

**Uniflex+®**

**System V, Release 2**

---

**Volume 2**

**User Manual, Sections 2-8**



## **COPYRIGHT**

Copyright © 1985 by UniSoft Systems. Portions of this material have been previously copyrighted by AT&T Bell Laboratories, Western Electric Company, and Regents of the University of California. Holders of a UNIX and UniPlus<sup>+</sup> 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.

## **DISCLAIMER**

While UniSoft Systems has endeavored to exercise care in the preparation of this guide, it nevertheless makes no warranties of any kind with regard to the documentation contained herein, including no warranty of merchantability or fitness for a particular purpose. In no event shall UniSoft be liable for incidental or consequential damages in connection with or arising out of the furnishing, performance, or use of any of this documentation.

## **TRADEMARKS**

UNIX is a trademark of AT&T Bell Laboratories. UniPlus<sup>+</sup> and UniSoft are registered trademarks of UniSoft Systems.

Adapted to UniPlus<sup>+</sup> by Heather Allen of UniSoft Systems.

# INTRODUCTION

This manual describes the features of System V UniPlus<sup>+</sup>, a UNIX operating system. All commands, features, and facilities described in this manual are available on UniPlus<sup>+</sup>.

This manual is divided into two volumes containing a total of six sections, some divided into subsections.

1. **Commands and Application Programs:**
  1. **General-Purpose Commands.**
  - 1C. **Communications Commands.**
  - 1G. **Graphics Commands.**
  - 1N. **Networking Commands.**
2. **System Calls.**
  - 2N. **Networking Calls.**
3. **Subroutines:**
  - 3C. **C and Assembler Library Routines.**
  - 3F. **FORTRAN Library Routines.**
  - 3M. **Mathematical Library Routines.**
  - 3N. **Networking Routines.**
  - 3S. **Standard I/O Library Routines.**
  - 3X. **Miscellaneous Routines.**
4. **File Formats.**
  - 4N. **Networking Formats.**
5. **Miscellaneous Facilities.**
  - 5F. **Protocol Family.**
  - 5P. **Protocol Descriptions.**
6. **Games.**

**Section 1** (*Commands and Application Programs*) describes programs invoked directly by the user or by command language procedures, as opposed to subroutines, which are called by the user's programs. Commands generally reside in the directory `/bin` (for **binary** programs). Some programs also reside in `/usr/bin`, to save space in `/bin`. These directories are searched automatically. Subsection 1C contains communication programs such as *cu*, *send*, *uucp*, etc.

**Section 2** (*System Calls*) describes the entries into the UNIX System kernel, including the C language interface.

## INTRODUCTION

**Section 3** (*Subroutines*) describes the available subroutines. Their binary versions reside in various system libraries in the directories `/lib` and `/usr/lib`. See *intro(3)* for descriptions of these libraries and the files in which they are stored.

**Section 4** (*File Formats*) documents the structure of particular kinds of files. Excluded are files used by only one command (for example, the assembler's intermediate files). In general, the C language `struct` declarations corresponding to these formats can be found in the directories `/usr/include` and `/usr/include/sys`.

**Section 5** (*Miscellaneous Facilities*) contains descriptions of character sets, macro packages, etc.

**Section 6** (*Games*) describes the games and educational programs that reside in the directory `/usr/games`.

Each section consists of several entries, each a page or so long. The name of the entry appears in the upper corners of its pages. Entries within each section are alphabetized, except the introduction that begins each section. The page numbers of each entry start at 1. Some entries may describe several routines, commands, etc. In such cases, the entry appears only once, alphabetized under its "major" name.

All entries are based on a common format, not all of whose parts always appear:

**NAME** gives the name(s) and a brief description of the entry.

**SYNOPSIS** summarizes the use of the program. A few conventions are used, particularly in Section 1 (*Commands*):

**Boldface** strings are typed just as they appear.

*Italic* strings usually represent substitutable argument prototypes (such as *filename*) which you are expected to substitute for the actual name. When an argument prototype is given as "name" or "file", it always refers to a *file* name.



## INTRODUCTION

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

Ellipses **...** show that the previous argument prototype may be repeated.

A final convention is used by the commands themselves. An argument beginning with a minus **-**, plus **+**, or equal sign **=** is often taken to a flag 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 **-**, **+**, or **=**.

**DESCRIPTION** discusses the program.

**EXAMPLE(S)** gives example(s) of usage.

**FILES** gives the file names that are built into the program.

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

**DIAGNOSTICS** discusses the diagnostic indications that may be produced. Self-explanatory messages are not listed.

**WARNINGS** points out potential pitfalls.

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

At the front of each volume there is a a table of contents and a permuted index. The permuted index lists the commands by the information in the **NAME** part of each entry in the User and Administrator Manual. The permuted index contains three columns. The center column is an alphabetic list of keywords. The last column is the entry that the keyword in the center column refers to. This entry is followed by the appropriate section number in parentheses. The first column contains the remaining information from the **NAME** part that either precedes or follows the keyword.

## INTRODUCTION

For example, to look for a text editor, scan the center column for the word "editor." There are several index lines containing an "editor reference, i.e.:

```
ed, red: text      editor ..... ed(1)
files. ld: link   editor for common object ..... ld(1)
```

You can then turn to the entries listed in the last column, *ed(1)* and *ld(1)*, to find information on the editor.

On most systems, all entries are available on-line via the *man(1)* command.

## TABLE OF CONTENTS

### 2. System Calls

intro	introduction to system calls and error numbers
accept	accept a connection on a socket
access	determine accessibility of a file
acct	enable or disable process accounting
alarm	set a process's alarm clock
bind	bind a name to a socket
brk	change data segment space allocation
chdir	change working directory
chmod	change mode of file
chown	change owner and group of a file
chroot	change root directory
close	close a file descriptor
connect	initiate a connection on a socket
creat	create a new file or rewrite an existing one
exec	execute a file
exit	terminate process
fcntl	file control
fork	create a new process
gethostid	get/set unique identifier of current host
gethostname	get/set name of current host
getpeername	get name of connected peer
getpid	get process, process group, and parent process IDs
getsockname	get socket name
getsockopt	get and set options on sockets
getuid	get real user, effective user, real group, and effective group IDs
ioctl	control device
kill	send a signal to a process or a group of processes
link	link to a file
listen	listen for connections on a socket
locking	provide exclusive file regions for reading or writing
lseek	move read/write file pointer
mknod	make a directory, or a special or ordinary file
mount	mount a file system
msgctl	message control operations
msgget	get message queue
msgop	message operations
nice	change priority of a process
open	open for reading or writing
pause	suspend process until signal
phys	allow a process to access physical addresses
pipe	create an interprocess channel
plock	lock process, text, or data in memory
profil	execution time profile
ptrace	process trace
read	read from file
reboot	reboot the system
recv	receive a message from a socket
select	synchronous i/o multiplexing
semctl	semaphore control operations
semget	get set of semaphores
semop	semaphore operations
send	send a message from a socket
setpgpr	set process group ID

Table of Contents

setregid	set real and effective group ID
setreuid	set real and effective user ID's
setuid	set user and group IDs
shmctl	shared memory control operations
shmget	get shared memory segment
shmop	shared memory operations
shutdown	shut down part of a full-duplex connection
signal	specify what to do upon receipt of a signal
socket	create an endpoint for communication
stat	get file status
stime	set time
sync	update super-block
time	get time
times	get process and child process times
ulimit	get and set user limits
umask	set and get file creation mask
umount	unmount a file system
uname	get name of current UNIX system
unlink	remove directory entry
ustat	get file system statistics
utime	set file access and modification times
uvar	returns system-specific configuration information
wait	wait for child process to stop or terminate
wait3	wait for child process to stop or terminate

3. Subroutines

intro	introduction to subroutines and libraries
a64l	convert between long integer and base-64 ASCII string
abort	generate an IOT fault
abort	terminate Fortran program
abs	return integer absolute value
abs	Fortran absolute value
acos	Fortran arccosine intrinsic function
aimag	Fortran imaginary part of complex argument
aint	Fortran integer part intrinsic function
asin	Fortran arcsine intrinsic function
assert	verify program assertion
atan	Fortran arctangent intrinsic function
atan2	Fortran arctangent intrinsic function
atof	convert ASCII string to floating-point number
bessel	Bessel functions
bft	block transfer data
bool	Fortran bitwise boolean functions
bsearch	binary search a sorted table
bstring	bit and byte string operations
byteorder	convert values between host and network byte order
clock	report CPU time used
conjg	Fortran complex conjugate intrinsic function
conv	translate characters
cos	Fortran cosine intrinsic function
cosh	Fortran hyperbolic cosine intrinsic function
crypt	generate DES encryption
ctermid	generate filename for terminal
ctime	convert date and time to string
ctype	classify characters
curses	CRT screen handling and optimization package

cuserid	get character login name of the user
dial	establish an out-going terminal line connection
dim	positive difference intrinsic functions
directory	flexible length directory operations
dprod	double precision product intrinsic function
drand48	generate uniformly distributed pseudo-random numbers
dup	duplicate a descriptor
dup2	duplicate a descriptor
ecvt	convert floating-point number to string
end	last locations in program
erf	error function and complementary error function
exp	Fortran exponential intrinsic function
exp	exponential, logarithm, power, square root functions
fclose	close or flush a stream
ferror	stream status inquiries
floor	floor, ceiling, remainder, absolute value functions
fopen	open a stream
fread	binary input/output
frexp	manipulate parts of floating-point numbers
fseek	reposition a file pointer in a stream
ftw	walk a file tree
ftype	explicit Fortran type conversion
gamma	log gamma function
getarg	return Fortran command-line argument
getc	get character or word from a stream
getcwd	get pathname of current working directory
getdtablesize	get descriptor table size
getenv	return value for environment name
getenv	return Fortran environment variable
getgrent	obtain group file entry from a group file
gethostent	get network host entry
getlogin	get login name
getnetent	get network entry
getopt	get option letter from argument vector
getpass	read a password
getprotoent	get protocol entry
getpw	get name from UID
getpwent	get password file entry
gets	get a string from a stream
getservent	get service entry
getut	access utmp file entry
hsearch	manage hash search tables
hypot	Euclidean distance function
largc	
index	return location of Fortran substring
inet	Internet address manipulation routines
insque	insert/remove element from a queue
killpg	send signal to a process group
l3tol	convert between 3-byte integers and long integers
ldahread	read the archive header of a member of an archive file
ldclose	close a common object file
ldfhead	read the file header of a common object file
ldgetname	retrieve symbol name for object file
ldlread	manipulate line number entries of a common object file function
ldlseek	seek to line number entries of a section of a common object file
ldohseek	seek to the optional file header of a common object file
ldopen	open a common object file for reading

