tracing information; if it is false then the information is not printed.

There is no restriction on the amount of information that can be traced.

stop if condition

stop at source-line-number [if condition]

stop in procedure/function [if condition]

stop variable [if condition]

Stop execution when the given line is reached, procedure or function called, variable changed, or condition true.

delete command-number

The trace or stop corresponding to the given number is removed. The numbers associated with traces and stops are printed by the **status** command.

status [> filename]

Print out the currently active **trace** and **stop** commands.

cont Continue execution from where it stopped. This can only be done when the program was stopped by an interrupt or through use of the **stop** command.

step Execute one source line.

next Execute up to the next source line. The difference between this and **step** is that if the line contains a call to a procedure or function the **step** command will stop at the beginning of that block, while the **next** command will not.

print expression [, expression ...]

Print out the values of the Pascal expressions. Variables declared in an outer block but having the same identifier as one in the current block may be referenced as "*block-name . variable*".

whatis identifier

Print the declaration of the given identifier.

which identifier

Print the full qualification of the given identifier, i.e. the outer blocks that the identifier is associated with.

assign variable expression

Assign the value of the expression to the variable.

call procedure(parameters)

Execute the object code associated with the named procedure or function.

help Print out a synopsis of *pdx* commands.

gripe Invokes a mail program to send a message to the person in charge of *pdx*.

where Print out a list of the active procedures and functions and the respective source line where they are called.

source filename

Read *pdx* commands from the given *filename*. Especially useful when the *filename* has been created by redirecting a **status** command from an earlier debugging session.

dump [> filename]

Print the names and values of all active data.

list [*source-line-number* [, *source-line-number*]]

list *procedure/function*

List the lines in the current source file from the first line number to the second inclusive. As in the editor "\$" can be used to refer to the last line. If no lines are specified, the entire file is listed. If the name of a procedure or function is given lines *n-k* to *n+k* are listed where *n* is the first statement in the procedure or function and *k* is small.

file [*filename*]

Change the current source file name to *filename*. If none is specified then the current source file name is printed.

edit [*filename*]

edit *procedure/function-name*

Invoke an editor on *filename* or the current source file if none is specified. If a *procedure* or *function* name is specified, the editor is invoked on the file that contains it. Which editor is invoked by default depends on the installation. The default can be overridden by setting the environment variable EDITOR to the name of the desired editor.

pi Recompile the program and read in the new symbol table information.

sh *command-line*

Pass the command line to the shell for execution. The SHELL environment variable determines which shell is used.

alias *new-command-name old-command-name*

This command makes *pdx* respond to *new-command-name* the way it used to respond to *old-command-name*.

quit Exit *pdx*.

The following commands deal with the program at the *px* instruction level rather than source level. They are not intended for general use.

tracei [*address*] [*if cond*]

tracei [*variable*] [*at address*] [*if cond*]

stopi [*address*] [*if cond*]

stopi [*at*] [*address*] [*if cond*]

Turn on tracing or set a stop using a *px* machine instruction addresses.

xi *address* [, *address*]

Print the instructions starting at the first *address*. Instructions up to the second *address* are printed.

xd *address* [, *address*]

Print in octal the specified data location(s).

FILES

obj	Pascal object file
.pdxinit	<i>Pdx</i> initialization file

SEE ALSO

pi(1), px(1)

*An Introduction to Pdx***BUGS**

Pdx does not understand sets, and provides no information about files.

The *whatis* command doesn't quite work for variant records.

Bad things will happen if a procedure invoked with the **call** command does a non-local goto.

The commands **step** and **next** should be able to take a *count* that specifies how many lines to execute.

There should be commands **stepi** and **nexti** that correspond to **step** and **next** but work at the instruction level.

There should be a way to get an address associated with a line number, procedure or function, and variable.

Most of the command names are too long.

The alias facility is quite weak.

A *cs**h*-like history capability would improve the situation.

NAME

`pi` - Pascal interpreter code translator

SYNOPSIS

`pi` [**option**] [`-i name ...`] `name.p`

DESCRIPTION

Pi translates the program in the file *name.p* leaving interpreter code in the file *obj* in the current directory. The interpreter code can be executed using *px*. *Pix* performs the functions of *pi* and *px* for 'load and go' Pascal.

The following flags are interpreted by *pi*; the associated options can also be controlled in comments within the program as described in the *Berkeley Pascal User's Manual*.

- b** Block buffer the file *output*.
- i** Enable the listing for any specified procedures and functions and while processing any specified **include** files.
- l** Make a program listing during translation.
- n** Begin each listed **include** file on a new page with a banner line.
- p** Suppress the post-mortem control flow backtrace if an error occurs; suppress statement limit counting.
- s** Accept standard Pascal only; non-standard constructs cause warning diagnostics.
- t** Suppress runtime tests of subrange variables and treat **assert** statements as comments.
- u** Card image mode; only the first 72 characters of input lines are used.
- w** Suppress warning diagnostics.
- z** Allow execution profiling with *pxp* by generating statement counters, and arranging for the creation of the profile data file *pmon.out* when the resulting object is executed.

FILES

<code>file.p</code>	input file
<code>file.i</code>	include file(s)
<code>/usr/lib/pi2.*strings</code>	text of the error messages
<code>/usr/lib/how_pi*</code>	basic usage explanation
<code>obj</code>	interpreter code output

SEE ALSO

Berkeley Pascal User's Manual
`pix(1)`, `px(1)`, `pxp(1)`, `pxref(1)`

DIAGNOSTICS

For a basic explanation do

pi

In the diagnostic output of the translator, lines containing syntax errors are listed with a flag indicating the point of error. Diagnostic messages indicate the action which the recovery mechanism took in order to be able to continue parsing. Some diagnostics indicate only that the input is 'malformed.' This occurs if the recovery can find no simple correction to make the input syntactically valid.

Semantic error diagnostics indicate a line in the source text near the point of error. Some errors evoke more than one diagnostic to help pinpoint the error; the follow-up messages begin with an ellipsis '...'.
The first character of each error message indicates its class:

E	Fatal error; no code will be generated.
e	Non-fatal error.
w	Warning - a potential problem.
s	Non-standard Pascal construct warning.

If a severe error occurs which inhibits further processing, the translator will give a diagnostic and then 'QUIT'.

AUTHORS

Charles B. Haley, William N. Joy, and Ken Thompson
Ported to VAX-11 by Peter Kessler

BUGS

The keyword **packed** is recognized but has no effect.

For clarity, semantic errors should be flagged at an appropriate place in the source text, and multiple instances of the 'same' semantic error should be summarized at the end of a **procedure** or **function** rather than evoking many diagnostics.

When **include** files are present, diagnostics relating to the last procedure in one file may appear after the beginning of the listing of the next.

NAME

`pix` - Pascal interpreter and executor

SYNOPSIS

`pix` [`-blnpstuwz`] [`-i name ...`] `name.p` [`argument ...`]

DESCRIPTION

Pix is a 'load and go' version of Pascal which combines the functions of the interpreter code translator *pi* and the executor *px*. It uses *pi* to translate the program in the file *name.p* and, if there were no fatal errors during translation, causes the resulting interpreter code to be executed by *px* with the specified arguments. A temporary file is used for the object code; the file *obj* is neither created nor destroyed.

FILES

<code>/usr/ucb/pi</code>	Pascal translator
<code>/usr/ucb/px</code>	Pascal executor
<code>/tmp/pix*</code>	temporary
<code>/usr/lib/how_pix</code>	basic explanation

SEE ALSO

Berkeley Pascal User's Manual
`pi(1)`, `px(1)`

DIAGNOSTICS

For a basic explanation do

`pix`

AUTHORS

Susan L. Graham and William N. Joy

NAME

plot - graphics filters

SYNOPSIS

plot [-Tterminal [raster]]

DESCRIPTION

These commands read plotting instructions (see *plot(5)*) from the standard input, and in general produce plotting instructions suitable for a particular *terminal* on the standard output.

If no *terminal* type is specified, the environment parameter \$TERM (see *environ(7)*) is used. Known *terminals* are:

4014 Tektronix 4014 storage scope.

450 DASI Hyterm 450 terminal (Diablo mechanism).

300 DASI 300 or GSI terminal (Diablo mechanism).

300S DASI 300S terminal (Diablo mechanism).

ver Versatec D1200A printer-plotter. This version of *plot* places a scan-converted image in '/usr/tmp/raster' and sends the result directly to the plotter device rather than to the standard output. The optional argument causes a previously scan-converted file *raster* to be sent to the plotter.

FILES

/usr/bin/tek
/usr/bin/t450
/usr/bin/t300
/usr/bin/t300s
/usr/bin/vplot
/usr/tmp/raster

SEE ALSO

plot(3X), plot(5)

BUGS

There is no lockout protection for /usr/tmp/raster.

NAME

pmerge - pascal file merger

SYNOPSIS

pmerge name.p ...

DESCRIPTION

Pmerge assembles the named Pascal files into a single standard Pascal program. The resulting program is listed on the standard output. It is intended to be used to merge a collection of separately compiled modules so that they can be run through **pi**, or exported to other sites.

FILES

/usr/tmp/MG* default temporary files

SEE ALSO

pc(1), pi(1),
Auxiliary documentation *Berkeley Pascal User's Manual*.

AUTHOR

M. Kirk McKusick

BUGS

Very minimal error checking is done, so incorrect programs will produce unpredictable results. Block comments should be placed after the keyword to which they refer or they are likely to end up in bizarre places.

NAME

pr - print file

SYNOPSIS

pr [option] ... [file] ...

DESCRIPTION

Pr produces a printed listing of one or more *files*. The output is separated into pages headed by a date, the name of the file or a specified header, and the page number. If there are no file arguments, *pr* prints its standard input.

Options apply to all following files but may be reset between files:

- n* Produce *n*-column output.
 - +*n* Begin printing with page *n*.
 - h* Take the next argument as a page header.
 - wn* For purposes of multi-column output, take the width of the page to be *n* characters instead of the default 72.
 - f* Use formfeeds instead of newlines to separate pages. A formfeed is assumed to use up two blank lines at the top of a page. (Thus this option does not affect the effective page length.)
 - ln* Take the length of the page to be *n* lines instead of the default 66.
 - t* Do not print the 5-line header or the 5-line trailer normally supplied for each page.
 - sc* Separate columns by the single character *c* instead of by the appropriate amount of white space. A missing *c* is taken to be a tab.
 - m* Print all *files* simultaneously, each in one column,
- Inter-terminal messages via *write(1)* are forbidden during a *pr*.

FILES

/dev/tty? to suspend messages.

SEE ALSO

cat(1)

DIAGNOSTICS

There are no diagnostics when *pr* is printing on a terminal.

NAME

print - pr to the printer

SYNOPSIS

print file ...

DESCRIPTION

Print pr's a copy of each named file on the printer. It is a one line shell script:

lpr -p \$*

SEE ALSO

lpr(1), pr(1)

NAME

`printenv` - print out the environment

SYNOPSIS

`printenv` [*name*]

DESCRIPTION

Printenv prints out the values of the variables in the environment. If a *name* is specified, only its value is printed.

If a *name* is specified and it is not defined in the environment, *printenv* returns exit status 1, else it returns status 0.

SEE ALSO

`sh(1)`, `environ(7)`, `csh(1)`

NAME

prmail - print out mail in the post office.

SYNOPSIS

prmail [user ...]

DESCRIPTION

Prmail prints the mail which waits for you, or the specified user, in the post office. The mail is not disturbed.

FILES

/usr/spool/mail/* post office

SEE ALSO

biff(1), mail(1), from(1), binmail(1)

NAME

prof - display profile data

SYNOPSIS

prof [-a] [-l] [-n] [-z] [-s] [-v [-low [-high]]] [a.out [mon.out ...]]

DESCRIPTION

Prof interprets the file produced by the *monitor* subroutine. Under default modes, the symbol table in the named object file (*a.out* default) is read and correlated with the profile file (*mon.out* default). For each external symbol, the percentage of time spent executing between that symbol and the next is printed (in decreasing order), together with the number of times that routine was called and the number of milliseconds per call. If more than one profile file is specified, the output represents the sum of the profiles.

In order for the number of calls to a routine to be tallied, the **-p** option of *cc*, *f77* or *pc* must have been given when the file containing the routine was compiled. This option also arranges for the profile file to be produced automatically.

Options are:

- a all symbols are reported rather than just external symbols.
- l the output is sorted by symbol value.
- n the output is sorted by number of calls
- s a summary profile file is produced in *mon.sum*. This is really only useful when more than one profile file is specified.
- v all printing is suppressed and a graphic version of the profile is produced on the standard output for display by the *plot(1)* filters. When plotting, the numbers *low* and *high*, by default 0 and 100, may be given to cause a selected percentage of the profile to be plotted with accordingly higher resolution.
- z routines which have zero usage (as indicated by call counts and accumulated time) are nevertheless printed in the output.

FILES

mon.out for profile
 a.out for namelist
 mon.sum for summary profile

SEE ALSO

monitor(3), profil(2), cc(1), plot(1G)

BUGS

Beware of quantization errors.

Is confused by *f77* which puts the entry points at the bottom of subroutines and functions.

NAME

ps - process status

SYNOPSIS

ps [**a****c****e****g****k****l****s****t****u****v****w****x****#**]

DESCRIPTION

Ps prints information about processes. Normally, only your processes are candidates to be printed by *ps*; specifying **a** causes other users processes to be candidates to be printed; specifying **x** includes processes without control terminals in the candidate pool.

All output formats include, for each process, the process id PID, control terminal of the process TT, cpu time used by the process TIME (this includes both user and system time), the state STAT of the process, and an indication of the COMMAND which is running. The state is given by a sequence of four letters, e.g. "RWNA". The first letter indicates the runnability of the process: R for runnable processes, T for stopped processes, P for processes in page wait, D for those in disk (or other short term) waits, S for those sleeping for less than about 20 seconds, and I for idle (sleeping longer than about 20 seconds) processes. The second letter indicates whether a process is swapped out, showing W if it is, or a blank if it is loaded (in-core); a process which has specified a soft limit on memory requirements and which is exceeding that limit shows >; such a process is (necessarily) not swapped. The third letter indicates whether a process is running with altered CPU scheduling priority (nice); if the process priority is reduced, an N is shown, if the process priority has been artificially raised then a '<' is shown; processes running without special treatment have just a blank. The final letter is not used in this version.

Here are the options:

- a** asks for information about all processes with terminals (ordinarily only one's own processes are displayed).
- c** prints the command name, as stored internally in the system for purposes of accounting, rather than the command arguments, which are kept in the process' address space. This is more reliable, if less informative, since the process is free to destroy the latter information.
- e** Asks for the environment to be printed as well as the arguments to the command.
- g** Asks for all processes. Without this option, *ps* only prints "interesting" processes. Processes are deemed to be uninteresting if they are process group leaders. This normally eliminates top-level command interpreters and processes waiting for users to login on free terminals.
- k** causes the file */vmcore* is used in place of */dev/kmem* and */dev/mem*. This is used for postmortem system debugging.
- l** asks for a long listing, with fields PPID, CP, PRI, NI, ADDR, SIZE, RSS and WCHAN as described below.
- s** Adds the size SSIZ of the kernel stack of each process (for use by system maintainers) to the basic output format.
- tx** restricts output to processes whose controlling tty is *x* (which should be specified as printed by *ps*, e.g. *t3* for tty3, *tco* for console, *td0* for ttyd0, *t?* for processes with no tty, *t* for processes at the current tty, etc). This option must be the last one given.
- u** A user oriented output is produced. This includes fields USER, %CPU, NICE, SIZE, and RSS as described below.

- v A version of the output containing virtual memory statistics is output. This includes fields RE, SL, PAGEIN, SIZE, RSS, LIM, TSIZ, TRS, %CPU and %MEM, described below.
- w Use a wide output format (132 columns rather than 80); if repeated, e.g. ww, use arbitrarily wide output. This information is used to decide how much of long commands to print.
- x asks even about processes with no terminal.
- # A process number may be given, (indicated here by #), in which case the output is restricted to that process. This option must also be last.

A second argument is taken to be the file containing the system's namelist. Otherwise, /vunix is used. A third argument tells *ps* where to look for *core* if the *k* option is given, instead of /vmcore. If a fourth argument is given, it is taken to be the name of a swap directory to use instead of the default.

Fields which are not common to all output formats:

USER	name of the owner of the process
%CPU	cpu utilization of the process; this is a decaying average over up to a minute of previous (real) time. Since the time base over which this is computed varies (since processes may be very young) it is possible for the sum of all %CPU fields to exceed 100%.
NICE	(or NI) process scheduling increment (see <i>setpriority(2)</i>)
SIZE	virtual size of the process (in 1024 byte units)
RSS	real memory (resident set) size of the process (in 1024 byte units)
LIM	soft limit on memory used, specified via a call to <i>setrlimit(2)</i> ; if no limit has been specified then shown as <i>xx</i>
TSIZ	size of text (shared program) image
TRS	size of resident (real memory) set of text
%MEM	percentage of real memory used by this process.
RE	residency time of the process (seconds in core)
SL	sleep time of the process (seconds blocked)
PAGEIN	number of disk i/o's resulting from references by the process to pages not loaded in core.
UID	numerical user-id of process owner
PPID	numerical id of parent of process
CP	short-term cpu utilization factor (used in scheduling)
PRI	process priority (non-positive when in non-interruptible wait)
ADDR	swap address of the process
WCHAN	event on which process is waiting (an address in the system).

F flags associated with process as in `<sys/proc.h>`:

SLOAD	0000001 in core
SSYS	0000002 swapper or pager process
SLOCK	0000004 process being swapped out
STRC	0000008 process is being traced
SWTED	0000010 another tracing flag
SULOCK	0000020 user settable lock in core
SOMASK	0000040 restore old mask after signal
SWEXIT	0000080 working on exiting
SVFORK	0000100 process resulted from <i>vfork()</i>
SOWEUPC	0000200 owe process an <i>addupc()</i> at next <i>ast</i>
SLOGIN	0000400 login process (legit child of <i>init</i>)

SSUPER 0000800 process has supervisor state
STHRASH 0001000 is a thrasher
STWIN 0002000 is a twin process
SDELTA 0004000 has delta segments
SDELLOCK 0008000 locked while mucking with segments
SINTR 0010000 is an interactive process
SSEL 0400000 selecting; wakeup/waiting danger

A process that has exited and has a parent, but has not yet been waited for by the parent is marked <defunct>; a process which is blocked trying to exit is marked <exiting>; *Ps* makes an educated guess as to the file name and arguments given when the process was created by examining memory or the swap area. The method is inherently somewhat unreliable and in any event a process is entitled to destroy this information, so the names cannot be counted on too much.

FILES

/vmunix system namelist
/dev/kmem kernel memory
/dev/swap swap directory
/vmcore core file
/dev searched to find tty names

SEE ALSO

kill(1), w(1)

BUGS

Things can change while *ps* is running; the picture it gives is only a close approximation to reality.

NAME

pti – phototypesetter interpreter

SYNOPSIS

pti [file ...]

DESCRIPTION

Pti shows the commands in a stream from the standard output of *troff*(1) using *troff*'s *-t* option, interpreting them as they would act on the typesetter. Horizontal motions shows as counts in internal units and are marked with '<' and '>' indicating left and right motion. Vertical space is called *lead* and is also indicated.

SEE ALSO

troff(1)

BUGS

Too cryptic for normal users, who should use “*troff -a ...*”.

NAME

ptx - permuted index

SYNOPSIS

ptx [option] ... [input [output]]

DESCRIPTION

Ptx generates a permuted index to file *input* on file *output* (standard input and output default). It has three phases: the first does the permutation, generating one line for each keyword in an input line. The keyword is rotated to the front. The permuted file is then sorted. Finally, the sorted lines are rotated so the keyword comes at the middle of the page. *Ptx* produces output in the form:

```
.xx "tail" "before keyword" "keyword and after" "head"
```

where *.xx* may be an *nroff* or *troff*(1) macro for user-defined formatting. The *before keyword* and *keyword and after* fields incorporate as much of the line as will fit around the keyword when it is printed at the middle of the page. *Tail* and *head*, at least one of which is an empty string "", are wrapped-around pieces small enough to fit in the unused space at the opposite end of the line. When original text must be discarded, '/' marks the spot.

The following options can be applied:

- f** Fold upper and lower case letters for sorting.
- t** Prepare the output for the phototypesetter; the default line length is 100 characters.
- w n** Use the next argument, *n*, as the width of the output line. The default line length is 72 characters.
- g n** Use the next argument, *n*, as the number of characters to allow for each gap among the four parts of the line as finally printed. The default gap is 3 characters.
- o only**
Use as keywords only the words given in the *only* file.
- i ignore**
Do not use as keywords any words given in the *ignore* file. If the **-i** and **-o** options are missing, use */usr/lib/eign* as the *ignore* file.
- b break**
Use the characters in the *break* file to separate words. In any case, tab, newline, and space characters are always used as break characters.
- r** Take any leading nonblank characters of each input line to be a reference identifier (as to a page or chapter) separate from the text of the line. Attach that identifier as a 5th field on each output line.

The index for this manual was generated using *ptx*.

FILES

/usr/bin/sort
/usr/lib/eign

BUGS

Line length counts do not account for overstriking or proportional spacing.

NAME

`pwd` - working directory name

SYNOPSIS

`pwd`

DESCRIPTION

Pwd prints the pathname of the working (current) directory.

SEE ALSO

`cd(1)`, `csh(1)`, `getwd(3)`

BUGS

In *csh(1)* the command *dirs* is always faster (although it can give a different answer in the rare case that the current directory or a containing directory was moved after the shell descended into it).

NAME

px - Pascal interpreter

SYNOPSIS

px [obj [argument ...]]

DESCRIPTION

Px interprets the abstract machine code generated by *pi*. The first argument is the file to be interpreted, and defaults to *obj*; remaining arguments are available to the Pascal program using the built-ins *argv* and *argc*. *Px* is also invoked by *pix* when running 'load and go'.

If the program terminates abnormally an error message and a control flow backtrace are printed. The number of statements executed and total execution time are printed after normal termination. The *p* option of *pi* suppresses all of this except the message indicating the cause of abnormal termination.

FILES

obj	default object file
pmon.out	profile data file

SEE ALSO

Berkeley Pascal User's Manual
pi(1), pix(1)

DIAGNOSTICS

Most run-time error messages are self-explanatory. Some of the more unusual ones are:

Reference to an inactive file

A file other than *input* or *output* was used before a call to *reset* or *rewrite*.

Statement count limit exceeded

The limit of 500,000 executed statements (which prevents excessive looping or recursion) has been exceeded.

Bad data found on integer read

Bad data found on real read

Usually, non-numeric input was found for a number. For reals, Pascal requires digits before and after the decimal point so that numbers like '.1' or '21.' evoke the second diagnostic.

panic: *Some message*

Indicates a internal inconsistency detected in *px* probably due to a Pascal system bug.

AUTHORS

Charles B. Haley, William Joy, and Ken Thompson
VAX-11 version by Kirk McKusick

BUGS

Post-mortem traceback is not limited; infinite recursion leads to almost infinite traceback.

NAME

pxp - Pascal execution profiler

SYNOPSIS

pxp [**-acdefjnstuw_**] [**-23456789**] [**-z** [name ...]] name.p

DESCRIPTION

Pxp can be used to obtain execution profiles of Pascal programs or as a pretty-printer. To produce an execution profile all that is necessary is to translate the program specifying the **z** option to *pi* or *pix*, to execute the program, and to then issue the command

```
pxp -z name.p
```

A reformatted listing is output if none of the **c**, **t**, or **z** options are specified; thus

```
pxp old.p > new.p
```

places a pretty-printed version of the program in 'old.p' in the file 'new.p'.

The use of the following options of *pxp* is discussed in sections 2.6, 5.4, 5.5 and 5.10 of the *Berkeley Pascal User's Manual*.

- a** Print the bodies of all procedures and functions in the profile; even those which were never executed.
- c** Extract profile data from the file *core*.
- d** Include declaration parts in a profile.
- e** Eliminate **include** directives when reformatting a file; the **include** is replaced by the reformatted contents of the specified file.
- f** Fully parenthesize expressions.
- j** Left justify all procedures and functions.
- n** Eject a new page as each file is included; in profiles, print a blank line at the top of the page.
- s** Strip comments from the input text.
- t** Print a table summarizing **procedure** and **function** call counts.
- u** Card image mode; only the first 72 characters of input lines are used.
- w** Suppress warning diagnostics.
- z** Generate an execution profile. If no *names*, are given the profile is of the entire program. If a list of names is given, then only any specified **procedures** or **functions** and the contents of any specified **include** files will appear in the profile.
- _** Underline keywords.
- d** With *d* a digit, $2 \leq d \leq 9$, causes *pxp* to use *d* spaces as the basic indenting unit. The default is 4.

FILES

name.p	input file
name.i	include file(s)
pmon.out	profile data
core	profile data source with -c
/usr/lib/how_pxp	information on basic usage

SEE ALSO

Berkeley Pascal User's Manual
pi(1), px(1)

DIAGNOSTICS

For a basic explanation do

pxp

Error diagnostics include 'No profile data in file' with the **c** option if the **z** option was not enabled to *pi*; 'Not a Pascal system core file' if the core is not from a *px* execution; 'Program and count data do not correspond' if the program was changed after compilation, before profiling; or if the wrong program is specified.

AUTHOR

William Joy

BUGS

Does not place multiple statements per line.

NAME

pxref - Pascal cross-reference program

SYNOPSIS

pxref [-] name

DESCRIPTION

Pxref makes a line numbered listing and a cross-reference of identifier usage for the program in *name*. The optional '-' argument suppresses the listing. The keywords **goto** and **label** are treated as identifiers for the purpose of the cross-reference. **Include** directives are not processed, but cause the placement of an entry indexed by '#include' in the cross-reference.

SEE ALSO

Berkeley Pascal User's Manual

AUTHOR

Niklaus Wirth

BUGS

Identifiers are trimmed to 10 characters.

NAME

`quota` - display disc usage and limits

SYNOPSIS

`quota` [`-qv`] [`user`]

DESCRIPTION

Quota displays users' disc usage and limits. Only the super-user may use the optional *user* argument to view the limits of users other than himself.

The `-q` flag prints a more terse message, containing only information on file systems where usage is over quota.

If a `-v` flag is supplied, *quota* will also display user's quotas on file systems where no storage is allocated.

Quota reports only on file systems which have disc quotas. If *quota* exits with a non-zero status, one or more file systems are over quota.

SEE ALSO

`quota(2)`, `quotaon(8)`

NAME

ranlib - convert archives to random libraries

SYNOPSIS

ranlib archive ...

DESCRIPTION

Ranlib converts each *archive* to a form which the loader can load more rapidly. *Ranlib* does this by adding a table of contents called **__SYMDEF** to the beginning of the archive. *Ranlib* uses *ar(1)* to reconstruct the archive, so that sufficient temporary file space must be available in the file system which contains the current directory.

SEE ALSO

ld(1), ar(1), lorder(1)

BUGS

Because generation of a library by *ar* and randomization of the library by *ranlib* are separate processes, phase errors are possible. The loader, *ld*, warns when the modification date of a library is more recent than the creation date of its dictionary; but this means that you get the warning even if you only copy the library.

NAME

ratfor – rational Fortran dialect

SYNOPSIS

ratfor [option ...] [filename ...]

DESCRIPTION

Ratfor converts a rational dialect of Fortran into ordinary irrational Fortran. *Ratfor* provides control flow constructs essentially identical to those in C:

statement grouping:

```
{ statement; statement; statement }
```

decision-making:

```
if (condition) statement [ else statement ]
switch (integer value) {
    case integer: statement
    ...
    [ default: ] statement
}
```

loops: while (condition) statement
 for (expression; condition; expression) statement
 do limits statement
 repeat statement [until (condition)]
 break
 next

and some syntactic sugar to make programs easier to read and write:

free form input:

multiple statements/line; automatic continuation

comments:

```
# this is a comment
```

translation of relationals:

>, >=, etc., become .GT., .GE., etc.

return (expression)

returns expression to caller from function

define: define name replacement

include:

```
include filename
```

Ratfor is best used with *f77(1)*.

SEE ALSO

f77(1)

B. W. Kernighan and P. J. Plauger, *Software Tools*, Addison-Wesley, 1976.

NAME

rcp - remote file copy

SYNOPSIS

```
rcp file1 file2
rcp [-r ] file ... directory
```

DESCRIPTION

Rcp copies files between machines. Each *file* or *directory* argument is either a remote file name of the form "rhost:path", or a local file name (containing no ':' characters, or a '/' before any '.'s.)

If the *-r* is specified and any of the source files are directories, *rcp* copies each subtree rooted at that name; in this case the destination must be a directory.

If *path* is not a full path name, it is interpreted relative to your login directory on *rhost*. A *path* on a remote host may be quoted (using \, ", or `) so that the metacharacters are interpreted remotely.

Rcp does not prompt for passwords; your current local user name must exist on *rhost* and allow remote command execution via *rsh*(1C).

Rcp handles third party copies, where neither source nor target files are on the current machine. Hostnames may also take the form "rhost.rname" to use *rname* rather than the current user name on the remote host.

SEE ALSO

ftp(1C), rsh(1C), rlogin(1C)

BUGS

Doesn't detect all cases where the target of a copy might be a file in cases where only a directory should be legal.

Is confused by any output generated by commands in a .login, .profile, or .cshrc file on the remote host.

NAME

`rcs` - change RCS file attributes

SYNOPSIS

`rcs` [options] file ...

DESCRIPTION

Rcs creates new RCS files or changes attributes of existing ones. An RCS file contains multiple revisions of text, an access list, a change log, descriptive text, and some control attributes. For *rcs* to work, the caller's login name must be on the access list, except if the access list is empty, the caller is the owner of the file or the superuser, or the `-i` option is present.

Files ending in `,v` are RCS files, all others are working files. If a working file is given, *rcs* tries to find the corresponding RCS file first in directory `./RCS` and then in the current directory, as explained in *co* (1).

- `-i` creates and initializes a new RCS file, but does not deposit any revision. If the RCS file has no path prefix, *rcs* tries to place it first into the subdirectory `./RCS`, and then into the current directory. If the RCS file already exists, an error message is printed.
- `-alogins` appends the login names appearing in the comma-separated list *logins* to the access list of the RCS file.
- `-Aoldfile` appends the access list of *oldfile* to the access list of the RCS file.
- `-e[logins]` erases the login names appearing in the comma-separated list *logins* from the access list of the RCS file. If *logins* is omitted, the entire access list is erased.
- `-cstring` sets the comment leader to *string*. The comment leader is printed before every log message line generated by the keyword `Log` during checkout (see *co*). This is useful for programming languages without multi-line comments. During *rcs -i* or initial *ci*, the comment leader is guessed from the suffix of the working file.
- `-l[rev]` locks the revision with number *rev*. If a branch is given, the latest revision on that branch is locked. If *rev* is omitted, the latest revision on the trunk is locked. Locking prevents overlapping changes. A lock is removed with *ci* or *rcs -u* (see below).
- `-u[rev]` unlocks the revision with number *rev*. If a branch is given, the latest revision on that branch is unlocked. If *rev* is omitted, the latest lock held by the caller is removed. Normally, only the locker of a revision may unlock it. Somebody else unlocking a revision breaks the lock. This causes a mail message to be sent to the original locker. The message contains a commentary solicited from the breaker. The commentary is terminated with a line containing a single `'` or control-D.
- `-L` sets locking to *strict*. Strict locking means that the owner of an RCS file is not exempt from locking for checkin. This option should be used for files that are shared.
- `-U` sets locking to non-strict. Non-strict locking means that the owner of a file need not lock a revision for checkin. This option should NOT be used for files that are shared. The default (`-L` or `-U`) is determined by your system administrator.
- `-nname[:rev]` associates the symbolic name *name* with the branch or revision *rev*. *Rcs* prints an error message if *name* is already associated with another number. If *rev* is omitted, the symbolic name is deleted.

- Nname[:rev]** same as **-n**, except that it overrides a previous assignment of *name*.
- orange** deletes ("outdates") the revisions given by *range*. A range consisting of a single revision number means that revision. A range consisting of a branch number means the latest revision on that branch. A range of the form *rev1-rev2* means revisions *rev1* to *rev2* on the same branch, **-rev** means from the beginning of the branch containing *rev* up to and including *rev*, and *rev-* means from revision *rev* to the end of the branch containing *rev*. None of the outdated revisions may have branches or locks.
- q** quiet mode; diagnostics are not printed.
- sstate[:rev]** sets the state attribute of the revision *rev* to *state*. If *rev* is omitted, the latest revision on the trunk is assumed; If *rev* is a branch number, the latest revision on that branch is assumed. Any identifier is acceptable for *state*. A useful set of states is *Exp* (for experimental), *Stab* (for stable), and *Rel* (for released). By default, *ci* sets the state of a revision to *Exp*.
- t[txtfile]** writes descriptive text into the RCS file (deletes the existing text). If *txtfile* is omitted, *rcs* prompts the user for text supplied from the std. input, terminated with a line containing a single '.' or control-D. Otherwise, the descriptive text is copied from the file *txtfile*. If the **-i** option is present, descriptive text is requested even if **-t** is not given. The prompt is suppressed if the std. input is not a terminal.

DIAGNOSTICS

The RCS file name and the revisions outdated are written to the diagnostic output. The exit status always refers to the last RCS file operated upon, and is 0 if the operation was successful, 1 otherwise.

FILES

The caller of the command must have read/write permission for the directory containing the RCS file and read permission for the RCS file itself. *Rcs* creates a semaphore file in the same directory as the RCS file to prevent simultaneous update. For changes, *rcs* always creates a new file. On successful completion, *rcs* deletes the old one and renames the new one. This strategy makes links to RCS files useless.

IDENTIFICATION

Author: Walter F. Tichy, Purdue University, West Lafayette, IN, 47907.
 Revision Number: 3.1 ; Release Date: 83/04/04 .
 Copyright © 1982 by Walter F. Tichy.

SEE ALSO

co (1), *ci* (1), *ident*(1), *rcsdiff* (1), *rcsintro* (1), *rcsmerge* (1), *rlog* (1), *rcsfile* (5), *scsstorcs* (8).
 Walter F. Tichy, "Design, Implementation, and Evaluation of a Revision Control System," in *Proceedings of the 6th International Conference on Software Engineering*, IEEE, Tokyo, Sept. 1982.

BUGS

NAME

`rcsdiff` - compare RCS revisions

SYNOPSIS

`rcsdiff` [`-biwt`] [`-cefh`] [`-rrev1`] [`-rrev2`] file ...

DESCRIPTION

Rcsdiff runs *diff* (1) to compare two revisions of each RCS file given. A file name ending in `,v` is an RCS file name, otherwise a working file name. *Rcsdiff* derives the working file name from the RCS file name and vice versa, as explained in *co* (1). Pairs consisting of both an RCS and a working file name may also be specified.

All options except `-r` have the same effect as described in *diff*(1).

If both *rev1* and *rev2* are omitted, *rcsdiff* compares the latest revision on the trunk with the contents of the corresponding working file. This is useful for determining what you changed since the last checkin.

If *rev1* is given, but *rev2* is omitted, *rcsdiff* compares revision *rev1* of the RCS file with the contents of the corresponding working file.

If both *rev1* and *rev2* are given, *rcsdiff* compares revisions *rev1* and *rev2* of the RCS file.

Both *rev1* and *rev2* may be given numerically or symbolically.

EXAMPLES

The command

```
rcsdiff f.c
```

runs *diff* on the latest trunk revision of RCS file `f.c,v` and the contents of working file `f.c`.

IDENTIFICATION

Author: Walter F. Tichy, Purdue University, West Lafayette, IN, 47907.

Revision Number: 1.2 ; Release Date: 86/05/19 .

Copyright © 1982 by Walter F. Tichy.

SEE ALSO

ci (1), *co* (1), *diff* (1), *ident* (1), *rcs* (1), *rcsintro* (1), *rcsmerge* (1), *rlog* (1), *rcsfile* (5).

Walter F. Tichy, "Design, Implementation, and Evaluation of a Revision Control System," in *Proceedings of the 6th International Conference on Software Engineering, IEEE, Tokyo, Sept. 1982*.

BUGS

NAME

rcsintro - introduction to RCS commands

DESCRIPTION

The Revision Control System (RCS) manages multiple revisions of text files. RCS automates the storing, retrieval, logging, identification, and merging of revisions. RCS is useful for text that is revised frequently, for example programs, documentation, graphics, papers, form letters, etc.

The basic user interface is extremely simple. The novice only needs to learn two commands: *ci* and *co*. *Ci*, short for "checkin", deposits the contents of a text file into an archival file called an RCS file. An RCS file contains all revisions of a particular text file. *Co*, short for "checkout", retrieves revisions from an RCS file.

SEE ALSO

ci(1), *co*(1), *ident*(1), *merge*(1), *rcs*(1), *rcsdiff*(1), *rcsmerge*(1), *rlog*(1), *rcsfile*(5).

Walter F. Tichy, "An Introduction to the Revision Control System", Programmer Supplementary Documents, Volume 1 (PS1), #13

NAME

rcsmerge - merge RCS revisions

SYNOPSIS

rcsmerge -rrev1 [-rrev2] [-p] file

DESCRIPTION

Rcsmerge incorporates the changes between *rev1* and *rev2* of an RCS file into the corresponding working file. If **-p** is given, the result is printed on the std. output, otherwise the result overwrites the working file.

A file name ending in *'v'* is an RCS file name, otherwise a working file name. *Merge* derives the working file name from the RCS file name and vice versa, as explained in *co* (1). A pair consisting of both an RCS and a working file name may also be specified.

Rev1 may not be omitted. If *rev2* is omitted, the latest revision on the trunk is assumed. Both *rev1* and *rev2* may be given numerically or symbolically.

Rcsmerge prints a warning if there are overlaps, and delimits the overlapping regions as explained in *co -j*. The command is useful for incorporating changes into a checked-out revision.

EXAMPLES

Suppose you have released revision 2.8 of *f.c*. Assume furthermore that you just completed revision 3.4, when you receive updates to release 2.8 from someone else. To combine the updates to 2.8 and your changes between 2.8 and 3.4, put the updates to 2.8 into file *f.c* and execute

```
rcsmerge -p -r2.8 -r3.4 f.c >f.merged.c
```

Then examine *f.merged.c*. Alternatively, if you want to save the updates to 2.8 in the RCS file, check them in as revision 2.8.1.1 and execute *co -j*:

```
ci -r2.8.1.1 f.c
co -r3.4 -j2.8:2.8.1.1 f.c
```

As another example, the following command undoes the changes between revision 2.4 and 2.8 in your currently checked out revision in *f.c*.

```
rcsmerge -r2.8 -r2.4 f.c
```

Note the order of the arguments, and that *f.c* will be overwritten.

IDENTIFICATION

Author: Walter F. Tichy, Purdue University, West Lafayette, IN, 47907.

Revision Number: 3.0 ; Release Date: 83/01/15 .

Copyright © 1982 by Walter F. Tichy.

SEE ALSO

ci (1), *co* (1), *merge* (1), *ident* (1), *rcs* (1), *rcsdiff* (1), *rlog* (1), *rcsfile* (5).

Walter F. Tichy, "Design, Implementation, and Evaluation of a Revision Control System," in *Proceedings of the 6th International Conference on Software Engineering*, IEEE, Tokyo, Sept. 1982.

BUGS

Rcsmerge does not work for files that contain lines with a single '.'.

NAME

rdist - remote file distribution program

SYNOPSIS

rdist [*-nqbRhivwy*] [*-f distfile*] [*-d var=value*] [*-m host*] [*name ...*]

rdist [*-nqbRhivwy*] *-c name ...* [*login@*]*host[:dest]*

DESCRIPTION

Rdist is a program to maintain identical copies of files over multiple hosts. It preserves the owner, group, mode, and mtime of files if possible and can update programs that are executing. *Rdist* reads commands from *distfile* to direct the updating of files and/or directories. If *distfile* is '-', the standard input is used. If no *-f* option is present, the program looks first for 'distfile', then 'Distfile' to use as the input. If no names are specified on the command line, *rdist* will update all of the files and directories listed in *distfile*. Otherwise, the argument is taken to be the name of a file to be updated or the label of a command to execute. If label and file names conflict, it is assumed to be a label. These may be used together to update specific files using specific commands.

The *-c* option forces *rdist* to interpret the remaining arguments as a small *distfile*. The equivalent *distfile* is as follows.

```
( name ... ) -> [login@]host
                install [dest];
```

Other options:

- d** Define *var* to have *value*. The *-d* option is used to define or override variable definitions in the *distfile*. *Value* can be the empty string, one name, or a list of names surrounded by parentheses and separated by tabs and/or spaces.
- m** Limit which machines are to be updated. Multiple *-m* arguments can be given to limit updates to a subset of the hosts listed the *distfile*.
- n** Print the commands without executing them. This option is useful for debugging *distfile*.
- q** Quiet mode. Files that are being modified are normally printed on standard output. The *-q* option suppresses this.
- R** Remove extraneous files. If a directory is being updated, any files that exist on the remote host that do not exist in the master directory are removed. This is useful for maintaining truly identical copies of directories.
- h** Follow symbolic links. Copy the file that the link points to rather than the link itself.
- i** Ignore unresolved links. *Rdist* will normally try to maintain the link structure of files being transferred and warn the user if all the links cannot be found.
- v** Verify that the files are up to date on all the hosts. Any files that are out of date will be displayed but no files will be changed nor any mail sent.
- w** Whole mode. The whole file name is appended to the destination directory name. Normally, only the last component of a name is used when renaming files. This will preserve the directory structure of the files being copied instead of flattening the directory structure. For example, renaming a list of files such as (*dir1/f1 dir2/f2*) to *dir3* would create files *dir3/dir1/f1* and *dir3/dir2/f2* instead of *dir3/f1* and *dir3/f2*.

- y Younger mode. Files are normally updated if their *mtime* and *size* (see *stat(2)*) disagree. The **-y** option causes *rdist* not to update files that are younger than the master copy. This can be used to prevent newer copies on other hosts from being replaced. A warning message is printed for files which are newer than the master copy.
- b Binary comparison. Perform a binary comparison and update files if they differ rather than comparing dates and sizes.

Distfile contains a sequence of entries that specify the files to be copied, the destination hosts, and what operations to perform to do the updating. Each entry has one of the following formats.

```
<variable name> '=' <name list>
[ label: ] <source list> '->' <destination list> <command list>
[ label: ] <source list> '::' <time_stamp file> <command list>
```

The first format is used for defining variables. The second format is used for distributing files to other hosts. The third format is used for making lists of files that have been changed since some given date. The *source list* specifies a list of files and/or directories on the local host which are to be used as the master copy for distribution. The *destination list* is the list of hosts to which these files are to be copied. Each file in the source list is added to a list of changes if the file is out of date on the host which is being updated (second format) or the file is newer than the time stamp file (third format).

Labels are optional. They are used to identify a command for partial updates.

Newlines, tabs, and blanks are only used as separators and are otherwise ignored. Comments begin with **#** and end with a newline.

Variables to be expanded begin with **\$** followed by one character or a name enclosed in curly braces (see the examples at the end).

The source and destination lists have the following format:

```
<name>
or
(' <zero or more names separated by white-space> ')
```

The shell meta-characters **[**, **]**, **{**, **}**, *****, and **?** are recognized and expanded (on the local host only) in the same way as *cs**h*(1). They can be escaped with a backslash. The **~** character is also expanded in the same way as *cs**h* but is expanded separately on the local and destination hosts. When the **-w** option is used with a file name that begins with **~**, everything except the home directory is appended to the destination name. File names which do not begin with **/** or **~** use the destination user's home directory as the root directory for the rest of the file name.

The command list consists of zero or more commands of the following format.

```
'install' <options> opt_dest_name ';'
'notify' <name list> ';'
'except' <name list> ';'
'except_pat' <pattern list> ';'
'special' <name list> string ';'

```

The *install* command is used to copy out of date files and/or directories. Each source file is copied to each host in the destination list. Directories are recursively copied in the same way. *Opt_dest_name* is an optional parameter to rename files. If no *install* command appears in the command list or the destination name is not specified, the source file name is used. Directories in the path name will be created if they do not exist on the remote host. To help prevent disasters, a non-empty directory on a target host will never be replaced with a regular file or a symbolic link. However, under the '-R' option a non-empty directory will be removed if the corresponding filename is completely absent on the master host. The *options* are '-R', '-h', '-i', '-v', '-w', '-y', and '-b' and have the same semantics as options on the command line except they only apply to the files in the source list. The login name used on the destination host is the same as the local host unless the destination name is of the format "login@host".

The *notify* command is used to mail the list of files updated (and any errors that may have occurred) to the listed names. If no '@' appears in the name, the destination host is appended to the name (e.g., name1@host, name2@host, ...).

The *except* command is used to update all of the files in the source list **except** for the files listed in *name list*. This is usually used to copy everything in a directory except certain files.

The *except_pat* command is like the *except* command except that *pattern list* is a list of regular expressions (see *ed(1)* for details). If one of the patterns matches some string within a file name, that file will be ignored. Note that since '\' is a quote character, it must be doubled to become part of the regular expression. Variables are expanded in *pattern list* but not shell file pattern matching characters. To include a '\$', it must be escaped with '\\$'.

The *special* command is used to specify *sh(1)* commands that are to be executed on the remote host after the file in *name list* is updated or installed. If the *name list* is omitted then the shell commands will be executed for every file updated or installed. The shell variable 'FILE' is set to the current filename before executing the commands in *string*. *String* starts and ends with '"' and can cross multiple lines in *distfile*. Multiple commands to the shell should be separated by ';'. Commands are executed in the user's home directory on the host being updated. The *special* command can be used to rebuild private databases, etc. after a program has been updated.

The following is a small example.

```
HOSTS = ( matisse root@arpa)

FILES = ( /bin /lib /usr/bin /usr/games
          /usr/include/{*.h,{stand,sys,vax*,pascal,machine}/*.h}
          /usr/lib /usr/man/man? /usr/ucb /usr/local/rdist )

EXLIB = ( Mail.rc aliases aliases.dir aliases.pag crontab dshrc
          sendmail.cf sendmail.fc sendmail.hf sendmail.st uucp vfont )

${FILES} -> ${HOSTS}
install -R ;
except /usr/lib/${EXLIB} ;
except /usr/games/lib ;
special /usr/lib/sendmail "/usr/lib/sendmail -bz" ;

srcs:
/usr/src/bin -> arpa
except_pat ( \\.\o\$ /SCCS\$ ) ;
```

```
IMAGEN = (ips dviimp catdvi)
```

```
imagen:  
/usr/local/${IMAGEN} -> arpa  
install /usr/local/lib ;  
notify ralph ;
```

```
${FILES} :: stamp.cory  
notify root@cory ;
```

FILES

distfile	input command file
/tmp/rdist*	temporary file for update lists

SEE ALSO

sh(1), csh(1), stat(2)

DIAGNOSTICS

A complaint about mismatch of rdist version numbers may really stem from some problem with starting your shell, e.g., you are in too many groups.

BUGS

Source files must reside on the local host where rdist is executed.

There is no easy way to have a special command executed after all files in a directory have been updated.

Variable expansion only works for name lists; there should be a general macro facility.

Rdist aborts on files which have a negative mtime (before Jan 1, 1970).

There should be a 'force' option to allow replacement of non-empty directories by regular files or symlinks. A means of updating file modes and owners of otherwise identical files is also needed.

NAME

refer - find and insert literature references in documents

SYNOPSIS

```
refer [-a] [-b] [-c] [-e] [-fn] [-kx] [-lm,n] [-n] [-p bib] [-skeys] [-Bl.m] [-P] [-S] [file ...]
```

DESCRIPTION

Refer is a preprocessor for *nroff* or *troff*(1) that finds and formats references for footnotes or endnotes. It is also the base for a series of programs designed to index, search, sort, and print stand-alone bibliographies, or other data entered in the appropriate form.

Given an incomplete citation with sufficiently precise keywords, *refer* will search a bibliographic database for references containing these keywords anywhere in the title, author, journal, etc. The input file (or standard input) is copied to standard output, except for lines between `.[` and `.]` delimiters, which are assumed to contain keywords, and are replaced by information from the bibliographic database. The user may also search different databases, override particular fields, or add new fields. The reference data, from whatever source, are assigned to a set of *troff* strings. Macro packages such as *ms*(7) print the finished reference text from these strings. By default references are flagged by footnote numbers.

The following options are available:

- ar Reverse the first *n* author names (Jones, J. A. instead of J. A. Jones). If *n* is omitted all author names are reversed.
- b Bare mode: do not put any flags in text (neither numbers nor labels).
- ckeys Capitalize (with CAPS SMALL CAPS) the fields whose key-letters are in *keys*.
- e Instead of leaving the references where encountered, accumulate them until a sequence of the form


```
.[
  $LIST$
.]
```

 is encountered, and then write out all references collected so far. Collapse references to same source.
- fn Set the footnote number to *n* instead of the default of 1 (one). With labels rather than numbers, this flag is a no-op.
- kx Instead of numbering references, use labels as specified in a reference data line beginning $\%x$; by default *x* is L.
- lm,n Instead of numbering references, use labels made from the senior author's last name and the year of publication. Only the first *m* letters of the last name and the last *n* digits of the date are used. If either *m* or *n* is omitted the entire name or date respectively is used.
- n Do not search the default file `/usr/dict/papers/Ind`. If there is a REFER environment variable, the specified file will be searched instead of the default file; in this case the `-n` flag has no effect.
- p *bib* Take the next argument *bib* as a file of references to be searched. The default file is searched last.

-skeys

Sort references by fields whose key-letters are in the *keys* string; permute reference numbers in text accordingly. Implies **-e**. The key-letters in *keys* may be followed by a number to indicate how many such fields are used, with **+** taken as a very large number. The default is **AD** which sorts on the senior author and then date; to sort, for example, on all authors and then title use **-sA+T**.

-Bl.m Bibliography mode. Take a file composed of records separated by blank lines, and turn them into *troff* input. Label *l* will be turned into the macro *.m* with *l* defaulting to **%X** and *.m* defaulting to **.AP** (annotation paragraph).

-P Place punctuation marks **.,;?!** after the reference signal, rather than before. (Periods and commas used to be done with strings.)

-S Produce references in the Natural or Social Science format.

To use your own references, put them in the format described below. They can be searched more rapidly by running *indxib(1)* on them before using *refer*; failure to index results in a linear search. When *refer* is used with the *eqn*, *neqn* or *tbl* preprocessors *refer* should be first, to minimize the volume of data passed through pipes.

The *refer* preprocessor and associated programs expect input from a file of references composed of records separated by blank lines. A record is a set of lines (fields), each containing one kind of information. Fields start on a line beginning with a "**%**", followed by a key-letter, then a blank, and finally the contents of the field, and continue until the next line starting with "**%**". The output ordering and formatting of fields is controlled by the macros specified for *nroff/troff* (for footnotes and endnotes) or *roffbib* (for stand-alone bibliographies). For a list of the most common key-letters and their corresponding fields, see *addbib(1)*. An example of a *refer* entry is given below.

EXAMPLE

```
%A M. E. Lesk
%T Some Applications of Inverted Indexes on the UNIX System
%B UNIX Programmer's Manual
%V 2b
%I Bell Laboratories
%C Murray Hill, NJ
%D 1978
```

FILES

```
/usr/dict/papers  directory of default publication lists
/usr/lib/refer    directory of companion programs
```

SEE ALSO

addbib(1), *sortbib(1)*, *roffbib(1)*, *indxib(1)*, *lookbib(1)*

AUTHOR

Mike Lesk

BUGS

Blank spaces at the end of lines in bibliography fields will cause the records to sort and reverse incorrectly. Sorting large numbers of references causes a core dump.

NAME

reset - reset the teletype bits to a sensible state

SYNOPSIS

reset

DESCRIPTION

Reset sets the terminal to cooked mode, turns off cbreak and raw modes, turns on nl, and restores special characters that are undefined to their default values.

This is most useful after a program dies leaving a terminal in a funny state; you have to type "<LF>reset<LF>" to get it to work then to the shell, as <CR> often doesn't work; often none of this will echo.

It is a good idea to follow *reset* with *tset*(1)

SEE ALSO

stty(1), tset(1)

BUGS

Doesn't set tabs properly; it can't intuit personal choices for interrupt and line kill characters, so it leaves these set to the local system standards.

NAME

rev - reverse lines of a file

SYNOPSIS

rev [file] ...

DESCRIPTION

Rev copies the named files to the standard output, reversing the order of characters in every line. If no file is specified, the standard input is copied.

NAME

rlogin - remote login

SYNOPSIS

```
rlogin rhost [ -e c ] [ -l username ]  
rhost [ -ec ] [ -l username ]
```

DESCRIPTION

Rlogin connects your terminal on the current local host system *lhost* to the remote host system *rhost*.

Each host has a file */etc/hosts.equiv* which contains a list of *rhost*'s with which it shares account names. (The host names must be the standard names as described in *rsh(1C)*.) When you *rlogin* as the same user on an equivalent host, you don't need to give a password. Each user may also have a private equivalence list in a file *.rhosts* in his login directory. Each line in this file should contain a *rhost* and a *username* separated by a space, giving additional cases where logins without passwords are to be permitted. If the originating user is not equivalent to the remote user, then a login and password will be prompted for on the remote machine as in *login(1)*. To avoid some security problems, the *.rhosts* file must be owned by either the remote user or root and may not be a symbolic link.

Your remote terminal type is the same as your local terminal type (as given in your environment *TERM* variable). All echoing takes place at the remote site, so that (except for delays) the *rlogin* is transparent. Flow control via *^S* and *^Q* and flushing of input and output on interrupts are handled properly. A line of the form "*~*." disconnects from the remote host, where "*~*" is the escape character. A different escape character may be specified by the *-e* option. There is no space separating this option flag and the argument character.

SEE ALSO

rsh(1C)

FILES

*/usr/hosts/** for *rhost* version of the command

BUGS

More terminal characteristics should be propagated.

NAME

`rm`, `rmdir` - remove (unlink) files or directories

SYNOPSIS

`rm` [`-f`] [`-r`] [`-i`] [`-`] file ...

`rmdir` dir ...

DESCRIPTION

Rm removes the entries for one or more files from a directory. If an entry was the last link to the file, the file is destroyed. Removal of a file requires write permission in its directory, but neither read nor write permission on the file itself.

If a file has no write permission and the standard input is a terminal, its permissions are printed and a line is read from the standard input. If that line begins with 'y' the file is deleted, otherwise the file remains. No questions are asked and no errors are reported when the `-f` (force) option is given.

If a designated file is a directory, an error comment is printed unless the optional argument `-r` has been used. In that case, *rm* recursively deletes the entire contents of the specified directory, and the directory itself.

If the `-i` (interactive) option is in effect, *rm* asks whether to delete each file, and, under `-r`, whether to examine each directory.

The null option `-` indicates that all the arguments following it are to be treated as file names. This allows the specification of file names starting with a minus.

Rmdir removes entries for the named directories, which must be empty.

SEE ALSO

`rm(1)`, `unlink(2)`, `rmdir(2)`

NAME

rmail - handle remote mail received via *uucp*

SYNOPSIS

rmail user ...

DESCRIPTION

Rmail interprets incoming mail received via *uucp*(1C), collapsing "From" lines in the form generated by *binmail*(1) into a single line of the form "return-path!sender", and passing the processed mail on to *sendmail*(8).

Rmail is explicitly designed for use with *uucp* and *sendmail*.

SEE ALSO

binmail(1), *uucp*(1C), *sendmail*(8)

BUGS

Rmail should not reside in /bin.

NAME

`rmdir`, `rm` - remove (unlink) directories or files

SYNOPSIS

`rmdir` dir ...

`rm` [-f] [-r] [-i] [-] file ...

DESCRIPTION

Rmdir removes entries for the named directories, which must be empty.

Rm removes the entries for one or more files from a directory. If an entry was the last link to the file, the file is destroyed. Removal of a file requires write permission in its directory, but neither read nor write permission on the file itself.

If a file has no write permission and the standard input is a terminal, its permissions are printed and a line is read from the standard input. If that line begins with 'y' the file is deleted, otherwise the file remains. No questions are asked and no errors are reported when the `-f` (force) option is given.

If a designated file is a directory, an error comment is printed unless the optional argument `-r` has been used. In that case, *rm* recursively deletes the entire contents of the specified directory, and the directory itself.

If the `-i` (interactive) option is in effect, *rm* asks whether to delete each file, and, under `-r`, whether to examine each directory.

The null option `-` indicates that all the arguments following it are to be treated as file names. This allows the specification of file names starting with a minus.

SEE ALSO

`rm(1)`, `unlink(2)`, `rmdir(2)`

NAME

roffbib - run off bibliographic database

SYNOPSIS

roffbib [-e] [-h] [-n] [-o] [-r] [-s] [-T*term*] [-x] [-m *mac*] [-V] [-Q] [file ...]

DESCRIPTION

Roffbib prints out all records in a bibliographic database, in bibliography format rather than as footnotes or endnotes. Generally it is used in conjunction with *sortbib*:

```
sortbib database | roffbib
```

Roffbib accepts most of the options understood by *nroff*(1), most importantly the **-T** flag to specify terminal type.

If abstracts or comments are entered following the %X field key, *roffbib* will format them into paragraphs for an annotated bibliography. Several %X fields may be given if several annotation paragraphs are desired. The **-x** flag will suppress the printing of these abstracts.

A user-defined set of macros may be specified after the **-m** option. There should be a space between the **-m** and the macro filename. This set of macros will replace the ones defined in /usr/lib/tmac/tmac.bib. The **-V** flag will send output to the Versatec; the **-Q** flag will queue output for the phototypesetter.

Four command-line registers control formatting style of the bibliography, much like the number registers of *ms*(7). The command-line argument **-rN1** will number the references starting at one (1). The flag **-rV2** will double space the bibliography, while **-rV1** will double space references but single space annotation paragraphs. The line length can be changed from the default 6.5 inches to 6 inches with the **-rL6i** argument, and the page offset can be set from the default of 0 to one inch by specifying **-rO1i** (capital O, not zero). Note: with the **-V** and **-Q** flags the default page offset is already one inch.

FILES

/usr/lib/tmac/tmac.bib file of macros used by *nroff*/*troff*

SEE ALSO

refer(1), addbib(1), sortbib(1), indxbib(1), lookbib(1)

AUTHORS

Greg Shenaut, Bill Tuthill

BUGS

Users have to rewrite macros to create customized formats.

NAME

rsh - remote shell

SYNOPSIS

```
rsh host [-l username] [-n] command
host [-l username] [-n] command
```

DESCRIPTION

Rsh connects to the specified *host*, and executes the specified *command*. *Rsh* copies its standard input to the remote command, the standard output of the remote command to its standard output, and the standard error of the remote command to its standard error. Interrupt, quit and terminate signals are propagated to the remote command; *rsh* normally terminates when the remote command does.

The remote username used is the same as your local username, unless you specify a different remote name with the *-l* option. This remote name must be equivalent (in the sense of *rlogin*(1C)) to the originating account; no provision is made for specifying a password with a command.

If you omit *command*, then instead of executing a single command, you will be logged in on the remote host using *rlogin*(1C).

Shell metacharacters which are not quoted are interpreted on local machine, while quoted metacharacters are interpreted on the remote machine. Thus the command

```
rsh otherhost cat remotefile >> localfile
```

appends the remote file *remotefile* to the localfile *localfile*, while

```
rsh otherhost cat remotefile ">>" otherremotefile
```

appends *remotefile* to *otherremotefile*.

Host names are given in the file */etc/hosts*. Each host has one standard name (the first name given in the file), which is rather long and unambiguous, and optionally one or more nicknames. The host names for local machines are also commands in the directory */usr/hosts*: if you put this directory in your search path then the **rsh** can be omitted.

FILES

```
/etc/hosts
/usr/hosts/*
```

SEE ALSO

rlogin(1C)

BUGS

If you are using *cs*(1) and put a *rsh*(1C) in the background without redirecting its input away from the terminal, it will block even if no reads are posted by the remote command. If no input is desired you should redirect the input of *rsh* to */dev/null* using the *-n* option.

You cannot run an interactive command (like *rogue*(6) or *vi*(1)); use *rlogin*(1C).

Stop signals stop the local *rsh* process only; this is arguably wrong, but currently hard to fix for reasons too complicated to explain here.

NAME

`ruptime` - show host status of local machines

SYNOPSIS

`ruptime` [**-a**] [**-l**] [**-t**] [**-u**]

DESCRIPTION

Ruptime gives a status line like *uptime* for each machine on the local network; these are formed from packets broadcast by each host on the network once a minute.

Machines for which no status report has been received for 5 minutes are shown as being down.

Users idle an hour or more are not counted unless the **-a** flag is given.

Normally, the listing is sorted by host name. The **-l**, **-t**, and **-u** flags specify sorting by load average, uptime, and number of users, respectively.

FILES

`/usr/spool/rwho/whod.*` data files

SEE ALSO

`rwho(1C)`

NAME

rwho - who's logged in on local machines

SYNOPSIS

rwho [**-a**]

DESCRIPTION

The *rwho* command produces output similar to *who*, but for all machines on the local network. If no report has been received from a machine for 5 minutes then *rwho* assumes the machine is down, and does not report users last known to be logged into that machine.

If a users hasn't typed to the system for a minute or more, then *rwho* reports this idle time. If a user hasn't typed to the system for an hour or more, then the user will be omitted from the output of *rwho* unless the **-a** flag is given.

FILES

/usr/spool/rwho/whod.* information about other machines

SEE ALSO

ruptime(1C), rwhod(8C)

BUGS

This is unwieldy when the number of machines on the local net is large.

NAME

`sccstorcs` - build RCS file from SCCS file

SYNOPSIS

`sccstorcs` [-t] [-v] *s.file* ...

DESCRIPTION

Sccstorcs builds an RCS file from each SCCS file argument. The deltas and comments for each delta are preserved and installed into the new RCS file in order. Also preserved are the user access list and descriptive text, if any, from the SCCS file.

The following flags are meaningful:

- t Trace only. Prints detailed information about the SCCS file and lists the commands that would be executed to produce the RCS file. No commands are actually executed and no RCS file is made.
- v Verbose. Prints each command that is run while it is building the RCS file.

FILES

For each *s.somefile*, *Sccstorcs* writes the files *somefile* and *somefile,v* which should not already exist. *Sccstorcs* will abort, rather than overwrite those files if they do exist.

SEE ALSO

`ci` (1), `co` (1), `rcs` (1).

Walter F. Tichy, "Design, Implementation, and Evaluation of a Revision Control System," in *Proceedings of the 6th International Conference on Software Engineering*, IEEE, Tokyo, Sept. 1982.

DIAGNOSTICS

All diagnostics are written to `stderr`. Non-zero exit status on error.

BUGS

Sccstorcs does not preserve all SCCS options specified in the SCCS file. Most notably, it does not preserve removed deltas, MR numbers, and cutoff points.

AUTHOR

Ken Greer

Copyright © 1983 by Kenneth L. Greer

NAME

`script` - make typescript of terminal session

SYNOPSIS

`script` [`-a`] [`file`]

DESCRIPTION

Script makes a typescript of everything printed on your terminal. The typescript is written to *file*, or appended to *file* if the `-a` option is given. It can be sent to the line printer later with *lpr*. If no file name is given, the typescript is saved in the file *typescript*.

The script ends when the forked shell exits.

This program is useful when using a crt and a hard-copy record of the dialog is desired, as for a student handing in a program that was developed on a crt when hard-copy terminals are in short supply.

BUGS

Script places **everything** in the log file. This is not what the naive user expects.

NAME

`sed` - stream editor

SYNOPSIS

`sed` [`-n`] [`-e script`] [`-f sfile`] [`file`] ...

DESCRIPTION

Sed copies the named *files* (standard input default) to the standard output, edited according to a script of commands. The `-f` option causes the script to be taken from file *sfile*; these options accumulate. If there is just one `-e` option and no `-f`'s, the flag `-e` may be omitted. The `-n` option suppresses the default output.

A script consists of editing commands, one per line, of the following form:

[address [, address]] function [arguments]

In normal operation *sed* cyclically copies a line of input into a *pattern space* (unless there is something left after a 'D' command), applies in sequence all commands whose *addresses* select that pattern space, and at the end of the script copies the pattern space to the standard output (except under `-n`) and deletes the pattern space.

An *address* is either a decimal number that counts input lines cumulatively across files, a '\$' that addresses the last line of input, or a context address, '/regular expression/', in the style of *ed*(1) modified thus:

The escape sequence '\n' matches a newline embedded in the pattern space.

A command line with no addresses selects every pattern space.

A command line with one address selects each pattern space that matches the address.

A command line with two addresses selects the inclusive range from the first pattern space that matches the first address through the next pattern space that matches the second. (If the second address is a number less than or equal to the line number first selected, only one line is selected.) Thereafter the process is repeated, looking again for the first address.

Editing commands can be applied only to non-selected pattern spaces by use of the negation function '!' (below).

In the following list of functions the maximum number of permissible addresses for each function is indicated in parentheses.

An argument denoted *text* consists of one or more lines, all but the last of which end with '\ ' to hide the newline. Backslashes in *text* are treated like backslashes in the replacement string of an 's' command, and may be used to protect initial blanks and tabs against the stripping that is done on every script line.

An argument denoted *rfile* or *wfile* must terminate the command line and must be preceded by exactly one blank. Each *wfile* is created before processing begins. There can be at most 10 distinct *wfile* arguments.

(1) a\
text

Append. Place *text* on the output before reading the next input line.

(2) b *label*

Branch to the ':' command bearing the *label*. If *label* is empty, branch to the end of the script.

(2) c\
text

Change. Delete the pattern space. With 0 or 1 address or at the end of a 2-address range, place *text* on the output. Start the next cycle.

(2) d Delete the pattern space. Start the next cycle.

(2) D Delete the initial segment of the pattern space through the first newline. Start the next cycle.

(2) g Replace the contents of the pattern space by the contents of the hold space.

(2) G Append the contents of the hold space to the pattern space.

(2) h Replace the contents of the hold space by the contents of the pattern space.

(2) H Append the contents of the pattern space to the hold space.

(1) i\
text

Insert. Place *text* on the standard output.

(2) n Copy the pattern space to the standard output. Replace the pattern space with the next line of input.

(2) N Append the next line of input to the pattern space with an embedded newline. (The current line number changes.)

(2) p Print. Copy the pattern space to the standard output.

(2) P Copy the initial segment of the pattern space through the first newline to the standard output.

(1) q Quit. Branch to the end of the script. Do not start a new cycle.

(2) r *rfile*

Read the contents of *rfile*. Place them on the output before reading the next input line.

(2) s/*regular expression/replacement/flags*

Substitute the *replacement* string for instances of the *regular expression* in the pattern space. Any character may be used instead of '/'. For a fuller description see *ed(1)*. *Flags* is zero or more of

g Global. Substitute for all nonoverlapping instances of the *regular expression* rather than just the first one.

p Print the pattern space if a replacement was made.

w *wfile*

Write. Append the pattern space to *wfile* if a replacement was made.

(2) t *label*

Test. Branch to the '.' command bearing the *label* if any substitutions have been made since the most recent reading of an input line or execution of a 't'. If *label* is empty, branch to the end of the script.

(2) w *wfile*

Write. Append the pattern space to *wfile*.

(2) x Exchange the contents of the pattern and hold spaces.

(2) y/*string1/string2/*

Transform. Replace all occurrences of characters in *string1* with the corresponding character in *string2*. The lengths of *string1* and *string2* must be equal.

(2)! *function*

Don't. Apply the *function* (or group, if *function* is '{') only to lines *not* selected by the address(es).

(0) : *label*

This command does nothing; it bears a *label* for 'b' and 't' commands to branch to.

(1) = Place the current line number on the standard output as a line.

(2) { Execute the following commands through a matching '}' only when the pattern space is selected.

(0) An empty command is ignored.

SEE ALSO

ed(1), grep(1), awk(1), lex(1)

NAME

sendbug - mail a system bug report to 4bsd-bugs

SYNOPSIS

sendbug [address]

DESCRIPTION

Bug reports sent to '4bsd-bugs@BERKELEY' are intercepted by a program which expects bug reports to conform to a standard format. *Sendbug* is a shell script to help the user compose and mail bug reports in the correct format. *Sendbug* works by invoking *vi*(1) on a temporary copy of the bug report format outline. The user must fill in the appropriate fields and exit *vi*. *Sendbug* then mails the completed report to '4bsd-bugs@BERKELEY' or the address specified on the command line.

FILES

/usr/ucb/bugformat contains the bug report outline

SEE ALSO

vi(1), *sendmail*(8)

NAME

`sfdate` - set the time/date of a file

SYNOPSIS

`sfdate yymmddhhmm [.ss] file ...`

DESCRIPTION

The "accessed" and "updated" times for the specified files are set to the specified date. *yy* is the last two digits of the year; the first *mm* is the month number; *dd* is the day number in the month; *hh* is the hour number (24 hour system); the second *mm* is the minute number; *.ss* is optional and is the seconds. For example:

```
date 10080045
```

sets the date to Oct 8, 12:45 AM. The year, month and day may be omitted, the current values being the defaults.

EXAMPLE

To set all the files in the directory hierarchy "dir" to midnight on April 1st, 1986:

```
find dir -exec sfdate 8604010000 {} \;
```


NAME

sh, for, case, if, while, :, ., break, continue, cd, eval, exec, exit, export, login, read, readonly, set, shift, times, trap, umask, wait - command language

SYNOPSIS

sh [**-ceiknrstuvx**] [arg] ...

DESCRIPTION

Sh is a command programming language that executes commands read from a terminal or a file. See **invocation** for the meaning of arguments to the shell.

Commands.

A *simple-command* is a sequence of non blank *words* separated by blanks (a blank is a **tab** or a **space**). 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 *execve(2)*). The *value* of a simple-command is its exit status if it terminates normally or 200+*status* if it terminates abnormally (see *sigvec(2)* for a list of status values).

A *pipeline* is a sequence of one or more *commands* separated by **|**. The standard output of each command but the last is connected by a *pipe(2)* to the standard input of the next command. Each command is run as a separate process; the shell waits for the last command to terminate.

A *list* is a sequence of one or more *pipelines* separated by **;**, **&**, **&&** or **||** and optionally terminated by **;** or **&**. **;** and **&** have equal precedence which is lower than that of **&&** and **||**, **&&** and **||** also have equal precedence. A semicolon causes sequential execution; an ampersand causes the preceding *pipeline* to be executed without waiting for it to finish. The symbol **&&** (**||**) causes the *list* following to be executed only if the preceding *pipeline* returns a zero (non zero) value. Newlines may appear in a *list*, instead of semicolons, to delimit commands.

A *command* is either a simple-command or one of the following. The value returned by a command is that of the last simple-command executed in the command.

for *name* [**in** *word* ...] **do** *list* **done**

Each time a **for** command is executed *name* is set to the next word in the **for** word list. If **in** *word* ... is omitted, **in "\$@"** is assumed. Execution ends when there are no more words in the list.

case *word* **in** [*pattern* [**|** *pattern*] ...] *list* **;;** ... **esac**

A **case** command executes the *list* associated with the first pattern that matches *word*. The form of the patterns is the same as that used for file name generation.

if *list* **then** *list* [**elif** *list* **then** *list*] ... [**else** *list*] **fi**

The *list* following **if** is executed and if it returns zero the *list* following **then** is executed. Otherwise, the *list* following **elif** is executed and if its value is zero the *list* following **then** is executed. Failing that the **else** *list* is executed.

while *list* [**do** *list*] **done**

A **while** command repeatedly executes the **while** *list* and if its value is zero executes the **do** *list*; otherwise the loop terminates. The value returned by a **while** command is that of the last executed command in the **do** *list*. **until** may be used in place of **while** to negate the loop termination test.

(*list*) Execute *list* in a subshell.

{ *list* } *list* is simply executed.

The following words are only recognized as the first word of a command and when not quoted.

if then else elif fi case in esac for while until do done { }

Command substitution.

The standard output from a command enclosed in a pair of back quotes (``) may be used as part or all of a word; trailing newlines are removed.

Parameter substitution.

The character \$ is used to introduce substitutable parameters. Positional parameters may be assigned values by **set**. Variables may be set by writing

name=value [*name=value*] ...

`\${parameter}

A *parameter* is a sequence of letters, digits or underscores (a *name*), a digit, or any of the characters * @ # ? - \$!. The value, if any, of the parameter is substituted. The braces are required only when *parameter* is followed by a letter, digit, or underscore that is not to be interpreted as part of its name. If *parameter* is a digit, it is a positional parameter. If *parameter* is * or @ then all the positional parameters, starting with \$1, are substituted separated by spaces. \$0 is set from argument zero when the shell is invoked.

`\${parameter-word}

If *parameter* is set, substitute its value; otherwise substitute *word*.

`\${parameter=word}

If *parameter* is not set, set it to *word*; the value of the parameter is then substituted. Positional parameters may not be assigned to in this way.

`\${parameter?word}

If *parameter* is set, 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, substitute *word*; otherwise substitute nothing.

In the above *word* is not evaluated unless it is to be used as the substituted string. (So that, for example, echo `\${d-pwd}` will only execute *pwd* if *d* is unset.)

The following *parameters* are automatically set by the shell.

- # The number of positional parameters in decimal.
- Options supplied to the shell on invocation or by **set**.
- ? The value returned by the last executed command in decimal.
- \$ The process number of this shell.
- ! The process number of the last background command invoked.

The following *parameters* are used but not set by the shell.

- HOME The default argument (home directory) for the **cd** command.
- PATH The search path for commands (see **execution**).
- MAIL If this variable is set to the name of a mail file, the shell informs the user of the arrival of mail in the specified file.
- PS1 Primary prompt string, by default '\$ '.
- PS2 Secondary prompt string, by default '> '.
- IFS Internal field separators, normally **space**, **tab**, and **newline**.

Blank interpretation.

After parameter and command substitution, any results of substitution are scanned for internal field separator characters (those found in \$IFS) and split into distinct arguments where

such characters are found. Explicit null arguments ("" or ``) are retained. Implicit null arguments (those resulting from *parameters* that have no values) are removed.

File name generation.

Following substitution, each command word is scanned for the characters *, ? and [. If one of these characters appears, the word is regarded as a pattern. The word is replaced with alphabetically sorted file names that match the pattern. If no file name is found that matches the pattern, the word is left unchanged. The character . at the start of a file name or immediately following a /, and the character /, must be matched explicitly.

- * Matches any string, including the null string.
- ? Matches any single character.
- [...] Matches any one of the characters enclosed. A pair of characters separated by - matches any character lexically between the pair.

Quoting.

The following characters have a special meaning to the shell and cause termination of a word unless quoted.

; & () | < > **newline space tab**

A character may be *quoted* by preceding it with a \. \bnewline is ignored. All characters enclosed between a pair of quote marks (` `), except a single quote, are quoted. Inside double quotes (" ") parameter and command substitution occurs and \ quotes the characters ` ` and \$.

"\$*" is equivalent to "\$1 \$2 ..." whereas

"\$@" is equivalent to "\$1" "\$2"

Prompting.

When used interactively, the shell prompts with the value of PS1 before reading a command. If at any time a newline is typed and further input is needed to complete a command, the secondary prompt (\$PS2) is issued.

Input output.

Before a command is executed its input and output may be redirected using a special notation interpreted by the shell. The following may appear anywhere in a simple-command or may precede or follow a *command* and are not passed on to the invoked command. Substitution occurs before *word* or *digit* is used.

< *word*

Use file *word* as standard input (file descriptor 0).

> *word*

Use file *word* as standard output (file descriptor 1). If the file does not exist, it is created; otherwise it is truncated to zero length.

>> *word*

Use file *word* as standard output. If the file exists, output is appended (by seeking to the end); otherwise the file is created.

<< *word*

The shell input is read up to a line the same as *word*, or end of file. The resulting document becomes the standard input. If any character of *word* is quoted, no interpretation is placed upon the characters of the document; otherwise, parameter and command substitution occurs, \bnewline is ignored, and \ is used to quote the characters ` ` \$ ` and the first character of *word*.

< & *digit*

The standard input is duplicated from file descriptor *digit*; see *dup(2)*. Similarly for the standard output using > .

< &- The standard input is closed. Similarly for the standard output using > .

If one of the above is preceded by a digit, the file descriptor created is that specified by the digit (instead of the default 0 or 1). For example,

```
... 2>&1
```

creates file descriptor 2 to be a duplicate of file descriptor 1.

If a command is followed by & then the default standard input for the command is the empty file (/dev/null). Otherwise, the environment for the execution of a command contains the file descriptors of the invoking shell as modified by input output specifications.

Environment.

The environment is a list of name-value pairs that is passed to an executed program in the same way as a normal argument list; see *execve(2)* and *environ(7)*. The shell interacts with the environment in several ways. On invocation, the shell scans the environment and creates a *parameter* for each name found, giving it the corresponding value. Executed commands inherit the same environment. If the user modifies the values of these *parameters* or creates new ones, none of these affects the environment unless the **export** command is used to bind the shell's *parameter* to the environment. The environment seen by any executed command is thus composed of any unmodified name-value pairs originally inherited by the shell, plus any modifications or additions, all of which must be noted in **export** commands.

The environment for any *simple-command* may be augmented by prefixing it with one or more assignments to *parameters*. Thus these two lines are equivalent

```
TERM=450 cmd args
(export TERM; TERM=450; cmd args)
```

If the **-k** flag is set, *all* keyword arguments are placed in the environment, even if they occur after the command name. The following prints 'a=b c' and 'c':

```
echo a=b c
set -k
echo a=b c
```

Signals.

The INTERRUPT and QUIT signals for an invoked command are ignored if the command is followed by &; otherwise signals have the values inherited by the shell from its parent. (But see also **trap**.)

Execution.

Each time a command is executed the above substitutions are carried out. Except for the 'special commands' listed below a new process is created and an attempt is made to execute the command via an *execve(2)*.

The shell parameter **\$PATH** defines the search path for the directory containing the command. Each alternative directory name is separated by a colon (:). The default path is **:/bin:/usr/bin**. If the command name contains a /, the search path is not used. Otherwise, each directory in the path is searched for an executable file. If the file has execute permission but is not an *a.out* file, it is assumed to be a file containing shell commands. A subshell (i.e., a separate process) is spawned to read it. A parenthesized command is also executed in a subshell.

Special commands.

The following commands are executed in the shell process and except where specified no input output redirection is permitted for such commands.

- : No effect; the command does nothing.
- . *file* Read and execute commands from *file* and return. The search path **\$PATH** is used to find the directory containing *file*.

- break** [*n*]
Exit from the enclosing **for** or **while** loop, if any. If *n* is specified, break *n* levels.
- continue** [*n*]
Resume the next iteration of the enclosing **for** or **while** loop. If *n* is specified, resume at the *n*-th enclosing loop.
- cd** [*arg*]
Change the current directory to *arg*. The shell parameter **\$HOME** is the default *arg*.
- eval** [*arg ...*]
The arguments are read as input to the shell and the resulting command(s) executed.
- exec** [*arg ...*]
The command specified by the arguments is executed in place of this shell without creating a new process. Input output arguments may appear and if no other arguments are given cause the shell input output to be modified.
- exit** [*n*]
Causes a non interactive 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 will also exit from the shell.)
- export** [*name ...*]
The given names are marked for automatic export to the *environment* of subsequently-executed commands. If no arguments are given, a list of exportable names is printed.
- login** [*arg ...*]
Equivalent to 'exec login arg ...'.
- read** *name ...*
One line is read from the standard input; successive words of the input are assigned to the variables *name* in order, with leftover words to the last variable. The return code is 0 unless the end-of-file is encountered.
- readonly** [*name ...*]
The given names are marked readonly and the values of the these names may not be changed by subsequent assignment. If no arguments are given, a list of all readonly names is printed.
- set** [-*eknptuvx*] [*arg ...*]
 - e If non interactive, exit immediately if a command fails.
 - k All keyword arguments are placed in the environment for a command, not just those that precede the command name.
 - n Read commands but do not execute them.
 - t Exit after reading and executing one command.
 - u Treat unset variables as an error when substituting.
 - v Print shell input lines as they are read.
 - x Print commands and their arguments as they are executed.
 - Turn off the -x and -v options.

These flags can also be used upon invocation of the shell. The current set of flags may be found in **\$-**.

Remaining arguments are positional parameters and are assigned, in order, to **\$1**, **\$2**, etc. If no arguments are given, the values of all names are printed.

shift The positional parameters from **\$2...** are renamed **\$1...**

times Print the accumulated user and system times for processes run from the shell.

trap [*arg*] [*n*] ...
Arg is a command to be read and executed when the shell receives signal(s) *n*. (Note that *arg* is scanned once when the trap is set and once when the trap is taken.) Trap commands are executed in order of signal number. If *arg* is absent, all trap(s) *n* are

reset to their original values. If *arg* is the null string, this signal is ignored by the shell and by invoked commands. If *n* is 0, the command *arg* is executed on exit from the shell, otherwise upon receipt of signal *n* as numbered in *sigvec(2)*. *Trap* with no arguments prints a list of commands associated with each signal number.

umask [*nnn*]

The user file creation mask is set to the octal value *nnn* (see *umask(2)*). If *nnn* is omitted, the current value of the mask is printed.

wait [*n*]

Wait for the specified process and report its termination status. If *n* is not given, all currently active child processes are waited for. The return code from this command is that of the process waited for.

Invocation.

If the first character of argument zero is -, commands are read from **\$HOME/.profile**, if such a file exists. Commands are then read as described below. The following flags are interpreted by the shell when it is invoked.

- c *string* If the -c flag is present, commands are read from *string*.
- s If the -s flag is present or if no arguments remain then commands are read from the standard input. Shell output is written to file descriptor 2.
- i If the -i flag is present or if the shell input and output are attached to a terminal (as told by *gtty*) then this shell is *interactive*. In this case the terminate signal SIGTERM (see *sigvec(2)*) is ignored (so that 'kill 0' does not kill an interactive shell) and the interrupt signal SIGINT is caught and ignored (so that **wait** is interruptible). In all cases SIGQUIT is ignored by the shell.

The remaining flags and arguments are described under the **set** command.

FILES

\$HOME/.profile
 /tmp/sh*
 /dev/null

SEE ALSO

csh(1), test(1), execve(2), environ(7)

DIAGNOSTICS

Errors detected by the shell, such as syntax errors cause the shell to return a non zero exit status. If the shell is being used non interactively then execution of the shell file is abandoned. Otherwise, the shell returns the exit status of the last command executed (see also **exit**).

BUGS

If << is used to provide standard input to an asynchronous process invoked by &, the shell gets mixed up about naming the input document. A garbage file /tmp/sh* is created, and the shell complains about not being able to find the file by another name.

SIZE (1)

USER COMMANDS

SIZE (1)

NAME

size - size of an object file

SYNOPSIS

size [object ...]

DESCRIPTION

Size prints the (decimal) number of bytes required by the text, data, and bss portions, and their sum in hex and decimal, of each object-file argument. If no file is specified, **a.out** is used.

SEE ALSO

a.out(5)

NAME

sleep - suspend execution for an interval

SYNOPSIS

sleep time

DESCRIPTION

Sleep suspends execution for *time* seconds. It is used to execute a command after a certain amount of time as in:

```
(sleep 105; command)&
```

or to execute a command every so often, as in:

```
while true
do
    command
    sleep 37
done
```

SEE ALSO

setitimer(2), alarm(3C), sleep(3)

BUGS

Time must be less than 2,147,483,647 seconds.

NAME

soelim - eliminate .so's from *nroff* input

SYNOPSIS

soelim [file ...]

DESCRIPTION

Soelim reads the specified files or the standard input and performs the textual inclusion implied by the *nroff* directives of the form

`.so somefile`

when they appear at the beginning of input lines. This is useful since programs such as *tbl* do not normally do this; it allows the placement of individual tables in separate files to be run as a part of a large document.

An argument consisting of a single minus (-) is taken to be a file name corresponding to the standard input.

Note that inclusion can be suppressed by using 'ˆ' instead of '.', i.e.

`ˆso /usr/lib/tmac.s`

A sample usage of *soelim* would be

`soelim exum?.n | tbl | nroff -ms | col | lpr`

SEE ALSO

colcrt(1), *more(1)*

AUTHOR

William Joy

BUGS

The format of the source commands must involve no strangeness - exactly one blank must precede and no blanks follow the file name.

NAME

sort – sort or merge files

SYNOPSIS

sort [**-mubdfinrtx**] [**+pos1** [**-pos2**]] ... [**-o** name] [**-T** directory] [name] ...

DESCRIPTION

Sort sorts lines of all the named files together and writes the result on the standard output. The name '-' means the standard input. If no input files are named, the standard input is sorted.

The default sort key is an entire line. Default ordering is lexicographic by bytes in machine collating sequence. The ordering is affected globally by the following options, one or more of which may appear.

- b** Ignore leading blanks (spaces and tabs) in field comparisons.
- d** 'Dictionary' order: only letters, digits and blanks are significant in comparisons.
- f** Fold upper case letters onto lower case.
- i** Ignore characters outside the ASCII range 040-0176 in nonnumeric comparisons.
- n** An initial numeric string, consisting of optional blanks, optional minus sign, and zero or more digits with optional decimal point, is sorted by arithmetic value. Option **n** implies option **b**.
- r** Reverse the sense of comparisons.
- tx** 'Tab character' separating fields is *x*.

The notation *+pos1 -pos2* restricts a sort key to a field beginning at *pos1* and ending just before *pos2*. *Pos1* and *pos2* each have the form *m.n*, optionally followed by one or more of the flags **bdfinr**, where *m* tells a number of fields to skip from the beginning of the line and *n* tells a number of characters to skip further. If any flags are present they override all the global ordering options for this key. If the **b** option is in effect *n* is counted from the first non-blank in the field; **b** is attached independently to *pos2*. A missing *.n* means *.0*; a missing *-pos2* means the end of the line. Under the **-tx** option, fields are strings separated by *x*; otherwise fields are nonempty nonblank strings separated by blanks.

When there are multiple sort keys, later keys are compared only after all earlier keys compare equal. Lines that otherwise compare equal are ordered with all bytes significant.

These option arguments are also understood:

- c** Check that the input file is sorted according to the ordering rules; give no output unless the file is out of sort.
- m** Merge only, the input files are already sorted.
- o** The next argument is the name of an output file to use instead of the standard output. This file may be the same as one of the inputs.
- T** The next argument is the name of a directory in which temporary files should be made.
- u** Suppress all but one in each set of equal lines. Ignored bytes and bytes outside keys do not participate in this comparison.

EXAMPLES

Print in alphabetical order all the unique spellings in a list of words. Capitalized words differ from uncapitalized.

```
sort -u +0f +0 list
```

Print the password file (*passwd*(5)) sorted by user id number (the 3rd colon-separated field).

```
sort -t: +2n /etc/passwd
```

Print the first instance of each month in an already sorted file of (month day) entries. The options **-um** with just one input file make the choice of a unique representative from a set of equal lines predictable.

```
sort -um +0 -1 dates
```

FILES

/usr/tmp/stm*, /tmp/* first and second tries for temporary files

SEE ALSO

uniq(1), comm(1), rev(1), join(1)

DIAGNOSTICS

Comments and exits with nonzero status for various trouble conditions and for disorder discovered under option **-c**.

BUGS

Very long lines are silently truncated.

NAME

sortbib - sort bibliographic database

SYNOPSIS

sortbib [-sKEYS] database ...

DESCRIPTION

Sortbib sorts files of records containing *refer* key-letters by user-specified keys. Records may be separated by blank lines, or by `.[` and `.]` delimiters, but the two styles may not be mixed together. This program reads through each *database* and pulls out key fields, which are sorted separately. The sorted key fields contain the file pointer, byte offset, and length of corresponding records. These records are delivered using disk seeks and reads, so *sortbib* may not be used in a pipeline to read standard input.

By default, *sortbib* alphabetizes by the first `%A` and the `%D` fields, which contain the senior author and date. The `-s` option is used to specify new *KEYS*. For instance, `-sATD` will sort by author, title, and date, while `-sA+D` will sort by all authors, and date. Sort keys past the fourth are not meaningful. No more than 16 databases may be sorted together at one time. Records longer than 4096 characters will be truncated.

Sortbib sorts on the last word on the `%A` line, which is assumed to be the author's last name. A word in the final position, such as "jr." or "ed.", will be ignored if the name beforehand ends with a comma. Authors with two-word last names or unusual constructions can be sorted correctly by using the *nroff* convention "`\0`" in place of a blank. A `%Q` field is considered to be the same as `%A`, except sorting begins with the first, not the last, word. *Sortbib* sorts on the last word of the `%D` line, usually the year. It also ignores leading articles (like "A" or "The") when sorting by titles in the `%T` or `%J` fields; it will ignore articles of any modern European language. If a sort-significant field is absent from a record, *sortbib* places that record before other records containing that field.

SEE ALSO

refer(1), addbib(1), roffb(1), indxbib(1), lookbib(1)

AUTHORS

Greg Shenaut, Bill Tuthill

BUGS

Records with missing author fields should probably be sorted by title.

NAME

spell, spellin, spellout - find spelling errors

SYNOPSIS

```
spell [-v] [-b] [-x] [-d hlist] [-s hstop] [-h spellhist] [file] ...
spellin [list]
spellout [-d] list
```

DESCRIPTION

Spell collects words from the named documents, and looks them up in a spelling list. Words that neither occur among nor are derivable (by applying certain inflections, prefixes or suffixes) from words in the spelling list are printed on the standard output. If no files are named, words are collected from the standard input.

Spell ignores most *troff*, *tbl* and *eqn(1)* constructions.

Under the *-v* option, all words not literally in the spelling list are printed, and plausible derivations from spelling list words are indicated.

Under the *-b* option, British spelling is checked. Besides preferring *centre*, *colour*, *speciality*, *travelled*, etc., this option insists upon *-ise* in words like *standardise*, Fowler and the OED to the contrary notwithstanding.

Under the *-x* option, every plausible stem is printed with '=' for each word.

The spelling list is based on many sources. While it is more haphazard than an ordinary dictionary, it is also more effective with proper names and popular technical words. Coverage of the specialized vocabularies of biology, medicine and chemistry is light.

The auxiliary files used for the spelling list, stop list, and history file may be specified by arguments following the *-d*, *-s*, and *-h* options. The default files are indicated below. Copies of all output may be accumulated in the history file. The stop list filters out misspellings (e.g. *thier=thy-y+ier*) that would otherwise pass.

Two routines help maintain the hash lists used by *spell*. Both expect a set of words, one per line, from the standard input. *Spellin* combines the words from the standard input and the preexisting *list* file and places a new list on the standard output. If no *list* file is specified, the new list is created from scratch. *Spellout* looks up each word from the standard input and prints on the standard output those that are missing from (or present on, with option *-d*) the hashed *list* file. For example, to verify that *hookey* is not on the default spelling list, add it to your own private list, and then use it with *spell*,

```
echo hookey | spellout /usr/dict/hlista
echo hookey | spellin /usr/dict/hlista > myhlist
spell -d myhlist huckfinn
```

FILES

/usr/dict/hlist[ab]	hashed spelling lists, American & British, default for <i>-d</i>
/usr/dict/hstop	hashed stop list, default for <i>-s</i>
/dev/null	history file, default for <i>-h</i>
/tmp/spell.\$\$*	temporary files
/usr/lib/spell	

SEE ALSO

deroff(1), sort(1), tee(1), sed(1)

BUGS

The spelling list's coverage is uneven; new installations will probably wish to monitor the output for several months to gather local additions.

British spelling was done by an American.

NAME

spline - interpolate smooth curve

SYNOPSIS

spline [option] ...

DESCRIPTION

Spline takes pairs of numbers from the standard input as abscissas and ordinates of a function. It produces a similar set, which is approximately equally spaced and includes the input set, on the standard output. The cubic spline output (R. W. Hamming, *Numerical Methods for Scientists and Engineers*, 2nd ed., 349ff) has two continuous derivatives, and sufficiently many points to look smooth when plotted, for example by *graph*(1G).

The following options are recognized, each as a separate argument.

- a Supply abscissas automatically (they are missing from the input); spacing is given by the next argument, or is assumed to be 1 if next argument is not a number.
- k The constant k used in the boundary value computation

$$y'_0 = ky'_1, \quad y'_n = ky'_{n-1}$$

is set by the next argument. By default $k = 0$.

- n Space output points so that approximately n intervals occur between the lower and upper x limits. (Default $n = 100$.)
- p Make output periodic, i.e. match derivatives at ends. First and last input values should normally agree.
- x Next 1 (or 2) arguments are lower (and upper) x limits. Normally these limits are calculated from the data. Automatic abscissas start at lower limit (default 0).

SEE ALSO

graph(1G), plot(1G)

DIAGNOSTICS

When data is not strictly monotone in x , *spline* reproduces the input without interpolating extra points.

BUGS

A limit of 1000 input points is enforced silently.

NAME

`split` - split a file into pieces

SYNOPSIS

`split` [*-n*] [file [name]]

DESCRIPTION

Split reads *file* and writes it in *n*-line pieces (default 1000), as many as necessary, onto a set of output files. The name of the first output file is *name* with **aa** appended, and so on lexicographically. If no output name is given, **x** is default.

If no input file is given, or if **-** is given in its stead, then the standard input file is used.

NAME

strip - remove symbols and relocation bits

SYNOPSIS

strip name ...

DESCRIPTION

Strip removes the symbol table and relocation bits ordinarily attached to the output of the assembler and loader. This is useful to save space after a program has been debugged.

The effect of *strip* is the same as use of the *-s* option of *ld*.

FILES

/tmp/stm? temporary file

SEE ALSO

ld(1)

NAME

strings - find the printable strings in a object, or other binary, file

SYNOPSIS

strings [-] [-o] [-number] file ...

DESCRIPTION

Strings looks for ascii strings in a binary file. A string is any sequence of 4 or more printing characters ending with a newline or a null. Unless the - flag is given, *strings* only looks in the initialized data space of object files. If the -o flag is given, then each string is preceded by its offset in the file (in octal). If the -number flag is given then number is used as the minimum string length rather than 4.