Table of Contents

ldrseek	seek to relocation entries of a section of a common object file
ldhread	read an indexed/named section header of a common object file
ldsseek	seek to an indexed/named section of a common object file
ldtbindx	compute the index of a symbol table entry of a common object file
ldtbread	read an indexed symbol table entry of a common object file
ldtbseek	seek to the symbol table of a common object file
len	return length of Fortran string
lockf	record locking on files
log	Fortran natural logarithm intrinsic function
log10	Fortran common logarithm intrinsic function
logname	return login name of user
lsearch	linear search and update
malloc	main memory allocator
malloca	fast main memory allocator
matherr	error-handling function
max	Fortran maximum-value functions
mclock	return Fortran time accounting
memory	memory operations
min	Fortran minimum-value functions
mktemp	make a unique filename
mod	Fortran remaindering intrinsic functions
monitor	prepare execution profile
nlist	get entries from name list
perror	system error messages
plot	graphics interface subroutines
popen	initiate pipe to/from a process
printf	print formatted output
putc	put character or word on a stream
putenv	change or add value to environment
putpwent	write password file entry
puts	put a string on a stream
qsort	quicker sort
rand	simple random-number generator
rand	Fortran uniform random-number generator
rcmd	routines for returning a stream to a remote command
readv	read from file
regcmp	compile and execute a regular expression
rexec	return stream to a remote command
round	Fortran nearest integer functions
scanf	convert formatted input
setbuf	assign buffering to a stream
setjmp	non-local goto
sign	Fortran transfer-of-sign intrinsic function
signal	specify Fortran action on receipt of a system signal
sin	Fortran sine intrinsic function
sinh	Fortran hyperbolic sine intrinsic function
sinh	hyperbolic functions
sleep	suspend execution for interval
sputl	access long integer data in a machine independent fashion
sqrt	Fortran square root intrinsic function
ssignal	software signals
stdio	standard buffered input/output package
stdipc	standard interprocess communication package
strcmp	string comparison intrinsic functions
string	string operations
strtod	convert string to double-precision number
strtol	convert string to integer

swab . . . . .	swap bytes
system . . . . .	issue a shell command from Fortran
system . . . . .	issue a shell command
tan . . . . .	Fortran tangent intrinsic function
tanh . . . . .	Fortran hyperbolic tangent intrinsic function
termcap . . . . .	terminal independent operation routines
tmpfile . . . . .	create a temporary file
tmpnam . . . . .	create a name for a temporary file
trig . . . . .	trigonometric functions
tsearch . . . . .	manage binary search trees
ttyname . . . . .	find name of a terminal
ttyslot . . . . .	find the slot in the utmp file of the current user
ungetc . . . . .	push character back into input stream
vprintf . . . . .	print formatted output of a varargs argument list
vprintf . . . . .	print formatted output of a varargs argument list
write . . . . .	write on a file
writev . . . . .	write on a file

#### 4. File Formats

intro . . . . .	introduction to file formats
a.out . . . . .	common assembler and link editor output
a.out5.0 . . . . .	assembler and link editor output
acct . . . . .	per-process accounting file format
altblk . . . . .	alternate block information for bad block handling
aouthdr . . . . .	aouthdr.h - a.out header for common object files
ar . . . . .	common archive file format
ar5.0 . . . . .	archive (library) file format
checklist . . . . .	list of file systems processed by fsck
core . . . . .	format of core image file
cpio . . . . .	format of cpio archive
dir . . . . .	format of directories
errfile . . . . .	error-log file format
filehdr . . . . .	file header for common object files
fs . . . . .	format of system volume
fspec . . . . .	format specification in text files
gettydefs . . . . .	speed and terminal settings used by getty
group . . . . .	group file
hosts . . . . .	host name data base
inittab . . . . .	script for the init process
inode . . . . .	format of an inode
issue . . . . .	issue identification file
ldfcn . . . . .	common object file access routines
linenum . . . . .	line number entries in a common object file
master . . . . .	master device information table
mnttab . . . . .	mounted file system table
networks . . . . .	network name data base
passwd . . . . .	password file
plot . . . . .	graphics interface
pnch . . . . .	file format for card images
profile . . . . .	setting up an environment at login time
protocols . . . . .	protocol name data base
reloc . . . . .	relocation information for a common object file
scsfile . . . . .	format of SCCS file
scnhdr . . . . .	section header for a common object file
services . . . . .	service name data base
syms . . . . .	common object file symbol table format

## Table of Contents

term	format of compiled term file.
terminfo	terminal capability data base
ttytype	data base of terminal types by port
utmp	utmp and wtmp entry formats

### 5. Miscellaneous Facilities

intro	introduction to miscellany
arp	Address Resolution Protocol
ascii	map of ASCII character set
environ	user environment
eqnchar	special character definitions for eqn and neqn
fcntl	file control options
greek	graphics for the extended TTY-37 type-box
inet	Internet protocol family
intro	introduction to networking facilities
ip	Internet Protocol
lo	software loopback network interface
man	macros for formatting entries in this manual
math	math functions and constants
mm	the MM macro package for formatting documents
mosd	the OSDD adapter macro package for formatting documents
mptx	the macro package for formatting a permuted index
mv	a troff macro package for typesetting view graphs and slides
prof	profile within a function
pty	pseudo terminal driver
regexp	regular expression compile and match routines
stat	data returned by stat system call
tcp	Internet Transmission Control Protocol
term	conventional names for terminals
termcap	terminal capability data base
types	primitive system data types
udp	Internet User Datagram Protocol
values	machine-dependent values
varargs	handle variable argument list

### 6. Games

intro	introduction to games
adventure	an exploration game
aliens	The alien invaders attack the earth
arithmetic	provide drill in number facts
autorobots	Escape from the automatic robots
back	the game of backgammon
bcd	convert to antique media
bj	the game of black jack
chase	Try to escape the killer robots
craps	the game of craps
cribbage	the card game cribbage
fish	play "Go Fish"
fortune	print a random, hopefully interesting, adage
hangman	guess the word
life	play the game of life
maze	generate a maze
moo	guessing game
number	convert Arabic numerals to English
quiz	test your knowledge
rain	animated raindrops display



*Table of Contents*

<b>robots</b>	Escape from the robots
<b>trek</b>	trekkie game
<b>ttt</b>	tic-tac-toe
<b>twinkle</b>	twinkle stars on the screen
<b>worm</b>	Play the growing worm game
<b>worms</b>	animate worms on a display terminal
<b>wump</b>	the game of hunt-the-wumpus



## COMMANDS

**300(1)**.....handle special functions of DASI 300 terminal  
**300s** (See *300(1)*).....handle special functions of DASI 300s terminal  
**4014(1)**.....paginator for the Tektronix 4014 terminal  
**450(1)**.....handle special functions of the DASI 450 terminal  
**\_exit** (See *exit(2)*).....terminate process  
**\_tolower** (See *conv(3C)*).....translate characters  
**\_toupper** (See *conv(3C)*).....translate characters  
**a.out(4)**.....common assembler and link editor output  
**a.out5.0(4)**.....assembler and link editor output (System V a.out format only)  
**a64i(3C)**.....convert between long integer and base-64 ASCII string  
**abort(3C)**.....generate an IOT fault  
**abort(3F)**.....terminate Fortran program  
**abs(3C)**.....return integer absolute value  
**abs(3F)**.....Fortran absolute value  
**accept(1M)**.....allow LP requests  
**accept(2N)**.....accept a connection on a socket  
**access(2)**.....determine accessibility of a file  
**acct(1M)**.....overview of accounting  
**acct(2)**.....enable or disable process accounting  
**acct(4)**.....per-process accounting file format  
**acctcms(1M)**.....command summary from per-process accounting records  
**acctcom(1)**.....search and print process accounting file(s)  
**acctcom(1M)**.....connect-time accounting  
**acctcon1** (See *acctcon(1M)*).....connect-time accounting  
**acctcon2** (See *acctcon(1M)*).....connect-time accounting  
**acctdisk** (See *acct(1M)*).....miscellaneous accounting command  
**acctdusg** (See *acct(1M)*).....miscellaneous accounting command  
**acctmerg(1M)**.....merge or add total accounting files  
**accton** (See *acct(1M)*).....miscellaneous accounting command  
**acctpre(1M)**.....process accounting  
**acctpre1** (See *acctpre(1M)*).....process accounting  
**acctpre2** (See *acctpre(1M)*).....process accounting  
**acctsh(1M)**.....shell procedures for accounting  
**acctwtmp** (See *acct(1M)*).....miscellaneous accounting command  
**acos** (See *trig(3M)*).....trigonometric function  
**acos(3F)**.....Fortran arccosine intrinsic function  
**admin(1)**.....create and administer SCCS files