Strings is useful for identifying random object files and many other things.

SEE ALSO

od(1)

BUGS

The algorithm for identifying strings is extremely primitive

NAME

struct – structure Fortran programs

SYNOPSIS

struct [option] ... file

DESCRIPTION

Struct translates the Fortran program specified by *file* (standard input default) into a Ratfor program. Wherever possible, Ratfor control constructs replace the original Fortran. Statement numbers appear only where still necessary. Cosmetic changes are made, including changing Hollerith strings into quoted strings and relational operators into symbols (e.g. ".GT." into ">"). The output is appropriately indented.

The following options may occur in any order.

- s** Input is accepted in standard format, i.e. comments are specified by a c, C, or * in column 1, and continuation lines are specified by a nonzero, nonblank character in column 6. Normally input is in the form accepted by *f77(1)*
- i** Do not turn computed goto statements into switches. (Ratfor does not turn switches back into computed goto statements.)
- a** Turn sequences of else ifs into a non-Ratfor switch of the form

```
switch
    {
        case pred1: code
        case pred2: code
        case pred3: code
        default: code
    }
```

The case predicates are tested in order; the code appropriate to only one case is executed. This generalized form of switch statement does not occur in Ratfor.

- b** Generate goto's instead of multilevel break statements.
- n** Generate goto's instead of multilevel next statements.
- tn** Make the nonzero integer *n* the lowest valued label in the output program (default 10).
- cn** Increment successive labels in the output program by the nonzero integer *n* (default 1).
- en** If *n* is 0 (default), place code within a loop only if it can lead to an iteration of the loop. If *n* is nonzero, admit a small code segments to a loop if otherwise the loop would have exits to several places including the segment, and the segment can be reached only from the loop. 'Small' is close to, but not equal to, the number of statements in the code segment. Values of *n* under 10 are suggested.

FILES

/tmp/struct*
/usr/lib/struct/*

SEE ALSO

f77(1)

BUGS

Struct knows Fortran 66 syntax, but not full Fortran 77.

If an input Fortran program contains identifiers which are reserved words in Ratfor, the structured version of the program will not be a valid Ratfor program.

The labels generated cannot go above 32767.

If you get a goto without a target, try **-e**.

NAME

stty - set terminal options

SYNOPSIS

stty [option ...]

DESCRIPTION

Stty sets certain I/O options on the current output terminal, placing its output on the diagnostic output. With no argument, it reports the speed of the terminal and the settings of the options which are different from their defaults. With the argument "all", all normally used option settings are reported. With the argument "everything", everything *stty* knows about is printed. The option strings are selected from the following set:

even	allow even parity input
-even	disallow even parity input
odd	allow odd parity input
-odd	disallow odd parity input
raw	raw mode input (no input processing (erase, kill, interrupt, ...); parity bit passed back)
-raw	negate raw mode
cooked	same as '-raw'
cbreak	make each character available to <i>read</i> (2) as received; no erase and kill processing, but all other processing (interrupt, suspend, ...) is performed
-cbreak	make characters available to <i>read</i> only when newline is received
-nl	allow carriage return for new-line, and output CR-LF for carriage return or new-line
nl	accept only new-line to end lines
echo	echo back every character typed
-echo	do not echo characters
lcase	map upper case to lower case
-lcase	do not map case
tandem	enable flow control, so that the system sends out the stop character when its internal queue is in danger of overflowing on input, and sends the start character when it is ready to accept further input
-tandem	disable flow control
hh	enable hardware handshaking, so that the system handshakes with RTS/CTS hardware protocol
-hh	disable hardware handshaking
-tabs	replace tabs by spaces when printing
tabs	preserve tabs
ek	set erase and kill characters to # and @

For the following commands which take a character argument *c*, you may also specify *c* as the "u" or "undef", to set the value to be undefined. A value of "^x", a 2 character sequence, is also interpreted as a control character, with "^?" representing delete.

erase c	set erase character to <i>c</i> (default '#', but often reset to ^H.)
kill c	set kill character to <i>c</i> (default '@', but often reset to ^U.)
intr c	set interrupt character to <i>c</i> (default DEL or ^? (delete), but often reset to ^C.)
quit c	set quit character to <i>c</i> (default control \.)
start c	set start character to <i>c</i> (default control Q.)
stop c	set stop character to <i>c</i> (default control S.)
eof c	set end of file character to <i>c</i> (default control D.)

brk c set break character to *c* (default undefined.) This character is an extra wakeup causing character.

cr0 cr1 cr2 cr3
select style of delay for carriage return (see *ioctl(2)*)

nl0 nl1 nl2 nl3
select style of delay for linefeed

tab0 tab1 tab2 tab3
select style of delay for tab

ff0 ff1 select style of delay for form feed

bs0 bs1 select style of delay for backspace

tty33 set all modes suitable for the Teletype Corporation Model 33 terminal.

tty37 set all modes suitable for the Teletype Corporation Model 37 terminal.

vt05 set all modes suitable for Digital Equipment Corp. VT05 terminal

dec set all modes suitable for Digital Equipment Corp. operating systems users; (erase, kill, and interrupt characters to ^? , ^U , and ^C , *decctlq* and "newcrt".)

tn300 set all modes suitable for a General Electric TermiNet 300

ti700 set all modes suitable for Texas Instruments 700 series terminal

tek set all modes suitable for Tektronix 4014 terminal

0 hang up phone line immediately

50 75 110 134 150 200 300 600 1200 1800 2400 4800 9600 exta extb rs422
Set terminal baud rate to the number given, if possible.

A teletype driver which supports the job control processing of *csH(1)* and more functionality than the basic driver is fully described in *tty(4)*. The following options apply only to it.

new Use new driver (switching flushes typeahead).

crt Set options for a CRT (*crtbs*, *ctlecho* and, if ≥ 1200 baud, *crterase* and *crtkill*.)

crtbs Echo backspaces on erase characters.

prterase For printing terminal echo erased characters backwards within "\" and "/".

crterase Wipe out erased characters with "backspace-space-backspace."

-crterase Leave erased characters visible; just backspace.

crtkill Wipe out input on like kill ala *crterase*.

-crtkill Just echo line kill character and a newline on line kill.

ctlecho Echo control characters as " ^x " (and delete as " ^? ".) Print two backspaces following the EOT character (control D).

-ctlecho Control characters echo as themselves; in cooked mode EOT (control-D) is not echoed.

decctlq After output is suspended (normally by ^S), only a start character (normally ^Q) will restart it. This is compatible with DEC's vendor supplied systems.

-decctlq After output is suspended, any character typed will restart it; the start character will restart output without providing any input. (This is the default.)

tostop Background jobs stop if they attempt terminal output.

-tostop Output from background jobs to the terminal is allowed.

tilde Convert "~" to "^" on output (for Hazeltine terminals).

-tilde Leave poor "~" alone.

flusho Output is being discarded usually because user hit control O (internal state bit).

-flusho Output is not being discarded.

pendin Input is pending after a switch from *cbreak* to cooked and will be re-input when a read becomes pending or more input arrives (internal state bit).

-pendin Input is not pending.

intrup Send a signal (SIGTINT) to the terminal control process group whenever an input record (line in cooked mode, character in *cbreak* or raw mode) is available for

- reading.
- intrup** Don't send input available interrupts.
 - mdmbuf** Start/stop output on carrier transitions (not implemented).
 - mdmbuf** Return error if write attempted after carrier drops.
 - litout** Send output characters without any processing.
 - litout** Do normal output processing, inserting delays, etc.
 - nohang** Don't send hangup signal if carrier drops.
 - nohang** Send hangup signal to control process group when carrier drops.
 - etxack** Diablo style etx/ack handshaking (not implemented).

The following special characters are applicable only to the new teletype driver and are not normally changed.

- susp c** set suspend process character to *c* (default control Z).
- dsusp c** set delayed suspend process character to *c* (default control Y).
- rprnt c** set reprint line character to *c* (default control R).
- flush c** set flush output character to *c* (default control O).
- werase c** set word erase character to *c* (default control W).
- lnext c** set literal next character to *c* (default control V).

SEE ALSO

ioctl(2), tabs(1), tset(1), tty(4)

NAME

style - analyze surface characteristics of a document

SYNOPSIS

style [**-ml**] [**-mm**] [**-a**] [**-e**] [**-l num**] [**-r num**] [**-p**] [**-P**] file ...

DESCRIPTION

Style analyzes the surface characteristics of the writing style of a document. It reports on readability, sentence length and structure, word length and usage, verb type, and sentence openers. Because *style* runs *deroff* before looking at the text, formatting header files should be included as part of the input. The default macro package **-ms** may be overridden with the flag **-mm**. The flag **-ml**, which causes **deroff** to skip lists, should be used if the document contains many lists of non-sentences. The other options are used to locate sentences with certain characteristics.

- a** print all sentences with their length and readability index.
- e** print all sentences that begin with an expletive.
- p** print all sentences that contain a passive verb.
- l num** print all sentences longer than *num*.
- r num** print all sentences whose readability index is greater than *num*.
- P** print parts of speech of the words in the document.

SEE ALSO

deroff(1), *diction*(1)

BUGS

Use of non-standard formatting macros may cause incorrect sentence breaks.

NAME

su - substitute user id temporarily

SYNOPSIS

su [userid]

DESCRIPTION

Su demands the password of the specified *userid*, and if it is given, changes to that *userid* and invokes the Shell *sh*(1) without changing the current directory. The user environment is unchanged except for HOME and SHELL, which are taken from the password file for the user being substituted (see *environ*(7)). The new user ID stays in force until the Shell exits.

If no *userid* is specified, 'root' is assumed. To remind the super-user of his responsibilities, the Shell substitutes '#' for its usual prompt.

SEE ALSO

sh(1)

BUGS

Local administrative rules cause restrictions to be placed on who can *su* to 'root', even with the root password. These rules vary from site to site.

NAME

sum - sum and count blocks in a file

SYNOPSIS

sum file

DESCRIPTION

Sum calculates and prints a 16-bit checksum for the named file, and also prints the number of blocks in the file. It is typically used to look for bad spots, or to validate a file communicated over some transmission line.

SEE ALSO

wc(1)

DIAGNOSTICS

'Read error' is indistinguishable from end of file on most devices; check the block count.

NAME

symorder - rearrange name list

SYNOPSIS

symorder orderlist symbolfile

DESCRIPTION

Orderlist is a file containing symbols to be found in symbolfile, 1 symbol per line.

Symbolfile is updated in place to put the requested symbols first in the symbol table, in the order specified. This is done by swapping the old symbols in the required spots with the new ones. If all of the order symbols are not found, an error is generated.

This program was specifically designed to cut down on the overhead of getting symbols from /vmunix.

SEE ALSO

nlist(3)

NAME

sysline - display system status on status line of a terminal

SYNOPSIS

sysline [**-bcdehDilmpqrsj**] [**-H remote**] [**+N**]

DESCRIPTION

Sysline runs in the background and periodically displays system status information on the status line of the terminal. Not all terminals contain a status line. Those that do include the h19, concept 108, Ann Arbor Ambassador, vt100, Televideo 925/950 and Freedom 100. If no flags are given, *sysline* displays the time of day, the current load average, the change in load average in the last 5 minutes, the number of users (followed by a 'u'), the number of runnable process, the number of suspended processes, and the users who have logged on and off since the last status report. Finally, if new mail has arrived, a summary of it is printed. If there is unread mail in your mailbox, an asterisk will appear after the display of the number of users. The display is normally in reverse video (if your terminal supports this in the status line) and is right justified to reduce distraction. Every fifth display is done in normal video to give the screen a chance to rest.

If you have a file named *.who* in your home directory, then the contents of that file is printed first. One common use of this feature is to alias *chdir*, *pushd*, and *popd* to place the current directory stack in *~/who* after it changes the new directory.

The following flags may be given on the command line.

- b** Beep once every half hour and twice every hour, just like those obnoxious watches you keep hearing.
- c** Clear the status line for 5 seconds before each redisplay.
- d** Debug mode -- print status line data in human readable format
- D** Print out the current day/date before the time.
- e** Print out only the information. Do not print out the control commands necessary to put the information on the bottom line. This option is useful for putting the output of *sysline* onto the mode line of an emacs window.
- H remote** Print the load average on the remote host *remote*. If the host is down, or is not sending out *rwhod* packets, then the down time is printed instead.
- h** Print out the host machine's name after the time.
- l** Don't print the names of people who log in and out.
- m** Don't check for mail.
- p** Don't report the number of process which are runnable and suspended.
- r** Don't display in reverse video.
- +N** Update the status line every N seconds. The default is 60 seconds.
- q** Don't print out diagnostic messages if something goes wrong when starting up.
- i** Print out the process id of the *sysline* process onto standard output upon startup. With this information you can send the alarm signal to the *sysline* process to cause it to update immediately. *sysline* writes to the standard error, so you can redirect the standard output into a file to catch the process id.

- s Print "short" form of line by left-justifying *iff* escapes are not allowed in the status line. Some terminals (the Televideos and Freedom 100 for example) do not allow cursor movement (or other "intelligent" operations) in the status line. For these terminals, *sysline* normally uses blanks to cause right-justification. This flag will disable the adding of the blanks.
- j Force the *sysline* output to be left justified even on terminals capable of cursor movement on the status line.

If you have a file *.syslinelock* in your home directory, then *sysline* will not update its statistics and write on your screen, it will just go to sleep for a minute. This is useful if you want to momentarily disable *sysline*. Note that it may take a few seconds from the time the lock file is created until you are guaranteed that *sysline* will not write on the screen.

FILES

<i>/etc/utmp</i>	names of people who are logged in
<i>/dev/kmem</i>	contains process table
<i>/usr/spool/rwho/whod.*</i>	who/uptime information for remote hosts
<i>\${HOME}/.who</i>	information to print on bottom line
<i>\${HOME}/.syslinelock</i>	when it exists, <i>sysline</i> will not print

AUTHORS

John Foderaro
 Tom Ferrin converted it to use termcap.
 Mark Horton added terminfo capability.

BUGS

If you interrupt the display then you may find your cursor missing or stuck on the status line. The best thing to do is reset the terminal.
 If there is too much for one line, the excess is thrown away.

NAME

tabs - set terminal tabs

SYNOPSIS

tabs [-n] [terminal]

DESCRIPTION

Tabs sets the tabs on a variety of terminals. Various terminal names given in *term(7)* are recognized; the default is, however, suitable for most 300 baud terminals. If the *-n* flag is present then the left margin is not indented as is normal.

SEE ALSO

stty(1), term(7)

BUGS

It's much better to use *tset(1)*.

NAME

`tac` - concatenate and print files in reverse order

SYNOPSIS

`tac` [*-string*] [*+string*] [*file ...*]

DESCRIPTION

Tac reads each *file* in sequence and writes it on the standard output, reversed by the file segments delimited by *string*. *-string* specifies segments bounded on the left by *string*, while *+string* specifies right-bounded segments. The default is *+\n* (print lines in reverse order). If no input *file* is given, or if the argument '-' is encountered, *tac* reads from the standard input. Note that in this case *tac* stores the entire standard input in a temporary file before it outputs anything, so for large input it is slow.

EXAMPLES

```
tac '-\nFrom ' /usr/spool/mail/$USER
```

prints out one's mail messages, most recent first.

```
tac file
```

prints the file in reverse, line by line.

```
tac /usr/adm/messages | egrep 'hp.*hard'
```

prints out the hard errors on MASSBUS disk drives, most recent first.

SEE ALSO

`cat(1)`, `rev(1)`, `tail(1)`, `tmail(1)`

BUGS

Tac doesn't handle multiple argument files exactly right, and it is unclear in which order they should be processed.

If invoked as '`tac < file`', *tac* uses a temp file but it doesn't have to.

NAME

tail - deliver the last part of a file

SYNOPSIS

tail [\pm number[lbc][fr]] [file]

DESCRIPTION

Tail copies the named file to the standard output beginning at a designated place. If no file is named, the standard input is used.

Copying begins at distance *+number* from the beginning, or *-number* from the end of the input. *Number* is counted in units of lines, blocks or characters, according to the appended option **l**, **b** or **c**. When no units are specified, counting is by lines.

Specifying **r** causes tail to print lines from the end of the file in reverse order. The default for **r** is to print the entire file this way. Specifying **f** causes *tail* to not quit at end of file, but rather wait and try to read repeatedly in hopes that the file will grow.

SEE ALSO

dd(1)

BUGS

Tails relative to the end of the file are treasured up in a buffer, and thus are limited in length.

Various kinds of anomalous behavior may happen with character special files.

NAME

talk - talk to another user

SYNOPSIS

talk person [ttyname]

DESCRIPTION

Talk is a visual communication program which copies lines from your terminal to that of another user.

If you wish to talk to someone on you own machine, then *person* is just the person's login name. If you wish to talk to a user on another host, then *person* is of the form :

host!user or
host.user or
host:user or
user@host

though *host@user* is perhaps preferred.

If you want to talk to a user who is logged in more than once, the *ttyname* argument may be used to indicate the appropriate terminal name.

When first called, it sends the message

```
Message from TalkDaemon@his_machine...
talk: connection requested by your_name@your_machine.
talk: respond with: talk your_name@your_machine
```

to the user you wish to talk to. At this point, the recipient of the message should reply by typing

```
talk your_name@your_machine
```

It doesn't matter from which machine the recipient replies, as long as his login-name is the same. Once communication is established, the two parties may type simultaneously, with their output appearing in separate windows. Typing control L will cause the screen to be reprinted, while your erase, kill, and word kill characters will work in talk as normal. To exit, just type your interrupt character; *talk* then moves the cursor to the bottom of the screen and restores the terminal.

Permission to talk may be denied or granted by use of the *mesg* command. At the outset talking is allowed. Certain commands, in particular *nroff* and *pr(1)* disallow messages in order to prevent messy output.

FILES

```
/etc/hosts    to find the recipient's machine
/etc/utmp    to find the recipient's tty
```

SEE ALSO

mesg(1), who(1), mail(1), write(1)

NAME

tar - tape archiver

SYNOPSIS

tar [*key*] [*name ...*]

DESCRIPTION

Tar saves and restores multiple files on a single file (usually a magnetic tape, but it can be any file). *Tar*'s actions are controlled by the *key* argument. The *key* is a string of characters containing at most one function letter and possibly one or more function modifiers. Other arguments to *tar* are file or directory names specifying which files to dump or restore. In all cases, appearance of a directory name refers to the files and (recursively) subdirectories of that directory.

The function portion of the key is specified by one of the following letters:

- r** The named files are written on the end of the tape. The **c** function implies this.
- x** The named files are extracted from the tape. If the named file matches a directory whose contents had been written onto the tape, this directory is (recursively) extracted. The owner, modification time, and mode are restored (if possible). If no file argument is given, the entire content of the tape is extracted. Note that if multiple entries specifying the same file are on the tape, the last one overwrites all earlier.
- t** The names of the specified files are listed each time they occur on the tape. If no file argument is given, all of the names on the tape are listed.
- u** The named files are added to the tape if either they are not already there or have been modified since last put on the tape.
- c** Create a new tape; writing begins on the beginning of the tape instead of after the last file. This command implies **r**.
- o** On output, *tar* normally places information specifying owner and modes of directories in the archive. Former versions of *tar*, when encountering this information will give error message of the form
 "<name>/: cannot create".
 This option will suppress the directory information.
- p** This option says to restore files to their original modes, ignoring the present *umask*(2). Setuid and sticky information will also be restored to the super-user.

The following characters may be used in addition to the letter which selects the function desired.

- 0, ..., 9** This modifier selects an alternate drive on which the tape is mounted. The default is drive 0 at 1600 bpi, which is normally /dev/rmt8.
- v** Normally *tar* does its work silently. The **v** (verbose) option make *tar* type the name of each file it treats preceded by the function letter. With the **t** function, the verbose option gives more information about the tape entries than just their names.
- w** *Tar* prints the action to be taken followed by file name, then wait for user confirmation. If a word beginning with 'y' is given, the action is done. Any other input means don't do it.
- f** *Tar* uses the next argument as the name of the archive instead of /dev/rmt?. If the name of the file is '-', *tar* writes to standard output or reads from standard input, whichever is appropriate. Thus, *tar* can be used as the head or tail of a filter

chain. *Tar* can also be used to move hierarchies with the command
`cd fromdir; tar cf - . | (cd todir; tar xf -)`

- b** *Tar* uses the next argument as the blocking factor for tape records. The default is 20 (the maximum). This option should only be used with raw magnetic tape archives (See **f** above). The block size is determined automatically when reading tapes (key letters 'x' and 't').
- l** tells *tar* to complain if it cannot resolve all of the links to the files dumped. If this is not specified, no error messages are printed.
- m** tells *tar* not to restore the modification times. The modification time will be the time of extraction.
- h** Force *tar* to follow symbolic links as if they were normal files or directories. Normally, *tar* does not follow symbolic links.
- B** Forces input and output blocking to 20 blocks per record. This option was added so that *tar* can work across a communications channel where the blocking may not be maintained.

If a file name is preceded by **-C**, then *tar* will perform a *chdir(2)* to that file name. This allows multiple directories not related by a close common parent to be archived using short relative path names. For example, to archive files from `/usr/include` and from `/etc`, one might use

```
tar c -C /usr include -C / etc
```

Previous restrictions dealing with *tar*'s inability to properly handle blocked archives have been lifted.

FILES

```
/dev/rmt?  
/tmp/tar*
```

DIAGNOSTICS

Complaints about bad key characters and tape read/write errors.
 Complaints if enough memory is not available to hold the link tables.

BUGS

There is no way to ask for the *n*-th occurrence of a file.
 Tape errors are handled ungracefully.
 The **u** option can be slow.
 The current limit on file name length is 100 characters.
 There is no way to selectively follow symbolic links.

NAME

`tbl` - format tables for `nroff` or `troff`

SYNOPSIS

`tbl` [files] ...

DESCRIPTION

`Tbl` is a preprocessor for formatting tables for `nroff` or `troff(1)`. The input files are copied to the standard output, except for lines between and are reformatted. Details are given in the `tbl(1)` reference manual.

EXAMPLE

As an example, letting `\t` represent a tab (which should be typed as a genuine tab) the input

```
.TS
c s s
c c s
c c c
l n n.
Household Population
Town\tHouseholds
\tNumber\tSize
Bedminster\t789\t3.26
Bernards Twp.\t3087\t3.74
Bernardsville\t2018\t3.30
Bound Brook\t3425\t3.04
Branchburg\t1644\t3.49
Bridgewater\t7897\t3.81
Far Hills\t240\t3.19
.TE
```

yields

Town	Household Population	
	Households Number	Size
Bedminster	789	3.26
Bernards Twp.	3087	3.74
Bernardsville	2018	3.30
Bound Brook	3425	3.04
Branchburg	1644	3.49
Bridgewater	7897	3.81
Far Hills	240	3.19

If no arguments are given, `tbl` reads the standard input, so it may be used as a filter. When `tbl` is used with `eqn` or `neqn` the `tbl` command should be first, to minimize the volume of data passed through pipes.

SEE ALSO

`troff(1)`, `eqn(1)`
M. E. Lesk, *TBL*.

NAME

tc - phototypesetter simulator

SYNOPSIS

tc [-t] [-sN] [-pL] [file]

DESCRIPTION

Tc interprets its input (standard input default) as device codes for a Graphic Systems phototypesetter (cat). The standard output of *tc* is intended for a Tektronix 4015 (a 4014 terminal with ASCII and APL character sets). The sixteen typesetter sizes are mapped into the 4014's four sizes; the entire TROFF character set is drawn using the 4014's character generator, using overstruck combinations where necessary. Typical usage:

```
troff -t file | tc
```

At the end of each page *tc* waits for a newline (empty line) from the keyboard before continuing on to the next page. In this wait state, the command *e* will suppress the screen erase before the next page; *sN* will cause the next *N* pages to be skipped; and *!line* will send line to the shell.

The command line options are:

-t Don't wait between pages; for directing output into a file.

-sN Skip the first *N* pages.

-pL Set page length to *L*. *L* may include the scale factors *p* (points), *i* (inches), *c* (centimeters), and *P* (picas); default is picas.

'-l w' Multiply the default aspect ratio, 1.5, of a displayed page by *l/w*.

SEE ALSO

troff(1), plot(1G)

BUGS

Font distinctions are lost.

tc's character set is limited to ASCII in just one size.

The aspect ratio option is unbelievable.

NAME

tcopy - copy a mag tape

SYNOPSIS

***tcopy* src [dest]**

DESCRIPTION

Tcopy is designed to copy magnetic tapes. The only assumption made about the tape is that there are two tape marks at the end. *Tcopy* with only a source tape specified will print information about the sizes of records and tape files. If a destination is specified, then, a copy will be made of the source tape. The blocking on the destination tape will be identical to that used on the source tape. Copying a tape will yield the same output as if just printing the sizes.

SEE ALSO

mtio(4)

NAME

TEACHJOVE - learn how to use the JOVE editor

SYNOPSIS

teachjove

DESCRIPTION

TEACHJOVE is a simple program that calls up the JOVE editor on a special file that is an interactive tutorial for the JOVE editor. Once in JOVE all you do is follow the instructions and by doing so you will learn all about JOVE!

FILES

/usr/new/lib/jove/teach-jove -- THE special file.

SEE ALSO

JOVE(1) - to learn about JOVE in general.

AUTHOR

Jonathan Payne

NAME

tee - pipe fitting

SYNOPSIS

tee [-i] [-a] [file] ...

DESCRIPTION

Tee transcribes the standard input to the standard output and makes copies in the *files*. Option **-i** ignores interrupts; option **-a** causes the output to be appended to the *files* rather than overwriting them.

NAME

telnet – user interface to the TELNET protocol

SYNOPSIS

telnet [*host* [*port*]]

DESCRIPTION

Telnet is used to communicate with another host using the TELNET protocol. If *telnet* is invoked without arguments, it enters command mode, indicated by its prompt (“*telnet*>”). In this mode, it accepts and executes the commands listed below. If it is invoked with arguments, it performs an *open* command (see below) with those arguments.

Once a connection has been opened, *telnet* enters input mode. In this mode, text typed is sent to the remote host. To issue *telnet* commands when in input mode, precede them with the *telnet* “escape character” (initially “^”). When in command mode, the normal terminal editing conventions are available.

The following commands are available. Only enough of each command to uniquely identify it need be typed.

open *host* [*port*]

Open a connection to the named host. If the no port number is specified, *telnet* will attempt to contact a TELNET server at the default port. The host specification may be either a host name (see *hosts*(5)) or an Internet address specified in the “dot notation”.

close Close a TELNET session and return to command mode.

quit Close any open TELNET session and exit *telnet*.

z Suspend *telnet*. This command only works when the user is using the *cs*(1).

escape [*escape-char*]

Set the *telnet* “escape character”. Control characters may be specified as “^” followed by a single letter; e.g. “control-X” is “^X”.

status Show the current status of *telnet*. This includes the peer one is connected to, as well as the state of debugging.

options

Toggle viewing of TELNET options processing. When options viewing is enabled, all TELNET option negotiations will be displayed. Options sent by *telnet* are displayed as “SENT”, while options received from the TELNET server are displayed as “RCVD”.

crmod

Toggle carriage return mode. When this mode is enabled any carriage return characters received from the remote host will be mapped into a carriage return and a line feed. This mode does not affect those characters typed by the user, only those received. This mode is not very useful, but is required for some hosts that like to ask the user to do local echoing.

? [*command*]

Get help. With no arguments, *telnet* prints a help summary. If a command is specified, *telnet* will print the help information available about the command only.

BUGS

This implementation is very simple because *rlogin*(1C) is the standard mechanism used to communicate locally with hosts.

NAME

test - condition command

SYNOPSIS

test expr

DESCRIPTION

test evaluates the expression *expr*, and if its value is true then returns zero exit status; otherwise, a non zero exit status is returned. *test* returns a non zero exit if there are no arguments.

The following primitives are used to construct *expr*.

-r file true if the file exists and is readable.

-w file true if the file exists and is writable.

-f file true if the file exists and is not a directory.

-d file true if the file exists exists and is a directory.

-s file true if the file exists and has a size greater than zero.

-t [fildes]

true if the open file whose file descriptor number is *fildes* (1 by default) is associated with a terminal device.

-z s1 true if the length of string *s1* is zero.

-n s1 true if the length of the string *s1* is nonzero.

s1 = s2 true if the strings *s1* and *s2* are equal.

s1 != s2 true if the strings *s1* and *s2* are not equal.

s1 true if *s1* is not the null string.

n1 -eq n2

true if the integers *n1* and *n2* are algebraically equal. Any of the comparisons **-ne**, **-gt**, **-ge**, **-lt**, or **-le** may be used in place of **-eq**.

These primaries may be combined with the following operators:

! unary negation operator

-a binary *and* operator

-o binary *or* operator

(expr)

parentheses for grouping.

-a has higher precedence than **-o**. Notice that all the operators and flags are separate arguments to *test*. Notice also that parentheses are meaningful to the Shell and must be escaped.

SEE ALSO

sh(1), find(1)

NAME

tftp – trivial file transfer program

SYNOPSIS

tftp [*host*]

DESCRIPTION

Tftp is the user interface to the Internet TFTP (Trivial File Transfer Protocol), which allows users to transfer files to and from a remote machine. The remote *host* may be specified on the command line, in which case *tftp* uses *host* as the default host for future transfers (see the **connect** command below).

COMMANDS

Once *tftp* is running, it issues the prompt **tftp>** and recognizes the following commands:

connect *host-name* [*port*]

Set the *host* (and optionally *port*) for transfers. Note that the TFTP protocol, unlike the FTP protocol, does not maintain connections between transfers; thus, the *connect* command does not actually create a connection, but merely remembers what host is to be used for transfers. You do not have to use the *connect* command; the remote host can be specified as part of the *get* or *put* commands.

mode *transfer-mode*

Set the mode for transfers; *transfer-mode* may be one of *ascii* or *binary*. The default is *ascii*.

put *file*

put *localfile remotefile*

put *file1 file2 ... fileN remote-directory*

Put a file or set of files to the specified remote file or directory. The destination can be in one of two forms: a filename on the remote host, if the host has already been specified, or a string of the form *host:filename* to specify both a host and filename at the same time. If the latter form is used, the hostname specified becomes the default for future transfers. If the remote-directory form is used, the remote host is assumed to be a *UNIX* machine.

get *filename*

get *remotename localname*

get *file1 file2 ... fileN*

Get a file or set of files from the specified *sources*. *Source* can be in one of two forms: a filename on the remote host, if the host has already been specified, or a string of the form *host:filename* to specify both a host and filename at the same time. If the latter form is used, the last hostname specified becomes the default for future transfers.

quit Exit *tftp*. An end of file also exits.

verbose

Toggle verbose mode.

trace Toggle packet tracing.

status Show current status.

rexmt *retransmission-timeout*

Set the per-packet retransmission timeout, in seconds.

timeout *total-transmission-timeout*

Set the total transmission timeout, in seconds.

ascii Shorthand for "mode ascii"

binary
Shorthand for "mode binary"

? [*command-name ...*]
Print help information.

BUGS

Because there is no user-login or validation within the *TFTP* protocol, the remote site will probably have some sort of file-access restrictions in place. The exact methods are specific to each site and therefore difficult to document here.

NAME

time - time a command

SYNOPSIS

time command

DESCRIPTION

The given command is executed; after it is complete, *time* prints the elapsed time during the command, the time spent in the system, and the time spent in execution of the command. Times are reported in seconds.

The times are printed on the diagnostic output stream.

Time is built in to *cs**h*(1), using a different output format.

BUGS

Elapsed time is accurate to the second, while the CPU times are measured to the 100th second. Thus the sum of the CPU times can be up to a second larger than the elapsed time.

Time is a built-in command to *cs**h*(1), with a much different syntax. This command is available as "/bin/time" to *cs**h* users.

NAME

tip, *cu* - connect to a remote system

SYNOPSIS

```
tip [-v] [-speed] system-name
tip [-v] [-speed] phone-number
cu phone-number [-t] [-s speed] [-a acu] [-l line] [-#]
```

DESCRIPTION

Tip and *cu* establish a full-duplex connection to another machine, giving the appearance of being logged in directly on the remote cpu. It goes without saying that you must have a login on the machine (or equivalent) to which you wish to connect. The preferred interface is *tip*. The *cu* interface is included for those people attached to the "call UNIX" command of version 7. This manual page describes only *tip*.

Typed characters are normally transmitted directly to the remote machine (which does the echoing as well). A tilde ("~") appearing as the first character of a line is an escape signal; the following are recognized:

- ~^D ~. Drop the connection and exit (you may still be logged in on the remote machine).
- ~c [name] Change directory to name (no argument implies change to your home directory).
- ~! Escape to a shell (exiting the shell will return you to *tip*).
- ~> Copy file from local to remote. *Tip* prompts for the name of a local file to transmit.
- ~< Copy file from remote to local. *Tip* prompts first for the name of the file to be sent, then for a command to be executed on the remote machine.
- ~p from [to]
Send a file to a remote UNIX host. The put command causes the remote UNIX system to run the command string "cat > 'to'", while *tip* sends it the "from" file. If the "to" file isn't specified the "from" file name is used. This command is actually a UNIX specific version of the "~>" command.
- ~t from [to]
Take a file from a remote UNIX host. As in the put command the "to" file defaults to the "from" file name if it isn't specified. The remote host executes the command string "cat 'from';echo ^A" to send the file to *tip*.
- ~| Pipe the output from a remote command to a local UNIX process. The command string sent to the local UNIX system is processed by the shell.
- ~# Send a BREAK to the remote system. For systems which don't support the necessary *ioctl* call the break is simulated by a sequence of line speed changes and DEL characters.
- ~s Set a variable (see the discussion below).
- ~^Z Stop *tip* (only available with job control).
- ~? Get a summary of the tilde escapes

Tip uses the file */etc/remote* to find how to reach a particular system and to find out how it should operate while talking to the system; refer to *remote(5)* for a full description. Each system has a default baud rate with which to establish a connection. If this value is not suitable, the baud rate to be used may be specified on the command line, e.g. "*tip* -300 mds".

When *tip* establishes a connection it sends out a connection message to the remote system; the default value, if any, is defined in `/etc/remote`.

When *tip* prompts for an argument (e.g. during setup of a file transfer) the line typed may be edited with the standard erase and kill characters. A null line in response to a prompt, or an interrupt, will abort the dialogue and return you to the remote machine.

Tip guards against multiple users connecting to a remote system by opening modems and terminal lines with exclusive access, and by honoring the locking protocol used by *uucp*(1C).

During file transfers *tip* provides a running count of the number of lines transferred. When using the `~>` and `~<` commands, the "eofread" and "eofwrite" variables are used to recognize end-of-file when reading, and specify end-of-file when writing (see below). File transfers normally depend on tandem mode for flow control. If the remote system does not support tandem mode, "echocheck" may be set to indicate *tip* should synchronize with the remote system on the echo of each transmitted character.

When *tip* must dial a phone number to connect to a system it will print various messages indicating its actions. *Tip* supports the DEC DN-11 and Racal-Vadic 831 auto-call-units; the DEC DF02 and DF03, Ventel 212+, Racal-Vadic 3451, and Bizcomp 1031 and 1032 integral call unit/modems.

VARIABLES

Tip maintains a set of *variables* which control its operation. Some of these variable are read-only to normal users (root is allowed to change anything of interest). Variables may be displayed and set through the "s" escape. The syntax for variables is patterned after *vi*(1) and *Mail*(1). Supplying "all" as an argument to the set command displays all variables readable by the user. Alternatively, the user may request display of a particular variable by attaching a "?" to the end. For example "escape?" displays the current escape character.

Variables are numeric, string, character, or boolean values. Boolean variables are set merely by specifying their name; they may be reset by prepending a '!' to the name. Other variable types are set by concatenating an '=' and the value. The entire assignment must not have any blanks in it. A single set command may be used to interrogate as well as set a number of variables. Variables may be initialized at run time by placing set commands (without the "~s" prefix in a file `.tiprc` in one's home directory). The `-v` option causes *tip* to display the sets as they are made. Certain common variables have abbreviations. The following is a list of common variables, their abbreviations, and their default values.

beautify

(bool) Discard unprintable characters when a session is being scripted; abbreviated *be*.

baudrate

(num) The baud rate at which the connection was established; abbreviated *ba*.

dialtimeout

(num) When dialing a phone number, the time (in seconds) to wait for a connection to be established; abbreviated *dial*.

echocheck

(bool) Synchronize with the remote host during file transfer by waiting for the echo of the last character transmitted; default is *off*.

eofread

(str) The set of characters which signify and end-of-tranmission during a `~<` file transfer command; abbreviated *eofr*.

eofwrite

(str) The string sent to indicate end-of-transmission during a `~>` file transfer command; abbreviated *eofw*.

eol

(str) The set of characters which indicate an end-of-line. *Tip* will recognize escape characters only after an end-of-line.

escape

(char) The command prefix (escape) character; abbreviated *es*; default value is '^'.

exceptions

(str) The set of characters which should not be discarded due to the beautification switch; abbreviated *ex*; default value is "\t\n\f\b".

force

(char) The character used to force literal data transmission; abbreviated *fo*; default value is '^P'.

framesize

(num) The amount of data (in bytes) to buffer between file system writes when receiving files; abbreviated *fr*.

host

(str) The name of the host to which you are connected; abbreviated *ho*.

prompt

(char) The character which indicates an end-of-line on the remote host; abbreviated *pr*; default value is '\n'. This value is used to synchronize during data transfers. The count of lines transferred during a file transfer command is based on receipt of this character.

raise

(bool) Upper case mapping mode; abbreviated *ra*; default value is *off*. When this mode is enabled, all lower case letters will be mapped to upper case by *tip* for transmission to the remote machine.

raisechar

(char) The input character used to toggle upper case mapping mode; abbreviated *rc*; default value is '^A'.

record

(str) The name of the file in which a session script is recorded; abbreviated *rec*; default value is "tip.record".

script

(bool) Session scripting mode; abbreviated *sc*; default is *off*. When *script* is *true*, *tip* will record everything transmitted by the remote machine in the script record file specified in *record*. If the *beautify* switch is on, only printable ASCII characters will be included in the script file (those characters between 040 and 0177). The variable *exceptions* is used to indicate characters which are an exception to the normal beautification rules.

tabexpand

(bool) Expand tabs to spaces during file transfers; abbreviated *tab*; default value is *false*. Each tab is expanded to 8 spaces.

verbose

(bool) Verbose mode; abbreviated *verb*; default is *true*. When verbose mode is enabled, *tip* prints messages while dialing, shows the current number of lines transferred during a file transfer operations, and more.

SHELL

(str) The name of the shell to use for the ~! command; default value is "/bin/sh", or taken from the environment.

HOME

(str) The home directory to use for the `~c` command; default value is taken from the environment.

FILES

<code>/etc/remote</code>	global system descriptions
<code>/etc/phones</code>	global phone number data base
<code>\${REMOTE}</code>	private system descriptions
<code>\${PHONES}</code>	private phone numbers
<code>~/tiprc</code>	initialization file.
<code>/usr/spool/uucp/LCK..*</code>	lock file to avoid conflicts with <i>uucp</i>

DIAGNOSTICS

Diagnostics are, hopefully, self explanatory.

SEE ALSO

`remote(5)`, `phones(5)`

BUGS

The full set of variables is undocumented and should, probably, be pared down.

NAME

`tk` - paginator for the Tektronix 4014

SYNOPSIS

`tk [-t] [-N] [-pL] [file]`

DESCRIPTION

The output of *tk* is intended for a Tektronix 4014 terminal. *Tk* arranges for 66 lines to fit on the screen, divides the screen into *N* columns, and contributes an eight space page offset in the (default) single-column case. Tabs, spaces, and backspaces are collected and plotted when necessary. Teletype Model 37 half- and reverse-line sequences are interpreted and plotted. At the end of each page *tk* waits for a newline (empty line) from the keyboard before continuing on to the next page. In this wait state, the command *!command* will send the *command* to the shell.

The command line options are:

- t Don't wait between pages; for directing output into a file.
- N Divide the screen into *N* columns and wait after the last column.
- pL Set page length to *L* lines.

SEE ALSO

`pr(1)`

NAME

tmail - print out mail messages, most recent first

SYNOPSIS

tmail [username] [mboxfile]

DESCRIPTION

Tmail prints mail messages in reverse order (most recent first). If no argument is given, *tmail* looks in your system maildrop (*/usr/spool/mail/\$USER*). An argument which is a valid *username* causes *tmail* to look in that person's maildrop; otherwise the argument should be the name of a "mailbox" file.

SEE ALSO

tac(1), *cat*(1).

BUGS

Should handle multiple arguments.

NAME

`touch` - update date last modified of a file

SYNOPSIS

`touch [-c] [-f] file ...`

DESCRIPTION

Touch attempts to set the modified date of each *file*. If a *file* exists, this is done by reading a character from the file and writing it back. If a *file* does not exist, an attempt will be made to create it unless the `-c` option is specified. The `-f` option will attempt to force the touch in spite of read and write permissions on a *file*.

SEE ALSO

`utimes(2)`

NAME

tr - translate characters

SYNOPSIS

tr [-cds] [string1 [string2]]

DESCRIPTION

Tr copies the standard input to the standard output with substitution or deletion of selected characters. Input characters found in *string1* are mapped into the corresponding characters of *string2*. When *string2* is short it is padded to the length of *string1* by duplicating its last character. Any combination of the options -cds may be used: -c complements the set of characters in *string1* with respect to the universe of characters whose ASCII codes are 01 through 0377 octal; -d deletes all input characters in *string1*; -s squeezes all strings of repeated output characters that are in *string2* to single characters.

In either string the notation *a-b* means a range of characters from *a* to *b* in increasing ASCII order. The character '\ ' followed by 1, 2 or 3 octal digits stands for the character whose ASCII code is given by those digits. A '\ ' followed by any other character stands for that character.

The following example creates a list of all the words in 'file1' one per line in 'file2', where a word is taken to be a maximal string of alphabetic characters. The second string is quoted to protect '\ ' from the Shell. 012 is the ASCII code for newline.

```
tr -cs A-Za-z '\012' <file1 >file2
```

SEE ALSO

ed(1), ascii(7), expand(1)

BUGS

Won't handle ASCII NUL in *string1* or *string2*; always deletes NUL from input.

NAME

`trman` - translate version 6 manual macros to version 7 macros

SYNOPSIS

`trman` [file]

DESCRIPTION

Trman reads the input file, which should be `nroff/troff` input and attempts to translate the version 6 manual sections therein to version 7 format. It is largely successful, but seems to have trouble with indented paragraphs and complicated font control. You should expect to have to fix up long sections by hand somewhat.

SEE ALSO

`man(7)`

BUGS

NAME

troff, **nroff** – text formatting and typesetting

SYNOPSIS

troff [option] ... [file] ...

nroff [option] ... [file] ...

DESCRIPTION

Troff formats text in the named *files* for printing on a Graphic Systems C/A/T phototypesetter; *nroff* is used for typewriter-like devices. Their capabilities are described in the *Nroff/Troff user's manual*.

If no *file* argument is present, the standard input is read. An argument consisting of a single minus (-) is taken to be a file name corresponding to the standard input. The options, which may appear in any order so long as they appear before the files, are:

-olist Print only pages whose page numbers appear in the comma-separated *list* of numbers and ranges. A range *N-M* means pages *N* through *M*; an initial *-N* means from the beginning to page *N*; and a final *N-* means from *N* to the end.

-nN Number first generated page *N*.

-sN Stop every *N* pages. *Nroff* will halt prior to every *N* pages (default *N=1*) to allow paper loading or changing, and will resume upon receipt of a newline. *Troff* will stop the phototypesetter every *N* pages, produce a trailer to allow changing cassettes, and resume when the typesetter's start button is pressed.

-mname Prepend the macro file */usr/lib/tmac/tmac.name* to the input *files*.

-raN Set register *a* (one-character) to *N*.

-i Read standard input after the input files are exhausted.

-q Invoke the simultaneous input-output mode of the **rd** request.

Troff only

-t Direct output to the standard output instead of the phototypesetter.

-f Refrain from feeding out paper and stopping phototypesetter at the end of the run.

-w Wait until phototypesetter is available, if currently busy.

-b Report whether the phototypesetter is busy or available. No text processing is done.

-a Send a printable ASCII approximation of the results to the standard output.

-pN Print all characters in point size *N* while retaining all prescribed spacings and motions, to reduce phototypesetter elapsed time.

-Ffontdir

The directory *fontdir* contains the font width tables */usr/lib/fonts*. This option can be used to produce output for devices besides the phototypesetter.

If the file */usr/adm/tracct* is writable, *troff* keeps phototypesetter accounting records there. The integrity of that file may be secured by making *troff* a 'set user-id' program.

FILES

*/tmp/ta** temporary file
*/usr/lib/tmac/tmac.** standard macro files

/usr/lib/term/*	terminal driving tables for <i>nroff</i>
/usr/lib/font/*	font width tables for <i>troff</i>
/dev/cat	phototypesetter
/usr/adm/tracct	accounting statistics for /dev/cat

SEE ALSO

J. F. Ossanna, *Nroff/Troff user's manual*

B. W. Kernighan, *A TROFF Tutorial*

eqn(1), tbl(1), ms(7), me(7), man(7), col(1)

NAME

true, false – provide truth values

SYNOPSIS

true

false

DESCRIPTION

True and *false* are usually used in a Bourne shell script. They test for the appropriate status "true" or "false" before running (or failing to run) a list of commands.

EXAMPLE

```
while true
do
    command list
done
```

SEE ALSO

csh(1), sh(1), false(1)

DIAGNOSTICS

True has exit status zero.

NAME

`tset` - terminal dependent initialization

SYNOPSIS

`tset` [options] [`-m` [ident][test baudrate]:type] ... [type]

`reset` ...

DESCRIPTION

Tset sets up your terminal when you first log in to a UNIX system. It does terminal dependent processing such as setting erase and kill characters, setting or resetting delays, sending any sequences needed to properly initialize the terminal, and the like. It first determines the *type* of terminal involved, and then does necessary initializations and mode settings. The *type* of terminal attached to each UNIX port is specified in the */etc/ttytype* database. Type names for terminals may be found in the *termcap*(5) database. If a port is not wired permanently to a specific terminal (not hardwired) it will be given an appropriate generic identifier such as *dialup*.

In the case where no arguments are specified, *tset* simply reads the terminal type out of the environment variable `TERM` and re-initializes the terminal. The rest of this manual concerns itself with mode and environment initialization, typically done once at login, and options used at initialization time to determine the terminal type and set up terminal modes.

When used in a startup script (*.profile* for *sh*(1) users or *.login* for *csh*(1) users) it is desirable to give information about the type of terminal you will usually use on ports which are not hardwired. These ports are identified in */etc/ttytype* as *dialup* or *plugboard* or *arpanet*, etc. To specify what terminal type you usually use on these ports, the `-m` (map) option flag is followed by the appropriate port type identifier, an optional baud rate specification, and the terminal type. (The effect is to "map" from some conditions to a terminal type, that is, to tell *tset* "If I'm on this kind of port, guess that I'm on that kind of terminal".) If more than one mapping is specified, the first applicable mapping prevails. A missing port type identifier matches all identifiers. Any of the alternate generic names given in *termcap* may be used for the identifier.

A *baudrate* is specified as with *stty*(1), and is compared with the speed of the diagnostic output (which should be the control terminal). The baud rate *test* may be any combination of: `>`, `@`, `<`, and `!`; `@` means "at" and `!` inverts the sense of the test. To avoid problems with metacharacters, it is best to place the entire argument to `-m` within `"` characters; users of *csh*(1) must also put a `\` before any `!` used here.

Thus

```
tset -m 'dialup>300:adm3a' -m dialup:dw2 -m 'plugboard:?adm3a'
```

causes the terminal type to be set to an *adm3a* if the port in use is a dialup at a speed greater than 300 baud; to a *dw2* if the port is (otherwise) a dialup (i.e. at 300 baud or less). (NOTE: the examples given here appear to take up more than one line, for text processing reasons. When you type in real *tset* commands, you must enter them entirely on one line.) If the *type* finally determined by *tset* begins with a question mark, the user is asked if s/he really wants that type. A null response means to use that type; otherwise, another type can be entered which will be used instead. Thus, in the above case, the user will be queried on a plugboard port as to whether they are actually using an *adm3a*.

If no mapping applies and a final *type* option, not preceded by a `-m`, is given on the command line then that type is used; otherwise the identifier found in the */etc/ttytype* database will be taken to be the terminal type. This should always be the case for hardwired ports.

It is usually desirable to return the terminal type, as finally determined by *tset*, and information about the terminal's capabilities to a shell's environment. This can be done using the *-o* option; using the Bourne shell, *sh*(1):

```
export TERM; TERM=`tset - options...`
```

or using the C shell, *cs*h(1):

```
setenv TERM `tset - options...`
```

With *cs*h it is convenient to make an alias in your *.cshrc*:

```
alias tset `setenv TERM `tset - \!*`
```

Either of these aliases allow the command

```
tset 2621
```

to be invoked at any time from your login *cs*h. **Note to Bourne Shell users:** It is **not** possible to get this aliasing effect with a shell script, because shell scripts cannot set the environment of their parent. (If a process could set its parent's environment, none of this nonsense would be necessary in the first place.)

These commands cause *tset* to place the name of your terminal in the variable *TERM* in the environment; see *environ*(7).

Once the terminal type is known, *tset* engages in terminal driver mode setting. This normally involves sending an initialization sequence to the terminal, setting the single character erase (and optionally the line-kill (full line erase)) characters, and setting special character delays. Tab and newline expansion are turned off during transmission of the terminal initialization sequence.

On terminals that can backspace but not overstrike (such as a CRT), and when the erase character is the default erase character ('#' on standard systems), the erase character is changed to *BACKSPACE* (Control-H).

The options are:

- ec* set the erase character to be the named character *c* on all terminals, the default being the backspace character on the terminal, usually *^H*. The character *c* can either be typed directly, or entered using the hat notation used here.
- kc* is similar to *-e* but for the line kill character rather than the erase character; *c* defaults to *^X* (for purely historical reasons). The kill character is left alone if *-k* is not specified. The hat notation can also be used for this option.
- The name of the terminal finally decided upon is output on the standard output. This is intended to be captured by the shell and placed in the environment variable *TERM*.
- n* On systems with the Berkeley 4BSD tty driver, specifies that the new tty driver modes should be initialized for this terminal. For a CRT, the *CRTERASE* and *CRTKILL* modes are set only if the baud rate is 1200 or greater. See *tty*(4) for more detail.
- I* suppresses transmitting terminal initialization strings.
- Q* suppresses printing the "Erase set to" and "Kill set to" messages.

If *tset* is invoked as *reset*, it will set cooked and echo modes, turn off *cbreak* and raw modes, turn on newline translation, and restore special characters to a sensible state before any terminal dependent processing is done. Any special character that is found to be *NULL* or "-1" is reset to its default value.

This is most useful after a program dies leaving a terminal in a funny state. You may have to type "*<LF>reset<LF>*" to get it to work since *<CR>* may not work in this state. Often none of this will echo.

EXAMPLES

These examples all assume the Bourne shell and use the `-` option. If you use `cs`, use one of the variations described above. Note that a typical use of `tset` in a `.profile` or `.login` will also use the `-e` and `-k` options, and often the `-n` or `-Q` options as well. These options have not been included here to keep the examples small. (NOTE: some of the examples given here appear to take up more than one line, for text processing reasons. When you type in real `tset` commands, you must enter them entirely on one line.)

At the moment, you are on a 2621. This is suitable for typing by hand but not for a `.profile`, unless you are *always* on a 2621.

```
export TERM; TERM=`tset - 2621`
```

You have an h19 at home which you dial up on, but your office terminal is hardwired and known in `/etc/ttytype`.

```
export TERM; TERM=`tset --m dialup:h19`
```

You have a switch which connects everything to everything, making it nearly impossible to key on what port you are coming in on. You use a vt100 in your office at 9600 baud, and dial up to switch ports at 1200 baud from home on a 2621. Sometimes you use someone else's terminal at work, so you want it to ask you to make sure what terminal type you have at high speeds, but at 1200 baud you are always on a 2621. Note the placement of the question mark, and the quotes to protect the greater than and question mark from interpretation by the shell.

```
export TERM; TERM=`tset --m 'switch>1200:vt100' -m 'switch<=1200:2621`
```

All of the above entries will fall back on the terminal type specified in `/etc/ttytype` if none of the conditions hold. The following entry is appropriate if you always dial up, always at the same baud rate, on many different kinds of terminals. Your most common terminal is an adm3a. It always asks you what kind of terminal you are on, defaulting to adm3a.

```
export TERM; TERM=`tset - ?adm3a`
```

If the file `/etc/ttytype` is not properly installed and you want to key entirely on the baud rate, the following can be used:

```
export TERM; TERM=`tset --m '>1200:vt100' 2621`
```

Here is a fancy example to illustrate the power of `tset` and to hopelessly confuse anyone who has made it this far. You dial up at 1200 baud or less on a concept100, sometimes over switch ports and sometimes over regular dialups. You use various terminals at speeds higher than 1200 over switch ports, most often the terminal in your office, which is a vt100. However, sometimes you log in from the university you used to go to, over the ARPANET; in this case you are on an ALTO emulating a dm2500. You also often log in on various hardwired ports, such as the console, all of which are properly entered in `/etc/ttytype`. You want your erase character set to control H, your kill character set to control U, and don't want `tset` to print the "Erase set to Backspace, Kill set to Control U" message.

```
export TERM; TERM=`tset -e -k^U -Q - -m 'switch<=1200:concept100' -m 'switch:vt100' -m dialup:concept100 -m arpanet:dm2500`
```

FILES

`/etc/ttytype` port name to terminal type mapping database
`/etc/termcap` terminal capability database

SEE ALSO

cs(1), sh(1), stty(1), ttytype(5), termcap(5), environ(7)

AUTHORS

Eric Allman
David Wasley
Mark Horton

BUGS

The *tset* command is one of the first commands a user must master when getting started on a UNIX system. Unfortunately, it is one of the most complex, largely because of the extra effort the user must go through to get the environment of the login shell set. Something needs to be done to make all this simpler, either the *login*(1) program should do this stuff, or a default shell alias should be made, or a way to set the environment of the parent should exist.

NAME

tsort - topological sort

SYNOPSIS

tsort [file]

DESCRIPTION

Tsort produces on the standard output a totally ordered list of items consistent with a partial ordering of items mentioned in the input *file*. If no *file* is specified, the standard input is understood.

The input consists of pairs of items (nonempty strings) separated by blanks. Pairs of different items indicate ordering. Pairs of identical items indicate presence, but not ordering.

SEE ALSO

lorder(1)

DIAGNOSTICS

Odd data: there is an odd number of fields in the input file.

BUGS

Uses a quadratic algorithm; not worth fixing for the typical use of ordering a library archive file.

NAME

`tty` - get terminal name

SYNOPSIS

`tty [-s]`

DESCRIPTION

Tty prints the pathname of the user's terminal unless the `-s` (silent) is given. In either case, the exit value is zero if the standard input is a terminal and one if it is not.

DIAGNOSTICS

'not a tty' if the standard input file is not a terminal.

NAME

ul - do underlining

SYNOPSIS

ul [-i] [-t *terminal*] [*name* ...]

DESCRIPTION

Ul reads the named files (or standard input if none are given) and translates occurrences of underscores to the sequence which indicates underlining for the terminal in use, as specified by the environment variable **TERM**. The **-t** option overrides the terminal kind specified in the environment. The file */etc/termcap* is read to determine the appropriate sequences for underlining. If the terminal is incapable of underlining, but is capable of a standout mode then that is used instead. If the terminal can overstrike, or handles underlining automatically, *ul* degenerates to *cat*(1). If the terminal cannot underline, underlining is ignored.

The **-i** option causes *ul* to indicate underlining onto by a separate line containing appropriate dashes '-'; this is useful when you want to look at the underlining which is present in an *nroff* output stream on a crt-terminal.

SEE ALSO

man(1), *nroff*(1), *colcrt*(1)

AUTHOR

Mark Horton wrote *ul*. The **-i** option was originally a option of the editor *ex*(1), then an *iul* command.

BUGS

Nroff usually outputs a series of backspaces and underlines intermixed with the text to indicate underlining. No attempt is made to optimize the backward motion.

NAME

`unifdef` - remove `ifdef`'ed lines

SYNOPSIS

`unifdef` [`-t -l -c -Dsym -Usym -idsym -iusym`] ... [`file`]

DESCRIPTION

Unifdef is useful for removing `ifdef`'ed lines from a file while otherwise leaving the file alone. *Unifdef* is like a stripped-down C preprocessor: it is smart enough to deal with the nested `ifdefs`, comments, single and double quotes of C syntax so that it can do its job, but it doesn't do any including or interpretation of macros. Neither does it strip out comments, though it recognizes and ignores them. You specify which symbols you want defined `-Dsym` or undefined `-Usym` and the lines inside those `ifdefs` will be copied to the output or removed as appropriate. The `ifdef`, `ifndef`, `else`, and `endif` lines associated with *sym* will also be removed. `ifdefs` involving symbols you don't specify are untouched and copied out along with their associated `ifdef`, `else`, and `endif` lines. If an `ifdef X` occurs nested inside another `ifdef X`, then the inside `ifdef` is treated as if it were an unrecognized symbol. If the same symbol appears in more than one argument, only the first occurrence is significant.

The `-l` option causes *unifdef* to replace removed lines with blank lines instead of deleting them.

If you use `ifdefs` to delimit non-C lines, such as comments or code which is under construction, then you must tell *unifdef* which symbols are used for that purpose so that it won't try to parse for quotes and comments in those `ifdef`'ed lines. You specify that you want the lines inside certain `ifdefs` to be ignored but copied out with `-idsym` and `-iusym` similar to `-Dsym` and `-Usym` above.

If you want to use *unifdef* for plain text (not C code), use the `-t` option. This makes *unifdef* refrain from attempting to recognize comments and single and double quotes.

Unifdef copies its output to *stdout* and will take its input from *stdin* if no *file* argument is given. If the `-c` argument is specified, then the operation of *unifdef* is complemented, i.e. the lines that would have been removed or blanked are retained and vice versa.

SEE ALSO

`diff(1)`

DIAGNOSTICS

Premature EOF, inappropriate `else` or `endif`.

Exit status is 0 if output is exact copy of input, 1 if not, 2 if trouble.

BUGS

Does not know how to deal with *cpp* constructs such as

```
#if defined(X) || defined(Y)
```

AUTHOR

Dave Yost

NAME

uniq - report repeated lines in a file

SYNOPSIS

uniq [**-udc** [**+n**] [**-n**]] [input [output]]

DESCRIPTION

Uniq reads the input file comparing adjacent lines. In the normal case, the second and succeeding copies of repeated lines are removed; the remainder is written on the output file. Note that repeated lines must be adjacent in order to be found; see *sort*(1). If the **-u** flag is used, just the lines that are not repeated in the original file are output. The **-d** option specifies that one copy of just the repeated lines is to be written. The normal mode output is the union of the **-u** and **-d** mode outputs.

The **-c** option supersedes **-u** and **-d** and generates an output report in default style but with each line preceded by a count of the number of times it occurred.

The *n* arguments specify skipping an initial portion of each line in the comparison:

- n** The first *n* fields together with any blanks before each are ignored. A field is defined as a string of non-space, non-tab characters separated by tabs and spaces from its neighbors.
- +n** The first *n* characters are ignored. Fields are skipped before characters.

SEE ALSO

sort(1), *comm*(1)

NAME

units - conversion program

SYNOPSIS

units

DESCRIPTION

Units converts quantities expressed in various standard scales to their equivalents in other scales. It works interactively in this fashion:

You have: inch

You want: cm

* 2.54000e+00

/ 3.93701e-01

A quantity is specified as a multiplicative combination of units optionally preceded by a numeric multiplier. Powers are indicated by suffixed positive integers, division by the usual sign:

You have: 15 pounds force/in²

You want: atm

* 1.02069e+00

/ 9.79790e-01

Units only does multiplicative scale changes. Thus it can convert Kelvin to Rankine, but not Centigrade to Fahrenheit. Most familiar units, abbreviations, and metric prefixes are recognized, together with a generous leavening of exotica and a few constants of nature including:

pi	ratio of circumference to diameter
c	speed of light
e	charge on an electron
g	acceleration of gravity
force	same as g
mole	Avogadro's number
water	pressure head per unit height of water
au	astronomical unit

'Pound' is a unit of mass. Compound names are run together, e.g. 'lightyear'. British units that differ from their US counterparts are prefixed thus: 'brgallon'. Currency is denoted 'belgiumfranc', 'britainpound', ...

For a complete list of units, 'cat /usr/lib/units'.

FILES

/usr/lib/units

BUGS

Don't base your financial plans on the currency conversions.

NAME

uptime - show how long system has been up

SYNOPSIS

uptime

DESCRIPTION

Uptime prints the current time, the length of time the system has been up, and the average number of jobs in the run queue over the last 1, 5 and 15 minutes. It is, essentially, the first line of a *w(1)* command.

FILES

/vmunix system name list

SEE ALSO

w(1)

NAME

users – compact list of users who are on the system

SYNOPSIS

users

DESCRIPTION

Users lists the login names of the users currently on the system in a compact, one-line format.

FILES

/etc/utmp

SEE ALSO

who(1)

NAME

uucp, **uulog** - unix to unix copy

SYNOPSIS

uucp [option] ... source-file ... destination-file
uulog [option] ...

DESCRIPTION

Uucp copies files named by the source-file arguments to the destination-file argument. A file name may be a path name on your machine, or may have the form

system-name!pathname

where 'system-name' is taken from a list of system names which *uucp* knows about. Shell metacharacters `?*[]` appearing in the pathname part will be expanded on the appropriate system.

Pathnames may be one of

- (1) a full pathname;
- (2) a pathname preceded by `~user`; where *user* is a userid on the specified system and is replaced by that user's login directory;
- (3) anything else is prefixed by the current directory.

If the result is an erroneous pathname for the remote system the copy will fail. If the destination-file is a directory, the last part of the source-file name is used.

Uucp preserves execute permissions across the transmission and gives 0666 read and write permissions (see *chmod(2)*).

The following options are interpreted by *uucp*.

- d** Make all necessary directories for the file copy.
- c** Use the source file when copying out rather than copying the file to the spool directory.
- m** Send mail to the requester when the copy is complete.

Uulog maintains a summary log of *uucp* and *uux(1C)* transactions in the file `/usr/spool/uucp/LOGFILE` by gathering information from partial log files named `/usr/spool/uucp/LOG.*.?`. It removes the partial log files.

The options cause *uulog* to print logging information:

- sys** Print information about work involving system *sys*.
- user** Print information about work done for the specified *user*.

FILES

`/usr/spool/uucp` - spool directory
`/usr/lib/uucp/*` - other data and program files

SEE ALSO

uux(1C), *mail(1)*
D. A. Nowitz, *Uucp Implementation Description*

WARNING

The domain of remotely accessible files can (and for obvious security reasons, usually should) be severely restricted. You will very likely not be able to fetch files by pathname; ask a responsible person on the remote system to send them to you. For the same reasons you will probably not be able to send files to arbitrary pathnames.

BUGS

All files received by *uucp* will be owned by *uucp*.

The `-m` option will only work sending files or receiving a single file. (Receiving multiple files specified by special shell characters `?*[]` will not activate the `-m` option.)

NAME

uuencode, uudecode – encode/decode a binary file for transmission via mail

SYNOPSIS

```
uuencode [ source ] remotetest | mail sys1!sys2!..!decode
uudecode [ file ]
```

DESCRIPTION

Uuencode and *uudecode* are used to send a binary file via uucp (or other) mail. This combination can be used over indirect mail links even when *uusend(1C)* is not available.

Uuencode takes the named source file (default standard input) and produces an encoded version on the standard output. The encoding uses only printing ASCII characters, and includes the mode of the file and the *remotedest* for recreation on the remote system.

Uudecode reads an encoded file, strips off any leading and trailing lines added by mailers, and recreates the original file with the specified mode and name.

The intent is that all mail to the user “decode” should be filtered through the *uudecode* program. This way the file is created automatically without human intervention. This is possible on the uucp network by either using *sendmail* or by making *rmail* be a link to *Mail* instead of *mail*. In each case, an alias must be created in a master file to get the automatic invocation of *uudecode*.

If these facilities are not available, the file can be sent to a user on the remote machine who can *uudecode* it manually.

The encode file has an ordinary text form and can be edited by any text editor to change the mode or remote name.

SEE ALSO

uuencode(5), uusend(1C), uucp(1C), uux(1C), mail(1)

AUTHOR

Mark Horton

BUGS

The file is expanded by 35% (3 bytes become 4 plus control information) causing it to take longer to transmit.

The user on the remote system who is invoking *uudecode* (often *uucp*) must have write permission on the specified file.

NAME

`uusend` - send a file to a remote host

SYNOPSIS

`uusend` [`-m mode`] sourcefile sys1!sys2!...!remotefile

DESCRIPTION

Uusend sends a file to a given location on a remote system. The system need not be directly connected to the local system, but a chain of *uucp*(1C) links needs to connect the two systems.

If the `-m` option is specified, the mode of the file on the remote end will be taken from the octal number given. Otherwise, the mode of the input file will be used.

The sourcefile can be "-", meaning to use the standard input. Both of these options are primarily intended for internal use of `uusend`.

The remotefile can include the ~userid syntax.

DIAGNOSTICS

If anything goes wrong any further away than the first system down the line, you will never hear about it.

SEE ALSO

`uux`(1C), `uucp`(1C), `uuencode`(1)

AUTHOR

Mark Horton

BUGS

This command shouldn't exist, since *uucp* should handle it.

All systems along the line must have the *uusend* command available and allow remote execution of it.

Some *uucp* systems have a bug where binary files cannot be the input to a *uux* command. If this bug exists in any system along the line, the file will show up severely munged.

NAME

`uux` - unix to unix command execution

SYNOPSIS

`uux [-] command-string`

DESCRIPTION

`Uux` will gather 0 or more files from various systems, execute a command on a specified system and send standard output to a file on a specified system.

The command-string is made up of one or more arguments that look like a shell command line, except that the command and file names may be prefixed by system-name!. A null system-name is interpreted as the local system.

File names may be one of

- (1) a full pathname;
- (2) a pathname preceded by `~xxx`; where `xxx` is a userid on the specified system and is replaced by that user's login directory;
- (3) anything else is prefixed by the current directory.

The `'-'` option will cause the standard input to the `uux` command to be the standard input to the command-string.

For example, the command

```
uux "!diff usg!/usr/dan/f1 pwba!/a4/dan/f1 > !fi.diff"
```

will get the `f1` files from the `usg` and `pwba` machines, execute a `diff` command and put the results in `f1.diff` in the local directory.

Any special shell characters such as `<>`;| should be quoted either by quoting the entire command-string, or quoting the special characters as individual arguments.

FILES

<code>/usr/spool/uucp</code>	spool directory
<code>/usr/lib/uucp/*</code>	other data and programs

SEE ALSO

`uucp(1C)`
 D. A. Nowitz, *Uucp Implementation Description*

WARNING

An installation may, and for security reasons generally will, limit the list of commands executable on behalf of an incoming request from `uux`. Typically, a restricted site will permit little other than the receipt of mail via `uux`.

BUGS

Only the first command of a shell pipeline may have a system-name!. All other commands are executed on the system of the first command.

The use of the shell metacharacter `*` will probably not do what you want it to do.

The shell tokens `<<` and `>>` are not implemented.

There is no notification of denial of execution on the remote machine.

NAME

`vers` - print version number of the kernel

SYNOPSIS

`vers` [`-v`]

DESCRIPTION

Vers prints the version number of the currently running UNIX kernel. It prints the same message that is seen at boot time. The `-v` (verbose) flag is used to determine the actual path-name of the kernel.

NAME

`vgrind` - grind nice listings of programs

SYNOPSIS

`vgrind` [**-f**] [**-**] [**-t**] [**-n**] [**-x**] [**-W**] [**-sn**] [**-h** header] [**-d** file] [**-llanguage**]
name ...

DESCRIPTION

Vgrind formats the program sources which are arguments in a nice style using *troff*(1). Comments are placed in italics, keywords in bold face, and the name of the current function is listed down the margin of each page as it is encountered.

Vgrind runs in two basic modes, filter mode or regular mode. In filter mode *vgrind* acts as a filter in a manner similar to *tbl*(1). The standard input is passed directly to the standard output except for lines bracketed by the *troff-like* macros:

`.vS` - starts processing

`.vE` - ends processing

These lines are formatted as described above. The output from this filter can be passed to *troff* for output. There need be no particular ordering with *eqn*(1) or *tbl*(1).

In regular mode *vgrind* accepts input files, processes them, and passes them to *troff*(1) for output.

In both modes *vgrind* passes any lines beginning with a decimal point without conversion.

The options are:

-f forces filter mode

- forces input to be taken from standard input (default if **-f** is specified)

-t similar to the same option in *troff* causing formatted text to go to the standard output

-n forces no keyword bolding

-x outputs the index file in a "pretty" format. The index file itself is produced whenever *vgrind* is run with a file called *index* in the current directory. The index of function definitions can then be run off by giving *vgrind* the **-x** option and the file *index* as argument.

-W forces output to the (wide) Versatec printer rather than the (narrow) Varian

-s specifies a point size to use on output (exactly the same as the argument of a `.ps`)

-h specifies a particular header to put on every output page (default is the file name)

-d specifies an alternate language definitions file (default is `/usr/lib/vgrindefs`)

-l specifies the language to use. Currently known are PASCAL (**-lp**), MODEL (**-lm**), C (**-lc** or the default), CSH (**-lcs**), SHELL (**-lsh**), RATFOR (**-lr**), and ICON (**-li**).

FILES

<code>index</code>	file where source for index is created
<code>/usr/lib/tmac/tmac.vgrind</code>	macro package
<code>/usr/lib/vfontedpr</code>	preprocessor
<code>/usr/lib/vgrindefs</code>	language descriptions

AUTHOR

Dave Presotto & William Joy

SEE ALSO

vlp(1), vtroff(1), vgrindefs(5)

BUGS

Vfontedpr assumes that a certain programming style is followed:

For **C** – function names can be preceded on a line only by spaces, tabs, or an asterisk. The parenthesized arguments must also be on the same line.

For **PASCAL** – function names need to appear on the same line as the keywords *function* or *procedure*.

For **MODEL** – function names need to appear on the same line as the keywords *is beginproc*.

If these conventions are not followed, the indexing and marginal function name comment mechanisms will fail.

More generally, arbitrary formatting styles for programs mostly look bad. The use of spaces to align source code fails miserably; if you plan to *vgrind* your program you should use tabs. This is somewhat inevitable since the font used by *vgrind* is variable width.

The mechanism of ctags in recognizing functions should be used here.

NAME

vi - screen oriented (visual) display editor based on *ex*

SYNOPSIS

vi [-t tag] [-r] [+command] [-l] [-wn] name ...

DESCRIPTION

Vi (visual) is a display oriented text editor based on *ex*(1). *Ex* and *vi* run the same code; it is possible to get to the command mode of *ex* from within *vi* and vice-versa.

The *Vi Quick Reference* card and the *Introduction to Display Editing with Vi* provide full details on using *vi*.

FILES

See *ex*(1).

SEE ALSO

ex (1), *edit* (1), "Vi Quick Reference" card, "An Introduction to Display Editing with Vi".

AUTHOR

William Joy

Mark Horton added macros to *visual* mode and is maintaining version 3

BUGS

Software tabs using **^T** work only immediately after the *autoindent*.

Left and right shifts on intelligent terminals don't make use of insert and delete character operations in the terminal.

The *wrapmargin* option can be fooled since it looks at output columns when blanks are typed. If a long word passes through the margin and onto the next line without a break, then the line won't be broken.

Insert/delete within a line can be slow if tabs are present on intelligent terminals, since the terminals need help in doing this correctly.

Saving text on deletes in the named buffers is somewhat inefficient.

The *source* command does not work when executed as **:source**; there is no way to use the **:append**, **:change**, and **:insert** commands, since it is not possible to give more than one line of input to a **:** escape. To use these on a **:global** you must **Q** to *ex* command mode, execute them, and then reenter the screen editor with *vi* or *open*.

NAME

`vmstat` - report virtual memory statistics

SYNOPSIS

`vmstat` [`-fs`] [`interval` [`count`]]

DESCRIPTION

`Vmstat` delves into the system and normally reports certain statistics kept about process, virtual memory, disk, trap and cpu activity. If given a `-f` argument, it instead reports on the number of *forks* and *vforks* since system startup and the number of pages of virtual memory involved in each kind of fork. If given a `-s` argument, it instead prints the contents of the *sum* structure, giving the total number of several kinds of paging related events which have occurred since boot.

If none of these options are given, `vmstat` will report in the first line a summary of the virtual memory activity since the system has been booted. If *interval* is specified, then successive lines are summaries over the last *interval* seconds. "`vmstat 5`" will print what the system is doing every five seconds; this is a good choice of printing interval since this is how often some of the statistics are sampled in the system; others vary every second, running the output for a while will make it apparent which are recomputed every second. If a *count* is given, the statistics are repeated *count* times. The format fields are:

Procs: information about numbers of processes in various states.

r	in run queue
b	blocked for resources (i/o, paging, etc.)
w	runnable or short sleeper (< 20 secs) but swapped

Memory: information about the usage of virtual and real memory. Virtual pages are considered active if they belong to processes which are running or have run in the last 20 seconds. A "page" here is 1024 bytes.

avm	active virtual pages
fre	size of the free list

Page: information about page faults and paging activity. These are averaged each five seconds, and given in units per second.

re	page reclaims (simulating reference bits)
pi	pages paged in
po	pages paged out
fr	pages freed per second
de	anticipated short term memory shortfall
sr	pages scanned by clock algorithm, per-second

up/hp/rk: Disk operations per second (this field is system dependent). Typically paging will be split across several of the available drives. The number under each of these is the unit number.

Faults: trap/interrupt rate averages per second over last 5 seconds.

in	(non clock) device interrupts per second
sy	system calls per second
cs	cpu context switch rate (switches/sec)



Cpu: breakdown of percentage usage of CPU time

us	user time for normal and low priority processes
sy	system time
id	cpu idle

FILES

/dev/kmem, /vmunix

SEE ALSO

The sections starting with "Interpreting system activity" in *Installing and Operating 4.2bsd*.

AUTHORS

William Joy and Ozalp Babaoglu

BUGS

There should be a screen oriented program which combines *vmstat* and *ps(1)* in real time as well as reporting on other system activity.

NAME

w - who is on and what they are doing

SYNOPSIS

w [-h] [-s] [user]

DESCRIPTION

W prints a summary of the current activity on the system, including what each user is doing. The heading line shows the current time of day, how long the system has been up, the number of users logged into the system, and the load averages. The load average numbers give the number of jobs in the run queue averaged over 1, 5 and 15 minutes.

The fields output are: the users login name, the name of the tty the user is on, the time of day the user logged on, the number of minutes since the user last typed anything, the CPU time used by all processes and their children on that terminal, the CPU time used by the currently active processes, the name and arguments of the current process.

The -h flag suppresses the heading. The -s flag asks for a short form of output. In the short form, the tty is abbreviated, the login time and cpu times are left off, as are the arguments to commands. -l gives the long output, which is the default.

If a user name is included, the output will be restricted to that user.

FILES

/etc/utmp
/dev/kmem
/dev/drum

SEE ALSO

who(1), finger(1), ps(1)

AUTHOR

Mark Horton

BUGS

The notion of the "current process" is muddy. The current algorithm is "the highest numbered process on the terminal that is not ignoring interrupts, or, if there is none, the highest numbered process on the terminal". This fails, for example, in critical sections of programs like the shell and editor, or when faulty programs running in the background fork and fail to ignore interrupts. (In cases where no process can be found, w prints "-".)

The CPU time is only an estimate, in particular, if someone leaves a background process running after logging out, the person currently on that terminal is "charged" with the time.

Background processes are not shown, even though they account for much of the load on the system.

Sometimes processes, typically those in the background, are printed with null or garbaged arguments. In these cases, the name of the command is printed in parentheses.

W does not know about the new conventions for detection of background jobs. It will sometimes find a background job instead of the right one.

NAME

wait - await completion of process

SYNOPSIS

wait

DESCRIPTION

Wait until all processes started with **&** have completed, and report on abnormal terminations.

Because the *wait(2)* system call must be executed in the parent process, the Shell itself executes *wait*, without creating a new process.

SEE ALSO

sh(1)

BUGS

Not all the processes of a 3- or more-stage pipeline are children of the Shell, and thus can't be waited for. (This bug does not apply to *cs*h(1).)

NAME

wall - write to all users

SYNOPSIS

wall

DESCRIPTION

Wall reads its standard input until an end-of-file. It then sends this message, preceded by 'Broadcast Message ...', to all logged in users.

The sender should be super-user to override any protections the users may have invoked.

FILES

/dev/tty?
/etc/utmp

SEE ALSO

mesg(1), write(1)

DIAGNOSTICS

'Cannot send to ...' when the open on a user's tty file fails.

NAME

`wc` - word count

SYNOPSIS

`wc` [`-lwc`] [name ...]

DESCRIPTION

`Wc` counts lines, words and characters in the named files, or in the standard input if no name appears. A word is a maximal string of characters delimited by spaces, tabs or newlines.

If an argument beginning with one of "lwc" is present, the specified counts (lines, words, or characters) are selected by the letters **l**, **w**, or **c**. The default is `-lwc`.

NAME

what - show what versions of object modules were used to construct a file

SYNOPSIS

what name ...

DESCRIPTION

What reads each file and searches for sequences of the form "@(#)" as inserted by the source code control system. It then prints the remainder of the string after this marker, up to a null character, newline, double quote, or ">" character.

BUGS

As SCCS is not licensed with UNIX/32V, this is a rewrite of the *what* command which is part of SCCS, and may not behave exactly the same as that command does.

NAME

whatis - describe what a command is

SYNOPSIS

whatis command ...

DESCRIPTION

Whatis looks up a given command and gives the header line from the manual section. You can then run the *man(1)* command to get more information. If the line starts 'name(section) ...' you can do 'man section name' to get the documentation for it. Try 'whatis ed' and then you should do 'man 1 ed' to get the manual.

Whatis is actually just the **-f** option to the *man(1)* command.

FILES

/usr/lib/whatis Data base

SEE ALSO

man(1), catman(8)

AUTHOR

William Joy

NAME

whereis - locate source, binary, and or manual for program

SYNOPSIS

whereis [**-sbm**] [**-u**] [**-SBM** dir ... **-f**] name ...

DESCRIPTION

Whereis locates source/binary and manuals sections for specified files. The supplied names are first stripped of leading pathname components and any (single) trailing extension of the form ".ext", e.g. ".c". Prefixes of "s." resulting from use of source code control are also dealt with. *Whereis* then attempts to locate the desired program in a list of standard places. If any of the **-b**, **-s** or **-m** flags are given then *whereis* searches only for binaries, sources or manual sections respectively (or any two thereof). The **-u** flag may be used to search for unusual entries. A file is said to be unusual if it does not have one entry of each requested type. Thus "whereis -m -u *" asks for those files in the current directory which have no documentation.

Finally, the **-B** **-M** and **-S** flags may be used to change or otherwise limit the places where *whereis* searches. The **-f** file flags is used to terminate the last such directory list and signal the start of file names.

EXAMPLE

The following finds all the files in /usr/bin which are not documented in /usr/man/man1 with source in /usr/src/cmd:

```
cd /usr/ucb
whereis -u -M /usr/man/man1 -S /usr/src/cmd -f *
```

FILES

```
/usr/src/*
/usr/{doc,man}/*
/lib, /etc, /usr/{lib,bin,ucb,old,new,local}
```

AUTHOR

William Joy

BUGS

Since the program uses *chdir*(2) to run faster, pathnames given with the **-M** **-S** and **-B** must be full; i.e. they must begin with a "/".

NAME

which - locate a program file including aliases and paths (*cs*h only)

SYNOPSIS

which [name] ...

DESCRIPTION

Which takes a list of names and looks for the files which would be executed had these names been given as commands. Each argument is expanded if it is aliased, and searched for along the user's path. Both aliases and path are taken from the user's *.cshrc* file.

FILES

~/cshrc source of aliases and path values

DIAGNOSTICS

A diagnostic is given for names which are aliased to more than a single word, or if an executable file with the argument name was not found in the path.

BUGS

Must be executed by a *cs*h, since only *cs*h's know about aliases.

NAME

who - who is on the system

SYNOPSIS

who [who-file] [**am I**]

DESCRIPTION

Who, without an argument, lists the login name, terminal name, and login time for each current UNIX user.

Without an argument, *who* examines the */etc/utmp* file to obtain its information. If a file is given, that file is examined. Typically the given file will be */usr/adm/wtmp*, which contains a record of all the logins since it was created. Then *who* lists logins, logouts, and crashes since the creation of the *wtmp* file. Each login is listed with user name, terminal name (with *'/dev/'* suppressed), and date and time. When an argument is given, logouts produce a similar line without a user name. Reboots produce a line with *'x'* in the place of the device name, and a fossil time indicative of when the system went down.

With two arguments, as in *'who am I'* (and also *'who are you'*), *who* tells who you are logged in as.

FILES

/etc/utmp

SEE ALSO

getuid(2), *utmp(5)*

NAME

whoami - print effective current user id

SYNOPSIS

whoami

DESCRIPTION

Whoami prints who you are. It works even if you are su'd, while 'who am i' does not since it uses /etc/utmp.

FILES

/etc/passwd Name data base

SEE ALSO

who(1)

NAME

whodos - display information about *dosc* users

SYNOPSIS

whodos

DESCRIPTION

Whodos is used to display information about *dosc* partition usage. It displays the Multi-link partition number, the *dosc* process PID, the tty line of the *dosc* user, the time the process started, and the user name for each available partition. The file */etc/mttys* tells the total number of available partitions.

FILES

/usr/spool/uucp/LCK..mtty, */etc/mttys*

SEE ALSO

dosc(1)

NOTE

If a *dosc* process is terminated other than by exiting, it may leave a lock file in */usr/spool/uucp*. This file must be deleted before the partition can be reused. If you suspect this has happened, verify that the PID for each partition reported as active by *whodos* is listed as an active process by *ps(1)*. If the PID does not exist, you may safely remove the lock file to allow the partition to be reused.

NAME

window - window environment

SYNOPSIS

window [**-t**] [**-f**] [**-d**] [**-e escape-char**] [**-c command**]

DESCRIPTION

Window implements a window environment on ASCII terminals.

A window is a rectangular portion of the physical terminal screen associated with a set of processes. Its size and position can be changed by the user at any time. Processes communicate with their window in the same way they normally interact with a terminal--through their standard input, output, and diagnostic file descriptors. The window program handles the details of redirecting input and output to and from the windows. At any one time, only one window can receive input from the keyboard, but all windows can simultaneously send output to the display.

Windows can overlap and are framed as necessary. Each window is named by one of the digits "1" to "9". This one character identifier, as well as a user definable label string, are displayed with the window on the top edge of its frame. A window can be designated to be in the *foreground*, in which case it will always be on top of all normal, non-foreground windows, and can be covered only by other foreground windows. A window need not be completely within the edges of the terminal screen. Thus a large window (possibly larger than the screen) may be positioned to show only a portion of its full size.

Each window has a cursor and a set of control functions. Most intelligent terminal operations such as line and character deletion and insertion are supported. Display modes such as underlining and reverse video are available if they are supported by the terminal. In addition, similar to terminals with multiple pages of memory, each window has a text buffer which can have more lines than the window itself.

OPTIONS

When *window* starts up, the commands (see long commands below) contained in the file *.windowrc* in the user's home directory are executed. If it does not exist, two equal sized windows spanning the terminal screen are created by default.

The command line options are

-t Turn on terse mode (see *terse* command below).

-f Fast. Don't perform any startup action.

-d Ignore *.windowrc* and create the two default windows instead.

-e escape-char

Set the escape character to *escape-char*. *Escape-char* can be a single character, or in the form *^X* where *X* is any character, meaning control-*X*.

-c command

Execute the string *command* as a long command (see below) before doing anything else.

PROCESS ENVIRONMENT

With each newly created window, a shell program is spawned with its process environment tailored to that window. Its standard input, output, and diagnostic file descriptors are bound to one end of either a pseudo-terminal (*pty* (4)) or a UNIX domain socket (*socketpair* (4)). If a pseudo-terminal is used, then its special characters and modes (see *stty* (1)) are copied from

the physical terminal. A *termcap* (5) entry tailored to this window is created and passed as environment (*environ* (5)) variable *TERMCAP*. The *termcap* entry contains the window's size and characteristics as well as information from the physical terminal, such as the existence of underline, reverse video, and other display modes, and the codes produced by the terminal's function keys, if any. In addition, the window size attributes of the pseudo-terminal are set to reflect the size of this window, and updated whenever it is changed by the user. In particular, the editor *vi* (1) uses this information to redraw its display.

OPERATION

During normal execution, *window* can be in one of two states: conversation mode and command mode. In conversation mode, the terminal's real cursor is placed at the cursor position of a particular window--called the current window--and input from the keyboard is sent to the process in that window. The current window is always on top of all other windows, except those in foreground. In addition, it is set apart by highlighting its identifier and label in reverse video.

Typing *window*'s escape character (normally ^P) in conversation mode switches it into command mode. In command mode, the top line of the terminal screen becomes the command prompt window, and *window* interprets input from the keyboard as commands to manipulate windows.

There are two types of commands: short commands are usually one or two key strokes; long commands are strings either typed by the user in the command window (see the ":" command below), or read from a file (see *source* below).

SHORT COMMANDS

Below, # represents one of the digits "1" to "9" corresponding to the windows 1 to 9. ^X means control-X, where X is any character. In particular, ^^ is control-^. *Escape* is the escape key, or ^[.

- # Select window # as the current window and return to conversation mode.
- %# Select window # but stay in command mode.
- ^^ Select the previous window and return to conversation mode. This is useful for toggling between two windows.

escape

- Return to conversation mode.
- ^P Return to conversation mode and write ^P to the current window. Thus, typing two ^P's in conversation mode sends one to the current window. If the *window* escape is changed to some other character, that character takes the place of ^P here.
- ? List a short summary of commands.
- ^L Redraw the screen.
- q Exit *window*. Confirmation is requested.
- ^Z Suspend *window*.
- w Create a new window. The user is prompted for the positions of the upper left and lower right corners of the window. The cursor is placed on the screen and the keys "h", "j", "k", and "l" move the cursor left, down, up, and right, respectively. The keys "H", "J", "K", and "L" move the cursor to the respective limits of the screen. Typing a number before the movement keys repeats the movement that number of times. Return enters the cursor position as the upper left corner of the window. The lower right corner is entered in the same manner. During this process, the placement

of the new window is indicated by a rectangular box drawn on the screen, corresponding to where the new window will be framed. Typing escape at any point cancels this command.

This window becomes the current window, and is given the first available ID. The default buffer size is used (see *nline* command below).

Only fully visible windows can be created this way.

- c#** Close window #. The process in the window is sent the hangup signal (see *kill* (1)). *Csh* (1) should handle this signal correctly and cause no problems.
- m#** Move window # to another location. A box in the shape of the window is drawn on the screen to indicate the new position of the window, and the same keys as those for the *w* command are used to position the box. The window can be moved partially off-screen.
- M#** Move window # to its previous position.
- s#** Change the size of window #. The user is prompted to enter the new lower right corner of the window. A box is drawn to indicate the new window size. The same keys used in *w* and *m* are used to enter the position.
- S#** Change window # to its previous size.
- ^Y** Scroll the current window up by one line.
- ^E** Scroll the current window down by one line.
- ^U** Scroll the current window up by half the window size.
- ^D** Scroll the current window down by half the window size.
- ^B** Scroll the current window up by the full window size.
- ^F** Scroll the current window down by the full window size.
- h** Move the cursor of the current window left by one column.
- j** Move the cursor of the current window down by one line.
- k** Move the cursor of the current window up by one line.
- l** Move the cursor of the current window right by one column.
- ^S** Stop output in the current window.
- ^Q** Start output in the current window.
- :** Enter a line to be executed as long commands. Normal line editing characters (erase character, erase word, erase line) are supported.

LONG COMMANDS

Long commands are a sequence of statements parsed much like a programming language, with a syntax similar to that of C. Numeric and string expressions and variables are supported, as well as conditional statements.

There are two data types: string and number. A string is a sequence of letters or digits beginning with a letter. “_” and “.” are considered letters. Alternately, non-alphanumeric characters can be included in strings by quoting them in “” or escaping them with “\”. In addition, the “\” sequences of C are supported, both inside and outside quotes (e.g., “\n” is a new line, “\r” a carriage return). For example, these are legal strings: abcde01234, “&#x26;#”, ab“\$#”cd, ab\\$\#cd, “/usr/ucb/window”.

A number is an integer value in one of three forms: a decimal number, an octal number preceded by "0", or a hexadecimal number preceded by "0x" or "0X". The natural machine integer size is used (i.e., the signed integer type of the C compiler). As in C, a non-zero number represents a boolean true.

The character "#" begins a comment which terminates at the end of the line.

A statement is either a conditional or an expression. Expression statements are terminated with a new line or ";". To continue an expression on the next line, terminate the first line with "\".

CONDITIONAL STATEMENT

Window has a single control structure: the fully bracketed if statement in the form

```
if <expr> then
    <statement>
    ...
elseif <expr> then
    <statement>
    ...
else
    <statement>
    ...
endif
```

The *else* and *elseif* parts are optional, and the latter can be repeated any number of times. *<Expr>* must be numeric.

EXPRESSIONS

Expressions in *window* are similar to those in the C language, with most C operators supported on numeric operands. In addition, some are overloaded to operate on strings.

When an expression is used as a statement, its value is discarded after evaluation. Therefore, only expressions with side effects (assignments and function calls) are useful as statements.

Single valued (no arrays) variables are supported, of both numeric and string values. Some variables are predefined. They are listed below.

The operators in order of increasing precedence:

<expr1> = <expr2>

Assignment. The variable of name *<expr1>*, which must be string valued, is assigned the result of *<expr2>*. Returns the value of *<expr2>*.

<expr1> ? <expr2> : <expr3>

Returns the value of *<expr2>* if *<expr1>* evaluates true (non-zero numeric value); returns the value of *<expr3>* otherwise. Only one of *<expr2>* and *<expr3>* is evaluated. *<Expr1>* must be numeric.

<expr1> || <expr2>

Logical or. Numeric values only. Short circuit evaluation is supported (i.e., if *<expr1>* evaluates true, then *<expr2>* is not evaluated).

<expr1> && <expr2>

Logical and with short circuit evaluation. Numeric values only.

<expr1> | <expr2>

Bitwise or. Numeric values only.

<expr1> ^ <expr2>

Bitwise exclusive or. Numeric values only.

<expr1> & <expr2>

Bitwise and. Numeric values only.

<expr1> == <expr2>, <expr1> != <expr2>

Comparison (equal and not equal, respectively). The boolean result (either 1 or 0) of the comparison is returned. The operands can be numeric or string valued. One string operand forces the other to be converted to a string in necessary.

<expr1> < <expr2>, <expr1> > <expr2>, <expr1> >= <expr2>, <expr1> <= <expr2>

Less than, greater than, less than or equal to, greater than or equal to. Both numeric and string values, with automatic conversion as above.

<expr1> << <expr2>, <expr1> >> <expr2>

If both operands are numbers, *<expr1>* is bit shifted left (or right) by *<expr2>* bits. If *<expr1>* is a string, then its first (or last) *<expr2>* characters are returned (if *<expr2>* is also a string, then its length is used in place of its value).

<expr1> + <expr2>, <expr1> - <expr2>

Addition and subtraction on numbers. For "+", if one argument is a string, then the other is converted to a string, and the result is the concatenation of the two strings.

<expr1> * <expr2>, <expr1> / <expr2>, <expr1> % <expr2>

Multiplication, division, modulo. Numbers only.

-<expr>, ~<expr>, !<expr>, \$<expr>, \$!<expr>

The first three are unary minus, bitwise complement and logical complement on numbers only. The operator, "\$", takes *<expr>* and returns the value of the variable of that name. If *<expr>* is numeric with value *n* and it appears within an alias macro (see below), then it refers to the *n*th argument of the alias invocation. "\$?" tests for the existence of the variable *<expr>*, and returns 1 if it exists or 0 otherwise.

<expr>(<arglist>)

Function call. *<Expr>* must be a string that is the unique prefix of the name of a builtin *window* function or the full name of a user defined alias macro. In the case of a builtin function, *<arglist>* can be in one of two forms:

<expr1>, <expr2>, ...

argname1 = <expr1>, argname2 = <expr2>, ...

The two forms can in fact be intermixed, but the result is unpredictable. Most arguments can be omitted; default values will be supplied for them. The *argnames* can be unique prefixes of the the argument names. The commas separating arguments are used only to disambiguate, and can usually be omitted.

Only the first argument form is valid for user defined aliases. Aliases are defined using the *alias* builtin function (see below). Arguments are accessed via a variant of the variable mechanism (see "\$" operator above).

Most functions return value, but some are used for side effect only and so must be used as statements. When a function or an alias is used as a statement, the parenthesis surrounding the argument list may be omitted. Aliases return no value.