## COMMANDS

<b>adventure(6)</b> .....	an exploration game
<b>almag(3F)</b> .....	Fortran imaginary part of complex argument
<b>aint(3F)</b> .....	Fortran integer part intrinsic function
<b>alarm(2)</b> .....	set a process's alarm clock
<b>aliases(7N)</b> .....	aliases file for delivermail
<b>aliens(6)</b> .....	the alien invaders attack the earth
<b>alog (See log(3F))</b> .....	Fortran natural logarithm intrinsic function
<b>alog10 (See log10(3F))</b> .....	Fortran common logarithm intrinsic function
<b>altblk(4)</b> .....	alternate block information for bad block handling
<b>amax0 (See max(3F))</b> .....	Fortran maximum-value function
<b>amax1 (See max(3F))</b> .....	Fortran maximum-value function
<b>amin0 (See min(3F))</b> .....	Fortran minimum-value function
<b>amin1 (See min(3F))</b> .....	Fortran minimum-value function
<b>amod (See mod(3F))</b> .....	Fortran remaindering intrinsic function
<b>and (See bool(3F))</b> .....	Fortran bitwise boolean function
<b>aint (See round(3F))</b> .....	Fortran nearest integer function
<b>aouthdr(4)</b> .....	a.out header for common object files
<b>ar(1)</b> .....	archive and library maintainer for portable archives
<b>ar(4)</b> .....	common archive file format
<b>ar5.0(1)</b> .....	archive and library maintainer (System V a.out format only)
<b>ar5.0(4)</b> .....	archive (library) file format (System V a.out format only)
<b>arithmetic(6)</b> .....	provide drill in number facts
<b>arp(5P)</b> .....	Address Resolution Protocol
<b>as(1)</b> .....	common assembler
<b>as5.0(1)</b> .....	assembler (System V a.out format only)
<b>asa(1)</b> .....	interpret ASA carriage control characters
<b>ascll(5)</b> .....	map of ASCII character set
<b>asctime (See ctime(3C))</b> .....	convert date and time to string
<b>asin (See trig(3M))</b> .....	trigonometric function
<b>asin(3F)</b> .....	Fortran arcsine intrinsic function
<b>assert(3X)</b> .....	verify program assertion
<b>at(1)</b> .....	execute commands at a later time
<b>atan (See trig(3M))</b> .....	trigonometric function
<b>atan(3F)</b> .....	Fortran arctangent intrinsic function
<b>atan2 (See trig(3M))</b> .....	trigonometric function
<b>atan2(3F)</b> .....	Fortran arctangent intrinsic function
<b>atof(3C)</b> .....	convert ASCII string to floating-point number
<b>atoi (See strtol(3C))</b> .....	convert string to integer
<b>atol (See strtol(3C))</b> .....	convert string to integer

## COMMANDS

**autorobots(6)** .....escape from the automatic robots  
**awk(1)** .....pattern scanning and processing language  
**back(6)** .....the game of backgammon  
**badblk(1M)** .....program to set or update bad block information  
**banner(1)** .....make posters  
**banner7(1)** .....print large banner on printer  
**basename(1)** .....deliver portions of path names  
**batch (See at(1))** .....execute commands at a later time  
**bc(1)** .....arbitrary-precision arithmetic language  
**bd(6)** .....convert to antique media  
**becheckrc (See brc(1M))** .....system initialization shell script  
**bcmp (See bstring(3N))** .....byte string operation  
**bcopy (See bstring(3N))** .....byte string operation  
**bcopy(1M)** .....interactive block copy  
**bdiff(1)** .....big diff  
**bessel(3M)** .....Bessel functions  
**bfs(1)** .....big file scanner  
**bind(2N)** .....bind a name to a socket  
**bj(6)** .....the game of black jack  
**blt(3C)** .....block transfer data  
**blt512 (See blt(3C))** .....block transfer data  
**bool(3F)** .....Fortran bitwise boolean functions  
**boot(8)** .....startup procedures  
**brc(1M)** .....system initialization shell script  
**brk(2)** .....change data segment space allocation  
**bs(1)** .....a compiler/interpreter for modest-sized programs  
**bsearch(3C)** .....binary search a sorted table  
**bstring(3N)** .....bit and byte string operations  
**byteorder(3N)** .....convert values between host and network byte order  
**bzero (See bstring(3N))** .....byte string operation  
**cabs (See abs(3F))** .....Fortran absolute value  
**cal(1)** .....print calendar  
**calendar(1)** .....reminder service  
**calloc (See malloc(3C))** .....main memory allocator  
**calloc (See malloc(3X))** .....fast main memory allocator  
**cancel (See lp(1))** .....cancel requests to an LP line printer  
**cat(1)** .....concatenate and print files  
**cb(1)** .....C program beautifier  
**cc(1)** .....C compiler

## COMMANDS

**cc5.0(1)** .....C compiler (System V a.out format only)  
**ccos** (*See cos(3F)*) .....Fortran cosine intrinsic function  
**cd(1)** .....change working directory  
**cdc(1)** .....change the delta commentary of an SCCS delta  
**ceil** (*See floor(3M)*) .....ceiling function  
**cexp** (*See exp(3F)*) .....Fortran exponential intrinsic function  
**cflow(1)** .....generate C flowgraph  
**char** (*See ftype(3F)*) .....explicit Fortran type conversion  
**chargefee** (*See acctsh(1M)*) .....shell procedure for accounting  
**chase(6)** .....try to escape the killer robots  
**chdir(2)** .....change working directory  
**checkall(1M)** .....faster file system checking procedure  
**checkcw** (*See cw(1)*) .....check text prepared with CW commands  
**checkeq** (*See eqn(1)*) .....check text prepared with eqn or neqn commands  
**checklist(4)** .....list of file systems processed by fsck  
**checkmm** (*See mm(1)*) .....check documents formatted with the MM macros  
**chmod(1M)** .....change current UNIX system nodename  
**chgrp** (*See chown(1)*) .....change group  
**chmod(1)** .....change mode  
**chmod(2)** .....change mode of file  
**chown(1)** .....change owner  
**chown(2)** .....change owner and group of a file  
**chroot(1M)** .....change root directory for a command  
**chroot(2)** .....change root directory  
**ckacct** (*See acctsh(1M)*) .....shell procedure for accounting  
**clear(1)** .....clear terminal screen  
**clearerr** (*See ferror(3S)*) .....stream status inquiry  
**clock(3C)** .....report CPU time used  
**clog** (*See log(3F)*) .....Fortran natural logarithm intrinsic function  
**close(2)** .....close a file descriptor  
**closedir** (*See directory(3X)*) .....flexible length directory operation  
**clr(1M)** .....clear inode  
**cmp(1)** .....compare two files  
**cmplx** (*See ftype(3F)*) .....explicit Fortran type conversion  
**col(1)** .....filter reverse line-feeds  
**comb(1)** .....combine SCCS deltas  
**comm(1)** .....select or reject lines common to two sorted files  
**config(1M)** .....configure system  
**conj(3F)** .....Fortran complex conjugate intrinsic function

## COMMANDS

**connect(2N)** .....initiate a connection on a socket  
**conv(1)** .....object file converter  
**conv(3C)** .....translate characters  
**core(4)** .....format of core image file  
**cos (See trig(3M))** .....trigonometric function  
**cos(3F)** .....Fortran cosine intrinsic function  
**cosh (See sinh(3M))** .....hyperbolic function  
**cosh(3F)** .....Fortran hyperbolic cosine intrinsic function  
**cp(1)** .....copy files  
**cpio(1)** .....copy file archives in and out  
**cpio(4)** .....format of cpio archive  
**cpp(1)** .....the C language preprocessor  
**cpp5.0(1)** .....the C language preprocessor (System V a.out format only)  
**cpset(1M)** .....install object files in binary directories  
**craps(6)** .....the game of craps  
**crash(8)** .....what to do when the system crashes  
**creat(2)** .....create a new file or rewrite an existing one  
**cribbage(6)** .....the card game cribbage  
**cron(1M)** .....clock daemon  
**crontab(1)** .....user crontab file  
**crypt(1)** .....encode/decode  
**crypt(3C)** .....generate DES encryption  
**esh(1)** .....a shell (command interpreter) with C-like syntax  
**csin (See sin(3F))** .....Fortran sine intrinsic function  
**csplit(1)** .....context split  
**csqrt (See sqrt(3F))** .....Fortran square root intrinsic function  
**ct(1C)** .....spawn getty to a remote terminal  
**ctags(1)** .....maintain a tags file for a C program  
**ctermid(3S)** .....generate filename for terminal  
**ctime(3C)** .....convert date and time to string  
**ctrace(1)** .....C program debugger  
**ctype(3C)** .....classify characters  
**cu(1C)** .....call another UNIX system  
**cubic (See m(6))** .....tic-tac-toe  
**curses(3X)** .....CRT screen handling and optimization package  
**userid(3S)** .....get character login name of the user  
**cut(1)** .....cut out selected fields of each line of a file  
**cw(1)** .....prepare constant-width text for troff  
**cxref(1)** .....generate C program cross-reference