BULTIN FUNCTIONS

The arguments are listed by name in their natural order. Optional arguments are in square brackets ("["]). Arguments that have no names are in angle brackets ("<>").

alias([<string>], [<string-list>])

If no argument is given, all currently defined alias macros are listed. Otherwise, *<string>* is defined as an alias, with expansion *<string-list>*. The previous

definition of *<string>*, if any, is returned. Default for *<string-list>* is no change.

close(<window-list>)

Close the windows specified in *<window-list>*. If *<window-list>* is the word *all*, then all windows are closed. No value is returned.

cursormodes([modes])

Set the window cursor to *modes*. *Modes* is the bitwise or of the mode bits defined as the variables *m_ul* (underline), *m_rev* (reverse video), *m_blk* (blinking), and *m_grp* (graphics, terminal dependent). Return value is the previous modes. Default is no change. For example, `cursor(m_revm_blk)` sets the window cursors to blinking reverse video.

echo([window], [<string-list>])

Write the list of strings, *<string-list>*, to *window*, separated by spaces and terminated with a new line. The strings are only displayed in the window, the processes in the window are not involved (see *write* below). No value is returned. Default is the current window.

escape([escapec])

Set the escape character to *escape-char*. Returns the old escape character as a one character string. Default is no change. *Escapec* can be a string of a single character, or in the form *^X*, meaning control-*X*.

foreground([window], [flag])

Move *window* in or out of foreground. *Flag* can be one of *on*, *off*, *yes*, *no*, *true*, or *false*, with obvious meanings, or it can be a numeric expression, in which case a non-zero value is true. Returns the old foreground flag as a number. Default for *window* is the current window, default for *flag* is no change.

label([window], [label])

Set the label of *window* to *label*. Returns the old label as a string. Default for *window* is the current window, default for *label* is no change. To turn off a label, set it to an empty string ("").

list() No arguments. List the identifiers and labels of all windows. No value is returned.

nline([nline])

Set the default buffer size to *nline*. Initially, it is 48 lines. Returns the old default buffer size. Default is no change. Using a very large buffer can slow the program down considerably.

select([window])

Make *window* the current window. The previous current window is returned. Default is no change.

shell([<string-list>])

Set the default window shell program to *<string-list>*. Returns the first string in the old shell setting. Default is no change. Initially, the default shell is taken from the environment variable *SHELL*.

source(filename)

Read and execute the long commands in *filename*. Returns -1 if the file cannot be read, 0 otherwise.

terse([flag])

Set terse mode to *flag*. In terse mode, the command window stays hidden even in command mode, and errors are reported by sounding the terminal's bell. *Flag* can take on the same values as in *foreground* above. Returns the old terse flag. Default is no change.

unalias(alias)

Undefine *alias*. Returns -1 if *alias* does not exist, 0 otherwise.

unset(variable)

Undefine *variable*. Returns -1 if *variable* does not exist, 0 otherwise.

variables()

No arguments. List all variables. No value is returned.

window([row], [column], [nrow], [ncol], [nline], [frame], [pty], [mapnl], [shell])

Open a window with upper left corner at *row*, *column* and size *nrow*, *ncol*. If *nline* is specified, then that many lines are allocated for the text buffer. Otherwise, the default buffer size is used. Default values for *row*, *column*, *nrow*, and *ncol* are, respectively, the upper, left-most, lower, or right-most extremes of the screen. *Frame*, *pty*, and *mapnl* are flag values interpreted in the same way as the argument to *foreground* (see above); they mean, respectively, put a frame around this window (default true), allocate pseudo-terminal for this window rather than socketpair (default true), and map new line characters in this window to carriage return and line feed (default true if socketpair is used, false otherwise). *Shell* is a list of strings that will be used as the shell program to place in the window (default is the program specified by *shell*, see below). The created window's identifier is returned as a number.

write([window], [<string-list>])

Send the list of strings, *<string-list>*, to *window*, separated by spaces but not terminated with a new line. The strings are actually given to the window as input. No value is returned. Default is the current window.

PREDEFINED VARIABLES

These variables are for information only. Redefining them does not affect the internal operation of *window*.

baud The baud rate as a number between 50 and 38400.

modes

The display modes (reverse video, underline, blinking, graphics) supported by the physical terminal. The value of *modes* is the bitwise or of some of the one bit values, *m_blk*, *m_grp*, *m_rev*, and *m_ul* (see below). These values are useful in setting the window cursors' modes (see *cursormodes* above).

m_blk

The blinking mode bit.

m_grp

The graphics mode bit (not very useful).

m_rev

The reverse video mode bit.

m_ul The underline mode bit.

ncol The number of columns on the physical screen.

nrow The number of rows on the physical screen.

term The terminal type. The standard name, found in the second name field of the terminal's *TERMCAP* entry, is used.

FILES

~/.windowrc
/dev/[pt]ty[pq]?

startup command file.
pseudo-terminal devices.

DIAGNOSTICS

Should be self explanatory.

NAME

write - write to another user

SYNOPSIS

write user [ttyname]

DESCRIPTION

Write copies lines from your terminal to that of another user. When first called, it sends the message

Message from yoursystem!yourname yourttyname...

The recipient of the message should write back at this point. Communication continues until an end of file is read from the terminal or an interrupt is sent. At that point *write* writes 'EOT' on the other terminal and exits.

If you want to write to a user who is logged in more than once, the *ttyname* argument may be used to indicate the appropriate terminal name.

Permission to write may be denied or granted by use of the *mesg* command. At the outset writing is allowed. Certain commands, in particular *nroff* and *pr(1)* disallow messages in order to prevent messy output.

If the character '!' is found at the beginning of a line, *write* calls the shell to execute the rest of the line as a command.

The following protocol is suggested for using *write*: when you first write to another user, wait for him to write back before starting to send. Each party should end each message with a distinctive signal—(o) for 'over' is conventional—that the other may reply. (oo) for 'over and out' is suggested when conversation is about to be terminated.

FILES

/etc/utmp to find user
/bin/sh to execute '!'

SEE ALSO

mesg(1), who(1), mail(1)

NAME

xsend, *xget*, *enroll* – secret mail

SYNOPSIS

xsend person
xget
enroll

DESCRIPTION

These commands implement a secure communication channel; it is like *mail(1)*, but no one can read the messages except the intended recipient. The method embodies a public-key cryptosystem using knapsacks.

To receive messages, use *enroll*; it asks you for a password that you must subsequently quote in order to receive secret mail.

To receive secret mail, use *xget*. It asks for your password, then gives you the messages.

To send secret mail, use *xsend* in the same manner as the ordinary mail command. (However, it will accept only one target). A message announcing the receipt of secret mail is also sent by ordinary mail.

FILES

/usr/spool/secretmail/*.key: keys
/usr/spool/secretmail/*.0-9: messages

SEE ALSO

mail (1)

BUGS

It should be integrated with ordinary mail. The announcement of secret mail makes traffic analysis possible.

NAME

xstr - extract strings from C programs to implement shared strings

SYNOPSIS

xstr [-c] [-] [file]

DESCRIPTION

Xstr maintains a file *strings* into which strings in component parts of a large program are hashed. These strings are replaced with references to this common area. This serves to implement shared constant strings, most useful if they are also read-only.

The command

xstr -c name

will extract the strings from the C source in *name*, replacing string references by expressions of the form (&xstr[number]) for some number. An appropriate declaration of *xstr* is prepended to the file. The resulting C text is placed in the file *x.c*, to then be compiled. The strings from this file are placed in the *strings* data base if they are not there already. Repeated strings and strings which are suffices of existing strings do not cause changes to the data base.

After all components of a large program have been compiled a file *xs.c* declaring the common *xstr* space can be created by a command of the form

xstr

This *xs.c* file should then be compiled and loaded with the rest of the program. If possible, the array can be made read-only (shared) saving space and swap overhead.

Xstr can also be used on a single file. A command

xstr name

creates files *x.c* and *xs.c* as before, without using or affecting any *strings* file in the same directory.

It may be useful to run *xstr* after the C preprocessor if any macro definitions yield strings or if there is conditional code which contains strings which may not, in fact, be needed. *Xstr* reads from its standard input when the argument '-' is given. An appropriate command sequence for running *xstr* after the C preprocessor is:

```
cc -E name.c | xstr -c -
cc -c x.c
mv x.o name.o
```

Xstr does not touch the file *strings* unless new items are added, thus *make* can avoid remaking *xs.o* unless truly necessary.

FILES

<i>strings</i>	Data base of strings
<i>x.c</i>	Massaged C source
<i>xs.c</i>	C source for definition of array 'xstr'
<i>/tmp/xs*</i>	Temp file when 'xstr name' doesn't touch <i>strings</i>

SEE ALSO

mkstr(1)

AUTHOR

William Joy

BUGS

If a string is a suffix of another string in the data base, but the shorter string is seen first by *xstr* both strings will be placed in the data base, when just placing the longer one there will do.

NAME

yacc - yet another compiler-compiler

SYNOPSIS

yacc [-vd] grammar

DESCRIPTION

Yacc converts a context-free grammar into a set of tables for a simple automaton which executes an LR(1) parsing algorithm. The grammar may be ambiguous; specified precedence rules are used to break ambiguities.

The output file, *y.tab.c*, must be compiled by the C compiler to produce a program *yyparse*. This program must be loaded with the lexical analyzer program, *yylex*, as well as *main* and *yyerror*, an error handling routine. These routines must be supplied by the user; *Lex(1)* is useful for creating lexical analyzers usable by *yacc*.

If the *-v* flag is given, the file *y.output* is prepared, which contains a description of the parsing tables and a report on conflicts generated by ambiguities in the grammar.

If the *-d* flag is used, the file *y.tab.h* is generated with the *define* statements that associate the *yacc*-assigned 'token codes' with the user-declared 'token names'. This allows source files other than *y.tab.c* to access the token codes.

FILES

y.output
y.tab.c
y.tab.h defines for token names
yacc.tmp, *yacc.acts* temporary files
/usr/lib/yaccpar parser prototype for C programs

SEE ALSO

lex(1)
LR Parsing by A. V. Aho and S. C. Johnson, Computing Surveys, June, 1974.
YACC - Yet Another Compiler Compiler by S. C. Johnson.

DIAGNOSTICS

The number of reduce-reduce and shift-reduce conflicts is reported on the standard output; a more detailed report is found in the *y.output* file. Similarly, if some rules are not reachable from the start symbol, this is also reported.

BUGS

Because file names are fixed, at most one *yacc* process can be active in a given directory at a time.

NAME

yes - be repetitively affirmative

SYNOPSIS

yes [*expletive*]

DESCRIPTION

Yes repeatedly outputs "y", or if *expletive* is given, that is output repeatedly. Termination is by rubout.

**ICON/UXB
OPERATING
SYSTEM
CALLS**



NAME

intro - introduction to system calls and error numbers

SYNOPSIS

```
#include <errno.h>
```

DESCRIPTION

This section describes all of the system calls. Most of these calls have one or more error returns. An error condition is indicated by an otherwise impossible return value. This is almost always -1; the individual descriptions specify the details.

As with normal arguments, all return codes and values from functions are of type integer unless otherwise noted. An error number is also made available in the external variable *errno*, which is not cleared on successful calls. Thus *errno* should be tested only after an error has occurred.

The following is a complete list of the errors and their names as given in *<errno.h>*.

- 0 "Error 0
Unused.
- 1 EPERM Not owner
Typically this error indicates an attempt to modify a file in some way forbidden except to its owner or super-user. It is also returned for attempts by ordinary users to do things allowed only to the super-user.
- 2 ENOENT No such file or directory
This error occurs when a file name is specified and the file should exist but doesn't, or when one of the directories in a path name does not exist.
- 3 ESRCH No such process
The process whose number was given to *kill* and *ptrace* does not exist, or is already dead.
- 4 EINTR Interrupted system call
An asynchronous signal (such as interrupt or quit), which the user has elected to catch, occurred during a system call. If execution is resumed after processing the signal, it will appear as if the interrupted system call returned this error condition.
- 5 EIO I/O error
Some physical I/O error occurred during a *read* or *write*. This error may in some cases occur on a call following the one to which it actually applies.
- 6 ENXIO No such device or address
I/O on a special file refers to a subdevice which does not exist, or beyond the limits of the device. It may also occur when, for example, an illegal tape drive unit number is selected or a disk pack is not loaded on a drive.
- 7 E2BIG Arg list too long
An argument list longer than 10240 bytes is presented to *execve*.
- 8 ENOEXEC Exec format error
A request is made to execute a file which, although it has the appropriate permissions, does not start with a valid magic number, see *a.out(5)*.
- 9 EBADF Bad file number
Either a file descriptor refers to no open file, or a read (resp. write) request is made to a file which is open only for writing (resp. reading).



- 26 **ETXTBSY** Text file busy
An attempt to execute a pure-procedure program which is currently open for writing (or reading!). Also an attempt to open for writing a pure-procedure program that is being executed.
- 27 **EFBIG** File too large
The size of a file exceeded the maximum (about 10^9 bytes).
- 28 **ENOSPC** No space left on device
During a *write* to an ordinary file, there is no free space left on the device.
- 29 **ESPIPE** Illegal seek
An *lseek* was issued to a pipe. This error may also be issued for other non-seekable devices.
- 30 **EROFS** Read-only file system
An attempt to modify a file or directory was made on a device mounted read-only.
- 31 **EMLINK** Too many links
An attempt to make more than 32767 hard links to a file.
- 32 **EPIPE** Broken pipe
A write on a pipe or socket for which there is no process to read the data. This condition normally generates a signal; the error is returned if the signal is ignored.
- 33 **EDOM** Math argument
The argument of a function in the math package (3M) is out of the domain of the function.
- 34 **ERANGE** Result too large
The value of a function in the math package (3M) is unrepresentable within machine precision.
- 35 **EWouldBLOCK** Operation would block
An operation which would cause a process to block was attempted on a object in non-blocking mode (see *ioctl* (2)).
- 36 **EINPROGRESS** Operation now in progress
An operation which takes a long time to complete (such as a *connect* (2)) was attempted on a non-blocking object (see *ioctl* (2)).
- 37 **EALREADY** Operation already in progress
An operation was attempted on a non-blocking object which already had an operation in progress.
- 38 **ENOTSOCK** Socket operation on non-socket
Self-explanatory.
- 39 **EDESTADDRREQ** Destination address required
A required address was omitted from an operation on a socket.
- 40 **EMSGSIZE** Message too long
A message sent on a socket was larger than the internal message buffer.
- 41 **EPROTOTYPE** Protocol wrong type for socket
A protocol was specified which does not support the semantics of the socket type requested. For example you cannot use the ARPA Internet UDP protocol with type `SOCK_STREAM`.

- 10 ECHILD No children
 Wait and the process has no living or unwaited-for children.
- 11 EAGAIN No more processes
 In a *fork*, the system's process table is full or the user is not allowed to create any more processes.
- 12 ENOMEM Not enough core
 During an *execve* or *break*, a program asks for more core or swap space than the system is able to supply. A lack of swap space is normally a temporary condition, however a lack of core is not a temporary condition; the maximum size of the text, data, and stack segments is a system parameter.
- 13 EACCES Permission denied
 An attempt was made to access a file in a way forbidden by the protection system.
- 14 EFAULT Bad address
 The system encountered a hardware fault in attempting to access the arguments of a system call.
- 15 ENOTBLK Block device required
 A plain file was mentioned where a block device was required, e.g. in *mount*.
- 16 EBUSY Mount device busy
 An attempt to mount a device that was already mounted or an attempt was made to dismount a device on which there is an active file directory. (open file, current directory, mounted-on file, active text segment).
- 17 EEXIST File exists
 An existing file was mentioned in an inappropriate context, e.g. *link*.
- 18 EXDEV Cross-device link
 A hard link to a file on another device was attempted.
- 19 ENODEV No such device
 An attempt was made to apply an inappropriate system call to a device; e.g. read a write-only device.
- 20 ENOTDIR Not a directory
 A non-directory was specified where a directory is required, for example in a path name or as an argument to *chdir*.
- 21 EISDIR Is a directory
 An attempt to write on a directory.
- 22 EINVAL Invalid argument
 Some invalid argument: dismounting a non-mounted device, mentioning an unknown signal in *signal*, reading or writing a file for which *seek* has generated a negative pointer. Also set by math functions, see *intro*(3).
- 23 ENFILE File table overflow
 The system's table of open files is full, and temporarily no more *opens* can be accepted.
- 24 EMFILE Too many open files
 Customary configuration limit is 20 per process.
- 25 ENOTTY Not a typewriter
 The file mentioned in an *ioctl* is not a terminal or one of the other devices to which these calls apply.

- 59 *unused*
- 60 ETIMEDOUT Connection timed out
A *connect* request failed because the connected party did not properly respond after a period of time. (The timeout period is dependent on the communication protocol.)
- 61 ECONNREFUSED Connection refused
No connection could be made because the target machine actively refused it. This usually results from trying to connect to a service which is inactive on the foreign host.
- 62 ELOOP Too many levels of symbolic links
A path name lookup involved more than 8 symbolic links.
- 63 ENAMETOOLONG File name too long
A component of a path name exceeded 255 characters, or an entire path name exceeded 1023 characters.
- 64 EHOSTDOWN Host is down
A socket operation failed because the destination host was down.
- 65 EHOSTUNREACH Host is unreachable
A socket operation was attempted to an unreachable host.
- 66 ENOTEMPTY Directory not empty
A directory with entries other than "." and ".." was supplied to a remove directory or rename call.
- 69 EDQUOT Disc quota exceeded
- 70 ENOTREG (used internally)
- 71 ECTNRDY Cassette tape not ready
- 72 ENOMSG No message of desired type
- 73 EIDRM Identifier removed
- 74 EDEADLK Deadlock condition if locked

DEFINITIONS

Process ID

Each active process in the system is uniquely identified by a positive integer called a process ID. The range of this ID is from 0 to {PROC_MAX}.

Parent process ID

A new process is created by a currently active process; see *fork(2)*. The parent process ID of a process is the process ID of its creator.

Process Group ID

Each active process is a member of a process group that is identified by a positive integer called the process group ID. This is the process ID of the group leader. This grouping permits the signalling of related processes (see *killpg(2)*) and the job control mechanisms of *cs(1)*.

Tty Group ID

Each active process can be a member of a terminal group that is identified by a positive integer called the tty group ID. This grouping is used to arbitrate between multiple jobs contending for the same terminal; see *cs(1)*, and *tty(4)*.

- 42 ENOPROTOOPT Bad protocol option
A bad option was specified in a *getsockopt(2)* or *setsockopt(2)* call.
- 43 EPROTONOSUPPORT Protocol not supported
The protocol has not been configured into the system or no implementation for it exists.
- 44 ESOCKTNOSUPPORT Socket type not supported
The support for the socket type has not been configured into the system or no implementation for it exists.
- 45 EOPNOTSUPP Operation not supported on socket
For example, trying to *accept* a connection on a datagram socket.
- 46 EPFNOSUPPORT Protocol family not supported
The protocol family has not been configured into the system or no implementation for it exists.
- 47 EAFNOSUPPORT Address family not supported by protocol family
An address incompatible with the requested protocol was used. For example, you shouldn't necessarily expect to be able to use PUP Internet addresses with ARPA Internet protocols.
- 48 EADDRINUSE Address already in use
Only one usage of each address is normally permitted.
- 49 EADDRNOTAVAIL Can't assign requested address
Normally results from an attempt to create a socket with an address not on this machine.
- 50 ENETDOWN Network is down
A socket operation encountered a dead network.
- 51 ENETUNREACH Network is unreachable
A socket operation was attempted to an unreachable network.
- 52 ENETRESET Network dropped connection on reset
The host you were connected to crashed and rebooted.
- 53 ECONNABORTED Software caused connection abort
A connection abort was caused internal to your host machine.
- 54 ECONNRESET Connection reset by peer
A connection was forcibly closed by a peer. This normally results from the peer executing a *shutdown(2)* call.
- 55 ENOBUFS No buffer space available
An operation on a socket or pipe was not performed because the system lacked sufficient buffer space.
- 56 EISCONN Socket is already connected
A *connect* request was made on an already connected socket; or, a *sendto* or *sendmsg* request on a connected socket specified a destination other than the connected party.
- 57 ENOTCONN Socket is not connected
An request to send or receive data was disallowed because the socket is not connected.
- 58 ESHUTDOWN Can't send after socket shutdown
A request to send data was disallowed because the socket had already been shut down with a previous *shutdown(2)* call.

Root Directory and Current Working Directory

Each process has associated with it a concept of a root directory and a current working directory for the purpose of resolving path name searches. A process's root directory need not be the root directory of the root file system.

File Access Permissions

Every file in the file system has a set of access permissions. These permissions are used in determining whether a process may perform a requested operation on the file (such as opening a file for writing). Access permissions are established at the time a file is created. They may be changed at some later time through the *chmod(2)* call.

File access is broken down according to whether a file may be: read, written, or executed. Directory files use the execute permission to control if the directory may be searched.

File access permissions are interpreted by the system as they apply to three different classes of users: the owner of the file, those users in the file's group, anyone else. Every file has an independent set of access permissions for each of these classes. When an access check is made, the system decides if permission should be granted by checking the access information applicable to the caller.

Read, write, and execute/search permissions on a file are granted to a process if:

The process's effective user ID is that of the super-user.

The process's effective user ID matches the user ID of the owner of the file and the owner permissions allow the access.

The process's effective user ID does not match the user ID of the owner of the file, and either the process's effective group ID matches the group ID of the file, or the group ID of the file is in the process's group access list, and the group permissions allow the access.

Neither the effective user ID nor effective group ID and group access list of the process match the corresponding user ID and group ID of the file, but the permissions for "other users" allow access.

Otherwise, permission is denied.

Sockets and Address Families

A socket is an endpoint for communication between processes. Each socket has queues for sending and receiving data.

Sockets are typed according to their communications properties. These properties include whether messages sent and received at a socket require the name of the partner, whether communication is reliable, the format used in naming message recipients, etc.

Each instance of the system supports some collection of socket types; consult *socket(2)* for more information about the types available and their properties.

Each instance of the system supports some number of sets of communications protocols. Each protocol set supports addresses of a certain format. An Address Family is the set of addresses for a specific group of protocols. Each socket has an address chosen from the address family in which the socket was created.

SEE ALSO

intro(3), perror(3)

Real User ID and Real Group ID

Each user on the system is identified by a positive integer termed the real user ID.

Each user is also a member of one or more groups. One of these groups is distinguished from others and used in implementing accounting facilities. The positive integer corresponding to this distinguished group is termed the real group ID.

All processes have a real user ID and real group ID. These are initialized from the equivalent attributes of the process which created it.

Effective User Id, Effective Group Id, and Access Groups

Access to system resources is governed by three values: the effective user ID, the effective group ID, and the group access list.

The effective user ID and effective group ID are initially the process's real user ID and real group ID respectively. Either may be modified through execution of a set-user-ID or set-group-ID file (possibly by one its ancestors); see *execve(2)*.

The group access list is an additional set of group ID's used only in determining resource accessibility. Access checks are performed as described below in "File Access Permissions".

Super-user

A process is recognized as a *super-user* process and is granted special privileges if its effective user ID is 0.

Special Processes

The processes with a process ID's of 0, 1, and 2 are special. Process 0 is the scheduler. Process 1 is the initialization process *init*, and is the ancestor of every other process in the system. It is used to control the process structure.

Descriptor

An integer assigned by the system when a file is referenced by *open(2)*, *dup(2)*, or *pipe(2)* or a socket is referenced by *socket(2)* or *socketpair(2)* which uniquely identifies an access path to that file or socket from a given process or any of its children.

File Name

Names consisting of up to {FILENAME_MAX} characters may be used to name an ordinary file, special file, or directory.

These characters may be selected from the set of all ASCII character excluding 0 (null) and the ASCII code for / (slash). (The parity bit, bit 8, must be 0.)

Note that it is generally unwise to use *, ?, [or] as part of file names because of the special meaning attached to these characters by the shell.

Path Name

A path name is a null-terminated character string starting with an optional slash (/), followed by zero or more directory names separated by slashes, optionally followed by a file name. The total length of a path name must be less than {PATHNAME_MAX} characters.

If a path name begins with a slash, the path search begins at the *root* directory. Otherwise, the search begins from the current working directory. A slash by itself names the *root* directory. A null pathname refers to the current directory.

Directory

A directory is a special type of file which contains entries which are references to other files. Directory entries are called links. By convention, a directory contains at least two links, . and .., referred to as *dot* and *dot-dot* respectively. Dot refers to the directory itself and dot-dot refers to its parent directory.

NAME

accept – accept a connection on a socket

SYNOPSIS

```
#include <sys/types.h>
#include <sys/socket.h>

ns = accept(s, addr, addrlen)
int ns, s;
struct sockaddr *addr;
int *addrlen;
```

DESCRIPTION

The argument *s* is a socket which has been created with *socket(2)*, bound to an address with *bind(2)*, and is listening for connections after a *listen(2)*. *Accept* extracts the first connection on the queue of pending connections, creates a new socket with the same properties of *s* and allocates a new file descriptor, *ns*, for the socket. If no pending connections are present on the queue, and the socket is not marked as non-blocking, *accept* blocks the caller until a connection is present. If the socket is marked non-blocking and no pending connections are present on the queue, *accept* returns an error as described below. The accepted socket, *ns*, may not be used to accept more connections. The original socket *s* remains open.

The argument *addr* is a result parameter which is filled in with the address of the connecting entity, as known to the communications layer. The exact format of the *addr* parameter is determined by the domain in which the communication is occurring. The *addrlen* is a value-result parameter; it should initially contain the amount of space pointed to by *addr*; on return it will contain the actual length (in bytes) of the address returned. This call is used with connection-based socket types, currently with *SOCK_STREAM*.

It is possible to *select(2)* a socket for the purposes of doing an *accept* by selecting it for read.

RETURN VALUE

The call returns *-1* on error. If it succeeds it returns a non-negative integer which is a descriptor for the accepted socket.

ERRORS

The *accept* will fail if:

- | | |
|---------------|--|
| [EBADF] | The descriptor is invalid. |
| [ENOTSOCK] | The descriptor references a file, not a socket. |
| [EOPNOTSUPP] | The referenced socket is not of type <i>SOCK_STREAM</i> . |
| [EFAULT] | The <i>addr</i> parameter is not in a writable part of the user address space. |
| [EWOULDBLOCK] | The socket is marked non-blocking and no connections are present to be accepted. |

SEE ALSO

bind(2), *connect(2)*, *listen(2)*, *select(2)*, *socket(2)*



NAME

access – determine accessibility of file

SYNOPSIS

```
#include <sys/file.h>
#define R_OK  4  /* test for read permission */
#define W_OK  2  /* test for write permission */
#define X_OK  1  /* test for execute (search) permission */
#define F_OK  0  /* test for presence of file */

accessible = access(path, mode)
int accessible;
char *path;
int mode;
```

DESCRIPTION

Access checks the given file *path* for accessibility according to *mode*, which is an inclusive or of the bits R_OK, W_OK and X_OK. Specifying *mode* as F_OK (i.e. 0) tests whether the directories leading to the file can be searched and the file exists.

The real user ID and the group access list (including the real group ID) are used in verifying permission, so this call is useful to set-UID programs.

Notice that only access bits are checked. A directory may be indicated as writable by *access*, but an attempt to open it for writing will fail (although files may be created there); a file may look executable, but *execve* will fail unless it is in proper format.

RETURN VALUE

If *path* cannot be found or if any of the desired access modes would not be granted, then a -1 value is returned; otherwise a 0 value is returned.

ERRORS

Access to the file is denied if one or more of the following are true:

- | | |
|-----------|--|
| [ENOTDIR] | A component of the path prefix is not a directory. |
| [ENOENT] | The argument path name was too long. |
| [ENOENT] | Read, write, or execute (search) permission is requested for a null path name or the named file does not exist. |
| [EPERM] | The argument contains a byte with the high-order bit set. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |
| [EROFS] | Write access is requested for a file on a read-only file system. |
| [ETXTBSY] | Write access is requested for a pure procedure (shared text) file that is being executed. |
| [EACCES] | Permission bits of the file mode do not permit the requested access; or search permission is denied on a component of the path prefix. The owner of a file has permission checked with respect to the “owner” read, write, and execute mode bits, members of the file’s group other than the owner have permission checked with respect to the “group” mode bits, and all others have permissions checked with respect to the “other” mode bits. |
| [EFAULT] | <i>Path</i> points outside the process’s allocated address space. |

SEE ALSO

chmod(2), stat(2)

NAME

`acct` - turn accounting on or off

SYNOPSIS

```
acct(file)
char *file;
```

DESCRIPTION

The system is prepared to write a record in an accounting *file* for each process as it terminates. This call, with a null-terminated string naming an existing file as argument, turns on accounting; records for each terminating process are appended to *file*. An argument of 0 causes accounting to be turned off.

The accounting file format is given in *acct(5)*.

This call is permitted only to the super-user.

NOTES

Accounting is automatically disabled when the file system the accounting file resides on runs out of space; it is enabled when space once again becomes available.

RETURN VALUE

On error -1 is returned. The file must exist and the call may be exercised only by the super-user. It is erroneous to try to turn on accounting when it is already on.

ERRORS

Acct will fail if one of the following is true:

[EPERM]	The caller is not the super-user.
[EPERM]	The pathname contains a character with the high-order bit set.
[ENOTDIR]	A component of the path prefix is not a directory.
[ENOENT]	The named file does not exist.
[EISDIR]	The named file is a directory.
[EROFS]	The named file resides on a read-only file system.
[EFAULT]	<i>File</i> points outside the process's allocated address space.
[ELOOP]	Too many symbolic links were encountered in translating the pathname.
[EACCES]	The file is a character or block special file.

SEE ALSO

acct(5), *sa(8)*

BUGS

No accounting is produced for programs running when a crash occurs. In particular nonterminating programs are never accounted for.

NAME

bind – bind a name to a socket

SYNOPSIS

```
#include <sys/types.h>
#include <sys/socket.h>
bind(s, name, namelen)
int s;
struct sockaddr *name;
int namelen;
```

DESCRIPTION

Bind assigns a name to an unnamed socket. When a socket is created with *socket(2)* it exists in a name space (address family) but has no name assigned. *Bind* requests the *name*, be assigned to the socket.

NOTES

Binding a name in the UNIX domain creates a socket in the file system which must be deleted by the caller when it is no longer needed (using *unlink(2)*). The file created is a side-effect of the current implementation, and will not be created in future versions of the UNIX ipc domain.

The rules used in name binding vary between communication domains. Consult the manual entries in section 4 for detailed information.

RETURN VALUE

If the bind is successful, a 0 value is returned. A return value of -1 indicates an error, which is further specified in the global *errno*.

ERRORS

The *bind* call will fail if:

[EBADF]	<i>S</i> is not a valid descriptor.
[ENOTSOCK]	<i>S</i> is not a socket.
[EADDRNOTAVAIL]	The specified address is not available from the local machine.
[EADDRINUSE]	The specified address is already in use.
[EINVAL]	The socket is already bound to an address.
[EACCESS]	The requested address is protected, and the current user has inadequate permission to access it.
[EFAULT]	The <i>name</i> parameter is not in a valid part of the user address space.

SEE ALSO

connect(2), listen(2), socket(2), getsockname(2)

NAME

brk, *sbrk* – change data segment size

SYNOPSIS

```
caddr_t brk(caddr_t addr)  
caddr_t addr;  
caddr_t sbrk(incr)  
int incr;
```

DESCRIPTION

Brk sets the system's idea of the lowest data segment location not used by the program (called the break) to *addr* (rounded up to the next multiple of the system's page size). Locations greater than *addr* and below the stack pointer are not in the address space and will thus cause a memory violation if accessed.

In the alternate function *sbrk*, *incr* more bytes are added to the program's data space and a pointer to the start of the new area is returned.

When a program begins execution via *execve* the break is set at the highest location defined by the program and data storage areas. Ordinarily, therefore, only programs with growing data areas need to use *sbrk*.

The *getrlimit*(2) system call may be used to determine the maximum permissible size of the data segment; it will not be possible to set the break beyond the *rlim_max* value returned from a call to *getrlimit*, e.g. "*etext* + *rlp* → *rlim_max*." (See *end*(3) for the definition of *etext*.)

RETURN VALUE

Zero is returned if the *brk* could be set; -1 if the program requests more memory than the system limit. *Sbrk* returns -1 if the break could not be set.

ERRORS

Sbrk will fail and no additional memory will be allocated if one of the following are true:

- [ENOMEM] The limit, as set by *setrlimit*(2), was exceeded.
- [ENOMEM] The maximum possible size of a data segment (compiled into the system) was exceeded.
- [ENOMEM] Insufficient space existed in the swap area to support the expansion.

SEE ALSO

execve(2), *getrlimit*(2), *malloc*(3), *end*(3)

BUGS

Setting the break may fail due to a temporary lack of swap space. It is not possible to distinguish this from a failure caused by exceeding the maximum size of the data segment without consulting *getrlimit*.

NAME

`chdir` - change current working directory

SYNOPSIS

```
chdir(path)
char *path;
```

DESCRIPTION

Path is the pathname of a directory. *Chdir* causes this directory to become the current working directory, the starting point for path names not beginning with “/”.

In order for a directory to become the current directory, a process must have execute (search) access to the directory.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

ERRORS

Chdir will fail and the current working directory will be unchanged if one or more of the following are true:

- | | |
|-----------|---|
| [ENOTDIR] | A component of the pathname is not a directory. |
| [ENOENT] | The named directory does not exist. |
| [ENOENT] | The argument path name was too long. |
| [EPERM] | The argument contains a byte with the high-order bit set. |
| [EACCES] | Search permission is denied for any component of the path name. |
| [EFAULT] | <i>Path</i> points outside the process's allocated address space. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |

SEE ALSO

`chroot(2)`

NAME

chmod - change mode of file

SYNOPSIS

```

chmod(path, mode)
char *path;
int mode;

fchmod(fd, mode)
int fd, mode;

```

DESCRIPTION

The file whose name is given by *path* or referenced by the descriptor *fd* has its mode changed to *mode*. Modes are constructed by *or*'ing together some combination of the following:

```

04000 set user ID on execution
02000 set group ID on execution
01000 save text image after execution
00400 read by owner
00200 write by owner
00100 execute (search on directory) by owner
00070 read, write, execute (search) by group
00007 read, write, execute (search) by others

```

If an executable file is set up for sharing (this is the default) then mode 1000 prevents the system from abandoning the swap-space image of the program-text portion of the file when its last user terminates. Ability to set this bit is restricted to the super-user.

Only the owner of a file (or the super-user) may change the mode.

Writing or changing the owner of a file turns off the set-user-id and set-group-id bits. This makes the system somewhat more secure by protecting set-user-id (set-group-id) files from remaining set-user-id (set-group-id) if they are modified, at the expense of a degree of compatibility.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

ERRORS

Chmod will fail and the file mode will be unchanged if:

[EPERM]	The argument contains a byte with the high-order bit set.
[ENOTDIR]	A component of the path prefix is not a directory.
[ENOENT]	The pathname was too long.
[ENOENT]	The named file does not exist.
[EACCES]	Search permission is denied on a component of the path prefix.
[EPERM]	The effective user ID does not match the owner of the file and the effective user ID is not the super-user.
[EROFS]	The named file resides on a read-only file system.
[EFAULT]	<i>Path</i> points outside the process's allocated address space.
[ELOOP]	Too many symbolic links were encountered in translating the pathname.

Fchmod will fail if:

- [EBADF] The descriptor is not valid.
- [EINVAL] *Fd* refers to a socket, not to a file.
- [EROFS] The file resides on a read-only file system.

SEE ALSO

open(2), chown(2)

NAME

`chown` – change owner and group of a file

SYNOPSIS

```
chown(path, owner, group)
char *path;
int owner, group;

fchown(fd, owner, group)
int fd, owner, group;
```

DESCRIPTION

The file which is named by *path* or referenced by *fd* has its *owner* and *group* changed as specified. Only the super-user may execute this call, because if users were able to give files away, they could defeat the file-space accounting procedures.

On some systems, *chown* clears the set-user-id and set-group-id bits on the file to prevent accidental creation of set-user-id and set-group-id programs owned by the super-user.

Fchown is particularly useful when used in conjunction with the file locking primitives (see *flock*(2)).

Only one of the owner and group id's may be set by specifying the other as -1.

RETURN VALUE

Zero is returned if the operation was successful; -1 is returned if an error occurs, with a more specific error code being placed in the global variable *errno*.

ERRORS

Chown will fail and the file will be unchanged if:

- | | |
|-----------|---|
| [EINVAL] | The argument <i>path</i> does not refer to a file. |
| [ENOTDIR] | A component of the path prefix is not a directory. |
| [ENOENT] | The argument pathname is too long. |
| [EPERM] | The argument contains a byte with the high-order bit set. |
| [ENOENT] | The named file does not exist. |
| [EACCES] | Search permission is denied on a component of the path prefix. |
| [EPERM] | The effective user ID does not match the owner of the file and the effective user ID is not the super-user. |
| [EROFS] | The named file resides on a read-only file system. |
| [EFAULT] | <i>Path</i> points outside the process's allocated address space. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |

Fchown will fail if:

- | | |
|----------|---|
| [EBADF] | <i>Fd</i> does not refer to a valid descriptor. |
| [EINVAL] | <i>Fd</i> refers to a socket, not a file. |

SEE ALSO

chmod(2), *flock*(2)

NAME

chroot – change root directory

SYNOPSIS

```
chroot(dirname)
char *dirname;
```

DESCRIPTION

Dirname is the address of the pathname of a directory, terminated by a null byte. *Chroot* causes this directory to become the root directory, the starting point for path names beginning with “/”.

In order for a directory to become the root directory a process must have execute (search) access to the directory.

This call is restricted to the super-user.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate an error.

ERRORS

Chroot will fail and the root directory will be unchanged if one or more of the following are true:

[ENOTDIR]	A component of the path name is not a directory.
[ENOENT]	The pathname was too long.
[EPERM]	The argument contains a byte with the high-order bit set.
[ENOENT]	The named directory does not exist.
[EACCES]	Search permission is denied for any component of the path name.
[EFAULT]	<i>Path</i> points outside the process's allocated address space.
[ELOOP]	Too many symbolic links were encountered in translating the pathname.

SEE ALSO

chdir(2)

NAME

close – delete a descriptor

SYNOPSIS

close(d)

int d;

DESCRIPTION

The *close* call deletes a descriptor from the per-process object reference table. If this is the last reference to the underlying object, then it will be deactivated. For example, on the last close of a file the current *seek* pointer associated with the file is lost; on the last close of a *socket(2)* associated naming information and queued data are discarded; on the last close of a file holding an advisory lock the lock is released; see further *flock(2)*.

A close of all of a process's descriptors is automatic on *exit*, but since there is a limit on the number of active descriptors per process, *close* is necessary for programs which deal with many descriptors.

When a process forks (see *fork(2)*), all descriptors for the new child process reference the same objects as they did in the parent before the fork. If a new process is then to be run using *execve(2)*, the process would normally inherit these descriptors. Most of the descriptors can be rearranged with *dup2(2)* or deleted with *close* before the *execve* is attempted, but if some of these descriptors will still be needed if the *execve* fails, it is necessary to arrange for them to be closed if the *execve* succeeds. For this reason, the call “*fcntl(d, F_SETFD, 1)*” is provided which arranges that a descriptor will be closed after a successful *execve*; the call “*fcntl(d, F_SETFD, 0)*” restores the default, which is to not close the descriptor.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and the global integer variable *errno* is set to indicate the error.

ERRORS

Close will fail if:

[EBADF] *D* is not an active descriptor.

SEE ALSO

accept(2), *flock(2)*, *open(2)*, *pipe(2)*, *socket(2)*, *socketpair(2)*, *execve(2)*, *fcntl(2)*

NAME

connect – initiate a connection on a socket

SYNOPSIS

```
#include <sys/types.h>
#include <sys/socket.h>
connect(s, name, namelen)
int s;
struct sockaddr *name;
int namelen;
```

DESCRIPTION

The parameter *s* is a socket. If it is of type `SOCK_DGRAM`, then this call permanently specifies the peer to which datagrams are to be sent; if it is of type `SOCK_STREAM`, then this call attempts to make a connection to another socket. The other socket is specified by *name* which is an address in the communications space of the socket. Each communications space interprets the *name* parameter in its own way.

RETURN VALUE

If the connection or binding succeeds, then 0 is returned. Otherwise a -1 is returned, and a more specific error code is stored in *errno*.

ERRORS

The call fails if:

[EBADF]	<i>S</i> is not a valid descriptor.
[ENOTSOCK]	<i>S</i> is a descriptor for a file, not a socket.
[EADDRNOTAVAIL]	The specified address is not available on this machine.
[EAFNOSUPPORT]	Addresses in the specified address family cannot be used with this socket.
[EISCONN]	The socket is already connected.
[ETIMEDOUT]	Connection establishment timed out without establishing a connection.
[ECONNREFUSED]	The attempt to connect was forcefully rejected.
[ENETUNREACH]	The network isn't reachable from this host.
[EADDRINUSE]	The address is already in use.
[EFAULT]	The <i>name</i> parameter specifies an area outside the process address space.
[EWOULDBLOCK]	The socket is non-blocking and the and the connection cannot be completed immediately. It is possible to <i>select(2)</i> the socket while it is connecting by selecting it for writing.

SEE ALSO

accept(2), select(2), socket(2), getsockname(2)

NAME

`creat` – create a new file

SYNOPSIS

```
creat(name, mode)
char *name;
```

DESCRIPTION

This interface is obsoleted by `open(2)`.

Creat creates a new file or prepares to rewrite an existing file called *name*, given as the address of a null-terminated string. If the file did not exist, it is given mode *mode*, as modified by the process's mode mask (see *umask(2)*). Also see *chmod(2)* for the construction of the *mode* argument.

If the file did exist, its mode and owner remain unchanged but it is truncated to 0 length.

The file is also opened for writing, and its file descriptor is returned.

NOTES

The *mode* given is arbitrary; it need not allow writing. This feature has been used in the past by programs to construct a simple exclusive locking mechanism. It is replaced by the `O_EXCL` open mode, or *flock(2)* facility.

RETURN VALUE

The value `-1` is returned if an error occurs. Otherwise, the call returns a non-negative descriptor which only permits writing.

ERRORS

Creat will fail and the file will not be created or truncated if one of the following occur:

[EPERM]	The argument contains a byte with the high-order bit set.
[ENOTDIR]	A component of the path prefix is not a directory.
[EACCES]	A needed directory does not have search permission.
[EACCES]	The file does not exist and the directory in which it is to be created is not writable.
[EACCES]	The file exists, but it is unwritable.
[EISDIR]	The file is a directory.
[EMFILE]	There are already too many files open.
[EROFS]	The named file resides on a read-only file system.
[ENXIO]	The file is a character special or block special file, and the associated device does not exist.
[ETXTBSY]	The file is a pure procedure (shared text) file that is being executed.
[EFAULT]	<i>Name</i> points outside the process's allocated address space.
[ELOOP]	Too many symbolic links were encountered in translating the pathname.
[EOPNOTSUPP]	The file was a socket (not currently implemented).

SEE ALSO

`open(2)`, `write(2)`, `close(2)`, `chmod(2)`, `umask(2)`

NAME

`dup`, `dup2` - duplicate a descriptor

SYNOPSIS

```
newd = dup(oldd)  
int newd, oldd;  
dup2(oldd, newd)  
int oldd, newd;
```

DESCRIPTION

Dup duplicates an existing object descriptor. The argument *oldd* is a small non-negative integer index in the per-process descriptor table. The value must be less than the size of the table, which is returned by *getdtablesize(2)*. The new descriptor *newd* returned by the call is the lowest numbered descriptor which is not currently in use by the process.

The object referenced by the descriptor does not distinguish between references using *oldd* and *newd* in any way. Thus if *newd* and *oldd* are duplicate references to an open file, *read(2)*, *write(2)* and *lseek(2)* calls all move a single pointer into the file. If a separate pointer into the file is desired, a different object reference to the file must be obtained by issuing an additional *open(2)* call.

In the second form of the call, the value of *newd* desired is specified. If this descriptor is already in use, the descriptor is first deallocated as if a *close(2)* call had been done first.

RETURN VALUE

The value -1 is returned if an error occurs in either call. The external variable *errno* indicates the cause of the error.

ERRORS

Dup and *dup2* fail if:

[EBADF]	<i>Oldd</i> or <i>newd</i> is not a valid active descriptor
[EMFILE]	Too many descriptors are active.

SEE ALSO

accept(2), *open(2)*, *close(2)*, *pipe(2)*, *socket(2)*, *socketpair(2)*, *getdtablesize(2)*

NAME

`execve` – execute a file

SYNOPSIS

```
execve(name, argv, envp)
char *name, *argv[], *envp[];
```

DESCRIPTION

Execve transforms the calling process into a new process. The new process is constructed from an ordinary file called the *new process file*. This file is either an executable object file, or a file of data for an interpreter. An executable object file consists of an identifying header, followed by pages of data representing the initial program (text) and initialized data pages. Additional pages may be specified by the header to be initialize with zero data. See *a.out*(5).

An interpreter file begins with a line of the form “*#! interpreter*”; When an interpreter file is *execve*’d, the system *execve*’s the specified *interpreter*, giving it the name of the originally *exec*’d file as an argument, shifting over the rest of the original arguments.

There can be no return from a successful *execve* because the calling core image is lost. This is the mechanism whereby different process images become active.

The argument *argv* is an array of character pointers to null-terminated character strings. These strings constitute the argument list to be made available to the new process. By convention, at least one argument must be present in this array, and the first element of this array should be the name of the executed program (i.e. the last component of *name*).

The argument *envp* is also an array of character pointers to null-terminated strings. These strings pass information to the new process which are not directly arguments to the command, see *environ*(7).

Descriptors open in the calling process remain open in the new process, except for those for which the close-on-exec flag is set; see *close*(2). Descriptors which remain open are unaffected by *execve*.

Ignored signals remain ignored across an *execve*, but signals that are caught are reset to their default values. The signal stack is reset to be undefined; see *sigvec*(2) for more information.

Each process has *real* user and group IDs and a *effective* user and group IDs. The *real* ID identifies the person using the system; the *effective* ID determines his access privileges. *Execve* changes the effective user and group ID to the owner of the executed file if the file has the “set-user-ID” or “set-group-ID” modes. The *real* user ID is not affected.

The new process also inherits the following attributes from the calling process:

process ID	see <i>getpid</i> (2)
parent process ID	see <i>getppid</i> (2)
process group ID	see <i>getpgrp</i> (2)
access groups	see <i>getgroups</i> (2)
working directory	see <i>chdir</i> (2)
root directory	see <i>chroot</i> (2)
control terminal	see <i>tty</i> (4)
resource usages	see <i>getrusage</i> (2)
interval timers	see <i>getitimer</i> (2)
resource limits	see <i>getrlimit</i> (2)
file mode mask	see <i>umask</i> (2)
signal mask	see <i>sigvec</i> (2)

When the executed program begins, it is called as follows:

```
main(argc, argv, envp)
int argc;
char **argv, **envp;
```

where *argc* is the number of elements in *argv* (the “arg count”) and *argv* is the array of character pointers to the arguments themselves.

Envp is a pointer to an array of strings that constitute the *environment* of the process. A pointer to this array is also stored in the global variable “environ”. Each string consists of a name, an “=”, and a null-terminated value. The array of pointers is terminated by a null pointer. The shell *sh*(1) passes an environment entry for each global shell variable defined when the program is called. See *environ*(7) for some conventionally used names.

RETURN VALUE

If *execve* returns to the calling process an error has occurred; the return value will be -1 and the global variable *errno* will contain an error code.