## COMMANDS

<b>abs</b> (See <i>abs(3F)</i> )	Fortran absolute value
<b>acos</b> (See <i>acos(3F)</i> )	Fortran arccosine intrinsic function
<b>asin</b> (See <i>asin(3F)</i> )	Fortran arcsine intrinsic function
<b>atan</b> (See <i>atan(3F)</i> )	Fortran arctangent intrinsic function
<b>atan2</b> (See <i>atan2(3F)</i> )	Fortran arctangent intrinsic function
<b>date(1)</b>	print and set the date
<b>dble</b> (See <i>fype(3F)</i> )	explicit Fortran type conversion
<b>dc(1)</b>	desk calculator
<b>dcmplx</b> (See <i>fype(3F)</i> )	explicit Fortran type conversion
<b>dconjg</b> (See <i>conjg(3F)</i> )	Fortran complex conjugate intrinsic function
<b>dcopy(1M)</b>	copy file systems for optimal access time
<b>dcopy1b(1M)</b>	copy file systems for optimal access time
<b>dcos</b> (See <i>cos(3F)</i> )	Fortran cosine intrinsic function
<b>dcosh</b> (See <i>cosh(3F)</i> )	Fortran hyperbolic cosine intrinsic function
<b>dd(1)</b>	convert and copy a file
<b>ddim</b> (See <i>dim(3F)</i> )	positive difference intrinsic function
<b>delivermail(8N)</b>	deliver mail to arbitrary people
<b>delta(1)</b>	make a delta (change) to an SCCS file
<b>deroff(1)</b>	remove <i>nroff/troff</i> , <i>tbl</i> , and <i>eqn</i> constructs
<b>dexp</b> (See <i>exp(3F)</i> )	Fortran exponential intrinsic function
<b>devnm(1M)</b>	device name
<b>df(1M)</b>	report number of free disk blocks
<b>dfscck</b> (See <i>fsck(1M)</i> )	file system consistency check and interactive repair
<b>dlal(3C)</b>	establish an out-going terminal line connection
<b>diff(1)</b>	differential file comparator
<b>diff3(1)</b>	3-way differential file comparison
<b>diffdir(1)</b>	diff directories
<b>diffmk(1)</b>	mark differences between files
<b>dim(3F)</b>	positive difference intrinsic functions
<b>dimag</b> (See <i>aimag(3F)</i> )	Fortran imaginary part of complex argument
<b>dint</b> (See <i>aint(3F)</i> )	Fortran integer part intrinsic function
<b>dir(4)</b>	format of directories
<b>dircmp(1)</b>	directory comparison
<b>directory(3X)</b>	flexible length directory operations
<b>dirname</b> (See <i>basename(1)</i> )	deliver portions of path names
<b>dis(1)</b>	disassembler
<b>disable</b> (See <i>enable(1)</i> )	disable LP printers
<b>diskformat(1M)</b>	format a disk
<b>disktune(1M)</b>	tune floppy disk settling time parameters



## COMMANDS

**disksg(1M)** .....generate disk accounting data by user ID  
**dlog** (See *log(3F)*) .....Fortran natural logarithm intrinsic function  
**dlog10** (See *log10(3F)*) .....Fortran common logarithm intrinsic function  
**dmax1** (See *max(3F)*) .....Fortran maximum-value function  
**dmin1** (See *min(3F)*) .....Fortran minimum-value function  
**dmod** (See *mod(3F)*) .....Fortran remaindering intrinsic function  
**drint** (See *round(3F)*) .....Fortran nearest integer function  
**do**disk (See *acctsh(1M)*) .....shell procedure for accounting  
**dprod(3F)** .....double precision product intrinsic function  
**drand48(3C)** .....generate uniformly distributed pseudo-random numbers  
**dsign** (See *sign(3F)*) .....Fortran transfer-of-sign intrinsic function  
**dsla** (See *sin(3F)*) .....Fortran sine intrinsic function  
**dsinh** (See *sinh(3F)*) .....Fortran hyperbolic sine intrinsic function  
**dsqrt** (See *sqr(3F)*) .....Fortran square root intrinsic function  
**dtan** (See *tan(3F)*) .....Fortran tangent intrinsic function  
**dtanh** (See *tanh(3F)*) .....Fortran hyperbolic tangent intrinsic function  
**du(1)** .....summarize disk usage  
**dump(1)** .....dump selected parts of an object file  
**dup(3)** .....duplicate a descriptor  
**dup2(3N)** .....duplicate a descriptor  
**echo(1)** .....echo arguments  
**ecvt(3C)** .....convert floating-point number to string  
**ed(1)** .....text editor  
**edata** (See *end(3C)*) .....last locations in program  
**edit** (See *ex(1)*) .....text editor  
**efl(1)** .....Extended Fortran Language  
**egrep** (See *grep(1)*) .....search a file for a pattern  
**enable(1)** .....enable LP printers  
**encrypt** (See *crypt(3C)*) .....generate DES encryption  
**end(3C)** .....last locations in program  
**endgrent** (See *getgrent(3C)*) .....obtain group file entry from a group file  
**endhostent** (See *gethostent(3N)*) .....get network host entry  
**endnetent** (See *getnetent(3N)*) .....get network entry  
**endprotoent** (See *getprotoent(3N)*) .....get protocol entry  
**endpwent** (See *getpwent(3C)*) .....get password file entry  
**endservent** (See *getservent(3N)*) .....get service entry  
**endutent** (See *getut(3C)*) .....access utmp file entry  
**env(1)** .....set environment for command execution  
**environ(5)** .....user environment

## COMMANDS

**eqn(1)** .....format mathematical text for troff  
**eqnchar(5)** .....special character definitions for eqn and neqn  
**erand48 (See drand48(3C))** .....generate uniformly distributed pseudo-random numbers  
**erf(3M)** .....error function  
**erfc (See erf(3M))** .....complementary error function  
**errdead(1M)** .....extract error records from dump  
**errdemon(1M)** .....error-logging daemon  
**errfile(4)** .....error-log file format  
**errno (See perror(3C))** .....system error message  
**error(7)** .....error-logging interface  
**errpt(1M)** .....process a report of logged errors  
**errstop(1M)** .....terminate the error-logging daemon  
**etext (See end(3C))** .....last locations in program  
**ex(1)** .....text editor  
**exec(2)** .....execute a file  
**execl (See exec(2))** .....execute a file  
**execle (See exec(2))** .....execute a file  
**execlp (See exec(2))** .....execute a file  
**execv (See exec(2))** .....execute a file  
**execve (See exec(2))** .....execute a file  
**execvp (See exec(2))** .....execute a file  
**exit(2)** .....terminate process  
**exp(3F)** .....Fortran exponential intrinsic function  
**exp(3M)** .....exponential function  
**expr(1)** .....evaluate arguments as an expression  
**exterr(1)** .....turn on/off the extended errors in the specified device  
**f77(1)** .....Fortran 77 compiler  
**fabs (See floor(3M))** .....absolute value function  
**factor(1)** .....factor a number  
**false (See true(1))** .....provide truth values  
**fclose(3S)** .....close a stream  
**fcntl(2)** .....file control  
**fcntl(5)** .....file control options  
**fevt (See ecvt(3C))** .....convert floating-point number to string  
**fdopen (See fopen(3S))** .....open a stream  
**feof (See ferror(3S))** .....stream status inquiry  
**ferror(3S)** .....stream status inquiry  
**ff(1M)** .....list file names and statistics for a file system  
**flush (See fclose(3S))** .....flush a stream

## COMMANDS

**fts** (See *bstring(3N)*) ..... bit string operation

**fgetc** (See *getc(3S)*) ..... get character from a stream

**fgetgrent** (See *getgrent(3C)*) ..... obtain group file entry from a group file

**fgetpwent** (See *getpwent(3C)*) ..... get password file entry

**fgets** (See *gets(3S)*) ..... get a string from a stream

**fgrep** (See *grep(1)*) ..... search a file for a pattern

**file(1)** ..... determine file type

**filehdr(4)** ..... file header for common object files

**fileno** (See *ferror(3S)*) ..... stream status inquiry

**filesave(1M)** ..... daily/weekly UNIX file system backup

**find(1M)** ..... fast incremental backup

**find(1)** ..... find files

**fish(6)** ..... play "Go Fish"

**float** (See *ftype(3F)*) ..... explicit Fortran type conversion

**floor(3M)** ..... floor function

**fmod** (See *floor(3M)*) ..... remainder function

**fopen(3S)** ..... open a stream

**fork(2)** ..... create a new process

**fortune(6)** ..... print a random, hopefully interesting, adage

**printf** (See *printf(3S)*) ..... print formatted output

**fputc** (See *putc(3S)*) ..... put character on a stream

**fputs** (See *puts(3S)*) ..... put a string on a stream

**fread(3S)** ..... binary input

**free(1M)** ..... recover files from a backup tape

**free** (See *malloc(3C)*) ..... main memory allocator

**free** (See *malloc(3X)*) ..... fast main memory allocator

**freopen** (See *fopen(3S)*) ..... open a stream

**freq(1)** ..... report on character frequencies in a file

**frexp(3C)** ..... manipulate parts of floating-point numbers

**fs(4)** ..... format of system volume

**fsconf** (See *scanf(3S)*) ..... convert formatted input

**fsck(1M)** ..... file system consistency check and interactive repair

**fscv(1M)** ..... convert files between M68000 and VAX-11/780 processors

**fsdb(1M)** ..... file system debugger

**fseek(3S)** ..... reposition a file pointer in a stream

**fspec(4)** ..... format specification in text files

**fsplit(1)** ..... split f77, ratfor, or efl files

**fstat** (See *stat(2)*) ..... get file status

**ftell** (See *fseek(3S)*) ..... reposition a file pointer in a stream

## COMMANDS

<b>ftok</b> (See <i>stdipc(3C)</i> )	.....	standard interprocess communication package
<b>ftp(1N)</b>	.....	file transfer program
<b>ftpd(8N)</b>	.....	DARPA Internet File Transfer Protocol server
<b>ftw(3C)</b>	.....	walk a file tree
<b>ftype(3F)</b>	.....	explicit Fortran type conversion
<b>fuser(1M)</b>	.....	identify processes using a file or file structure
<b>fwrite</b> (See <i>fwread(3S)</i> )	.....	binary output
<b>fwtmp(1M)</b>	.....	manipulate connect accounting records
<b>gamma(3M)</b>	.....	log gamma function
<b>gcvt</b> (See <i>ecvt(3C)</i> )	.....	convert floating-point number to string
<b>get(1)</b>	.....	get a version of an SCCS file
<b>getarg(3F)</b>	.....	return Fortran command-line argument
<b>getc(3S)</b>	.....	get character from a stream
<b>getchar</b> (See <i>getc(3S)</i> )	.....	get character from a stream
<b>getcwd(3C)</b>	.....	get pathname of current working directory
<b>getdtablesize(3N)</b>	.....	get descriptor table size
<b>getgid</b> (See <i>getuid(2)</i> )	.....	get effective group ID
<b>getenv(3C)</b>	.....	return value for environment name
<b>getenv(3F)</b>	.....	return Fortran environment variable
<b>geteuid</b> (See <i>getuid(2)</i> )	.....	get effective user ID
<b>getgid</b> (See <i>getuid(2)</i> )	.....	get real group ID
<b>getgrent(3C)</b>	.....	obtain group file entry from a group file
<b>getgrgid</b> (See <i>getgrent(3C)</i> )	.....	obtain group file entry from a group file
<b>getgrnam</b> (See <i>getgrent(3C)</i> )	.....	obtain group file entry from a group file
<b>gethostbyaddr</b> (See <i>gethostent(3N)</i> )	.....	get network host entry
<b>gethostbyname</b> (See <i>gethostent(3N)</i> )	.....	get network host entry
<b>gethostent(3N)</b>	.....	get network host entry
<b>gethostid(2N)</b>	.....	get unique identifier of current host
<b>gethostname(2N)</b>	.....	get name of current host
<b>getlogin(3C)</b>	.....	get login name
<b>getnetbyaddr</b> (See <i>getnetent(3N)</i> )	.....	get network entry
<b>getnetbyname</b> (See <i>getnetent(3N)</i> )	.....	get network entry
<b>getnetent(3N)</b>	.....	get network entry
<b>getopt(1)</b>	.....	parse command options
<b>getopt(3C)</b>	.....	get option letter from argument vector
<b>getpass(3C)</b>	.....	read a password
<b>getpeername(2N)</b>	.....	get name of connected peer
<b>getppgrp</b> (See <i>getpid(2)</i> )	.....	get process group ID
<b>getpid(2)</b>	.....	get process ID

## COMMANDS

**getppid** (See *getpid(2)*) .....get parent process ID  
**getprotobyname** (See *getprotoent(3N)*) .....get protocol entry  
**getprotobyname** (See *getprotoent(3N)*) .....get protocol entry  
**getprotoent(3N)** .....get protocol entry  
**getpw(3C)** .....get name from UID  
**getpwent(3C)** .....get password file entry  
**getpwnam** (See *getpwent(3C)*) .....get password file entry  
**getpwuid** (See *getpwent(3C)*) .....get password file entry  
**gets(3S)** .....get a string from a stream  
**getservbyname** (See *getservent(3N)*) .....get service entry  
**getservbyport** (See *getservent(3N)*) .....get service entry  
**getservent(3N)** .....get service entry  
**getsockname(2N)** .....get socket name  
**getsockopt(2N)** .....get options on sockets  
**getty(1M)** .....set terminal type, modes, speed, and line discipline  
**gettydefs(4)** .....speed and terminal settings used by *getty*  
**getuid(2)** .....get real user ID  
**getut(3C)** .....access utmp file entry  
**getutent** (See *getut(3C)*) .....access utmp file entry  
**getutid** (See *getut(3C)*) .....access utmp file entry  
**getutline** (See *getut(3C)*) .....access utmp file entry  
**getw** (See *getc(3S)*) .....get word from a stream  
**gmtime** (See *ctime(3C)*) .....convert date and time to string  
**graph(1G)** .....draw a graph  
**greek(1)** .....select terminal filter  
**greek(5)** .....graphics for the extended TTY-37 type-box  
**grep(1)** .....search a file for a pattern  
**group(4)** .....group file  
**grpck** (See *pwck(1M)*) .....group file checker  
**gsignal** (See *ssignal(3C)*) .....software signal  
**hangman(6)** .....guess the word  
**hashcheck** (See *spell(1)*) .....work with the *spell* program's hash lists  
**hashmake spell(1)** .....work with the *spell* program's hash lists  
**hcreate** (See *hsearch(3C)*) .....manage hash search tables  
**hdestroy** (See *hsearch(3C)*) .....manage hash search tables  
**head(1)** .....give first few lines  
**help(1)** .....ask for help in using SCCS  
**hex(1)** .....translates object files  
**hostid(1N)** .....set or print identifier of current host system

## COMMANDS

**hostname(1N)** .....set or print name of current host system  
**hosts(4N)** .....host name data base  
**hsearch(3C)** .....manage hash search tables  
**htonl (See byteorder(3N))** .....convert values between host and network byte order  
**htons (See byteorder(3N))** .....convert values between host and network byte order  
**hyphen(1)** .....find hyphenated words  
**hypot(3M)** .....Euclidean distance function  
**iabs (See abs(3F))** .....Fortran absolute value  
**inrc(3F)** .....count command line arguments  
**ichar (See ftype(3F))** .....explicit Fortran type conversion  
**id(1)** .....print user and group IDs and names  
**idim (See dim(3F))** .....positive difference intrinsic function  
**idint (See ftype(3F))** .....explicit Fortran type conversion  
**idmint (See round(3F))** .....Fortran nearest integer function  
**ifconfig(8N)** .....configure network interface parameters  
**ifix (See ftype(3F))** .....explicit Fortran type conversion  
**index(3F)** .....return location of Fortran substring  
**inet(3N)** .....Internet address manipulation routines  
**inet(5F)** .....Internet protocol family  
**inet\_addr (See inet(3N))** .....Internet address manipulation routine  
**inet\_inaddr (See inet(3N))** .....Internet address manipulation routine  
**inet\_makeaddr (See inet(3N))** .....Internet address manipulation routine  
**inet\_netof (See inet(3N))** .....Internet address manipulation routine  
**inet\_network (See inet(3N))** .....Internet address manipulation routine  
**inet\_ntoa (See inet(3N))** .....Internet address manipulation routine  
**init(1M)** .....process control initialization  
**inittab(4)** .....script for the init process  
**inode(4)** .....format of an inode  
**insque(3N)** .....insert element from a queue  
**install(1M)** .....install commands  
**int (See ftype(3F))** .....explicit Fortran type conversion  
**ioctl(2)** .....control device  
**ip(5P)** .....Internet Protocol  
**lperm(1)** .....remove a message queue, semaphore set or shared memory id  
**lps(1)** .....report inter-process communication facilities status  
**irand (See rand(3F))** .....Fortran uniform random-number generator  
**isalnum (See ctype(3C))** .....classify characters  
**isalpha (See ctype(3C))** .....classify characters  
**isascii (See ctype(3C))** .....classify characters

## COMMANDS

<b>Isatty</b> ( <i>See ttyname(3C)</i> )	.....	find name of a terminal
<b>iscntrl</b> ( <i>See ctype(3C)</i> )	.....	classify characters
<b>isdigit</b> ( <i>See ctype(3C)</i> )	.....	classify characters
<b>isgraph</b> ( <i>See ctype(3C)</i> )	.....	classify characters
<b>isign</b> ( <i>See sign(3F)</i> )	.....	Fortran transfer-of-sign intrinsic function
<b>islower</b> ( <i>See ctype(3C)</i> )	.....	classify characters
<b>isprint</b> ( <i>See ctype(3C)</i> )	.....	classify characters
<b>ispunct</b> ( <i>See ctype(3C)</i> )	.....	classify characters
<b>isspace</b> ( <i>See ctype(3C)</i> )	.....	classify characters
<b>issue</b> (4)	.....	issue identification file
<b>isupper</b> ( <i>See ctype(3C)</i> )	.....	classify characters
<b>isxdigit</b> ( <i>See ctype(3C)</i> )	.....	classify characters
<b>j0</b> ( <i>See bessel(3M)</i> )	.....	Bessel function
<b>j1</b> ( <i>See bessel(3M)</i> )	.....	Bessel function
<b>jn</b> ( <i>See bessel(3M)</i> )	.....	Bessel function
<b>join</b> (1)	.....	relational database operator
<b>rand48</b> ( <i>See drand48(3C)</i> )	.....	generate uniformly distributed pseudo-random numbers
<b>kill</b> (1)	.....	terminate a process
<b>kill</b> (2)	.....	send a signal to a process or a group of processes
<b>killall</b> (1M)	.....	kill all active processes
<b>killpg</b> (3N)	.....	send signal to a process group
<b>kmem</b> ( <i>See mem(7)</i> )	.....	core memory
<b>l3tol</b> (3C)	.....	convert between 3-byte integers and long integers
<b>l64a</b> ( <i>See a64i(3C)</i> )	.....	convert between long integer and base-64 ASCII string
<b>labelit</b> ( <i>See volcopy(1M)</i> )	.....	copy file systems with label checking
<b>last</b> (1)	.....	indicate last logins of users and teletypes
<b>lastlogin</b> ( <i>See acctsh(1M)</i> )	.....	shell procedure for accounting
<b>lav</b> (1)	.....	print load average statistics
<b>lcong48</b> ( <i>See drand48(3C)</i> )	.....	generate uniformly distributed pseudo-random numbers
<b>ld</b> (1)	.....	link editor for common object files
<b>ld5.0</b> (1)	.....	link editor (System V a.out format only)
<b>ldaclose</b> ( <i>See ldclose(3X)</i> )	.....	close a common object file
<b>ldahread</b> (3X)	.....	read the archive header of a member of an archive file
<b>ldlopen</b> ( <i>See ldopen(3X)</i> )	.....	open a common object file for reading
<b>ldclose</b> (3X)	.....	close a common object file
<b>ldexp</b> ( <i>See frexp(3C)</i> )	.....	manipulate parts of floating-point numbers
<b>ldfcn</b> (4)	.....	common object file access routines
<b>ldhread</b> (3X)	.....	read the file header of a common object file
<b>ldgetname</b> (3X)	.....	retrieve symbol name for object file