ERRORS

Execve will fail and return to the calling process if one or more of the following are true:

- | | |
|-----------|--|
| [ENOENT] | One or more components of the new process file’s path name do not exist. |
| [ENOTDIR] | A component of the new process file is not a directory. |
| [EACCES] | Search permission is denied for a directory listed in the new process file’s path prefix. |
| [EACCES] | The new process file is not an ordinary file. |
| [EACCES] | The new process file mode denies execute permission. |
| [ENOEXEC] | The new process file has the appropriate access permission, but has an invalid magic number in its header. |
| [ETXTBSY] | The new process file is a pure procedure (shared text) file that is currently open for writing or reading by some process. |
| [ENOMEM] | The new process requires more virtual memory than is allowed by the imposed maximum (<i>getrlimit</i> (2)). |
| [E2BIG] | The number of bytes in the new process’s argument list is larger than the system-imposed limit of {ARG_MAX} bytes. |
| [EFAULT] | The new process file is not as long as indicated by the size values in its header. |
| [EFAULT] | <i>Path</i> , <i>argv</i> , or <i>envp</i> point to an illegal address. |

CAVEATS

If a program is *setuid* to a non-super-user, but is executed when the real *uid* is “root”, then the program has the powers of a super-user as well.

SEE ALSO

exit(2), *fork*(2), *execl*(3), *environ*(7)

NAME

`_exit` - terminate a process

SYNOPSIS

```
_exit(status)  
int status;
```

DESCRIPTION

`_exit` terminates a process with the following consequences:

All of the descriptors open in the calling process are closed.

If the parent process of the calling process is executing a *wait* or is interested in the SIGCHLD signal, then it is notified of the calling process's termination and the low-order eight bits of *status* are made available to it; see *wait(2)*.

The parent process ID of all of the calling process's existing child processes are also set to 1. This means that the initialization process (see *intro(2)*) inherits each of these processes as well.

Most C programs call the library routine *exit(3)* which performs cleanup actions in the standard i/o library before calling `_exit`.

RETURN VALUE

This call never returns.

SEE ALSO

fork(2), *wait(2)*, *exit(3)*

NAME

`fcntl` - file control

SYNOPSIS

```
#include <fcntl.h>
res = fcntl(fd, cmd, arg)
int res;
int fd, cmd, arg;
```

DESCRIPTION

Fcntl provides for control over descriptors. The argument *fd* is a descriptor to be operated on by *cmd* as follows:

- F_DUPFD** Return a new descriptor as follows:
- Lowest numbered available descriptor greater than or equal to *arg*.
 - Same object references as the original descriptor.
 - New descriptor shares the same file pointer if the object was a file.
 - Same access mode (read, write or read/write).
 - Same file status flags (i.e., both file descriptors share the same file status flags).
 - The close-on-exec flag associated with the new file descriptor is set to remain open across *execv(2)* system calls.
- F_GETFD** Get the close-on-exec flag associated with the file descriptor *fd*. If the low-order bit is 0, the file will remain open across *exec*, otherwise the file will be closed upon execution of *exec*.
- F_SETFD** Set the close-on-exec flag associated with *fd* to the low order bit of *arg* (0 or 1 as above).
- F_GETFL** Get descriptor status flags, as described below.
- F_SETFL** Set descriptor status flags. Only certain flags can be set; see *fcntl(5)*.
- F_GETOWN** Get the process ID or process group currently receiving SIGIO and SIGURG signals; process groups are returned as negative values.
- F_SETOWN** Set the process or process group to receive SIGIO and SIGURG signals; process groups are specified by supplying *arg* as negative, otherwise *arg* is interpreted as a process ID.
- F_GETLK** Get the first lock which blocks the lock description given by the variable of type *struct flock* pointed to by *arg*. The information retrieved overwrites the information passed to *fcntl* in the *flock* structure. If no lock is found that would prevent this lock from being created, then the structure is passed back unchanged except for the lock type which will be set to F_UNLCK.
- F_SETLK** Set or clear a file segment lock according to the variable of type *struct flock* pointed to by *arg* [see *fcntl(5)*]. The *cmd* F_SETLK is used to establish read (F_RDLCK) and write (F_WRLCK) locks, as well as remove either type of lock (F_UNLCK). If a read or write lock cannot be set, *fcntl* will return immediately with an error value of -1.

F_SETLKW This *cmd* is the same as **F_SETLK** except that if a read or write lock is blocked by other locks, the process will sleep until the segment is free to be locked.

The flags for the **F_GETFL** and **F_SETFL** flags are as follows:

FNDELAY Non-blocking I/O; if no data is available to a *read* call, or if a write operation would block, the call returns -1 with the error **EWOULDLOCK**.

FAPPEND Force each write to append at the end of file; corresponds to the **O_APPEND** flag of *open(2)*.

FASYNC Enable the **SIGIO** signal to be sent to the process group when I/O is possible, e.g. upon availability of data to be read.

A read lock prevents any process from write locking the protected area. More than one read lock may exist for a given segment of a file at a given time. The file descriptor on which a read lock is being placed must have been opened with read access.

A write lock prevents any process from read locking or write locking the protected area. Only one write lock may exist for a given segment of a file at a given time. The file descriptor on which a write lock is being placed must have been opened with write access.

The structure *flock* describes the type (**L_type**), starting offset (**L_whence**), relative offset (**L_start**), size (**L_len**), and process id (**L_pid**) of the segment of the file to be affected. The process id field is only used with the **F_GETLK** *cmd* to return the value for a block in lock. Locks may start and extend beyond the current end of a file, but may not be negative relative to the beginning of the file. A lock may be set to always extend to the end of file by setting **L_len** to zero (0). If such a lock also has **L_start** set to zero (0), the whole file will be locked. Changing or unlocking a segment from the middle of a larger locked segment leaves two smaller segments for either end. Locking a segment that is already locked by the calling process causes the old lock type to be removed and the new lock type to take affect. All locks associated with a file for a given process are removed when a file descriptor for that file is closed by that process or the process holding that file descriptor terminates. Locks are not inherited by a child process in a *fork(2)* system call.

RETURN VALUE

Upon successful completion, the value returned depends on *cmd* as follows:

F_DUPFD	A new file descriptor.
F_GETFD	Value of flag (only the low-order bit is defined).
F_GETFL	Value of flags.
F_GETLK	Value other than -1.
F_GETOWN	Value of file descriptor owner.
F_SETFD	Value other than -1.
F_SETFL	Value other than -1.
F_SETLK	Value other than -1.
F_SETLKW	Value other than -1.

Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

ERRORS

Fcntl will fail if one or more of the following are true:

[EBADF]	<i>Fildes</i> is not a valid open file descriptor.
[EMFILE]	<i>Cmd</i> is F_DUPFD and the maximum allowed number of file descriptors are currently open.

- [EINVAL] *Cmd* is F_DUPFD and *arg* is negative or greater the maximum allowable number (see *getdtablesize(2)*).
- [EINVAL] *Cmd* is F_GETLK, F_SETLK, or SETLKW and *arg* or the data it points to is not valid.
- [EACCESS] *Cmd* is F_SETLK, the type of lock (*Ltype*) is a read (F_RDLCK) or write (F_WRLCK) lock and the segment of a file to be locked is already write locked by another process or the type is a write lock and the segment of a file to be locked is already read or write locked by another process.
- [EMFILE] *Cmd* is F_SETLK or F_SETLKW, the type of lock is a read or write lock and there are no more file locking headers available (too many files have segments locked).
- [ENOSPC] *Cmd* is F_SETLK or F_SETLKW, the type of lock is a read or write lock and there are no more file locking headers available (too many files have segments locked) or there are no more record locks available (too many file segments locked).
- [EDEADLK] *Cmd* is F_SETLK, when the lock is blocked by some lock from another process and sleeping (waiting) for that lock to become free, this causes a deadlock situation.

SEE ALSO

close(2), execve(2), getdtablesize(2), open(2), sigvec(2), fcntl(5)

BUGS

The asynchronous I/O facilities of FNDELAY and FASYNC are currently available only for tty operations. No SIGIO signal is sent upon draining of output sufficiently for non-blocking writes to occur.

NAME

flock – apply or remove an advisory lock on an open file

SYNOPSIS

```
#include <sys/file.h>
#define LOCK_SH 1 /* shared lock */
#define LOCK_EX 2 /* exclusive lock */
#define LOCK_NB 4 /* don't block when locking */
#define LOCK_UN 8 /* unlock */

flock(fd, operation)
int fd, operation;
```

DESCRIPTION

Flock applies or removes an *advisory* lock on the file associated with the file descriptor *fd*. A lock is applied by specifying an *operation* parameter which is the inclusive or of `LOCK_SH` or `LOCK_EX` and, possibly, `LOCK_NB`. To unlock an existing lock *operation* should be `LOCK_UN`.

Advisory locks allow cooperating processes to perform consistent operations on files, but do not guarantee consistency (i.e. processes may still access files without using advisory locks possibly resulting in inconsistencies).

The locking mechanism allows two types of locks: *shared* locks and *exclusive* locks. At any time multiple shared locks may be applied to a file, but at no time are multiple exclusive, or both shared and exclusive, locks allowed simultaneously on a file.

A shared lock may be *upgraded* to an exclusive lock, and vice versa, simply by specifying the appropriate lock type; this results in the previous lock being released and the new lock applied (possibly after other processes have gained and released the lock).

Requesting a lock on an object which is already locked normally causes the caller to be blocked until the lock may be acquired. If `LOCK_NB` is included in *operation*, then this will not happen; instead the call will fail and the error `EWOULDBLOCK` will be returned.

NOTES

Locks are on files, not file descriptors. That is, file descriptors duplicated through `dup(2)` or `fork(2)` do not result in multiple instances of a lock, but rather multiple references to a single lock. If a process holding a lock on a file forks and the child explicitly unlocks the file, the parent will lose its lock.

Processes blocked awaiting a lock may be awakened by signals.

RETURN VALUE

Zero is returned if the operation was successful; on an error a `-1` is returned and an error code is left in the global location *errno*.

ERRORS

The *flock* call fails if:

[<code>EWOULDBLOCK</code>]	The file is locked and the <code>LOCK_NB</code> option was specified.
[<code>EBADF</code>]	The argument <i>fd</i> is an invalid descriptor.
[<code>EINVAL</code>]	The argument <i>fd</i> refers to an object other than a file.

SEE ALSO

open(2), close(2), dup(2), execve(2), fork(2)

NAME

fork - create a new process

SYNOPSIS

```
pid = fork()
int pid;
```

DESCRIPTION

Fork causes creation of a new process. The new process (child process) is an exact copy of the calling process except for the following:

The child process has a unique process ID.

The child process has a different parent process ID (i.e., the process ID of the parent process).

The child process has its own copy of the parent's descriptors. These descriptors reference the same underlying objects, so that, for instance, file pointers in file objects are shared between the child and the parent, so that a *lseek(2)* on a descriptor in the child process can affect a subsequent *read* or *write* by the parent. This descriptor copying is also used by the shell to establish standard input and output for newly created processes as well as to set up pipes.

The child processes resource utilizations are set to 0; see *setrlimit(2)*.

RETURN VALUE

Upon successful completion, *fork* returns a value of 0 to the child process and returns the process ID of the child process to the parent process. Otherwise, a value of -1 is returned to the parent process, no child process is created, and the global variable *errno* is set to indicate the error.

ERRORS

Fork will fail and no child process will be created if one or more of the following are true:

- [EAGAIN] The system-imposed limit {PROC_MAX} on the total number of processes under execution would be exceeded.
- [EAGAIN] The system-imposed limit {KID_MAX} on the total number of processes under execution by a single user would be exceeded.

SEE ALSO

execve(2), *wait(2)*



NAME

fsync - synchronize a file's in-core state with that on disk

SYNOPSIS

```
fsync(fd)
int fd;
```

DESCRIPTION

Fsync causes all modified data and attributes of *fd* to be moved to a permanent storage device. This normally results in all in-core modified copies of buffers for the associated file to be written to a disk.

Fsync should be used by programs which require a file to be in a known state; for example in building a simple transaction facility.

RETURN VALUE

A 0 value is returned on success. A -1 value indicates an error.

ERRORS

The *fsync* fails if:

[EBADF]	<i>Fd</i> is not a valid descriptor.
[EINVAL]	<i>Fd</i> refers to a socket, not to a file.

SEE ALSO

sync(2), *sync*(8), *update*(8)

BUGS

The current implementation of this call is expensive for large files.

NAME

getdtablesize - get descriptor table size

SYNOPSIS

```
nds = getdtablesize()  
int nds;
```

DESCRIPTION

Each process has a fixed size descriptor table which is guaranteed to have at least 20 slots. The entries in the descriptor table are numbered with small integers starting at 0. The call *getdtablesize* returns the size of this table.

SEE ALSO

close(2), *dup(2)*, *open(2)*

NAME

getgid, getegid - get group identity

SYNOPSIS

```
gid = getgid()
int gid;
egid = getegid()
int egid;
```

DESCRIPTION

Getgid returns the real group ID of the current process, *getegid* the effective group ID.

The real group ID is specified at login time.

The effective group ID is more transient, and determines additional access permission during execution of a "set-group-ID" process, and it is for such processes that *getgid* is most useful.

SEE ALSO

getuid(2), setregid(2), setgid(3)

NAME

getgroups - get group access list

SYNOPSIS

```
#include <sys/param.h>
getgroups(ngroups, gidset)
int *ngroups, *gidset;
```

DESCRIPTION

Getgroups gets the current group access list of the user process and stores it in the array *gidset*. The parameter *ngroups* indicates the number of entries which may be placed in *gidset* and is modified on return to indicate the actual number of groups returned. No more than NGRPS, as defined in *<sys/param.h>*, will ever be returned.

RETURN VALUE

A value of 0 indicates that the call succeeded, and that the number of elements of *gidset* and the set itself were returned. A value of -1 indicates that an error occurred, and the error code is stored in the global variable *errno*.

ERRORS

The possible errors for *getgroup* are:

[EFAULT] The arguments *ngroups* or *gidset* specify invalid addresses.

SEE ALSO

setgroups(2), initgroups(3)

NAME

gethostid, sethostid – get/set unique identifier of current host

SYNOPSIS

```
hostid = gethostid()
int hostid;
sethostid(hostid)
int hostid;
```

DESCRIPTION

Sethostid establishes a 32-bit identifier for the current processor which is intended to be unique among all UNIX systems in existence. This is normally a DARPA Internet address for the local machine. This call is allowed only to the super-user and is normally performed at boot time.

Gethostid returns the 32-bit identifier for the current processor.

SEE ALSO

hostid(1), gethostname(2)

BUGS

32 bits for the identifier is too small.

NAME

gethostname, sethostname – get/set name of current host

SYNOPSIS

```
gethostname(name, namelen)
char *name;
int namelen;

sethostname(name, namelen)
char *name;
int namelen;
```

DESCRIPTION

Gethostname returns the standard host name for the current processor, as previously set by *sethostname*. The parameter *namelen* specifies the size of the *name* array. The returned name is null-terminated unless insufficient space is provided.

Sethostname sets the name of the host machine to be *name*, which has length *namelen*. This call is restricted to the super-user and is normally used only when the system is bootstrapped.

RETURN VALUE

If the call succeeds a value of 0 is returned. If the call fails, then a value of -1 is returned and an error code is placed into the global location *errno*.

ERRORS

The following errors may be returned by these calls:

[EFAULT]	The <i>name</i> or <i>namelen</i> parameter gave an invalid address.
[EPERM]	The caller was not the super-user.

SEE ALSO

gethostid(2)

BUGS

Host names are limited to 255 characters.

NAME

getitimer, setitimer – get/set value of interval timer

SYNOPSIS

```
#include <sys/time.h>
#define ITIMER_REAL 0 /* real time intervals */
#define ITIMER_VIRTUAL 1 /* virtual time intervals */
#define ITIMER_PROF 2 /* user and system virtual time */
```

```
getitimer(which, value)
int which;
struct itimerval *value;

setitimer(which, value, ovalue)
int which;
struct itimerval *value, *ovalue;
```

DESCRIPTION

The system provides each process with three interval timers, defined in `<sys/time.h>`. The `getitimer` call returns the current value for the timer specified in `which`, while the `setitimer` call sets the value of a timer (optionally returning the previous value of the timer).

A timer value is defined by the `itimerval` structure:

```
struct itimerval {
    struct timeval it_interval; /* timer interval */
    struct timeval it_value; /* current value */
};
```

If `it_value` is non-zero, it indicates the time to the next timer expiration. If `it_interval` is non-zero, it specifies a value to be used in reloading `it_value` when the timer expires. Setting `it_value` to 0 disables a timer. Setting `it_interval` to 0 causes a timer to be disabled after its next expiration (assuming `it_value` is non-zero).

Time values smaller than the resolution of the system clock are rounded up to this resolution.

The `ITIMER_REAL` timer decrements in real time. A `SIGALRM` signal is delivered when this timer expires.

The `ITIMER_VIRTUAL` timer decrements in process virtual time. It runs only when the process is executing. A `SIGVTALRM` signal is delivered when it expires.

The `ITIMER_PROF` timer decrements both in process virtual time and when the system is running on behalf of the process. It is designed to be used by interpreters in statistically profiling the execution of interpreted programs. Each time the `ITIMER_PROF` timer expires, the `SIGPROF` signal is delivered. Because this signal may interrupt in-progress system calls, programs using this timer must be prepared to restart interrupted system calls.

NOTES

Three macros for manipulating time values are defined in `<sys/time.h>`. `Timerclear` sets a time value to zero, `timerisset` tests if a time value is non-zero, and `timercmp` compares two time values (beware that `>=` and `<=` do not work with this macro).

RETURN VALUE

If the calls succeed, a value of 0 is returned. If an error occurs, the value -1 is returned, and a more precise error code is placed in the global variable *errno*.

ERRORS

The possible errors are:

[EFAULT] The *value* structure specified a bad address.

[EINVAL] A *value* structure specified a time was too large to be handled.

SEE ALSO

sigvec(2), gettimeofday(2)

NAME

getpagesize – get system page size

SYNOPSIS

```
pagesize = getpagesize()
int pagesize;
```

DESCRIPTION

Getpagesize returns the number of bytes in a page. Page granularity is the granularity of many of the memory management calls.

The page size is a *system* page size and may not be the same as the underlying hardware page size.

SEE ALSO

sbrk(2), pagesize(1)

NAME

getpeername – get name of connected peer

SYNOPSIS

```
getpeername(s, name, namelen)
int s;
struct sockaddr *name;
int *namelen;
```

DESCRIPTION

Getpeername returns the name of the peer connected to socket *s*. The *namelen* parameter should be initialized to indicate the amount of space pointed to by *name*. On return it contains the actual size of the name returned (in bytes).

DIAGNOSTICS

A 0 is returned if the call succeeds, -1 if it fails.

ERRORS

The call succeeds unless:

- [EBADF] The argument *s* is not a valid descriptor.
- [ENOTSOCK] The argument *s* is a file, not a socket.
- [ENOTCONN] The socket is not connected.
- [ENOBUFS] Insufficient resources were available in the system to perform the operation.
- [EFAULT] The *name* parameter points to memory not in a valid part of the process address space.

SEE ALSO

bind(2), socket(2), getsockname(2)

BUGS

Names bound to sockets in the UNIX domain are inaccessible; *getpeername* returns a zero length name.

NAME

getpgrp - get process group

SYNOPSIS

```
pgrp = getpgrp(pid)
int pgrp;
int pid;
```

DESCRIPTION

The process group of the specified process is returned by *getpgrp*. If *pid* is zero, then the call applies to the current process.

Process groups are used for distribution of signals, and by terminals to arbitrate requests for their input: processes which have the same process group as the terminal are foreground and may read, while others will block with a signal if they attempt to read.

This call is thus used by programs such as *cs**h*(1) to create process groups in implementing job control. The *TIOCGPGRP* and *TIOCSPGRP* calls described in *tt**y*(4) are used to get/set the process group of the control terminal.

SEE ALSO

setpgrp(2), getuid(2), tt*y*(4)

NAME

getpid, getppid - get process identification

SYNOPSIS

```
pid = getpid()
long pid;
```

```
ppid = getppid()
long ppid;
```

DESCRIPTION

Getpid returns the process ID of the current process. Most often it is used with the host identifier *gethostid(2)* to generate uniquely-named temporary files.

Getppid returns the process ID of the parent of the current process.

SEE ALSO

gethostid(2)

NAME

getpriority, setpriority – get/set program scheduling priority

SYNOPSIS

```
#include <sys/resource.h>
#define PRIO_PROCESS 0    /* process */
#define PRIO_PGRP     1    /* process group */
#define PRIO_USER     2    /* user id */

prio = getpriority(which, who)
int prio, which, who;

setpriority(which, who, prio)
int which, who, prio;
```

DESCRIPTION

The scheduling priority of the process, process group, or user, as indicated by *which* and *who* is obtained with the *getpriority* call and set with the *setpriority* call. *Which* is one of PRIO_PROCESS, PRIO_PGRP, or PRIO_USER, and *who* is interpreted relative to *which* (a process identifier for PRIO_PROCESS, process group identifier for PRIO_PGRP, and a user ID for PRIO_USER). *Prio* is a value in the range -20 to 20. The default priority is 0; lower priorities cause more favorable scheduling.

The *getpriority* call returns the highest priority (lowest numerical value) enjoyed by any of the specified processes. The *setpriority* call sets the priorities of all of the specified processes to the specified value. Only the super-user may lower priorities.

RETURN VALUE

Since *getpriority* can legitimately return the value -1, it is necessary to clear the external variable *errno* prior to the call, then check it afterward to determine if a -1 is an error or a legitimate value. The *setpriority* call returns 0 if there is no error, or -1 if there is.

ERRORS

Getpriority and *setpriority* may return one of the following errors:

- [ESRCH] No process(es) were located using the *which* and *who* values specified.
- [EINVAL] *Which* was not one of PRIO_PROCESS, PRIO_PGRP, or PRIO_USER.

In addition to the errors indicated above, *setpriority* may fail with one of the following errors returned:

- [EACCES] A process was located, but neither its effective nor real user ID matched the effective user ID of the caller.
- [EACCES] A non super-user attempted to change a process priority to a negative value.

SEE ALSO

nice(1), fork(2), renice(8)

NAME

getrlimit, setrlimit – control maximum system resource consumption

SYNOPSIS

```
#include <sys/time.h>
#include <sys/resource.h>

getrlimit(resource, rlp)
int resource;
struct rlimit *rlp;

setrlimit(resource, rlp)
int resource;
struct rlimit *rlp;
```

DESCRIPTION

Limits on the consumption of system resources by the current process and each process it creates may be obtained with the *getrlimit* call, and set with the *setrlimit* call.

The *resource* parameter is one of the following:

RLIMIT_CPU	the maximum amount of cpu time (in milliseconds) to be used by each process.
RLIMIT_FSIZE	the largest size, in bytes, of any single file which may be created.
RLIMIT_DATA	the maximum size, in bytes, of the data segment for a process; this defines how far a program may extend its break with the <i>sbrk(2)</i> system call.
RLIMIT_STACK	the maximum size, in bytes, of the stack segment for a process; this defines how far a program's stack segment may be extended, either automatically by the system, or explicitly by a user with the <i>sbrk(2)</i> system call.
RLIMIT_CORE	the largest size, in bytes, of a <i>core</i> file which may be created.
RLIMIT_RSS	the maximum size, in bytes, a process's resident set size may grow to. This imposes a limit on the amount of physical memory to be given to a process; if memory is tight, the system will prefer to take memory from processes which are exceeding their declared resident set size.

A resource limit is specified as a soft limit and a hard limit. When a soft limit is exceeded a process may receive a signal (for example, if the cpu time is exceeded), but it will be allowed to continue execution until it reaches the hard limit (or modifies its resource limit). The *rlimit* structure is used to specify the hard and soft limits on a resource,

```
struct rlimit {
    int    rlim_cur;    /* current (soft) limit */
    int    rlim_max;    /* hard limit */
};
```

Only the super-user may raise the maximum limits. Other users may only alter *rlim_cur* within the range from 0 to *rlim_max* or (irreversibly) lower *rlim_max*.

An "infinite" value for a limit is defined as RLIMIT_INFINITY (0x7fffffff).

Because this information is stored in the per-process information, this system call must be executed directly by the shell if it is to affect all future processes created by the shell; *limit* is thus a built-in command to *cs(1)*.

The system refuses to extend the data or stack space when the limits would be exceeded in the normal way: a *break* call fails if the data space limit is reached, or the process is killed when the stack limit is reached (since the stack cannot be extended, there is no way to send a signal!).

A file i/o operation which would create a file which is too large will cause a signal SIGXFSZ to be generated, this normally terminates the process, but may be caught. When the soft cpu time limit is exceeded, a signal SIGXCPU is sent to the offending process.

RETURN VALUE

A 0 return value indicates that the call succeeded, changing or returning the resource limit. A return value of -1 indicates that an error occurred, and an error code is stored in the global location *errno*.

ERRORS

The possible errors are:

- [EFAULT] The address specified for *rlp* is invalid.
- [EPERM] The limit specified to *setrlimit* would have raised the maximum limit value, and the caller is not the super-user.

SEE ALSO

csh(1), *quota*(2)

BUGS

There should be *limit* and *unlimit* commands in *sh*(1) as well as in *csh*.

NAME

getrusage - get information about resource utilization

SYNOPSIS

```
#include <sys/time.h>
#include <sys/resource.h>
#define RUSAGE_SELF      0
    /* calling process */
#define RUSAGE_CHILDREN -1
    /* terminated child processes */

getrusage(who, rusage)
int who;
struct rusage *rusage;
```

DESCRIPTION

Getrusage returns information describing the resources utilized by the current process, or all its terminated child processes. The *who* parameter is one of `RUSAGE_SELF` and `RUSAGE_CHILDREN`. If *rusage* is non-zero, the buffer it points to will be filled in with the following structure:

```
struct rusage {
    struct timeval ru_utime;    /* user time used */
    struct timeval ru_stime;    /* system time used */
    int ru_maxrss;
    int ru_ixrss;              /* integral shared memory size */
    int ru_idrss;              /* integral unshared data size */
    int ru_isrss;              /* integral unshared stack size */
    int ru_minflt;            /* page reclaims */
    int ru_majflt;            /* page faults */
    int ru_nswap;              /* swaps */
    int ru_inblock;            /* block input operations */
    int ru_oublock;            /* block output operations */
    int ru_msgsnd;             /* messages sent */
    int ru_msrvcv;             /* messages received */
    int ru_nsignals;           /* signals received */
    int ru_nvcsw;              /* voluntary context switches */
    int ru_nivcsw;             /* involuntary context switches */
};
```

The fields are interpreted as follows:

<code>ru_utime</code>	the total amount of time spent executing in user mode.
<code>ru_stime</code>	the total amount of time spent in the system executing on behalf of the process(es).
<code>ru_maxrss</code>	the maximum resident set size utilized (in kilobytes).
<code>ru_ixrss</code>	an "integral" value indicating the amount of memory used which was also shared among other processes. This value is expressed in units of kilobytes * seconds-of-execution and is calculated by summing the number of shared memory pages in use each time the internal system clock ticks and then averaging over 1 second intervals.
<code>ru_idrss</code>	an integral value of the amount of unshared memory residing in the data

	segment of a process (expressed in units of kilobytes * seconds-of-execution).
<code>ru_isrss</code>	an integral value of the amount of unshared memory residing in the stack segment of a process (expressed in units of kilobytes * seconds-of-execution).
<code>ru_minflt</code>	the number of page faults serviced without any i/o activity; here i/o activity is avoided by "reclaiming" a page frame from the list of pages awaiting reallocation.
<code>ru_majflt</code>	the number of page faults serviced which required i/o activity.
<code>ru_nswap</code>	the number of times a process was "swapped" out of main memory.
<code>ru_inblock</code>	the number of times the file system had to perform input.
<code>ru_outblock</code>	the number of times the file system had to perform output.
<code>ru_msgsnd</code>	the number of ipc messages sent.
<code>ru_msgrcv</code>	the number of ipc messages received.
<code>ru_nsignals</code>	the number of signals delivered.
<code>ru_nvcsw</code>	the number of times a context switch resulted due to a process voluntarily giving up the processor before its time slice was completed (usually to await availability of a resource).
<code>ru_nivcsw</code>	the number of times a context switch resulted due to a higher priority process becoming runnable or because the current process exceeded its time slice.

NOTES

The numbers `ru_inblock` and `ru_outblock` account only for real i/o; data supplied by the caching mechanism is charged only to the first process to read or write the data.

SEE ALSO

`gettimeofday(2)`, `wait(2)`

BUGS

There is no way to obtain information about a child process which has not yet terminated.

NAME

`getsockname` - get socket name

SYNOPSIS

```
getsockname(s, name, namelen)
int s;
struct sockaddr *name;
int *namelen;
```

DESCRIPTION

Getsockname returns the current *name* for the specified socket. The *namelen* parameter should be initialized to indicate the amount of space pointed to by *name*. On return it contains the actual size of the name returned (in bytes).

DIAGNOSTICS

A 0 is returned if the call succeeds, -1 if it fails.

ERRORS

The call succeeds unless:

- | | |
|------------|--|
| [EBADF] | The argument <i>s</i> is not a valid descriptor. |
| [ENOTSOCK] | The argument <i>s</i> is a file, not a socket. |
| [ENOBUFS] | Insufficient resources were available in the system to perform the operation. |
| [EFAULT] | The <i>name</i> parameter points to memory not in a valid part of the process address space. |

SEE ALSO

`bind(2)`, `socket(2)`

BUGS

Names bound to sockets in the UNIX domain are inaccessible; *getsockname* returns a zero length name.

NAME

getsockopt, setsockopt – get and set options on sockets

SYNOPSIS

```
#include <sys/types.h>
#include <sys/socket.h>

getsockopt(s, level, optname, optval, optlen)
int s, level, optname;
char *optval;
int *optlen;

setsockopt(s, level, optname, optval, optlen)
int s, level, optname;
char *optval;
int optlen;
```

DESCRIPTION

Getsockopt and *setsockopt* manipulate *options* associated with a socket. Options may exist at multiple protocol levels; they are always present at the uppermost “socket” level.

When manipulating socket options the level at which the option resides and the name of the option must be specified. To manipulate options at the “socket” level, *level* is specified as SOL_SOCKET. To manipulate options at any other level the protocol number of the appropriate protocol controlling the option is supplied. For example, to indicate an option is to be interpreted by the TCP protocol, *level* should be set to the protocol number of TCP; see *getprotoent(3N)*.

The parameters *optval* and *optlen* are used to access option values for *setsockopt*. For *getsockopt* they identify a buffer in which the value for the requested option(s) are to be returned. For *getsockopt*, *optlen* is a value-result parameter, initially containing the size of the buffer pointed to by *optval*, and modified on return to indicate the actual size of the value returned. If no option value is to be supplied or returned, *optval* may be supplied as 0.

Optname and any specified options are passed uninterpreted to the appropriate protocol module for interpretation. The include file *<sys/socket.h>* contains definitions for “socket” level options; see *socket(2)*. Options at other protocol levels vary in format and name, consult the appropriate entries in (4P).

RETURN VALUE

A 0 is returned if the call succeeds, -1 if it fails.

ERRORS

The call succeeds unless:

[EBADF]	The argument <i>s</i> is not a valid descriptor.
[ENOTSOCK]	The argument <i>s</i> is a file, not a socket.
[ENOPROTOOPT]	The option is unknown.
[EFAULT]	The options are not in a valid part of the process address space.

SEE ALSO

socket(2), getprotoent(3N)

NAME

gettimeofday, settimeofday – get/set date and time

SYNOPSIS

```
#include <sys/time.h>
gettimeofday(tp, tzp)
struct timeval *tp;
struct timezone *tzp;
settimeofday(tp, tzp)
struct timeval *tp;
struct timezone *tzp;
```

DESCRIPTION

Gettimeofday returns the system's notion of the current Greenwich time and the current time zone. Time returned is expressed relative in seconds and microseconds since midnight January 1, 1970.

The structures pointed to by *tp* and *tzp* are defined in *<sys/time.h>* as:

```
struct timeval {
    u_long tv_sec;
        /* seconds since Jan. 1, 1970 */
    long tv_usec;
        /* and microseconds */
};

struct timezone {
    int tz_minuteswest;
        /* of Greenwich */
    int tz_dsttime;
        /* type of dst correction to apply */
};
```

The *timezone* structure indicates the local time zone (measured in minutes of time westward from Greenwich), and a flag that, if nonzero, indicates that Daylight Saving time applies locally during the appropriate part of the year.

Only the super-user may set the time of day.

RETURN

A 0 return value indicates that the call succeeded. A -1 return value indicates an error occurred, and in this case an error code is stored into the global variable *errno*.

ERRORS

The following error codes may be set in *errno*:

[EFAULT]	An argument address referenced invalid memory.
[EPERM]	A user other than the super-user attempted to set the time.

SEE ALSO

date(1), ctime(3)

BUGS

Time is never correct enough to believe the microsecond values. There should a mechanism by which, at least, local clusters of systems might synchronize their clocks to millisecond granularity.

NAME

getuid, geteuid - get user identity

SYNOPSIS

```
uid = getuid()
int uid;

euid = geteuid()
int euid;
```

DESCRIPTION

Getuid returns the real user ID of the current process, *geteuid* the effective user ID.

The real user ID identifies the person who is logged in. The effective user ID gives the process additional permissions during execution of "set-user-ID" mode processes, which use *getuid* to determine the real-user-id of the process which invoked them.

SEE ALSO

getgid(2), setreuid(2)

NAME

ioctl – control device

SYNOPSIS

```
#include <sys/ioctl.h>
ioctl(d, request, argp)
int d, request;
char *argp;
```

DESCRIPTION

Ioctl performs a variety of functions on open descriptors. In particular, many operating characteristics of character special files (e.g. terminals) may be controlled with *ioctl* requests. The writeups of various devices in section 4 discuss how *ioctl* applies to them.

An *ioctl request* has encoded in it whether the argument is an “in” parameter or “out” parameter, and the size of the argument *argp* in bytes. Macros and defines used in specifying an *ioctl request* are located in the file *<sys/ioctl.h>*.

RETURN VALUE

If an error has occurred, a value of -1 is returned and *errno* is set to indicate the error.

ERRORS

Ioctl will fail if one or more of the following are true:

- | | |
|----------|--|
| [EBADF] | <i>D</i> is not a valid descriptor. |
| [ENOTTY] | <i>D</i> is not associated with a character special device. |
| [ENOTTY] | The specified request does not apply to the kind of object which the descriptor <i>d</i> references. |
| [EINVAL] | <i>Request</i> or <i>argp</i> is not valid. |

SEE ALSO

execve(2), *fcntl(2)*, *mt(4)*, *tty(4)*, *intro(4N)*

NAME

kill - send signal to a process

SYNOPSIS

```
kill(pid, sig)
int pid, sig;
```

DESCRIPTION

Kill sends the signal *sig* to a process, specified by the process number *pid*. *Sig* may be one of the signals specified in *sigvec(2)*, or it may be 0, in which case error checking is performed but no signal is actually sent. This can be used to check the validity of *pid*.

The sending and receiving processes must have the same effective user ID, otherwise this call is restricted to the super-user. A single exception is the signal SIGCONT which may always be sent to any child or grandchild of the current process.

If the process number is 0, the signal is sent to all other processes in the sender's process group; this is a variant of *killpg(2)*.

If the process number is -1, and the user is the super-user, the signal is broadcast universally except to system processes and the process sending the signal.

Processes may send signals to themselves.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

ERRORS

Kill will fail and no signal will be sent if any of the following occur:

- | | |
|----------|--|
| [EINVAL] | <i>Sig</i> is not a valid signal number. |
| [ESRCH] | No process can be found corresponding to that specified by <i>pid</i> . |
| [EPERM] | The sending process is not the super-user and its effective user id does not match the effective user-id of the receiving process. |

SEE ALSO

getpid(2), getpgrp(2), killpg(2), sigvec(2)

NAME

killpg - send signal to a process group

SYNOPSIS

```
killpg(pgrp, sig)
int pgrp, sig;
```

DESCRIPTION

Killpg sends the signal *sig* to the process group *pgrp*. See *sigvec(2)* for a list of signals.

The sending process and members of the process group must have the same effective user ID, otherwise this call is restricted to the super-user. As a single special case the continue signal SIGCONT may be sent to any process which is a descendant of the current process.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and the global variable *errno* is set to indicate the error.

ERRORS

Killpg will fail and no signal will be sent if any of the following occur:

- | | |
|----------|--|
| [EINVAL] | <i>Sig</i> is not a valid signal number. |
| [ESRCH] | No process can be found corresponding to that specified by <i>pid</i> . |
| [EPERM] | The sending process is not the super-user and one or more of the target processes has an effective user ID different from that of the sending process. |

SEE ALSO

kill(2), getpgrp(2), sigvec(2)

NAME

link - make a hard link to a file

SYNOPSIS

```
link(name1, name2)
char *name1, *name2;
```

DESCRIPTION

A hard link to *name1* is created; the link has the name *name2*. *Name1* must exist.

With hard links, both *name1* and *name2* must be in the same file system. Unless the caller is the super-user, *name1* must not be a directory. Both the old and the new *link* share equal access and rights to the underlying object.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

ERRORS

Link will fail and no link will be created if one or more of the following are true:

- | | |
|-----------|--|
| [EPERM] | Either pathname contains a byte with the high-order bit set. |
| [ENOENT] | Either pathname was too long. |
| [ENOTDIR] | A component of either path prefix is not a directory. |
| [ENOENT] | A component of either path prefix does not exist. |
| [EACCES] | A component of either path prefix denies search permission. |
| [ENOENT] | The file named by <i>name1</i> does not exist. |
| [EEXIST] | The link named by <i>name2</i> does exist. |
| [EPERM] | The file named by <i>name1</i> is a directory and the effective user ID is not super-user. |
| [EXDEV] | The link named by <i>name2</i> and the file named by <i>name1</i> are on different file systems. |
| [EACCES] | The requested link requires writing in a directory with a mode that denies write permission. |
| [EROFS] | The requested link requires writing in a directory on a read-only file system. |
| [EFAULT] | One of the pathnames specified is outside the process's allocated address space. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |

SEE ALSO

symlink(2), unlink(2)

NAME

listen – listen for connections on a socket

SYNOPSIS

```
listen(s, backlog)  
int s, backlog;
```

DESCRIPTION

To accept connections, a socket is first created with *socket*(2), a backlog for incoming connections is specified with *listen*(2) and then the connections are accepted with *accept*(2). The *listen* call applies only to sockets of type SOCK_STREAM or SOCK_DGRAM.

The *backlog* parameter defines the maximum length the queue of pending connections may grow to. If a connection request arrives with the queue full the client will receive an error with an indication of ECONNREFUSED.

RETURN VALUE

A 0 return value indicates success; -1 indicates an error.

ERRORS

The call fails if:

[EBADF]	The argument <i>s</i> is not a valid descriptor.
[ENOTSOCK]	The argument <i>s</i> is not a socket.
[EOPNOTSUPP]	The socket is not of a type that supports the operation <i>listen</i> .

SEE ALSO

accept(2), *connect*(2), *socket*(2)

BUGS

The *backlog* is currently limited (silently) to 5.

NAME

`lseek` – move read/write pointer

SYNOPSIS

```
#define L_SET 0    /* set the seek pointer */
#define L_INCR 1  /* increment the seek pointer */
#define L_XTND 2  /* extend the file size */

pos = lseek(d, offset, whence)
int pos;
int d, offset, whence;
```

DESCRIPTION

The descriptor *d* refers to a file or device open for reading and/or writing. *Lseek* sets the file pointer of *d* as follows:

If *whence* is `L_SET`, the pointer is set to *offset* bytes.

If *whence* is `L_INCR`, the pointer is set to its current location plus *offset*.

If *whence* is `L_XTND`, the pointer is set to the size of the file plus *offset*.

Upon successful completion, the resulting pointer location as measured in bytes from beginning of the file is returned. Some devices are incapable of seeking. The value of the pointer associated with such a device is undefined.

NOTES

Seeking far beyond the end of a file, then writing, creates a gap or “hole”, which occupies no physical space and reads as zeros.

RETURN VALUE

Upon successful completion, a non-negative integer, the current file pointer value, is returned. Otherwise, a value of `-1` is returned and *errno* is set to indicate the error.

ERRORS

Lseek will fail and the file pointer will remain unchanged if:

- [EBADF] *Fildes* is not an open file descriptor.
- [ESPIPE] *Fildes* is associated with a pipe or a socket.
- [EINVAL] *Whence* is not a proper value.
- [EINVAL] The resulting file pointer would be negative.

SEE ALSO

`dup(2)`, `open(2)`

BUGS

This document’s use of *whence* is incorrect English, but maintained for historical reasons.

NAME

`mkdir` – make a directory file

SYNOPSIS

```
mkdir(path, mode)  
char *path;  
int mode;
```

DESCRIPTION

Mkdir creates a new directory file with name *path*. The mode of the new file is initialized from *mode*. (The protection part of the mode is modified by the process's mode mask; see *umask*(2)).

The directory's owner ID is set to the process's effective user ID. The directory's group ID is set to that of the parent directory in which it is created.

The low-order 9 bits of mode are modified by the process's file mode creation mask: all bits set in the process's file mode creation mask are cleared. See *umask*(2).

RETURN VALUE

A 0 return value indicates success. A -1 return value indicates an error, and an error code is stored in *errno*.

ERRORS

Mkdir will fail and no directory will be created if:

[EPERM]	The process's effective user ID is not super-user.
[EPERM]	The <i>path</i> argument contains a byte with the high-order bit set.
[ENOTDIR]	A component of the path prefix is not a directory.
[ENOENT]	A component of the path prefix does not exist.
[EROFS]	The named file resides on a read-only file system.
[EEXIST]	The named file exists.
[EFAULT]	<i>Path</i> points outside the process's allocated address space.
[ELOOP]	Too many symbolic links were encountered in translating the pathname.
[EIO]	An I/O error occurred while writing to the file system.

SEE ALSO

chmod(2), *stat*(2), *umask*(2)

NAME

`mknod` - make a special file

SYNOPSIS

```
mknod(path, mode, dev)  
char *path;  
int mode, dev;
```

DESCRIPTION

Mknod creates a new file whose name is *path*. The mode of the new file (including special file bits) is initialized from *mode*. (The protection part of the mode is modified by the process's mode mask; see *umask*(2)). The first block pointer of the i-node is initialized from *dev* and is used to specify which device the special file refers to.

If *mode* indicates a block or character special file, *dev* is a configuration dependent specification of a character or block I/O device. If *mode* does not indicate a block special or character special device, *dev* is ignored.

Mknod may be invoked only by the super-user.

RETURN VALUE

Upon successful completion a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

ERRORS

Mknod will fail and the file mode will be unchanged if:

[EPERM]	The process's effective user ID is not super-user.
[EPERM]	The pathname contains a character with the high-order bit set.
[ENOTDIR]	A component of the path prefix is not a directory.
[ENOENT]	A component of the path prefix does not exist.
[EROFS]	The named file resides on a read-only file system.
[EEXIST]	The named file exists.
[EFAULT]	<i>Path</i> points outside the process's allocated address space.
[ELOOP]	Too many symbolic links were encountered in translating the pathname.

SEE ALSO

chmod(2), *stat*(2), *umask*(2)

NAME

mount, umount – mount or remove file system

SYNOPSIS

```
mount(special, name, rwflag)
char *special, *name;
int rwflag;

umount(special)
char *special;
```

DESCRIPTION

Mount announces to the system that a removable file system has been mounted on the block-structured special file *special*; from now on, references to file *name* will refer to the root file on the newly mounted file system. *Special* and *name* are pointers to null-terminated strings containing the appropriate path names.

Name must exist already. *Name* must be a directory. Its old contents are inaccessible while the file system is mounted.

The *rwflag* argument determines whether the file system can be written on; if it is 0 writing is allowed, if non-zero no writing is done. Physically write-protected and magnetic tape file systems must be mounted read-only or errors will occur when access times are updated, whether or not any explicit write is attempted.

Unmount announces to the system that the *special* file is no longer to contain a removable file system. The associated file reverts to its ordinary interpretation.

RETURN VALUE

Mount returns 0 if the action occurred, -1 if *special* is inaccessible or not an appropriate file, if *name* does not exist, if *special* is already mounted, if *name* is in use, or if there are already too many file systems mounted.

Unmount returns 0 if the action occurred; -1 if the special file is inaccessible or does not have a mounted file system, or if there are active files in the mounted file system.

ERRORS

Mount will fail when one of the following occurs:

- | | |
|-----------|---|
| [NODEV] | The caller is not the super-user. |
| [NODEV] | <i>Special</i> does not exist. |
| [ENOTBLK] | <i>Special</i> is not a block device. |
| [ENXIO] | The major device number of <i>special</i> is out of range (this indicates no device driver exists for the associated hardware). |
| [EPERM] | The pathname contains a character with the high-order bit set. |
| [ENOTDIR] | A component of the path prefix in <i>name</i> is not a directory. |
| [EROFS] | <i>Name</i> resides on a read-only file system. |
| [EBUSY] | <i>Name</i> is not a directory, or another process currently holds a reference to it. |
| [EBUSY] | No space remains in the mount table. |
| [EBUSY] | The super block for the file system had a bad magic number or an out of range block size. |

- [EBUSY] Not enough **memory** was available to read the cylinder group information for the file system.
- [EBUSY] An i/o error occurred while reading the super block or cylinder group information.
- Umount* may fail with one of the following errors:
- [NODEV] The caller is not the super-user.
- [NODEV] *Special* does not exist.
- [ENOTBLK] *Special* is not a block device.
- [ENXIO] The major device number of *special* is out of range (this indicates no device driver exists for the associated hardware).
- [EINVAL] The requested device is not in the mount table.
- [EBUSY] A process is holding a reference to a file located on the file system.

SEE ALSO

mount(8), umount(8)

BUGS

The error codes are in a state of **disarray**; too many errors appear to the caller as one value.

NAME

`msgctl` – message control operations

SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/msg.h>

int msgctl (msqid, cmd, buf)
int msqid, cmd;
struct msqid_ds *buf;
```

DESCRIPTION

Msgctl provides a variety of message control operations as specified by *cmd*. The following *cmds* are available:

- IPC_STAT** Place the current value of each member of the data structure associated with *msqid* into the structure pointed to by *buf*. The contents of this structure are defined in *intro(2)*. {READ}
- IPC_SET** Set the value of the following members of the data structure associated with *msqid* to the corresponding value found in the structure pointed to by *buf*:
- `msg_perm.uid`
 - `msg_perm.gid`
 - `msg_perm.mode` /* only low 9 bits */
 - `msg_qbytes`
- This *cmd* can only be executed by a process that has an effective user ID equal to either that of super user or to the value of `msg_perm.uid` in the data structure associated with *msqid*. Only super user can raise the value of `msg_qbytes`.
- IPC_RMID** Remove the message queue identifier specified by *msqid* from the system and destroy the message queue and data structure associated with it. This *cmd* can only be executed by a process that has an effective user ID equal to either that of super user or to the value of `msg_perm.uid` in the data structure associated with *msqid*.

Msgctl will fail if one or more of the following are true:

- [EINVAL] *Msqid* is not a valid message queue identifier.
- [EINVAL] *Cmd* is not a valid command.
- [EACCES] *Cmd* is equal to **IPC_STAT** and {READ} operation permission is denied to the calling process (see *intro(2)*).
- [EPERM] *Cmd* is equal to **IPC_RMID** or **IPC_SET**. The effective user ID of the calling process is not equal to that of super user and it is not equal to the value of `msg_perm.uid` in the data structure associated with *msqid*.
- [EPERM] *Cmd* is equal to **IPC_SET**, an attempt is being made to increase to the value of `msg_qbytes`, and the effective user ID of the calling process is not equal to that of super user.
- [EFAULT] *Buf* points to an illegal address.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

SEE ALSO

intro(2), msgget(2), msgop(2).

NAME

`msgget` – get message queue

SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/msg.h>
int msgget (key, msgflg)
key_t key;
int msgflg;
```

DESCRIPTION

Msgget returns the message queue identifier associated with *key*.