## COMMANDS

**ldlinit** (See *ldhread(3X)*) ...manipulate line number entries of a common object file function  
**ldlitem** (See *ldhread(3X)*) ...manipulate line number entries of a common object file function  
**ldhread(3X)** .....manipulate line number entries of a common object file function  
**ldlseek(3X)** .....seek to line number entries of a section of a common object file  
**ldnrseek** (See *ldlseek(3X)*) seek to line number entries of a section of a common object file  
**ldnrseek** (See *ldrseek(3X)*) ...seek to relocation entries of a section of a common object file  
**ldnshread** (See *ldshread(3X)*) read an indexed/named section header of a common object file  
**ldnsseek** (See *ldsseek(3X)*) .....seek to an indexed/named section of a common object file  
**ldohseek(3X)** .....seek to the optional file header of a common object file  
**ldopen(3X)** .....open a common object file for reading  
**ldrseek(3X)** .....seek to relocation entries of a section of a common object file  
**ldshread(3X)** .....read an indexed/named section header of a common object file  
**ldsseek(3X)** .....seek to an indexed/named section of a common object file  
**ldtbindex(3X)** .....compute the index of a symbol table entry of a common object file  
**ldtbread(3X)** .....read an indexed symbol table entry of a common object file  
**ldtbseek(3X)** .....seek to the symbol table of a common object file  
**len(3F)** .....return length of Fortran string  
**lex(1)** .....generate programs for simple lexical tasks  
**lfind** (See *lsearch(3C)*) .....linear search and update  
**lge** (See *strcmp(3F)*) .....string comparison intrinsic function  
**lgt** (See *strcmp(3F)*) .....string comparison intrinsic function  
**life(6)** .....play the game of life  
**line(1)** .....read one line  
**linenum(4)** .....line number entries in a common object file  
**link(1M)** .....exercise link system call  
**link(2)** .....link to a file  
**lint(1)** .....a C program checker  
**listen(2N)** .....listen for connections on a socket  
**lle** (See *strcmp(3F)*) .....string comparison intrinsic function  
**lft** (See *strcmp(3F)*) .....string comparison intrinsic function  
**ln** (See *cp(1)*) .....link files  
**lo(5)** .....software loopback network interface  
**localtime** (See *ctime(3C)*) .....convert date and time to string  
**lockf(3C)** .....record locking on files  
**locking(2)** .....provide exclusive file regions for reading or writing  
**log** (See *exp(3M)*) .....logarithm function  
**log(3F)** .....Fortran natural logarithm intrinsic function  
**log10** (See *exp(3M)*) .....logarithm function  
**log10(3F)** .....Fortran common logarithm intrinsic function



## COMMANDS

**login(1)**.....sign on  
**logname(1)**.....get login name  
**logname(3X)**.....return login name of user  
**longjmp (See setjmp(3C))**.....non-local goto  
**lorder(1)**.....find ordering relation for an object library  
**lorder5.0(1)**.....find ordering relation for an object library (System V a.out format only)  
**lp(1)**.....send requests to an LP line printer  
**lpadmin(1M)**.....configure the LP spooling system  
**lpmove (See lpsched(1M))**.....move LP requests  
**lpsched(1M)**.....start the LP request scheduler  
**lpshut (See lpsched(1M))**.....stop the LP request scheduler  
**lpstat(1)**.....print LP status information  
**lrand48 (See drand48(3C))**.....generate uniformly distributed pseudo-random numbers  
**ls(1)**.....list contents of directory  
**lsearch(3C)**.....linear search and update  
**lseek(2)**.....move read/write file pointer  
**lshlft (See bool(3F))**.....Fortran bitwise boolean function  
**lto13 (See l3to1(3C))**.....convert between 3-byte integers and long integers  
**m4(1)**.....macro processor  
**m68k (See machid(1))**.....provide truth value about your processor type  
**machid(1)**.....provide truth value about your processor type  
**mail(1)**.....send mail to users or read mail  
**mailx(1)**.....interactive message processing system  
**make(1)**.....maintain, update, and regenerate groups of programs  
**makekey(1)**.....generate encryption key  
**mailinfo (See malloc(3X))**.....fast main memory allocator  
**malloc(3C)**.....main memory allocator  
**malloc(3X)**.....fast main memory allocator  
**mallopt (See malloc(3X))**.....fast main memory allocator  
**man(1)**.....print entries in this manual  
**man(5)**.....macros for formatting entries in this manual  
**master(4)**.....master device information table  
**math(5)**.....math functions and constants  
**matherr(3M)**.....error-handling function  
**max(3F)**.....Fortran maximum-value function  
**max0 (See max(3F))**.....Fortran maximum-value function  
**max1 (See max(3F))**.....Fortran maximum-value function  
**maze(6)**.....generate a maze  
**mc68cc(1)**.....C compiler

## COMMANDS

<b>mclock(3F)</b> .....	return Fortran time accounting
<b>mem(7)</b> .....	core memory
<b>memccpy (See memory(3C))</b> .....	memory operation
<b>memchr (See memory(3C))</b> .....	memory operation
<b>memcmp (See memory(3C))</b> .....	memory operation
<b>memcpy (See memory(3C))</b> .....	memory operation
<b>memory(3C)</b> .....	memory operation
<b>memset (See memory(3C))</b> .....	memory operation
<b>msg(1)</b> .....	permit or deny messages
<b>min(3F)</b> .....	Fortran minimum-value function
<b>min0 (See min(3F))</b> .....	Fortran minimum-value function
<b>min1 (See min(3F))</b> .....	Fortran minimum-value function
<b>mkdir(1)</b> .....	make a directory
<b>mkfs(1M)</b> .....	construct a file system
<b>mkfs1b(1M)</b> .....	construct a file system
<b>mklost+found(1M)</b> .....	make a lost+found directory for fsck
<b>mknod(1M)</b> .....	build special file
<b>mknod(2)</b> .....	make a directory, or a special or ordinary file
<b>mkstr(1)</b> .....	create an error message file by massaging C source
<b>mktemp(3C)</b> .....	make a unique filename
<b>mm(1)</b> .....	print documents formatted with the MM macros
<b>mm(5)</b> .....	the MM macro package for formatting documents
<b>mmt(1)</b> .....	typeset documents
<b>mnttab(4)</b> .....	mounted file system table
<b>mod(3F)</b> .....	Fortran remaindering intrinsic function
<b>modf (See frexp(3C))</b> .....	manipulate parts of floating-point numbers
<b>monacct (See acctsh(1M))</b> .....	shell procedure for accounting
<b>monitor(3C)</b> .....	prepare execution profile
<b>moo(6)</b> .....	guessing game
<b>more(1)</b> .....	file perusal filter for crt viewing
<b>mosd(5)</b> .....	the OSDD adapter macro package for formatting documents
<b>mount(1M)</b> .....	mount file system
<b>mount(2)</b> .....	mount a file system
<b>mpix(5)</b> .....	the macro package for formatting a permuted index
<b>mrnd48 (See drand48(3C))</b> .....	generate uniformly distributed pseudo-random numbers
<b>msgctl(2)</b> .....	message control operations
<b>msgget(2)</b> .....	get message queue
<b>msgop(2)</b> .....	message operations
<b>mv (See cp(1))</b> .....	move files

## COMMANDS

**mv(5)** .....a troff macro package for typesetting view graphs and slides  
**mvdir(1M)** .....move a directory  
**mvf (See mmt(1))** .....typeset view graphs and slides  
**ncheck(1M)** .....generate names from inumbers  
**neqn (See eqn(1))** .....format mathematical text for nroff  
**netmail(8N)** .....the B-NET network mail system  
**netmailer(8N)** .....deliver mail to B-NET  
**netstat(1N)** .....show network status  
**networks(4N)** .....network name data base  
**newform(1)** .....change the format of a text file  
**newgrp(1)** .....log in to a new group  
**news(1)** .....print news items  
**nice(1)** .....run a command at low priority  
**nice(2)** .....change priority of a process  
**nint (See round(3F))** .....Fortran nearest integer function  
**nl(1)** .....line numbering filter  
**nlist(3C)** .....get entries from name list  
**nm(1)** .....print name list of common object file  
**nm5.0(1)** .....print name list (System V a.out format only)  
**nohup(1)** .....run a command immune to hangups (sh only)  
**not (See bool(3F))** .....Fortran bitwise boolean function  
**nrand48 (See drand48(3C))** .....generate uniformly distributed pseudo-random numbers  
**nroff(1)** .....format text  
**ntohl (See byteorder(3N))** .....convert values between host and network byte order  
**ntohs (See byteorder(3N))** .....convert values between host and network byte order  
**null(7)** .....the null file  
**nulladm (See acctsh(1M))** .....shell procedure for accounting  
**number(6)** .....convert Arabic numerals to English  
**od(1)** .....octal dump  
**open(2)** .....open for reading or writing  
**opendir (See directory(3X))** .....flexible length directory operation  
**or (See bool(3F))** .....Fortran bitwise boolean function  
**osdd (See mm(1))** .....print documents formatted with the MM and OSDD macros  
**pack(1)** .....compress files  
**passwd(1)** .....change login password  
**passwd(4)** .....password file  
**paste(1)** .....merge same lines of several files or subsequent lines of one file  
**pause(2)** .....suspend process until signal  
**pcat (See pack(1))** .....expand compressed files

## COMMANDS

**pclose** (See *popen(3S)*) .....initiate pipe to/from a process  
**pdp11** (See *machid(1)*) .....provide truth value about your processor type  
**perror(3C)** .....system error message  
**pg(1)** .....file perusal filter for soft-copy terminals  
**phys(2)** .....allow a process to access physical addresses  
**pipe(2)** .....create an interprocess channel  
**plock(2)** .....lock process, text, or data in memory  
**plot(3X)** .....graphics interface subroutines  
**plot(4)** .....graphics interface  
**pnch(4)** .....file format for card images  
**popen(3S)** .....initiate pipe to/from a process  
**pow** (See *exp(3M)*) .....power function  
**powerfail** (See *brc(1M)*) .....system initialization shell script  
**pr(1)** .....print files  
**prctmp** (See *acctsh(1M)*) .....shell procedure for accounting  
**prdaily** (See *acctsh(1M)*) .....shell procedure for accounting  
**printenv(1)** .....print out the environment  
**printf(3S)** .....print formatted output  
**prof(1)** .....display profile data  
**prof(5)** .....profile within a function  
**profil(2)** .....execution time profile  
**profile(4)** .....setting up an environment at login time  
**protocols(4N)** .....protocol name data base  
**prs(1)** .....print an SCCS file  
**prtacct** (See *acctsh(1M)*) .....shell procedure for accounting  
**ps(1)** .....report process status  
**pstat(1M)** .....print system facts  
**ptrace(2)** .....process trace  
**ptx(1)** .....permuted index  
**pty(5)** .....pseudo terminal driver  
**put(1C)** .....puts a file onto a remote machine  
**putc(3S)** .....put character on a stream  
**putchar** (See *putc(3S)*) .....put character on a stream  
**putenv(3C)** .....change or add value to environment  
**putpwent(3C)** .....write password file entry  
**puts(3S)** .....put a string on a stream  
**pututline** (See *getut(3C)*) .....access utmp file entry  
**putw** (See *putc(3S)*) .....put word on a stream  
**pwck(1M)** .....password file checker

## COMMANDS

<b>pwd(1)</b> .....	working directory name
<b>qsort(3C)</b> .....	quicker sort
<b>quiz(6)</b> .....	test your knowledge
<b>rain(6)</b> .....	animated raindrops display
<b>rand(3C)</b> .....	simple random-number generator
<b>rand(3F)</b> .....	Fortran uniform random-number generator
<b>ratfor(1)</b> .....	rational Fortran dialect
<b>rc</b> (See <i>brc(1M)</i> ) .....	system initialization shell script
<b>remd(3N)</b> .....	routine for returning a stream to a remote command
<b>rcp(1N)</b> .....	remote file copy
<b>rcvhex(1)</b> .....	translates Motorola S-records from downloading into a file
<b>read(2)</b> .....	read from file
<b>readdir</b> (See <i>directory(3X)</i> ) .....	flexible length directory operation
<b>readv(3N)</b> .....	read from file
<b>real</b> (See <i>ftype(3F)</i> ) .....	explicit Fortran type conversion
<b>realloc</b> (See <i>malloc(3C)</i> ) .....	main memory allocator
<b>realloc</b> (See <i>malloc(3X)</i> ) .....	fast main memory allocator
<b>reboot(1M)</b> .....	reboot the system
<b>reboot(2)</b> .....	reboot the system
<b>recv(2N)</b> .....	receive a message from a socket
<b>recvfrom</b> (See <i>recv(2N)</i> ) .....	receive a message from a socket
<b>recvmsg</b> (See <i>recv(2N)</i> ) .....	receive a message from a socket
<b>red</b> (See <i>ed(1)</i> ) .....	text editor
<b>regcmp(1)</b> .....	regular expression compile
<b>regcmp(3X)</b> .....	compile a regular expression
<b>regex</b> (See <i>regcmp(3X)</i> ) .....	execute a regular expression
<b>regexp(5)</b> .....	regular expression compile and match routines
<b>reject</b> (See <i>accept(1M)</i> ) .....	prevent LP requests
<b>reloc(4)</b> .....	relocation information for a common object file
<b>remque</b> (See <i>insque(3N)</i> ) .....	remove element from a queue
<b>remsh(1N)</b> .....	remote shell
<b>remshd(8N)</b> .....	remote shell server
<b>reset</b> (See <i>iset(1)</i> ) .....	reset the teletype bits to a sensible state
<b>rewind</b> (See <i>seek(3S)</i> ) .....	reposition a file pointer in a stream
<b>rewinddir</b> (See <i>directory(3X)</i> ) .....	flexible length directory operation
<b>rexec(3N)</b> .....	return stream to a remote command
<b>rexecd(8N)</b> .....	remote execution server
<b>rlogin(1N)</b> .....	remote login
<b>rlogind(8N)</b> .....	remote login server

## COMMANDS

**rm** (1) .....remove files  
**rmail** (See *mail(1)*) .....send mail to users or read mail  
**rmdel** (1) .....remove a delta from an SCCS file  
**rmdir** (See *rm(1)*) .....remove directories  
**robots** (6) .....escape from the robots  
**round** (3F) .....Fortran nearest integer functions  
**route** (8N) .....manually manipulate the routing tables  
**routed** (8N) .....network routing daemon  
**rresvport** (See *rcmd(3N)*) .....routine for returning a stream to a remote command  
**rsh** (See *sh(1)*) .....shell, the restricted command programming language  
**rshift** (See *bool(3F)*) .....Fortran bitwise boolean function  
**runacct** (1M) .....run daily accounting  
**ruptime** (1N) .....show host status of local machines  
**ruserok** (See *rcmd(3N)*) .....routine for returning a stream to a remote command  
**rwho** (1N) .....who's logged in on local machines  
**rwhod** (8N) .....system status server  
**sa1** (See *sar(1M)*) .....system activity report package  
**sa2** (See *sar(1M)*) .....system activity report package  
**sact** (1) .....print current SCCS file editing activity  
**sade** (See *sar(1M)*) .....system activity report package  
**sag** (1G) .....system activity graph  
**sar** (1) .....system activity reporter  
**sar** (1M) .....system activity report package  
**sbrk** (See *brk(2)*) .....change data segment space allocation  
**scanf** (3S) .....convert formatted input  
**scsdiff** (1) .....compare two versions of an SCCS file  
**scsfile** (4) .....format of SCCS file  
**scnhdr** (4) .....section header for a common object file  
**script** (1) .....make typescript of terminal session  
**sdb** (1) .....symbolic debugger  
**sdiff** (1) .....side-by-side difference program  
**sed** (1) .....stream editor  
**seed48** (See *drand48(3C)*) .....generate uniformly distributed pseudo-random numbers  
**seekdir** (See *directory(3X)*) .....flexible length directory operation  
**select** (2N) .....synchronous I/O multiplexing  
**semctl** (2) .....semaphore control operations  
**semget** (2) .....get set of semaphores  
**semop** (2) .....semaphore operation  
**send** (2N) .....send a message from a socket

## COMMANDS

**sendmsg** (*See send(2N)*) ..... send a message from a socket  
**sendto** (*See send(2N)*) ..... send a message from a socket  
**services(4N)** ..... service name data base  
**setbuf(3S)** ..... assign buffering to a stream  
**setgid** (*See setuid(2)*) ..... set group ID  
**setgrent** (*See getgrent(3C)*) ..... obtain group file entry from a group file  
**sethostent** (*See gethostent(3N)*) ..... get network host entry  
**sethostid** (*See gethostid(2N)*) ..... set unique identifier of current host  
**sethostname** (*See gethostname(2N)*) ..... set name of current host  
**setjmp(3C)** ..... non-local goto  
**setkey** (*See crypt(3C)*) ..... generate DES encryption  
**setmnt(1M)** ..... establish mount table  
**setnetent** (*See getnetent(3N)*) ..... get network entry  
**setpggrp(2)** ..... set process group ID  
**setprotoent** (*See getprotoent(3N)*) ..... get protocol entry  
**setpwent** (*See getpwent(3C)*) ..... get password file entry  
**setregid(2)** ..... set real and effective group ID  
**setreuid(2)** ..... set real and effective user IDs  
**setservent** (*See getservent(3N)*) ..... get service entry  
**setsockopt** (*See getsockopt(2N)*) ..... set options on sockets  
**setuid(2)** ..... set user ID  
**setutent** (*See getut(3C)*) ..... access utmp file entry  
**setvbuf** (*See setbuf(3S)*) ..... assign buffering to a stream  
**sgctl** (*See spctl(3X)*) ..... access long integer data in a machine independent fashion  
**sh(1)** ..... shell, the standard command programming language  
**shl(1)** ..... shell layer manager  
**shmctl(2)** ..... shared memory control operations  
**shmget(2)** ..... get shared memory segment  
**shmop(2)** ..... shared memory operations  
**shutacct** (*See acctsh(1M)*) ..... shell procedure for accounting  
**shutdown(1M)** ..... terminate all processing  
**shutdown(2N)** ..... shut down part of a full-duplex connection  
**sign(3F)** ..... Fortran transfer-of-sign intrinsic function  
**signal(2)** ..... specify what to do upon receipt of a signal  
**signal(3F)** ..... specify Fortran action on receipt of a system signal  
**sin** (*See trig(3M)*) ..... trigonometric function  
**sin(3F)** ..... Fortran sine intrinsic function  
**sinh(3F)** ..... Fortran hyperbolic sine intrinsic function  
**sinh(3M)** ..... hyperbolic function