A message queue identifier and associated message queue and data structure (see *intro(2)*) are created for *key* if one of the following are true:

10 *Key* is equal to `IPC_PRIVATE`.

Key does not already have a message queue identifier associated with it, and $(msgflg \& IPC_CREAT)$ is “true”.

Upon creation, the data structure associated with the new message queue identifier is initialized as follows:

`Msg_perm.cuid`, `msg_perm.uid`, `msg_perm.cgid`, and `msg_perm.gid` are set equal to the effective user ID and effective group ID, respectively, of the calling process.

The low-order 9 bits of `msg_perm.mode` are set equal to the low-order 9 bits of *msgflg*.

`Msg_qnum`, `msg_lspid`, `msg_lrpid`, `msg_stime`, and `msg_rtime` are set equal to 0.

`Msg_ctime` is set equal to the current time.

`Msg_qbytes` is set equal to the system limit.

Msgget will fail if one or more of the following are true:

- | | |
|----------|--|
| [EACCES] | A message queue identifier exists for <i>key</i> , but operation permission (see <i>intro(2)</i>) as specified by the low-order 9 bits of <i>msgflg</i> would not be granted. |
| [ENOENT] | A message queue identifier does not exist for <i>key</i> and $(msgflg \& IPC_CREAT)$ is “false”. |
| [ENOSPC] | A message queue identifier is to be created but the system-imposed limit on the maximum number of allowed message queue identifiers system wide would be exceeded. |
| [EEXIST] | A message queue identifier exists for <i>key</i> but $((msgflg \& IPC_CREAT) \& (msgflg \& IPC_EXCL))$ is “true”. |

RETURN VALUE

Upon successful completion, a non-negative integer, namely a message queue identifier, is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

SEE ALSO

intro(2), *msgctl(2)*, *msgop(2)*.

NAME

msgop - message operations

SYNOPSIS

```

#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/msg.h>

int msgsnd (msqid, msgp, msgsz, msgflg)
int msqid;
struct msgbuf *msgp;
int msgsz, msgflg;

int msgrcv (msqid, msgp, msgsz, msgtyp, msgflg)
int msqid;
struct msgbuf *msgp;
int msgsz;
long msgtyp;
int msgflg;

```

DESCRIPTION

Msgsnd is used to send a message to the queue associated with the message queue identifier specified by *msqid*. {WRITE} *Msgp* points to a structure containing the message. This structure is composed of the following members:

```

long    mtype;    /* message type */
char    mtext[]; /* message text */

```

Mtype is a positive integer that can be used by the receiving process for message selection (see *msgrcv* below). *Mtext* is any text of length *msgsz* bytes. *Msgsz* can range from 0 to a system-imposed maximum.

Msgflg specifies the action to be taken if one or more of the following are true:

The number of bytes already on the queue is equal to **msg_qbytes** (see *intro(2)*).

The total number of messages on all queues system-wide is equal to the system-imposed limit.

These actions are as follows:

If (*msgflg* & IPC_NOWAIT) is "true", the message will not be sent and the calling process will return immediately.

If (*msgflg* & IPC_NOWAIT) is "false", the calling process will suspend execution until one of the following occurs:

The condition responsible for the suspension no longer exists, in which case the message is sent.

Msqid is removed from the system (see *msgctl(2)*). When this occurs, *errno* is set equal to EIDRM, and a value of -1 is returned.

The calling process receives a signal that is to be caught. In this case the message is not sent and the calling process resumes execution in the manner prescribed in *signal(2)*.

Msgsnd will fail and no message will be sent if one or more of the following are true:

[EINVAL]	<i>Msqid</i> is not a valid message queue identifier.
[EACCES]	Operation permission is denied to the calling process (see <i>intro(2)</i>).

- [EINVAL] *Mtype* is less than 1.
- [EAGAIN] The message cannot be sent for one of the reasons cited above and (*msgflg* & *IPC_NOWAIT*) is "true".
- [EINVAL] *Msgsz* is less than zero or greater than the system-imposed limit.
- [EFAULT] *Msgp* points to an illegal address.

Upon successful completion, the following actions are taken with respect to the data structure associated with *msqid* (see intro (2)).

Msg_qnum is incremented by 1.

Msg_lspid is set equal to the process ID of the calling process.

Msg_stime is set equal to the current time.

Msgrcv reads a message from the queue associated with the message queue identifier specified by *msqid* and places it in the structure pointed to by *msgp*. {READ} This structure is composed of the following members:

```

long    mtype;        /* message type */
char    mtext[];     /* message text */

```

Mtype is the received message's type as specified by the sending process. *Mtext* is the text of the message. *Msgsz* specifies the size in bytes of *mtext*. The received message is truncated to *msgsz* bytes if it is larger than *msgsz* and (*msgflg* & *MSG_NOERROR*) is "true". The truncated part of the message is lost and no indication of the truncation is given to the calling process.

Msgtyp specifies the type of message requested as follows:

If *msgtyp* is equal to 0, the first message on the queue is received.

If *msgtyp* is greater than 0, the first message of type *msgtyp* is received.

If *msgtyp* is less than 0, the first message of the lowest type that is less than or equal to the absolute value of *msgtyp* is received.

Msgflg specifies the action to be taken if a message of the desired type is not on the queue. These are as follows:

If (*msgflg* & *IPC_NOWAIT*) is "true", the calling process will return immediately with a return value of -1 and *errno* set to *ENOMSG*.

If (*msgflg* & *IPC_NOWAIT*) is "false", the calling process will suspend execution until one of the following occurs:

A message of the desired type is placed on the queue.

Msqid is removed from the system. When this occurs, *errno* is set equal to *EIDRM*, and a value of -1 is returned.

The calling process receives a signal that is to be caught. In this case a message is not received and the calling process resumes execution in the manner prescribed in *signal(2)*.

Msgrcv will fail and no message will be received if one or more of the following are true:

- [EINVAL] *Msqid* is not a valid message queue identifier.
- [EACCES] Operation permission is denied to the calling process.
- [EINVAL] *Msgsz* is less than 0.
- [E2BIG] *Mtext* is greater than *msgsz* and (*msgflg* & *MSG_NOERROR*) is "false".

[ENOMSG] The queue does not contain a message of the desired type and (*msgtyp* & *IPC_NOWAIT*) is "true".

[EFAULT] *Msgp* points to an illegal address.

Upon successful completion, the following actions are taken with respect to the data structure associated with *msqid* (see intro (2)).

Msg_qnum is decremented by 1.

Msg_lrpid is set equal to the process ID of the calling process.

Msg_rtime is set equal to the current time.

RETURN VALUES

If *msgsnd* or *msgrcv* return due to the receipt of a signal, a value of -1 is returned to the calling process and *errno* is set to EINTR. If they return due to removal of *msqid* from the system, a value of -1 is returned and *errno* is set to EIDRM.

Upon successful completion, the return value is as follows:

Msgsnd returns a value of 0.

Msgrcv returns a value equal to the number of bytes actually placed into *mtext*.

Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

SEE ALSO

intro(2), msgctl(2), msgget(2), signal(2).

NAME

`open` – open a file for reading or writing, or create a new file

SYNOPSIS

```
#include <sys/file.h>
open(path, flags, mode)
char *path;
int flags, mode;
```

DESCRIPTION

`Open` opens the file *path* for reading and/or writing, as specified by the *flags* argument and returns a descriptor for that file. The *flags* argument may indicate the file is to be created if it does not already exist (by specifying the `O_CREAT` flag), in which case the file is created with mode *mode* as described in `chmod(2)` and modified by the process' `umask` value (see `umask(2)`).

Path is the address of a string of ASCII characters representing a path name, terminated by a null character. The flags specified are formed by *or*'ing the following values

```
O_RDONLY  open for reading only
O_WRONLY  open for writing only
O_RDWR    open for reading and writing
O_NDELAY  do not block on open
O_APPEND  append on each write
O_CREAT   create file if it does not exist
O_TRUNC   truncate size to 0
O_EXCL    error if create and file exists
```

Opening a file with `O_APPEND` set causes each write on the file to be appended to the end. If `O_TRUNC` is specified and the file exists, the file is truncated to zero length. If `O_EXCL` is set with `O_CREAT`, then if the file already exists, the open returns an error. This can be used to implement a simple exclusive access locking mechanism. If the `O_NDELAY` flag is specified and the open call would result in the process being blocked for some reason (e.g. waiting for carrier on a dialup line), the open returns immediately. The first time the process attempts to perform i/o on the open file it will block (not currently implemented).

Upon successful completion a non-negative integer termed a file descriptor is returned. The file pointer used to mark the current position within the file is set to the beginning of the file.

The new descriptor is set to remain open across `execve` system calls; see `close(2)`.

No process may have more than `{OPEN_MAX}` file descriptors open simultaneously.

ERRORS

The named file is opened unless one or more of the following are true:

[EPERM]	The pathname contains a character with the high-order bit set.
[ENOTDIR]	A component of the path prefix is not a directory.
[ENOENT]	<code>O_CREAT</code> is not set and the named file does not exist.
[EACCES]	A component of the path prefix denies search permission.
[EACCES]	The required permissions (for reading and/or writing) are denied for the named flag.
[EISDIR]	The named file is a directory, and the arguments specify it is to be opened for writing.

- [EROFS] The named file resides on a read-only file system, and the file is to be modified.
- [EMFILE] {OPEN_MAX} file descriptors are currently open.
- [ENXIO] The named file is a character special or block special file, and the device associated with this special file does not exist.
- [ETXTBSY] The file is a pure procedure (shared text) file that is being executed and the *open* call requests write access.
- [EFAULT] *Path* points outside the process's allocated address space.
- [ELOOP] Too many symbolic links were encountered in translating the pathname.
- [EEXIST] O_EXCL was specified and the file exists.
- [ENXIO] The O_NDELAY flag is given, and the file is a communications device on which there is no carrier present.
- [EOPNOTSUPP] An attempt was made to open a socket (not currently implemented).

SEE ALSO

chmod(2), close(2), dup(2), lseek(2), read(2), write(2), umask(2)

NAME

pipe – create an interprocess communication channel

SYNOPSIS

```
pipe(fildes)
int fildes[2];
```

DESCRIPTION

The *pipe* system call creates an I/O mechanism called a pipe. The file descriptors returned can be used in read and write operations. When the pipe is written using the descriptor *fildes*[1] up to 4096 bytes of data are buffered before the writing process is suspended. A read using the descriptor *fildes*[0] will pick up the data.

It is assumed that after the pipe has been set up, two (or more) cooperating processes (created by subsequent *fork* calls) will pass data through the pipe with *read* and *write* calls.

The shell has a syntax to set up a linear array of processes connected by pipes.

Read calls on an empty pipe (no buffered data) with only one end (all write file descriptors closed) returns an end-of-file.

Pipes are really a special case of the *socketpair*(2) call and, in fact, are implemented as such in the system.

A signal is generated if a write on a pipe with only one end is attempted.

RETURN VALUE

The function value zero is returned if the pipe was created; -1 if an error occurred.

ERRORS

The *pipe* call will fail if:

- [EMFILE] Too many descriptors are active.
- [EFAULT] The *fildes* buffer is in an invalid area of the process's address space.

SEE ALSO

sh(1), read(2), write(2), fork(2), socketpair(2)

BUGS

Should more than 4096 bytes be necessary in any pipe among a loop of processes, deadlock will occur.

NAME

profil - execution time profile

SYNOPSIS

```
profil(buff, bufsiz, offset, scale)
char *buff;
int bufsiz, offset, scale;
```

DESCRIPTION

Buff points to an area of core whose length (in bytes) is given by *bufsiz*. After this call, the user's program counter (*pc*) is examined each clock tick (20 milliseconds); *offset* is subtracted from it, and the result multiplied by *scale*. If the resulting number corresponds to a word inside *buff*, that word is incremented.

The scale is interpreted as an unsigned, fixed-point fraction with binary point at the left: 0x10000 gives a 1-1 mapping of *pc*'s to words in *buff*; 0x8000 maps each pair of instruction words together. 0x2 maps all instructions onto the beginning of *buff* (producing a non-interrupting core clock).

Profiling is turned off by giving a *scale* of 0 or 1. It is rendered ineffective by giving a *bufsiz* of 0. Profiling is turned off when an *execve* is executed, but remains on in child and parent both after a *fork*. Profiling is turned off if an update in *buff* would cause a memory fault.

RETURN VALUE

A 0, indicating success, is always returned.

SEE ALSO

gprof(1), setitimer(2), monitor(3)

NAME

`ptrace` - process trace

SYNOPSIS

```
#include <signal.h>
ptrace(request, pid, addr, data)
int request, pid, *addr, data;
```

DESCRIPTION

Ptrace provides a means by which a parent process may control the execution of a child process, and examine and change its core image. Its primary use is for the implementation of breakpoint debugging. There are four arguments whose interpretation depends on a *request* argument. Generally, *pid* is the process ID of the traced process, which must be a child (no more distant descendant) of the tracing process. A process being traced behaves normally until it encounters some signal whether internally generated like "illegal instruction" or externally generated like "interrupt". See *sigvec(2)* for the list. Then the traced process enters a stopped state and its parent is notified via *wait(2)*. When the child is in the stopped state, its core image can be examined and modified using *ptrace*. If desired, another *ptrace* request can then cause the child either to terminate or to continue, possibly ignoring the signal.

The value of the *request* argument determines the precise action of the call:

- 0 This request is the only one used by the child process; it declares that the process is to be traced by its parent. All the other arguments are ignored. Peculiar results will ensue if the parent does not expect to trace the child.
- 1,2 The word in the child process's address space at *addr* is returned. If I and D space are separated (e.g. historically on a pdp-11), request 1 indicates I space, 2 D space. *Addr* must be even. The child must be stopped. The input *data* is ignored.
- 3 The word of the system's per-process data area corresponding to *addr* is returned. *Addr* must be even and less than 512. This space contains the registers and other information about the process; its layout corresponds to the *user* structure in the system.
- 4,5 The given *data* is written at the word in the process's address space corresponding to *addr*, which must be even. No useful value is returned. If I and D space are separated, request 4 indicates I space, 5 D space. Attempts to write in pure procedure fail if another process is executing the same file.
- 6 The process's system data is written, as it is read with request 3. Only a few locations can be written in this way: the general registers, the floating point status and registers, and certain bits of the processor status word.
- 7 The *data* argument is taken as a signal number and the child's execution continues at location *addr* as if it had incurred that signal. Normally the signal number will be either 0 to indicate that the signal that caused the stop should be ignored, or that value fetched out of the process's image indicating which signal caused the stop. If *addr* is (int *)1 then execution continues from where it stopped.
- 8 The traced process terminates.
- 9 Execution continues as in request 7; however, as soon as possible after execution of at least one instruction, execution stops again. The signal number from the stop is SIGTRAP. (On ICON systems the trace-bit is set and just one instruction is executed.) This is part of the mechanism for implementing breakpoints.

As indicated, these calls (except for request 0) can be used only when the subject process has stopped. The *wait* call is used to determine when a process stops; in such a case the "termination" status returned by *wait* has the value 0177 to indicate stoppage rather than genuine termination.

To forestall possible fraud, *ptrace* inhibits the set-user-id and set-group-id facilities on subsequent *execve(2)* calls. If a traced process calls *execve*, it will stop before executing the first instruction of the new image showing signal SIGTRAP.

On an ICON system, "word" also means a 32-bit integer, but the "even" restriction does not apply.

RETURN VALUE

A 0 value is returned if the call succeeds. If the call fails then a -1 is returned and the global variable *errno* is set to indicate the error.

ERRORS

[EINVAL]	The request code is invalid.
[EINVAL]	The specified process does not exist.
[EINVAL]	The given signal number is invalid.
[EFAULT]	The specified address is out of bounds.
[EPERM]	The specified process cannot be traced.

SEE ALSO

wait(2), *sigvec(2)*, *adb(1)*

BUGS

Ptrace is unique and arcane; it should be replaced with a special file which can be opened and read and written. The control functions could then be implemented with *ioctl(2)* calls on this file. This would be simpler to understand and have much higher performance.

The request 0 call should be able to specify signals which are to be treated normally and not cause a stop. In this way, for example, programs with simulated floating point (which use "illegal instruction" signals at a very high rate) could be efficiently debugged.

The error indication, -1, is a legitimate function value; *errno*, see *intro(2)*, can be used to disambiguate.

It should be possible to stop a process on occurrence of a system call; in this way a completely controlled environment could be provided.

NAME

quota – manipulate disk quotas

SYNOPSIS

```
#include <sys/quota.h>
quota(cmd, uid, arg, addr)
int cmd, uid, arg;
caddr_t addr;
```

DESCRIPTION

The *quota* call manipulates disk quotas for file systems which have had quotas enabled with *setquota(2)*. The *cmd* parameter indicates a command to be applied to the user ID *uid*. *Arg* is a command specific argument and *addr* is the address of an optional, command specific, data structure which is copied in or out of the system. The interpretation of *arg* and *addr* is given with each command below.

Q_SETDLIM

Set disc quota limits and current usage for the user with ID *uid*. *Arg* is a major-minor device indicating a particular file system. *Addr* is a pointer to a struct *dqblk* structure (defined in *<sys/quota.h>*). This call is restricted to the super-user.

Q_GETDLIM

Get disc quota limits and current usage for the user with ID *uid*. The remaining parameters are as for *Q_SETDLIM*.

Q_SETDUSE

Set disc usage limits for the user with ID *uid*. *Arg* is a major-minor device indicating a particular file system. *Addr* is a pointer to a struct *dqusage* structure (defined in *<sys/quota.h>*). This call is restricted to the super-user.

Q_SYNC

Update the on-disc copy of quota usages. The *uid*, *arg*, and *addr* parameters are ignored.

Q_SETUID

Change the calling process's quota limits to those of the user with ID *uid*. The *arg* and *addr* parameters are ignored. This call is restricted to the super-user.

Q_SETWARN

Alter the disc usage warning limits for the user with ID *uid*. *Arg* is a major-minor device indicating a particular file system. *Addr* is a pointer to a struct *dqwarn* structure (defined in *<sys/quota.h>*). This call is restricted to the super-user.

Q_DOWARN

Warn the user with user ID *uid* about excessive disc usage. This call causes the system to check its current disc usage information and print a message on the terminal of the caller for each file system on which the user is over quota. If the *arg* parameter is specified as *NODEV*, all file systems which have disc quotas will be checked. Otherwise, *arg* indicates a specific major-minor device to be checked. This call is restricted to the super-user.

RETURN VALUE

A successful call returns 0 and, possibly, more information specific to the *cmd* performed; when an error occurs, the value -1 is returned and *errno* is set to indicate the reason.

ERRORS

A *quota* call will fail when one of the following occurs:

- [EINVAL] *Cmd* is invalid.
- [ESRCH] No disc quota is found for the indicated user.
- [EPERM] The call is privileged and the caller was not the super-user.
- [EINVAL] The *arg* parameter is being interpreted as a major-minor device and it indicates an unmounted file system.
- [EFAULT] An invalid *addr* is supplied; the associated structure could not be copied in or out of the kernel.
- [EUSERS] The quota table is full.

SEE ALSO

setquota(2), quotaon(8), quotacheck(8)

BUGS

There should be some way to integrate this call with the resource limit interface provided by *setrlimit(2)* and *getrlimit(2)*.

The Australian spelling of *disk* is used throughout the quota facilities in honor of the implementors.

NAME

`read`, `readv` – read input

SYNOPSIS

```
cc = read(d, buf, nbytes)
```

```
int cc, d;
```

```
char *buf;
```

```
int nbytes;
```

```
#include <sys/types.h>
```

```
#include <sys/uio.h>
```

```
cc = readv(d, iov, iovcnt)
```

```
int cc, d;
```

```
struct iovec *iov;
```

```
int iovcnt;
```

DESCRIPTION

Read attempts to read *nbytes* of data from the object referenced by the descriptor *d* into the buffer pointed to by *buf*. *Readv* performs the same action, but scatters the input data into the *iovcnt* buffers specified by the members of the *iovec* array: *iov*[0], *iov*[1], ..., *iov*[*iovcnt* - 1].

For *readv*, the *iovec* structure is defined as

```
struct iovec {
    caddr_t    iov_base;
    int       iov_len;
};
```

Each *iovec* entry specifies the base address and length of an area in memory where data should be placed. *Readv* will always fill an area completely before proceeding to the next.

On objects capable of seeking, the *read* starts at a position given by the pointer associated with *d*, see *lseek*(2). Upon return from *read*, the pointer is incremented by the number of bytes actually read.

Objects that are not capable of seeking always read from the current position. The value of the pointer associated with such a object is undefined.

Upon successful completion, *read* and *readv* return the number of bytes actually read and placed in the buffer. The system guarantees to read the number of bytes requested if the descriptor references a file which has that many bytes left before the end-of-file, but in no other cases.

If the returned value is 0, then end-of-file has been reached.

RETURN VALUE

If successful, the number of bytes actually read is returned. Otherwise, a -1 is returned and the global variable *errno* is set to indicate the error.

ERRORS

Read and *readv* will fail if one or more of the following are true:

[EBADF] *Fildes* is not a valid file descriptor open for reading.

[EFAULT] *Buf* points outside the allocated address space.

[EINTR] A read from a slow device was interrupted before any data arrived by the delivery of a signal.

In addition, *readv* may return one of the following errors:

[EINVAL] *Iovcnt* was less than or equal to 0, or greater than 16.

[EINVAL] One of the *iov_len* values in the *iov* array was negative.

[EINVAL] The sum of the *iov_len* values in the *iov* array overflowed a 32-bit integer.

SEE ALSO

dup(2), *open(2)*, *pipe(2)*, *socket(2)*, *socketpair(2)*

NAME

readlink – read value of a symbolic link

SYNOPSIS

```
cc = readlink(path, buf, bufsiz)
int cc;
char *path, *buf;
int bufsiz;
```

DESCRIPTION

Readlink places the contents of the symbolic link *name* in the buffer *buf* which has size *bufsiz*. The contents of the link are not null terminated when returned.

RETURN VALUE

The call returns the count of characters placed in the buffer if it succeeds, or a -1 if an error occurs, placing the error code in the global variable *errno*.

ERRORS

Readlink will fail and the file mode will be unchanged if:

- | | |
|-----------|---|
| [EPERM] | The <i>path</i> argument contained a byte with the high-order bit set. |
| [ENOENT] | The pathname was too long. |
| [ENOTDIR] | A component of the path prefix is not a directory. |
| [ENOENT] | The named file does not exist. |
| [ENXIO] | The named file is not a symbolic link. |
| [EACCES] | Search permission is denied on a component of the path prefix. |
| [EPERM] | The effective user ID does not match the owner of the file and the effective user ID is not the super-user. |
| [EINVAL] | The named file is not a symbolic link. |
| [EFAULT] | <i>Buf</i> extends outside the process's allocated address space. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |

SEE ALSO

stat(2), lstat(2), symlink(2)

NAME

reboot – reboot system or halt processor

SYNOPSIS

```
#include <sys/reboot.h>
reboot(howto)
int howto;
```

DESCRIPTION

Reboot reboots the system, and is invoked automatically in the event of unrecoverable system failures. *Howto* is a mask of options passed to the bootstrap program. The system call interface permits only `RB_HALT` or `RB_AUTOBOOT` to be passed to the reboot program. When none of these options (e.g. `RB_AUTOBOOT`) is given, the system is rebooted from file “*vmunix*” in the root file system of unit 0 of a disk chosen in a processor specific way.

The bits of *howto* are:

RB_HALT

the processor is simply halted; no reboot takes place. `RB_HALT` should be used with caution. This switch is not available from the system call interface.

RB_AUTOBOOT

All media are synced, and the machine reboots. At reboot time you are offered the option to continue with a normal reboot, or to interact with the system loader(8).

Only the super-user may *reboot* a machine.

RETURN VALUES

If successful, this call never returns. Otherwise, a -1 is returned and an error is returned in the global variable *errno*.

ERRORS

[EPERM] The caller is not the super-user.

SEE ALSO

crash(8), halt(8), init(8), reboot(8)

NAME

recv, *recvfrom*, *recvmsg* – receive a message from a socket

SYNOPSIS

```
#include <sys/types.h>
#include <sys/socket.h>

cc = recv(s, buf, len, flags)
int cc, s;
char *buf;
int len, flags;

cc = recvfrom(s, buf, len, flags, from, fromlen)
int cc, s;
char *buf;
int len, flags;
struct sockaddr *from;
int *fromlen;

cc = recvmsg(s, msg, flags)
int cc, s;
struct msghdr msg[];
int flags;
```

DESCRIPTION

Recv, *recvfrom*, and *recvmsg* are used to receive messages from a socket.

The *recv* call may be used only on a *connected* socket (see *connect(2)*), while *recvfrom* and *recvmsg* may be used to receive data on a socket whether it is in a connected state or not.

If *from* is non-zero, the source address of the message is filled in. *Fromlen* is a value-result parameter, initialized to the size of the buffer associated with *from*, and modified on return to indicate the actual size of the address stored there. The length of the message is returned in *cc*. If a message is too long to fit in the supplied buffer, excess bytes may be discarded depending on the type of socket the message is received from; see *socket(2)*.

If no messages are available at the socket, the receive call waits for a message to arrive, unless the socket is nonblocking (see *ioctl(2)*) in which case a *cc* of -1 is returned with the external variable *errno* set to *EWOULDBLOCK*.

The *select(2)* call may be used to determine when more data arrives.

The *flags* argument to a send call is formed by *or*'ing one or more of the values,

```
#define MSG_PEEK 0x1 /* peek at incoming message */
#define MSG_OOB 0x2 /* process out-of-band data */
```

The *recvmsg* call uses a *msghdr* structure to minimize the number of directly supplied parameters. This structure has the following form, as defined in *<sys/socket.h>*:

```
struct msghdr {
    caddr_t      msg_name;          /* optional address */
    int         msg_namelen;       /* size of address */
    struct iovec *msg_iov;         /* scatter/gather array */
    int         msg_iovlen;        /* # elements in msg_iov */
    caddr_t      msg_accrightrights; /* access rights sent/received */
    int         msg_accrightrightslen;
};
```

Here *msg_name* and *msg_namelen* specify the destination address if the socket is unconnected; *msg_name* may be given as a null pointer if no names are desired or required. The *msg_iov* and *msg_iovlen* describe the scatter gather locations, as described in *read(2)*. Access rights to be sent along with the message are specified in *msg_accrights*, which has length *msg_accrightslen*.

RETURN VALUE

These calls return the number of bytes received, or -1 if an error occurred.

ERRORS

The calls fail if:

[EBADF]	The argument <i>s</i> is an invalid descriptor.
[ENOTSOCK]	The argument <i>s</i> is not a socket.
[EWOULDBLOCK]	The socket is marked non-blocking and the receive operation would block.
[EINTR]	The receive was interrupted by delivery of a signal before any data was available for the receive.
[EFAULT]	The data was specified to be received into a non-existent or protected part of the process address space.

SEE ALSO

read(2), *send(2)*, *socket(2)*

NAME

rename – change the name of a file

SYNOPSIS

```
rename(from, to)
char *from, *to;
```

DESCRIPTION

Rename causes the link named *from* to be renamed as *to*. If *to* exists, then it is first removed. Both *from* and *to* must be of the same type (that is, both directories or both non-directories), and must reside on the same file system.

Rename guarantees that an instance of *to* will always exist, even if the system should crash in the middle of the operation.

CAVEAT

The system can deadlock if a loop in the file system graph is present. This loop takes the form of an entry in directory “a”, say “a/foo”, being a hard link to directory “b”, and an entry in directory “b”, say “b/bar”, being a hard link to directory “a”. When such a loop exists and two separate processes attempt to perform “rename a/foo b/bar” and “rename b/bar a/foo”, respectively, the system may deadlock attempting to lock both directories for modification. Hard links to directories should be replaced by symbolic links by the system administrator.

RETURN VALUE

A 0 value is returned if the operation succeeds, otherwise *rename* returns -1 and the global variable *errno* indicates the reason for the failure.

ERRORS

Rename will fail and neither of the argument files will be affected if any of the following are true:

[ENOTDIR]	A component of either path prefix is not a directory.
[ENOENT]	A component of either path prefix does not exist.
[EACCES]	A component of either path prefix denies search permission.
[ENOENT]	The file named by <i>from</i> does not exist.
[EPERM]	The file named by <i>from</i> is a directory and the effective user ID is not super-user.
[EXDEV]	The link named by <i>to</i> and the file named by <i>from</i> are on different logical devices (file systems). Note that this error code will not be returned if the implementation permits cross-device links.
[EACCES]	The requested link requires writing in a directory with a mode that denies write permission.
[EROFS]	The requested link requires writing in a directory on a read-only file system.
[EFAULT]	<i>Path</i> points outside the process’s allocated address space.
[EINVAL]	<i>From</i> is a parent directory of <i>to</i> .

SEE ALSO

open(2)

NAME

`rmdir` – remove a directory file

SYNOPSIS

```
rmdir(path)
char *path;
```

DESCRIPTION

Rmdir removes a directory file whose name is given by *path*. The directory must not have any entries other than “.” and “..”.

RETURN VALUE

A 0 is returned if the remove succeeds; otherwise a -1 is returned and an error code is stored in the global location *errno*.

ERRORS

The named file is removed unless one or more of the following are true:

- [ENOTEMPTY] The named directory contains files other than “.” and “..” in it.
- [EPERM] The pathname contains a character with the high-order bit set.
- [ENOENT] The pathname was too long.
- [ENOTDIR] A component of the path prefix is not a directory.
- [ENOENT] The named file does not exist.
- [EACCES] A component of the path prefix denies search permission.
- [EACCES] Write permission is denied on the directory containing the link to be removed.
- [EBUSY] The directory to be removed is the mount point for a mounted file system.
- [EROFS] The directory entry to be removed resides on a read-only file system.
- [EFAULT] *Path* points outside the process’s allocated address space.
- [ELOOP] Too many symbolic links were encountered in translating the pathname.

SEE ALSO

`mkdir(2)`, `unlink(2)`

NAME

select – synchronous i/o multiplexing

SYNOPSIS

```
#include <sys/time.h>
nfound = select(nfds, readfds, writefds, exceptfds, timeout)
int nfound, nfds, *readfds, *writefds, *exceptfds;
struct timeval *timeout;
```

DESCRIPTION

Select examines the i/o descriptors specified by the bit masks *readfds*, *writefds*, and *exceptfds* to see if they are ready for reading, writing, or have an exceptional condition pending, respectively. File descriptor *f* is represented by the bit “1<<*f*” in the mask. *Nfds* descriptors are checked, i.e. the bits from 0 through *nfds*-1 in the masks are examined. *Select* returns, in place, a mask of those descriptors which are ready. The total number of ready descriptors is returned in *nfound*.

If *timeout* is a non-zero pointer, it specifies a maximum interval to wait for the selection to complete. If *timeout* is a zero pointer, the *select* blocks indefinitely. To affect a poll, the *timeout* argument should be non-zero, pointing to a zero valued *timeval* structure.

Any of *readfds*, *writefds*, and *exceptfds* may be given as 0 if no descriptors are of interest.

RETURN VALUE

Select returns the number of descriptors which are contained in the bit masks, or -1 if an error occurred. If the time limit expires then *select* returns 0.

ERRORS

An error return from *select* indicates:

[EBADF]	One of the bit masks specified an invalid descriptor.
[EINTR]	An signal was delivered before any of the selected for events occurred or the time limit expired.

SEE ALSO

accept(2), *connect(2)*, *read(2)*, *write(2)*, *recv(2)*, *send(2)*

BUGS

The descriptor masks are always modified on return, even if the call returns as the result of the timeout.

NAME

semctl – semaphore control operations

SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/sem.h>

int semctl (semid, semnum, cmd, arg)
int semid, cmd;
int semnum;
union semun {
    int val;
    struct semid_ds *buf;
    ushort *array;
} arg;
```

DESCRIPTION

Semctl provides a variety of semaphore control operations as specified by *cmd*.

The following *cmds* are executed with respect to the semaphore specified by *semid* and *semnum*:

GETVAL	Return the value of <i>semval</i> (see <i>intro(2)</i>). {READ}
SETVAL	Set the value of <i>semval</i> to <i>arg.val</i> . {ALTER} When this cmd is successfully executed, the <i>semadj</i> value corresponding to the specified semaphore in all processes is cleared.
GETPID	Return the value of <i>sempid</i> . {READ}
GETNCNT	Return the value of <i>semncnt</i> . {READ}
GETZCNT	Return the value of <i>semzcnt</i> . {READ}

The following *cmds* return and set, respectively, every *semval* in the set of semaphores.

GETALL	Place <i>semvals</i> into array pointed to by <i>arg.array</i> . {READ}
SETALL	Set <i>semvals</i> according to the array pointed to by <i>arg.array</i> . {ALTER} When this cmd is successfully executed the <i>semadj</i> values corresponding to each specified semaphore in all processes are cleared.

The following *cmds* are also available:

IPC_STAT	Place the current value of each member of the data structure associated with <i>semid</i> into the structure pointed to by <i>arg.buf</i> . The contents of this structure are defined in <i>intro(2)</i> . {READ}
IPC_SET	Set the value of the following members of the data structure associated with <i>semid</i> to the corresponding value found in the structure pointed to by <i>arg.buf</i> : sem_perm.uid sem_perm.gid sem_perm.mode /* only low 9 bits */

This cmd can only be executed by a process that has an effective user ID equal to either that of super-user or to the value of **sem_perm.uid** in the data structure associated with *semid*.

IPC_RMID Remove the semaphore identifier specified by *semid* from the system and destroy the set of semaphores and data structure associated with it. This cmd can only be executed by a process that has an effective user ID equal to either that of super-user or to the value of **sem_perm.uid** in the data structure associated with *semid*.

Semctl will fail if one or more of the following are true:

- [EINVAL] *Semid* is not a valid semaphore identifier.
- [EINVAL] *Semnum* is less than zero or greater than **sem_nsems**.
- [EINVAL] *Cmd* is not a valid command.
- [EACCES] Operation permission is denied to the calling process (see *intro(2)*).
- [ERANGE] *Cmd* is **SETVAL** or **SETALL** and the value to which *semval* is to be set is greater than the system imposed maximum.
- [EPERM] *Cmd* is equal to **IPC_RMID** or **IPC_SET** and the effective user ID of the calling process is not equal to that of super-user and it is not equal to the value of **sem_perm.uid** in the data structure associated with *semid*.
- [EFAULT] *Arg.buf* points to an illegal address.

RETURN VALUE

Upon successful completion, the value returned depends on *cmd* as follows:

- GETVAL** The value of *semval*.
- GETPID** The value of *sempid*.
- GETNCNT** The value of *semncnt*.
- GETZCNT** The value of *semzcnt*.
- All others A value of 0.

Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

SEE ALSO

intro(2), *semget(2)*, *semop(2)*.

NAME

`semget` – get set of semaphores

SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/sem.h>
int semget (key, nsems, semflg)
key_t key;
int nsems, semflg;
```

DESCRIPTION

Semget returns the semaphore identifier associated with *key*.

A semaphore identifier and associated data structure and set containing *nsems* semaphores (see *intro(2)*) are created for *key* if one of the following are true:

Key is equal to `IPC_PRIVATE`.

Key does not already have a semaphore identifier associated with it, and (*semflg* & `IPC_CREAT`) is "true".

Upon creation, the data structure associated with the new semaphore identifier is initialized as follows:

`Sem_perm.cuid`, `Sem_perm.uid`, `Sem_perm.cgid`, and `Sem_perm.gid` are set equal to the effective user ID and effective group ID, respectively, of the calling process.

The low-order 9 bits of `Sem_perm.mode` are set equal to the low-order 9 bits of *semflg*.

`Sem_nsems` is set equal to the value of *nsems*.

`Sem_otime` is set equal to 0 and `Sem_ctime` is set equal to the current time.

Semget will fail if one or more of the following are true:

- [EINVAL] *Nsems* is either less than or equal to zero or greater than the system-imposed limit.
- [EACCES] A semaphore identifier exists for *key*, but operation permission (see *intro(2)*) as specified by the low-order 9 bits of *semflg* would not be granted.
- [EINVAL] A semaphore identifier exists for *key*, but the number of semaphores in the set associated with it is less than *nsems* and *nsems* is not equal to zero.
- [ENOENT] A semaphore identifier does not exist for *key* and (*semflg* & `IPC_CREAT`) is "false".
- [ENOSPC] A semaphore identifier is to be created but the system-imposed limit on the maximum number of allowed semaphore identifiers system wide would be exceeded.
- [ENOSPC] A semaphore identifier is to be created but the system-imposed limit on the maximum number of allowed semaphores system wide would be exceeded.
- [EEXIST] A semaphore identifier exists for *key* but ((*semflg* & `IPC_CREAT`) and (*semflg* & `IPC_EXCL`)) is "true".

RETURN VALUE

Upon successful completion, a non-negative integer, namely a semaphore identifier, is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

SEE ALSO

intro(2), semctl(2), semop(2).

NAME

semop - semaphore operations

SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/sem.h>

int semop (semid, sops, nsops)
int semid;
struct sembuf **sops;
int nsops;
```

DESCRIPTION

Semop is used to automatically perform an array of semaphore operations on the set of semaphores associated with the semaphore identifier specified by *semid*. *Sops* is a pointer to the array of semaphore-operation structures. *Nsops* is the number of such structures in the array. The contents of each structure includes the following members:

```
short  sem_num; /* semaphore number */
short  sem_op;  /* semaphore operation */
short  sem_flg; /* operation flags */
```

Each semaphore operation specified by *sem_op* is performed on the corresponding semaphore specified by *semid* and *sem_num*.

Sem_op specifies one of three semaphore operations as follows:

If *sem_op* is a negative integer, one of the following will occur: {ALTER}

If *semval* (see *intro(2)*) is greater than or equal to the absolute value of *sem_op*, the absolute value of *sem_op* is subtracted from *semval*. Also, if (*sem_flg* & SEM_UNDO) is "true", the absolute value of *sem_op* is added to the calling process's *semadj* value (see *exit(2)*) for the specified semaphore.

If *semval* is less than the absolute value of *sem_op* and (*sem_flg* & IPC_NOWAIT) is "true", *semop* will return immediately.

If *semval* is less than the absolute value of *sem_op* and (*sem_flg* & IPC_NOWAIT) is "false", *semop* will increment the *semncnt* associated with the specified semaphore and suspend execution of the calling process until one of the following conditions occur.

Semval becomes greater than or equal to the absolute value of *sem_op*. When this occurs, the value of *semncnt* associated with the specified semaphore is decremented, the absolute value of *sem_op* is subtracted from *semval* and, if (*sem_flg* & SEM_UNDO) is "true", the absolute value of *sem_op* is added to the calling process's *semadj* value for the specified semaphore.

The *semid* for which the calling process is awaiting action is removed from the system (see *semctl(2)*). When this occurs, *errno* is set equal to EIDRM, and a value of -1 is returned.

The calling process receives a signal that is to be caught. When this occurs, the value of *semncnt* associated with the specified semaphore is decremented, and the calling process resumes execution in the manner prescribed in *signal(2)*.

If *sem_op* is a positive integer, the value of *sem_op* is added to *semval* and, if (*sem_flg* & SEM_UNDO) is "true", the value of *sem_op* is subtracted from the calling process's *semadj* value for the specified semaphore. {ALTER}

If *sem_op* is zero, one of the following will occur: {READ}

If *semval* is zero, *semop* will return immediately.

If *semval* is not equal to zero and (*sem_flg* & IPC_NOWAIT) is "true", *semop* will return immediately.

If *semval* is not equal to zero and (*sem_flg* & IPC_NOWAIT) is "false", *semop* will increment the *semzcnt* associated with the specified semaphore and suspend execution of the calling process until one of the following occurs:

semval becomes zero, at which time the value of *semzcnt* associated with the specified semaphore is decremented.

The *semid* for which the calling process is awaiting action is removed from the system. When this occurs, *errno* is set equal to EIDRM, and a value of -1 is returned.

The calling process receives a signal that is to be caught. When this occurs, the value of *semzcnt* associated with the specified semaphore is decremented, and the calling process resumes execution in the manner prescribed in *signal(2)*.

semop will fail if one or more of the following are true for any of the semaphore operations specified by *sops*:

- [EINVAL] *semid* is not a valid semaphore identifier.
- [EFBIG] *sem_num* is less than zero or greater than or equal to the number of semaphores in the set associated with *semid*.
- [E2BIG] *Nsops* is greater than the system-imposed maximum.
- [EACCES] Operation permission is denied to the calling process (see *intro(2)*).
- [EAGAIN] The operation would result in suspension of the calling process but (*sem_flg* & IPC_NOWAIT) is "true".
- [ENOSPC] The limit on the number of individual processes requesting an SEM_UNDO would be exceeded.
- [EINVAL] The number of individual semaphores for which the calling process requests a SEM_UNDO would exceed the limit.
- [ERANGE] An operation would cause a *semval* to overflow the system-imposed limit.
- [ERANGE] An operation would cause a *semadj* value to overflow the system-imposed limit.
- [EFAULT] *Sops* points to an illegal address.

Upon successful completion, the value of *sempid* for each semaphore specified in the array pointed to by *sops* is set equal to the process ID of the calling process.

RETURN VALUE

If *semop* returns due to the receipt of a signal, a value of -1 is returned to the calling process and *errno* is set to EINTR. If it returns due to the removal of a *semid* from the system, a value of -1 is returned and *errno* is set to EIDRM.

Upon successful completion, the value of `semval` at the time of the call for the last operation in the array pointed to by `sops` is returned. Otherwise, a value of `-1` is returned and `errno` is set to indicate the error.

SEE ALSO

`exec(2)`, `exit(2)`, `fork(2)`, `intro(2)`, `semctl(2)`, `semget(2)`.

NAME

`send`, `sendto`, `sendmsg` – send a message from a socket

SYNOPSIS

```
#include <sys/types.h>
#include <sys/socket.h>

cc = send(s, msg, len, flags)
int cc, s;
char *msg;
int len, flags;

cc = sendto(s, msg, len, flags, to, tolen)
int cc, s;
char *msg;
int len, flags;
struct sockaddr *to;
int tolen;

cc = sendmsg(s, msg, flags)
int cc, s;
struct msghdr msg[];
int flags;
```

DESCRIPTION

`Send`, `sendto`, and `sendmsg` are used to transmit a message to another socket. `Send` may be used only when the socket is in a *connected* state, while `sendto` and `sendmsg` may be used at any time.

The address of the target is given by `to` with `tolen` specifying its size. The length of the message is given by `len`. If the message is too long to pass atomically through the underlying protocol, then the error `EMSGSIZE` is returned, and the message is not transmitted.

No indication of failure to deliver is implicit in a `send`. Return values of `-1` indicate some locally detected errors.

If no messages space is available at the socket to hold the message to be transmitted, then `send` normally blocks, unless the socket has been placed in non-blocking i/o mode. The `select(2)` call may be used to determine when it is possible to send more data.

The `flags` parameter may be set to `SOF_OOB` to send “out-of-band” data on sockets which support this notion (e.g. `SOCK_STREAM`).

See `recv(2)` for a description of the `msghdr` structure.

RETURN VALUE

The call returns the number of characters sent, or `-1` if an error occurred.

ERRORS

[EBADF]	An invalid descriptor was specified.
[ENOTSOCK]	The argument <code>s</code> is not a socket.
[EFAULT]	An invalid user space address was specified for a parameter.

[EMSGSIZE]

The socket requires that message be sent atomically, and the size of the message to be sent made this impossible.

[EWOULDBLOCK]

The socket is marked non-blocking and the requested operation would block.

SEE ALSO

recv(2), socket(2)

NAME

setgroups – set group access list

SYNOPSIS

```
#include <sys/param.h>
setgroups(ngroups, gidset)
int ngroups, *gidset;
```

DESCRIPTION

Setgroups sets the group access list of the current user process according to the array *gidset*. The parameter *ngroups* indicates the number of entries in the array and must be no more than *NGRPS*, as defined in *<sys/param.h>*.

Only the super-user may set new groups.

RETURN VALUE

A 0 value is returned on success, -1 on error, with a error code stored in *errno*.

ERRORS

The *setgroups* call will fail if:

- [EPERM] The caller is not the super-user.
- [EFAULT] The address specified for *gidset* is outside the process address space.

SEE ALSO

getgroups(2), initgroups(3X)

NAME

setpgrp - set process group

SYNOPSIS

```
setpgrp(pid, pgrp)
int pid, pgrp;
```

DESCRIPTION

Setpgrp sets the process group of the specified process *pid* to the specified *pgrp*. If *pid* is zero, then the call applies to the current process.

If the invoker is not the super-user, then the affected process must have the same effective user-id as the invoker or be a descendant of the invoking process.

RETURN VALUE

Setpgrp returns when the operation was successful. If the request failed, -1 is returned and the global variable *errno* indicates the reason.

ERRORS

Setpgrp will fail and the process group will not be altered if one of the following occur:

- | | |
|---------|---|
| [ESRCH] | The requested process does not exist. |
| [EPERM] | The effective user ID of the requested process is different from that of the caller and the process is not a descendent of the calling process. |

SEE ALSO

getpgrp(2)

NAME

setquota - enable/disable quotas on a file system

SYNOPSIS

setquota(special, file)

char *special, *file;

DESCRIPTION

Disc quotas are enabled or disabled with the *setquota* call. *Special* indicates a block special device on which a mounted file system exists. If *file* is nonzero, it specifies a file in that file system from which to take the quotas. If *file* is 0, then quotas are disabled on the file system. The quota file must exist; it is normally created with the *checkquota(8)* program.

Only the super-user may turn quotas on or off.

SEE ALSO

quota(2), quotacheck(8), quotaon(8)

RETURN VALUE

A 0 return value indicates a successful call. A value of -1 is returned when an error occurs and *errno* is set to indicate the reason for failure.

ERRORS

Setquota will fail when one of the following occurs:

- | | |
|-----------|---|
| [NODEV] | The caller is not the super-user. |
| [NODEV] | <i>Special</i> does not exist. |
| [ENOTBLK] | <i>Special</i> is not a block device. |
| [ENXIO] | The major device number of <i>special</i> is out of range (this indicates no device driver exists for the associated hardware). |
| [EPERM] | The pathname contains a character with the high-order bit set. |
| [ENOTDIR] | A component of the path prefix in <i>file</i> is not a directory. |
| [EROFS] | <i>File</i> resides on a read-only file system. |
| [EACCES] | <i>File</i> resides on a file system different from <i>special</i> . |
| [EACCES] | <i>File</i> is not a plain file. |

BUGS

The error codes are in a state of disarray; too many errors appear to the caller as one value.

NAME

setregid – set real and effective group ID

SYNOPSIS

```
setregid(rgid, egid)
int rgid, egid;
```

DESCRIPTION

The real and effective group ID's of the current process are set to the arguments. Only the super-user may change the real group ID of a process. Unprivileged users may change the effective group ID to the real group ID, but to no other.

Supplying a value of -1 for either the real or effective group ID forces the system to substitute the current ID in place of the -1 parameter.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

ERRORS

[EPERM] The current process is not the super-user and a change other than changing the effective group-id to the real group-id was specified.

SEE ALSO

getgid(2), setreuid(2), setgid(3)

NAME

setreuid – set real and effective user ID's

SYNOPSIS

```
setreuid(ruid, euid)
int ruid, euid;
```

DESCRIPTION

The real and effective user ID's of the current process are set according to the arguments. If *ruid* or *euid* is -1, the current uid is filled in by the system. Only the super-user may modify the real uid of a process. Users other than the super-user may change the effective uid of a process only to the real uid.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

ERRORS

[EPERM] The current process is not the super-user and a change other than changing the effective user-id to the real user-id was specified.

SEE ALSO

getuid(2), setregid(2), setuid(3)

NAME

shmctl – shared memory control operations

SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/shm.h>

int shmctl (shmid, cmd, buf)
int shmid, cmd;
struct shmids *buf;
```

DESCRIPTION

Shmctl provides a variety of shared memory control operations as specified by *cmd*. The following *cmds* are available:

- IPC_STAT** Place the current value of each member of the data structure associated with *shmid* into the structure pointed to by *buf*. The contents of this structure are defined in [EINVAL] *intro*(2). {READ}
- IPC_SET** Set the value of the following members of the data structure associated with *shmid* to the corresponding value found in the structure pointed to by *buf*:
- ```
shm_perm.uid
shm_perm.gid
shm_perm.mode /* only low 9 bits */
```
- This *cmd* can only be executed by a process that has an effective user ID equal to either that of super user or to the value of **shm\_perm.uid** in the data structure associated with *shmid*.
- IPC\_RMID** Remove the shared memory identifier specified by *shmid* from the system and destroy the shared memory segment and data structure associated with it. This *cmd* can only be executed by a process that has an effective user ID equal to either that of super user or to the value of **shm\_perm.uid** in the data structure associated with *shmid*.
- SHM\_LOCK** Lock the shared memory segment specified by *shmid* in memory. This *cmd* can only be executed by a process that has an effective user ID equal to super user.
- SHM\_UNLOCK** Unlock the shared memory segment specified by *shmid*. This *cmd* can only be executed by a process that has an effective user ID equal to super user.

*Shmctl* will fail if one or more of the following are true:

*Shmid* is not a valid shared memory identifier. [EINVAL]

*Cmd* is not a valid command. [EINVAL]

*Cmd* is equal to **IPC\_STAT** and {READ} operation permission is denied to the calling process (see *intro*(2)). [EACCESS]

*Cmd* is equal to **IPC\_RMID** or **IPC\_SET** and the effective user ID of the calling process is not equal to that of super user and it is not equal to the value of **shm\_perm.uid** in the data structure associated with *shmid*. [EPERM]

**NAME**

setquota – enable/disable quotas on a file system

**SYNOPSIS**

```
setquota(special, file)
char *special, *file;
```

**DESCRIPTION**

Disc quotas are enabled or disabled with the *setquota* call. *Special* indicates a block special device on which a mounted file system exists. If *file* is nonzero, it specifies a file in that file system from which to take the quotas. If *file* is 0, then quotas are disabled on the file system. The quota file must exist; it is normally created with the *checkquota*(8) program.

Only the super-user may turn quotas on or off.

**SEE ALSO**

quota(2), quotacheck(8), quotaon(8)

**RETURN VALUE**

A 0 return value indicates a successful call. A value of -1 is returned when an error occurs and *errno* is set to indicate the reason for failure.

**ERRORS**

*Setquota* will fail when one of the following occurs:

- |           |                                                                                                                                 |
|-----------|---------------------------------------------------------------------------------------------------------------------------------|
| [NODEV]   | The caller is not the super-user.                                                                                               |
| [NODEV]   | <i>Special</i> does not exist.                                                                                                  |
| [ENOTBLK] | <i>Special</i> is not a block device.                                                                                           |
| [ENXIO]   | The major device number of <i>special</i> is out of range (this indicates no device driver exists for the associated hardware). |
| [EPERM]   | The pathname contains a character with the high-order bit set.                                                                  |
| [ENOTDIR] | A component of the path prefix in <i>file</i> is not a directory.                                                               |
| [EROFS]   | <i>File</i> resides on a read-only file system.                                                                                 |
| [EACCES]  | <i>File</i> resides on a file system different from <i>special</i> .                                                            |
| [EACCES]  | <i>File</i> is not a plain file.                                                                                                |

**BUGS**

The error codes are in a state of disarray; too many errors appear to the caller as one value.

**NAME**

setregid – set real and effective group ID

**SYNOPSIS**

```
setregid(rgid, egid)
int rgid, egid;
```

**DESCRIPTION**

The real and effective group ID's of the current process are set to the arguments. Only the super-user may change the real group ID of a process. Unprivileged users may change the effective group ID to the real group ID, but to no other.

Supplying a value of -1 for either the real or effective group ID forces the system to substitute the current ID in place of the -1 parameter.

**RETURN VALUE**

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

**ERRORS**

[EPERM] The current process is not the super-user and a change other than changing the effective group-id to the real group-id was specified.

**SEE ALSO**

getgid(2), setreuid(2), setgid(3)

**NAME**

setreuid - set real and effective user ID's

**SYNOPSIS**

```
setreuid(ruid, euid)
int ruid, euid;
```

**DESCRIPTION**

The real and effective user ID's of the current process are set according to the arguments. If *ruid* or *euid* is -1, the current uid is filled in by the system. Only the super-user may modify the real uid of a process. Users other than the super-user may change the effective uid of a process only to the real uid.

**RETURN VALUE**

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

**ERRORS**

[EPERM] The current process is not the super-user and a change other than changing the effective user-id to the real user-id was specified.

**SEE ALSO**

getuid(2), setregid(2), setuid(3)

**NAME**

shmctl – shared memory control operations

**SYNOPSIS**

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/shm.h>
int shmctl (shmid, cmd, buf)
int shmid, cmd;
struct shm_id *buf;
```

**DESCRIPTION**

*Shmctl* provides a variety of shared memory control operations as specified by *cmd*. The following *cmds* are available:

- IPC\_STAT** Place the current value of each member of the data structure associated with *shmid* into the structure pointed to by *buf*. The contents of this structure are defined in [EINVAL] *intro(2)*. {READ}
- IPC\_SET** Set the value of the following members of the data structure associated with *shmid* to the corresponding value found in the structure pointed to by *buf*:
- ```
shm_perm.uid
shm_perm.gid
shm_perm.mode /* only low 9 bits */
```
- This *cmd* can only be executed by a process that has an effective user ID equal to either that of super user or to the value of **shm_perm.uid** in the data structure associated with *shmid*.
- IPC_RMID** Remove the shared memory identifier specified by *shmid* from the system and destroy the shared memory segment and data structure associated with it. This *cmd* can only be executed by a process that has an effective user ID equal to either that of super user or to the value of **shm_perm.uid** in the data structure associated with *shmid*.
- SHM_LOCK** Lock the shared memory segment specified by *shmid* in memory. This *cmd* can only be executed by a process that has an effective user ID equal to super user.
- SHM_UNLOCK** Unlock the shared memory segment specified by *shmid*. This *cmd* can only be executed by a process that has an effective user ID equal to super user.

Shmctl will fail if one or more of the following are true:

Shmid is not a valid shared memory identifier. [EINVAL]

Cmd is not a valid command. [EINVAL]

Cmd is equal to **IPC_STAT** and {READ} operation permission is denied to the calling process (see *intro(2)*). [EACCES]

Cmd is equal to **IPC_RMID** or **IPC_SET** and the effective user ID of the calling process is not equal to that of super user and it is not equal to the value of **shm_perm.uid** in the data structure associated with *shmid*. [EPERM]

Cmd is equal to **SHM_LOCK** or **SHM_UNLOCK** and the effective user ID of the calling process is not equal to that of super user. [EPERM]

Cmd is equal to **SHM_UNLOCK** and the shared-memory segment specified by *shmid* is not locked in memory. [EINVAL] *Buf* points to an illegal address. [EFAULT]

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

SEE ALSO

shmget(2), shmop(2).

NAME

`shmget` – get shared memory segment

SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/shm.h>
int shmget (key, size, shmflg)
key_t key;
int size, shmflg;
```

DESCRIPTION

Shmget returns the shared memory identifier associated with *key*.

A shared memory identifier and associated data structure and shared memory segment of size *size* bytes (see *intro(2)*) are created for *key* if one of the following are true:

Key is equal to `IPC_PRIVATE`.

Key does not already have a shared memory identifier associated with it, and (*shmflg* & `IPC_CREAT`) is “true”.

Upon creation, the data structure associated with the new shared memory identifier is initialized as follows:

`Shm_perm.cuid`, `shm_perm.uid`, `shm_perm.cgid`, and `shm_perm.gid` are set equal to the effective user ID and effective group ID, respectively, of the calling process.

The low-order 9 bits of `shm_perm.mode` are set equal to the low-order 9 bits of *shmflg*. `Shm_segsz` is set equal to the value of *size*.

`Shm_lpid`, `shm_nattch`, `shm_atime`, and `shm_dtime` are set equal to 0.

`Shm_ctime` is set equal to the current time.

Shmget will fail if one or more of the following are true:

- [EINVAL] *Size* is less than the system-imposed minimum or greater than the system-imposed maximum.
- [EACCES] A shared memory identifier exists for *key* but operation permission (see *intro(2)*) as specified by the low-order 9 bits of *shmflg* would not be granted.
- [EINVAL] A shared memory identifier exists for *key* but the size of the segment associated with it is less than *size* and *size* is not equal to zero.
- [ENOENT] A shared memory identifier does not exist for *key* and (*shmflg* & `IPC_CREAT`) is “false”.
- [ENOSPC] A shared memory identifier is to be created but the system-imposed limit on the maximum number of allowed shared memory identifiers system wide would be exceeded.
- [ENOMEM] A shared memory identifier and associated shared memory segment are to be created but the amount of available physical memory is not sufficient to fill the request.
- [EEXIST] A shared memory identifier exists for *key* but ((*shmflg* & `IPC_CREAT`) and (*shmflg* & `IPC_EXCL`)) is “true”.

RETURN VALUE

Upon successful completion, a non-negative integer, namely a shared memory identifier is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

SEE ALSO

intro(2), shmctl(2), shmop(2).

NAME

shmop – shared memory operations

SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/shm.h>

char *shmat (shmid, shmaddr, shmflg)
int shmid;
char *shmaddr
int shmflg;

int shmdt (shmaddr)
char *shmaddr
```

DESCRIPTION

Shmat attaches the shared memory segment associated with the shared memory identifier specified by *shmid* to the data segment of the calling process. The segment is attached at the address specified by one of the following criteria:

If *shmaddr* is equal to zero, the segment is attached at the first available address as selected by the system.

If *shmaddr* is not equal to zero and (*shmflg* & SHM_RND) is “true”, the segment is attached at the address given by (*shmaddr* - (*shmaddr* modulus SHMLBA)).

If *shmaddr* is not equal to zero and (*shmflg* & SHM_RND) is “false”, the segment is attached at the address given by *shmaddr*.

The segment is attached for reading if (*shmflg* & SHM_RDONLY) is “true” {READ}, otherwise it is attached for reading and writing {READ/WRITE}.

Shmat will fail and not attach the shared memory segment if one or more of the following are true:

- | | |
|----------|---|
| [EINVAL] | <i>Shmid</i> is not a valid shared memory identifier. |
| [EACCES] | Operation permission is denied to the calling process (see <i>intro(2)</i>). |
| [ENOMEM] | The available data space is not large enough to accommodate the shared memory segment. |
| [EINVAL] | <i>Shmaddr</i> is not equal to zero, and the value of (<i>shmaddr</i> - (<i>shmaddr</i> modulus SHMLBA)) is an illegal address. |
| [EINVAL] | <i>Shmaddr</i> is not equal to zero, (<i>shmflg</i> & SHM_RND) is “false”, and the value of <i>shmaddr</i> is an illegal address. |
| [EMFILE] | The number of shared memory segments attached to the calling process would exceed the system-imposed limit. |
| [EINVAL] | <i>Shmdt</i> detaches from the calling process’s data segment the shared memory segment located at the address specified by <i>shmaddr</i> . |
| [EINVAL] | <i>Shmdt</i> will fail and not detach the shared memory segment if <i>shmaddr</i> is not the data segment start address of a shared memory segment. |

RETURN VALUES

Upon successful completion, the return value is as follows:

Shmat returns the data segment start address of the attached shared memory segment.

Shmdt returns a value of 0.

Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

SEE ALSO

exec(2), *exit(2)*, *fork(2)*, *intro(2)*, *shmctl(2)*, *shmget(2)*.

NAME

shutdown – shut down part of a full-duplex connection

SYNOPSIS

```
shutdown(s, how)
int s, how;
```

DESCRIPTION

The *shutdown* call causes all or part of a full-duplex connection on the socket associated with *s* to be shut down. If *how* is 0, then further receives will be disallowed. If *how* is 1, then further sends will be disallowed. If *how* is 2, then further sends and receives will be disallowed.

DIAGNOSTICS

A 0 is returned if the call succeeds, -1 if it fails.

ERRORS

The call succeeds unless:

- [EBADF] *S* is not a valid descriptor.
- [ENOTSOCK] *S* is a file, not a socket.
- [ENOTCONN] The specified socket is not connected.

SEE ALSO

connect(2), socket(2)

NAME

sigblock - block signals

SYNOPSIS

```
sigblock(mask);  
int mask;
```

DESCRIPTION

Sigblock causes the signals specified in *mask* to be added to the set of signals currently being blocked from delivery. Signal *i* is blocked if the *i*-th bit in *mask* is a 1.

It is not possible to block SIGKILL, SIGSTOP, or SIGCONT; this restriction is silently imposed by the system.

RETURN VALUE

The previous set of masked signals is returned.

SEE ALSO

kill(2), sigvec(2), sigsetmask(2),

NAME

`sigpause` - atomically release blocked signals and wait for interrupt

SYNOPSIS

```
sigpause(sigmask)  
int sigmask;
```

DESCRIPTION

Sigpause assigns *sigmask* to the set of masked signals and then waits for a signal to arrive; on return the set of masked signals is restored. *Sigmask* is usually 0 to indicate that no signals are now to be blocked. *Sigpause* always terminates by being interrupted, returning EINTR.

In normal usage, a signal is blocked using *sigblock(2)*, to begin a critical section, variables modified on the occurrence of the signal are examined to determine that there is no work to be done, and the process pauses awaiting work by using *sigpause* with the mask returned by *sigblock*.

SEE ALSO

`sigblock(2)`, `sigvec(2)`

NAME

sigsetmask - set current signal mask

SYNOPSIS

```
sigsetmask(mask);  
int mask;
```

DESCRIPTION

Sigsetmask sets the current signal mask (those signals which are blocked from delivery). Signal *i* is blocked if the *i*-th bit in *mask* is a 1.

The system quietly disallows SIGKILL, SIGSTOP, or SIGCONT to be blocked.

RETURN VALUE

The previous set of masked signals is returned.

SEE ALSO

kill(2), sigvec(2), sigblock(2), sigpause(2)

NAME

`sigstack` – set and/or get signal stack context

SYNOPSIS

```
#include <signal.h>
struct sigstack {
    caddr_t ss_sp;
    int     ss_onstack;
};
sigstack(ss, oss);
struct sigstack *ss, *oss;
```

DESCRIPTION

Sigstack allows users to define an alternate stack on which signals are to be processed. If *ss* is non-zero, it specifies a *signal stack* on which to deliver signals and tells the system if the process is currently executing on that stack. When a signal's action indicates its handler should execute on the signal stack (specified with a *sigvec*(2) call), the system checks to see if the process is currently executing on that stack. If the process is not currently executing on the signal stack, the system arranges a switch to the signal stack for the duration of the signal handler's execution. If *oss* is non-zero, the current signal stack state is returned.

NOTES

Signal stacks are not "grown" automatically, as is done for the normal stack. If the stack overflows unpredictable results may occur.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

ERRORS

Sigstack will fail and the signal stack context will remain unchanged if one of the following occurs.

[EFAULT] Either *ss* or *oss* points to memory which is not a valid part of the process address space.

SEE ALSO

sigvec(2), *setjmp*(3)

NAME

sigvec - software signal facilities

SYNOPSIS

```
#include <signal.h>
struct sigvec {
    int      (*sv_handler)();
    int      sv_mask;
    int      sv_onstack;
};
sigvec(sig, vec, ovec)
int sig;
struct sigvec *vec, *ovec;
```

DESCRIPTION

The system defines a set of signals that may be delivered to a process. Signal delivery resembles the occurrence of a hardware interrupt: the signal is blocked from further occurrence, the current process context is saved, and a new one is built. A process may specify a *handler* to which a signal is delivered, or specify that a signal is to be *blocked* or *ignored*. A process may also specify that a default action is to be taken by the system when a signal occurs. Normally, signal handlers execute on the current stack of the process. This may be changed, on a per-handler basis, so that signals are taken on a special *signal stack*.

All signals have the same *priority*. Signal routines execute with the signal that caused their invocation *blocked*, but other signals may yet occur. A global *signal mask* defines the set of signals currently blocked from delivery to a process. The signal mask for a process is initialized from that of its parent (normally 0). It may be changed with a *sigblock(2)* or *sigsetmask(2)* call, or when a signal is delivered to the process.

When a signal condition arises for a process, the signal is added to a set of signals pending for the process. If the signal is not currently *blocked* by the process then it is delivered to the process. When a signal is delivered, the current state of the process is saved, a new signal mask is calculated (as described below), and the signal handler is invoked. The call to the handler is arranged so that if the signal handling routine returns normally the process will resume execution in the context from before the signal's delivery. If the process wishes to resume in a different context, then it must arrange to restore the previous context itself.

When a signal is delivered to a process a new signal mask is installed for the duration of the process' signal handler (or until a *sigblock* or *sigsetmask* call is made). This mask is formed by taking the current signal mask, adding the signal to be delivered, and *or'ing* in the signal mask associated with the handler to be invoked.

Sigvec assigns a handler for a specific signal. If *vec* is non-zero, it specifies a handler routine and mask to be used when delivering the specified signal. Further, if *sv_onstack* is 1, the system will deliver the signal to the process on a *signal stack*, specified with *sigstack(2)*. If *ovec* is non-zero, the previous handling information for the signal is returned to the user.

The following is a list of all signals with names as in the include file *<signal.h>*:

SIGHUP	1	hangup
SIGINT	2	interrupt
SIGQUIT	3*	quit
SIGILL	4*	illegal instruction
SIGTRAP	5*	trace trap

SIGIOT	6*	IOT instruction
SIGEMT	7*	EMT instruction
SIGFPE	8*	floating point exception
SIGKILL	9	kill (cannot be caught, blocked, or ignored)
SIGBUS	10*	bus error
SIGSEGV	11*	segmentation violation
SIGSYS	12*	bad argument to system call
SIGPIPE	13	write on a pipe with no one to read it
SIGALRM	14	alarm clock
SIGTERM	15	software termination signal
SIGURG	16•	urgent condition present on socket
SIGSTOP	17†	stop (cannot be caught, blocked, or ignored)
SIGTSTP	18†	stop signal generated from keyboard
SIGCONT	19•	continue after stop (cannot be blocked)
SIGCHLD	20•	child status has changed
SIGTTIN	21†	background read attempted from control terminal
SIGTTOU	22†	background write attempted to control terminal
SIGIO	23•	i/o is possible on a descriptor (see <i>fcntl(2)</i>)
SIGXCPU	24	cpu time limit exceeded (see <i>setrlimit(2)</i>)
SIGXFSZ	25	file size limit exceeded (see <i>setrlimit(2)</i>)
SIGVTALRM	26	virtual time alarm (see <i>setitimer(2)</i>)
SIGPROF	27	profiling timer alarm (see <i>setitimer(2)</i>)

The starred signals in the list above cause a core image if not caught or ignored.

Once a signal handler is installed, it remains installed until another *sigvec* call is made, or an *execve(2)* is performed. The default action for a signal may be reinstated by setting *sv_handler* to SIG_DFL; this default is termination (with a core image for starred signals) except for signals marked with • or †. Signals marked with • are discarded if the action is SIG_DFL; signals marked with † cause the process to stop. If *sv_handler* is SIG_IGN the signal is subsequently ignored, and pending instances of the signal are discarded.

If a caught signal occurs during certain system calls, causing the call to terminate prematurely, the call is automatically restarted. In particular this can occur during a *read* or *write(2)* on a slow device (such as a terminal; but not a file) and during a *wait(2)*.

After a *fork(2)* or *vfork(2)* the child inherits all signals, the signal mask, and the signal stack.

Execve(2) resets all caught signals to default action; ignored signals remain ignored; the signal mask remains the same; the signal stack state is reset.

NOTES

The mask specified in *vec* is not allowed to block SIGKILL, SIGSTOP, or SIGCONT. This is done silently by the system.

RETURN VALUE

A 0 value indicated that the call succeeded. A -1 return value indicates an error occurred and *errno* is set to indicated the reason.

ERRORS

Sigvec will fail and no new signal handler will be installed if one of the following occurs:

- [EFAULT] Either *vec* or *ovec* points to memory which is not a valid part of the process address space.
- [EINVAL] *Sig* is not a valid signal number.

[EINVAL] An attempt is made to ignore or supply a handler for SIGKILL or SIGSTOP.

[EINVAL] An attempt is made to ignore SIGCONT (by default SIGCONT is ignored).

SEE ALSO

kill(1), ptrace(2), kill(2), sigblock(2), sigsetmask(2), sigpause(2), sigstack(2), sigvec(2), setjmp(3), tty(4)

BUGS

This manual page is confusing.

NAME

socket – create an endpoint for communication

SYNOPSIS

```
#include <sys/types.h>
#include <sys/socket.h>

s = socket(af, type, protocol)
int s, af, type, protocol;
```

DESCRIPTION

Socket creates an endpoint for communication and returns a descriptor.

The *af* parameter specifies an address format with which addresses specified in later operations using the socket should be interpreted. These formats are defined in the include file `<sys/socket.h>`. The currently understood formats are

AF_UNIX	(UNIX path names),
AF_INET	(ARPA Internet addresses),
AF_PUP	(Xerox PUP-I Internet addresses)
AF_IMPLINK	(IMP “host at IMP” addresses).

The socket has the indicated *type* which specifies the semantics of communication. Currently defined types are:

```
SOCK_STREAM
SOCK_DGRAM
SOCK_RAW
SOCK_SEQPACKET
SOCK_RDM
```

A `SOCK_STREAM` type provides sequenced, reliable, two-way connection based byte streams with an out-of-band data transmission mechanism. A `SOCK_DGRAM` socket supports datagrams (connectionless, unreliable messages of a fixed (typically small) maximum length). `SOCK_RAW` sockets provide access to internal network interfaces. The types `SOCK_RAW`, which is available only to the super-user, and `SOCK_SEQPACKET` and `SOCK_RDM`, which are planned, but not yet implemented, are not described here.

The *protocol* specifies a particular protocol to be used with the socket. Normally only a single protocol exists to support a particular socket type using a given address format. However, it is possible that many protocols may exist in which case a particular protocol must be specified in this manner. The protocol number to use is particular to the “communication domain” in which communication is to take place; see *services(3N)* and *protocols(3N)*.

Sockets of type `SOCK_STREAM` are full-duplex byte streams, similar to pipes. A stream socket must be in a *connected* state before any data may be sent or received on it. A connection to another socket is created with a *connect(2)* call. Once connected, data may be transferred using *read(2)* and *write(2)* calls or some variant of the *send(2)* and *recv(2)* calls. When a session has been completed a *close(2)* may be performed. Out-of-band data may also be transmitted as described in *send(2)* and received as described in *recv(2)*.

The communications protocols used to implement a `SOCK_STREAM` insure that data is not lost or duplicated. If a piece of data for which the peer protocol has buffer space cannot be successfully transmitted within a reasonable length of time, then the connection is considered broken and calls will indicate an error with `-1` returns and with `ETIMEDOUT` as the specific code in the global variable `errno`. The protocols optionally keep sockets “warm” by forcing transmissions roughly every minute in the absence of other activity. An error is then

indicated if no response can be elicited on an otherwise idle connection for a extended period (e.g. 5 minutes). A SIGPIPE signal is raised if a process sends on a broken stream; this causes naive processes, which do not handle the signal, to exit.

SOCK_DGRAM and SOCK_RAW sockets allow sending of datagrams to correspondents named in *send(2)* calls. It is also possible to receive datagrams at such a socket with *recv(2)*.

An *fcntl(2)* call can be used to specify a process group to receive a SIGURG signal when the out-of-band data arrives.

The operation of sockets is controlled by socket level *options*. These options are defined in the file *<sys/socket.h>* and explained below. *Setsockopt* and *getsockopt(2)* are used to set and get options, respectively.

SO_DEBUG	turn on recording of debugging information
SO_REUSEADDR	allow local address reuse
SO_KEEPALIVE	keep connections alive
SO_DONTROUTE	do not apply routing on outgoing messages
SO_LINGER	linger on close if data present
SO_DONTLINGER	do not linger on close

SO_DEBUG enables debugging in the underlying protocol modules. SO_REUSEADDR indicates the rules used in validating addresses supplied in a *bind(2)* call should allow reuse of local addresses. SO_KEEPALIVE enables the periodic transmission of messages on a connected socket. Should the connected party fail to respond to these messages, the connection is considered broken and processes using the socket are notified via a SIGPIPE signal. SO_DONTROUTE indicates that outgoing messages should bypass the standard routing facilities. Instead, messages are directed to the appropriate network interface according to the network portion of the destination address. SO_LINGER and SO_DONTLINGER control the actions taken when unsent messages are queued on socket and a *close(2)* is performed. If the socket promises reliable delivery of data and SO_LINGER is set, the system will block the process on the *close* attempt until it is able to transmit the data or until it decides it is unable to deliver the information (a timeout period, termed the linger interval, is specified in the *setsockopt* call when SO_LINGER is requested). If SO_DONTLINGER is specified and a *close* is issued, the system will process the close in a manner which allows the process to continue as quickly as possible.

RETURN VALUE

A -1 is returned if an error occurs, otherwise the return value is a descriptor referencing the socket.

ERRORS

The *socket* call fails if:

[EAFNOSUPPORT]	The specified address family is not supported in this version of the system.
[ESOCKTNOSUPPORT]	The specified socket type is not supported in this address family.
[EPROTONOSUPPORT]	The specified protocol is not supported.
[EMFILE]	The per-process descriptor table is full.
[ENOBUFS]	No buffer space is available. The socket cannot be created.

SEE ALSO

accept(2), bind(2), connect(2), getsockname(2), getsockopt(2), ioctl(2), listen(2), recv(2),
select(2), send(2), shutdown(2), socketpair(2)
"A 4.2BSD Interprocess Communication Primer".

BUGS

The use of keepalives is a questionable feature for this layer.

NAME

socketpair – create a pair of connected sockets

SYNOPSIS

```
#include <sys/types.h>
#include <sys/socket.h>
socketpair(d, type, protocol, sv)
int d, type, protocol;
int sv[2];
```

DESCRIPTION

The *socketpair* call creates an unnamed pair of connected sockets in the specified domain *d*, of the specified *type*, and using the optionally specified *protocol*. The descriptors used in referencing the new sockets are returned in *sv*[0] and *sv*[1]. The two sockets are indistinguishable.

DIAGNOSTICS

A 0 is returned if the call succeeds, -1 if it fails.

ERRORS

The call succeeds unless:

- [EMFILE] Too many descriptors are in use by this process.
- [EAFNOSUPPORT] The specified address family is not supported on this machine.
- [EPROTONOSUPPORT] The specified protocol is not supported on this machine.
- [EOPNOSUPPORT] The specified protocol does not support creation of socket pairs.
- [EFAULT] The address *sv* does not specify a valid part of the process address space.

SEE ALSO

read(2), write(2), pipe(2)

BUGS

This call is currently implemented only for the UNIX domain.

NAME

stat, lstat, fstat – get file status

SYNOPSIS

```
#include <sys/types.h>
```

```
#include <sys/stat.h>
```

```
stat(path, buf)
```

```
char *path;
```

```
struct stat *buf;
```

```
lstat(path, buf)
```

```
char *path;
```

```
struct stat *buf;
```

```
fstat(fd, buf)
```

```
int fd;
```

```
struct stat *buf;
```

DESCRIPTION

Stat obtains information about the file *path*. Read, write or execute permission of the named file is not required, but all directories listed in the path name leading to the file must be reachable.

Lstat is like *stat* except in the case where the named file is a symbolic link, in which case *lstat* returns information about the link, while *stat* returns information about the file the link references.

Fstat obtains the same information about an open file referenced by the argument descriptor, such as would be obtained by an *open* call.

Buf is a pointer to a *stat* structure into which information is placed concerning the file. The contents of the structure pointed to by *buf*

```
struct stat {
    dev_t      st_dev; /* device inode resides on */
    ino_t      st_ino; /* this inode's number */
    u_short    st_mode; /* protection */
    short      st_nlink; /* number or hard links to the file */
    short      st_uid; /* user-id of owner */
    short      st_gid; /* group-id of owner */
    dev_t      st_rdev; /* the device type, for inode that is device */
    off_t      st_size; /* total size of file */
    time_t     st_atime; /* file last access time */
    int st_spare1;
    time_t     st_mtime; /* file last modify time */
    int st_spare2;
    time_t     st_ctime; /* file last status change time */
    int st_spare3;
    long       st_blksize; /* optimal blocksize for file system i/o ops */
    long       st_blocks; /* actual number of blocks allocated */
    long       st_spare4[2];
};
```

- st_atime** Time when file data was last read or modified. Changed by the following system calls: *mknod(2)*, *utimes(2)*, *read(2)*, and *write(2)*. For reasons of efficiency, *st_atime* is not set when a directory is searched, although this would be more logical.
- st_mtime** Time when data was last modified. It is not set by changes of owner, group, link count, or mode. Changed by the following system calls: *mknod(2)*, *utimes(2)*, *write(2)*.
- st_ctime** Time when file status was last changed. It is set both both by writing and changing the i-node. Changed by the following system calls: *chmod(2)* *chown(2)*, *link(2)*, *mknod(2)*, *unlink(2)*, *utimes(2)*, *write(2)*.

The status information word *st_mode* has bits:

```

#define S_IFMT      0170000 /* type of file */
#define S_IFDIR     0040000 /* directory */
#define S_IFCHR     0020000 /* character special */
#define S_IFBLK    0060000 /* block special */
#define S_IFREG     0100000 /* regular */
#define S_IFLNK    0120000 /* symbolic link */
#define S_IFSOCK   0140000 /* socket */
#define S_ISUID    0004000 /* set user id on execution */
#define S_ISGID    0002000 /* set group id on execution */
#define S_ISVTX    0001000 /* save swapped text even after use */
#define S_IRREAD   0000400 /* read permission, owner */
#define S_IWWRITE  0000200 /* write permission, owner */
#define S_IXEXEC  0000100 /* execute/search permission, owner */

```

The mode bits 0000070 and 0000007 encode group and others permissions (see *chmod(2)*).

When *fd* is associated with a pipe, *fstat* reports an ordinary file with an i-node number, restricted permissions, and a not necessarily meaningful length.

RETURN VALUE

Upon successful completion a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

ERRORS

Stat and *lstat* will fail if one or more of the following are true:

- [ENOTDIR] A component of the path prefix is not a directory.
- [EPERM] The pathname contains a character with the high-order bit set.
- [ENOENT] The pathname was too long.
- [ENOENT] The named file does not exist.
- [EACCES] Search permission is denied for a component of the path prefix.
- [EFAULT] *Buf* or *name* points to an invalid address.

Fstat will fail if one or both of the following are true:

- [EBADF] *Fildes* is not a valid open file descriptor.
- [EFAULT] *Buf* points to an invalid address.
- [ELOOP] Too many symbolic links were encountered in translating the pathname.

CAVEAT

The fields in the `stat` structure currently marked `st_spare1`, `st_spare2`, and `st_spare3` are present in preparation for inode time stamps expanding to 64 bits. This, however, can break certain programs which depend on the time stamps being contiguous (in calls to `utimes(2)`).

SEE ALSO

`chmod(2)`, `chown(2)`, `utimes(2)`

BUGS

Applying `fstat` to a socket returns a zero'd buffer.

The list of calls which modify the various fields should be carefully checked with reality.

NAME

swapon - specify a swap directory

SYNOPSIS

```
swapon(directory)
char *directory;
```

DESCRIPTION

Swapon makes the directory *directory* available to the system for allocation for paging and swapping.

SEE ALSO

swapon(8)

NAME

symlink – make symbolic link to a file

SYNOPSIS

```
symlink(name1, name2)
char *name1, *name2;
```

DESCRIPTION

A symbolic link *name2* is created to *name1* (*name2* is the name of the file created, *name1* is the string used in creating the symbolic link). Either name may be an arbitrary path name; the files need not be on the same file system.

RETURN VALUE

Upon successful completion, a zero value is returned. If an error occurs, the error code is stored in *errno* and a -1 value is returned.

ERRORS

The symbolic link is made unless one or more of the following are true:

- | | |
|-----------|---|
| [EPERM] | Either <i>name1</i> or <i>name2</i> contains a character with the high-order bit set. |
| [ENOENT] | One of the pathnames specified was too long. |
| [ENOTDIR] | A component of the <i>name2</i> prefix is not a directory. |
| [EEXIST] | <i>Name2</i> already exists. |
| [EACCES] | A component of the <i>name2</i> path prefix denies search permission. |
| [EROFS] | The file <i>name2</i> would reside on a read-only file system. |
| [EFAULT] | <i>Name1</i> or <i>name2</i> points outside the process's allocated address space. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |

SEE ALSO

link(2), ln(1), unlink(2)

NAME

sync - update super-block

SYNOPSIS

sync()

DESCRIPTION

Sync causes all information in core memory that should be on disk to be written out. This includes modified super blocks, modified i-nodes, and delayed block I/O.

Sync should be used by programs which examine a file system, for example *fsck*, *df*, etc. *Sync* is mandatory before a boot.

SEE ALSO

fsync(2), sync(8), update(8)

BUGS

The writing, although scheduled, is not necessarily complete upon return from *sync*.

NAME

syscall – indirect system call

SYNOPSIS

syscall(**number, arg, ...**)

DESCRIPTION

Syscall performs the system call whose assembly language interface has the specified *number*, register arguments *d0* and *d1* and further arguments *arg*.

The *d0* value of the system call is returned.

DIAGNOSTICS

When the C-bit is set, *syscall* returns -1 and sets the external variable *errno* (see *intro(2)*).

BUGS

There is no way to simulate system calls such as *pipe(2)*, which return values in register *d1*.

NAME

`truncate` – truncate a file to a specified length

SYNOPSIS

```
truncate(path, length)
char *path;
int length;

ftruncate(fd, length)
int fd, length;
```

DESCRIPTION

Truncate causes the file named by *path* or referenced by *fd* to be truncated to at most *length* bytes in size. If the file previously was larger than this size, the extra data is lost. With *ftruncate*, the file must be open for writing.

RETURN VALUES

A value of 0 is returned if the call succeeds. If the call fails a -1 is returned, and the global variable *errno* specifies the error.

ERRORS

Truncate succeeds unless:

- [EPERM] The pathname contains a character with the high-order bit set.
- [ENOENT] The pathname was too long.
- [ENOTDIR] A component of the path prefix of *path* is not a directory.
- [ENOENT] The named file does not exist.
- [EACCES] A component of the *path* prefix denies search permission.
- [EISDIR] The named file is a directory.
- [EROFS] The named file resides on a read-only file system.
- [ETXTBSY] The file is a pure procedure (shared text) file that is being executed.
- [EFAULT] *Name* points outside the process's allocated address space.

Ftruncate succeeds unless:

- [EBADF] The *fd* is not a valid descriptor.
- [EINVAL] The *fd* references a socket, not a file.

SEE ALSO

`open(2)`

BUGS

Partial blocks discarded as the result of truncation are not zero filled; this can result in holes in files which do not read as zero.

These calls should be generalized to allow ranges of bytes in a file to be discarded.

NAME

umask - set file creation mode mask

SYNOPSIS

```
oumask = umask(numask)
int oumask, numask;
```

DESCRIPTION

Umask sets the process's file mode creation mask to *numask* and returns the previous value of the mask. The low-order 9 bits of *numask* are used whenever a file is created, clearing corresponding bits in the file mode (see *chmod(2)*). This clearing allows each user to restrict the default access to his files.

The value is initially 022 (write access for owner only). The mask is inherited by child processes.

RETURN VALUE

The previous value of the file mode mask is returned by the call.

SEE ALSO

chmod(2), mknod(2), open(2)

NAME

unlink - remove directory entry

SYNOPSIS

```
unlink(path)
char *path;
```

DESCRIPTION

Unlink removes the entry for the file *path* from its directory. If this entry was the last link to the file, and no process has the file open, then all resources associated with the file are reclaimed. If, however, the file was open in any process, the actual resource reclamation is delayed until it is closed, even though the directory entry has disappeared.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

ERRORS

The *unlink* succeeds unless:

- | | |
|-----------|---|
| [EPERM] | The path contains a character with the high-order bit set. |
| [ENOENT] | The path name is too long. |
| [ENOTDIR] | A component of the path prefix is not a directory. |
| [ENOENT] | The named file does not exist. |
| [EACCES] | Search permission is denied for a component of the path prefix. |
| [EACCES] | Write permission is denied on the directory containing the link to be removed. |
| [EPERM] | The named file is a directory and the effective user ID of the process is not the super-user. |
| [EBUSY] | The entry to be unlinked is the mount point for a mounted file system. |
| [EROFS] | The named file resides on a read-only file system. |
| [EFAULT] | <i>Path</i> points outside the process's allocated address space. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |

SEE ALSO

close(2), link(2), rmdir(2)

NAME

utimes - set file times

SYNOPSIS

```
#include <sys/time.h>
utimes(file, tvp)
char *file;
struct timeval *tvp[2];
```

DESCRIPTION

The *utimes* call uses the “accessed” and “updated” times in that order from the *tvp* vector to set the corresponding recorded times for *file*.

The caller must be the owner of the file or the super-user. The “inode-changed” time of the file is set to the current time.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

ERRORS

Utime will fail if one or more of the following are true:

- | | |
|-----------|--|
| [EPERM] | The pathname contained a character with the high-order bit set. |
| [ENOENT] | The pathname was too long. |
| [ENOENT] | The named file does not exist. |
| [ENOTDIR] | A component of the path prefix is not a directory. |
| [EACCES] | A component of the path prefix denies search permission. |
| [EPERM] | The process is not super-user and not the owner of the file. |
| [EACCES] | The effective user ID is not super-user and not the owner of the file and <i>times</i> is NULL and write access is denied. |
| [EROFS] | The file system containing the file is mounted read-only. |
| [EFAULT] | <i>Tvp</i> points outside the process's allocated address space. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |

SEE ALSO

stat(2)

NAME

vfork - spawn new process in a virtual memory efficient way

SYNOPSIS

```
pid = vfork()
int pid;
```

DESCRIPTION

Vfork can be used to create new processes without fully copying the address space of the old process, which is horrendously inefficient in a paged environment. It is useful when the purpose of *fork(2)* would have been to create a new system context for an *execve*. *Vfork* differs from *fork* in that the child borrows the parent's memory and thread of control until a call to *execve(2)* or an exit (either by a call to *exit(2)* or abnormally.) The parent process is suspended while the child is using its resources.

Vfork returns 0 in the child's context and (later) the pid of the child in the parent's context.

Vfork can normally be used just like *fork*. It does not work, however, to return while running in the child's context from the procedure which called *vfork* since the eventual return from *vfork* would then return to a no longer existent stack frame. Be careful, also, to call *_exit* rather than *exit* if you can't *execve*, since *exit* will flush and close standard I/O channels, and thereby mess up the parent process's standard I/O data structures. (Even with *fork* it is wrong to call *exit* since buffered data would then be flushed twice.)

SEE ALSO

fork(2), *execve(2)*, *sigvec(2)*, *wait(2)*,

DIAGNOSTICS

Same as for *fork*.

BUGS

This system call will be eliminated when proper system sharing mechanisms are implemented. Users should not depend on the memory sharing semantics of *vfork* as it will, in that case, be made synonymous to *fork*.

To avoid a possible deadlock situation, processes which are children in the middle of a *vfork* are never sent SIGTTOU or SIGTTIN signals; rather, output or *ioctl*s are allowed and input attempts result in an end-of-file indication.

NAME

vhangup - virtually "hangup" the current control terminal

SYNOPSIS

vhangup()

DESCRIPTION

Vhangup is used by the initialization process *init*(8) (among others) to arrange that users are given "clean" terminals at login, by revoking access of the previous users' processes to the terminal. To effect this, *vhangup* searches the system tables for references to the control terminal of the invoking process, revoking access permissions on each instance of the terminal which it finds. Further attempts to access the terminal by the affected processes will yield i/o errors (EBADF). Finally, a hangup signal (SIGHUP) is sent to the process group of the control terminal.

SEE ALSO

init (8)

BUGS

Access to the control terminal via */dev/tty* is still possible.

This call should be replaced by an automatic mechanism which takes place on process exit.

NAME

wait, *wait3* – wait for process to terminate

SYNOPSIS

```
#include <sys/wait.h>
pid = wait(status)
int pid;
union wait *status;

pid = wait(0)
int pid;

#include <sys/time.h>
#include <sys/resource.h>

pid = wait3(status, options, rusage)
int pid;
union wait *status;
int options;
struct rusage *rusage;
```

DESCRIPTION

Wait causes its caller to delay until a signal is received or one of its child processes terminates. If any child has died since the last *wait*, return is immediate, returning the process id and exit status of one of the terminated children. If there are no children, return is immediate with the value -1 returned.

On return from a successful *wait* call, *status* is nonzero, and the high byte of *status* contains the low byte of the argument to *exit* supplied by the child process; the low byte of *status* contains the termination status of the process. A more precise definition of the *status* word is given in *<sys/wait.h>*.

Wait3 provides an alternate interface for programs which must not block when collecting the status of child processes. The *status* parameter is defined as above. The *options* parameter is used to indicate the call should not block if there are no processes which wish to report status (WNOHANG), and/or that only children of the current process which are stopped due to a SIGTTIN, SIGTTOU, SIGTSTP, or SIGSTOP signal should have their status reported (WUNTRACED). If *rusage* is non-zero, a summary of the resources used by the terminated process and all its children is returned (this information is currently not available for stopped processes).

When the WNOHANG option is specified and no processes wish to report status, *wait3* returns a *pid* of 0. The WNOHANG and WUNTRACED options may be combined by *or*'ing the two values.

NOTES

See *sigvec(2)* for a list of termination statuses (signals); 0 status indicates normal termination. A special status (0177) is returned for a stopped process which has not terminated and can be restarted; see *ptrace(2)*. If the 0200 bit of the termination status is set, a core image of the process was produced by the system.

If the parent process terminates without waiting on its children, the initialization process (process ID = 1) inherits the children.

Wait and *wait9* are automatically restarted when a process receives a signal while awaiting termination of a child process.

RETURN VALUE

If *wait* returns due to a stopped or terminated child process, the process ID of the child is returned to the calling process. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

Wait9 returns -1 if there are no children not previously waited for; 0 is returned if WNOHANG is specified and there are no stopped or exited children.

ERRORS

Wait will fail and return immediately if one or more of the following are true:

- [ECHILD] The calling process has no existing unwaited-for child processes.
- [EFAULT] The *status* or *rusage* arguments point to an illegal address.

SEE ALSO

exit(2)

NAME

write, writev – write on a file

SYNOPSIS

```

write(d, buf, nbytes)
int d;
char *buf;
int nbytes;
#include <sys/types.h>
#include <sys/uio.h>
writev(d, iov, iovectlen)
int d;
struct iovec *iov;
int iovectlen;

```

DESCRIPTION

Write attempts to write *nbytes* of data to the object referenced by the descriptor *d* from the buffer pointed to by *buf*. *Writev* performs the same action, but gathers the output data from the *iovlen* buffers specified by the members of the *iovec* array: *iov*[0], *iov*[1], etc.

On objects capable of seeking, the *write* starts at a position given by the pointer associated with *d*, see *lseek*(2). Upon return from *write*, the pointer is incremented by the number of bytes actually written.

Objects that are not capable of seeking always write from the current position. The value of the pointer associated with such an object is undefined.

If the real user is not the super-user, then *write* clears the set-user-id bit on a file. This prevents penetration of system security by a user who “captures” a writable set-user-id file owned by the super-user.

RETURN VALUE

Upon successful completion the number of bytes actually written is returned. Otherwise a *-1* is returned and *errno* is set to indicate the error.

ERRORS

Write will fail and the file pointer will remain unchanged if one or more of the following are true:

- | | |
|----------|--|
| [EBADF] | <i>D</i> is not a valid descriptor open for writing. |
| [EPIPE] | An attempt is made to write to a pipe that is not open for reading by any process. |
| [EPIPE] | An attempt is made to write to a socket of type SOCK_STREAM which is not connected to a peer socket. |
| [EFBIG] | An attempt was made to write a file that exceeds the process’s file size limit or the maximum file size. |
| [EFAULT] | Part of <i>iov</i> or data to be written to the file points outside the process’s allocated address space. |

SEE ALSO

lseek(2), *open*(2), *pipe*(2)

C O M M E N T S

ICON/UXB REFERENCE MANUAL Volume 1A

P/N 172-022-001

Your comments and suggestions are appreciated and will help us to provide you with the very best in system and application documentation. Send your comments to the address at the bottom of this page. Users who respond will be entitled to free updates of this manual for one year.

1. How would you rate this manual for **COMPLETENESS?** (Please Circle)

Excellent Poor
5 ----- 4 ----- 3 ----- 2 ----- 1 ----- 0

2. Is there any information that you feel should be included or removed?

3. How would you rate this manual for **ACCURACY?** (Please Circle)

Excellent Poor
5 ----- 4 ----- 3 ----- 2 ----- 1 ----- 0

4. Indicate the page number and nature of any error(s) found in this manual.

5. How would you rate this manual for **USABILITY?** (Please Circle)

Excellent Poor
5 ----- 4 ----- 3 ----- 2 ----- 1 ----- 0

6. Describe any format or packaging problems you have experienced with this manual and/or binder.

7. Do you have any general comments or suggestions regarding this publication or future publications?

Your Name _____


Company _____

Address _____ Phone (____) _____

City & State _____ Zip Code _____

Job Function _____

Type of Equipment Installed: _____





Copyright © 1987 Icon International, Inc.

Printed in the U.S.A.

Rev B

172-022-001