## COMMANDS

**size(1)** .....print section sizes of common object files  
**size5.0(1)** .....size of an object file (System V a.out format only)  
**sleep(1)** .....suspend execution for an interval  
**sleep(3C)** .....suspend execution for interval  
**sngl (See ftype(3F))** .....explicit Fortran type conversion  
**sno(1)** .....SNOBOL interpreter  
**socket(2N)** .....create an endpoint for communication  
**sort(1)** .....sort and/or merge files  
**spell(1)** .....find spelling errors  
**spellin (See spell(1))** .....work with the spell program's hash lists  
**spline(1G)** .....interpolate smooth curve  
**split(1)** .....split a file into pieces  
**sprintf (See printf(3S))** .....print formatted output  
**sputl(3X)** .....access long integer data in a machine independent fashion  
**sqrt (See exp(3M))** .....square root function  
**sqrt(3F)** .....Fortran square root intrinsic function  
**srand (See rand(3C))** .....simple random-number generator  
**rand (See rand(3F))** .....Fortran uniform random-number generator  
**rand48 (See drand48(3C))** .....generate uniformly distributed pseudo-random numbers  
**sscanf (See scanf(3S))** .....convert formatted input  
**ssignal(3C)** .....software signal  
**ssp(1)** .....make output single spaced  
**startup (See acctsh(1M))** .....shell procedure for accounting  
**stat(2)** .....get file status  
**stat(5)** .....data returned by stat system call  
**stdio(3S)** .....standard buffered input/output package  
**stdipe(3C)** .....standard interprocess communication package  
**stime(2)** .....set time  
**streat (See string(3C))** .....string operation  
**strchr (See string(3C))** .....string operation  
**strcmp (See string(3C))** .....string operation  
**strcmp(3F)** .....string comparison intrinsic function  
**streqy (See string(3C))** .....string operation  
**strcsyn (See string(3C))** .....string operation  
**string(3C)** .....string operation  
**strings(1)** .....find the printable strings in an object, or other binary file  
**strip(1)** .....strip symbol and line number information from an object file  
**strip5.0(1)** .....remove symbols and relocation bits (System V a.out format only)  
**strlen (See string(3C))** .....string operation



## COMMANDS

**strncat** (See *string(3C)*) ..... string operation  
**strncmp** (See *string(3C)*) ..... string operation  
**strncpy** (See *string(3C)*) ..... string operation  
**strpbrk** (See *string(3C)*) ..... string operation  
**strrchr** (See *string(3C)*) ..... string operation  
**strspn** (See *string(3C)*) ..... string operation  
**strtod**(3C) ..... convert string to double-precision number  
**strtok** (See *string(3C)*) ..... string operation  
**strtol**(3C) ..... convert string to integer  
**stty**(1) ..... set the options for a terminal  
**su**(1) ..... become super-user or another user  
**sum**(1) ..... print checksum and block count of a file  
**sum7**(1) ..... sum and count blocks in a file  
**sumdir**(1) ..... sum and count characters in the files in the given directories  
**swab**(3C) ..... swap bytes  
**sxt**(7) ..... pseudo-device driver  
**syms**(4) ..... common object file symbol table format  
**sync**(1) ..... update the super block  
**sync**(2) ..... update super-block  
**sys\_errlist** (See *perorr(3C)*) ..... system error message  
**sys\_nerr** (See *perorr(3C)*) ..... system error message  
**sysdef**(1M) ..... system definition  
**system**(3F) ..... issue a shell command from Fortran  
**system**(3S) ..... issue a shell command  
**tabs**(1) ..... set tabs on a terminal  
**tall**(1) ..... deliver the last part of a file  
**take**(1C) ..... takes a file from a remote machine  
**talk**(1N) ..... talk to another user  
**tan** (See *trig(3M)*) ..... trigonometric function  
**tan**(3F) ..... Fortran tangent intrinsic function  
**tanh** (See *sinh(3M)*) ..... hyperbolic function  
**tanh**(3F) ..... Fortran hyperbolic tangent intrinsic function  
**tapesave** (See *filesave(1M)*) ..... daily/weekly UNIX file system backup  
**tar**(1) ..... tape file archiver  
**tbl**(1) ..... format tables for nroff or troff  
**tc**(1) ..... phototypesetter simulator  
**tcp**(5P) ..... Internet Transmission Control Protocol  
**tdelete** (See *isearch(3C)*) ..... manage binary search trees  
**tee**(1) ..... pipe fitting

## COMMANDS

**telinit** (See *init(1M)*) ..... process control initialization

**telldir** (See *directory(3X)*) ..... flexible length directory operation

**telnet(1N)** ..... user interface to the TELNET protocol

**telnetd(8N)** ..... DARPA TELNET protocol server

**tempnam** (See *tmpnam(3S)*) ..... create a name for a temporary file

**term(4)** ..... format of compiled term file.

**term(5)** ..... conventional names for terminals

**termcap(3X)** ..... terminal independent operation routines

**termcap(5)** ..... terminal capability data base

**terminfo(4)** ..... terminal capability data base

**termio(7)** ..... general terminal interface

**test(1)** ..... condition evaluation command

**tfind** (See *tsearch(3C)*) ..... manage binary search trees

**tftpd(8N)** ..... DARPA Trivial File Transfer Protocol server

**tgsetent** (See *termcap(3X)*) ..... terminal independent operation routine

**tggetflag** (See *termcap(3X)*) ..... terminal independent operation routine

**tggetnam** (See *termcap(3X)*) ..... terminal independent operation routine

**tggetstr** (See *termcap(3X)*) ..... terminal independent operation routine

**tgoto** (See *termcap(3X)*) ..... terminal independent operation routine

**tie(1M)** ..... terminfo compiler

**time(1)** ..... time a command

**time(2)** ..... get time

**times(2)** ..... get process and child process times

**timex(1)** ..... time a command; report process data and system activity

**tmpfile(3S)** ..... create a temporary file

**tmpnam(3S)** ..... create a name for a temporary file

**tonscii** (See *conv(3C)*) ..... translate characters

**tolower** (See *conv(3C)*) ..... translate characters

**touch(1)** ..... update access and modification times of a file

**toupper** (See *conv(3C)*) ..... translate characters

**tp(1)** ..... manipulate tape archive

**tplot(1G)** ..... graphics filters

**tput(1)** ..... query terminfo database

**tputs** (See *termcap(3X)*) ..... terminal independent operation routine

**tr(1)** ..... translate characters

**trek(6)** ..... trekkie game

**trig(3M)** ..... trigonometric functions

**troff(1)** ..... typeset text

**trpt(8N)** ..... transliterate protocol trace

## COMMANDS


**true(1)** ..... provide truth values  
**tsearch(3C)** ..... manage binary search trees  
**tset(1)** ..... set the teletype bits to a sensible state  
**tsort(1)** ..... topological sort  
**tt(6)** ..... tic-tac-toe  
**tty(1)** ..... get the terminal's name  
**tty(7)** ..... controlling terminal interface  
**ttynam(3C)** ..... find name of a terminal  
**ttyslot(3C)** ..... find the slot in the utmp file of the current user  
**ttyp(4)** ..... data base of terminal types by port  
**turnacct (See acctsh(1M))** ..... shell procedure for accounting  
**twalk (See tsearch(3C))** ..... manage binary search trees  
**twinkle(6)** ..... twinkle stars on the screen  
**types(5)** ..... primitive system data types  
**tzset (See ctim(3C))** ..... convert date and time to string  
**u3b (See machid(1))** ..... provide truth value about your processor type  
**u3b5 (See machid(1))** ..... provide truth value about your processor type  
**udp(5P)** ..... Internet User Datagram Protocol  
**ul(1)** ..... do underlining  
**ulimit(2)** ..... get and set user limits  
**umask(1)** ..... set file-creation mode mask  
**umask(2)** ..... set and get file creation mask  
**umount (See mount(1M))** ..... dismount file system  
**umount(2)** ..... unmount a file system  
**uname(1)** ..... print name of current UNIX system  
**uname(2)** ..... get name of current UNIX system  
**unget(1)** ..... undo a previous get of an SCCS file  
**ungetc(3S)** ..... push character back into input stream  
**uniq(1)** ..... report repeated lines in a file  
**units(1)** ..... conversion program  
**unlink (See link(1M))** ..... exercise unlink system call  
**unlink(2)** ..... remove directory entry  
**unpack (See pack(1))** ..... expand compressed files  
**updater(1)** ..... update files between two machines  
**updater(1M)** ..... update files between two machines  
**ustat(2)** ..... get file system statistics  
**utime(2)** ..... set file access and modification times  
**utmp(4)** ..... utmp entry format  
**utmpname (See getut(3C))** ..... access utmp file entry

## COMMANDS

**uuclean(1M)** .....UUCP spool directory clean-up  
**uucp(1C)** .....UNIX system to UNIX system copy  
**uulog** (*See uucp(1C)*) .....prints a summary log of UUCP and UUX transactions  
**uuname** (*See uucp(1C)*) .....lists the UUCP names of known systems  
**uuplex** (*See uuto(1C)*) .....public UNIX-to-UNIX system file copy  
**uustat(1C)** .....UUCP status inquiry and job control  
**uusub(1M)** .....monitor UUCP network  
**uuto(1C)** .....public UNIX-to-UNIX system file copy  
**uux(1C)** .....UNIX-to-UNIX system command execution  
**uvar(2)** .....returns system-specific configuration information  
**val(1)** .....validate SCCS file  
**values(5)** .....machine-dependent values  
**varargs(5)** .....handle variable argument list  
**vax** (*See machid(1)*) .....provide truth value about your processor type  
**vc(1)** .....version control  
**vchk(1M)** .....version checkup  
**vedit** (*See vi(1)*) .....screen-oriented (visual) display editor based on ex  
**version(1)** .....reports version number of files  
**vfprintf** (*See vprintf(3S)*) .....print formatted output of a varargs argument list  
**vfprintf** (*See vprintf(3X)*) .....print formatted output of a varargs argument list  
**vi(1)** .....screen-oriented (visual) display editor based on ex  
**view** (*See vi(1)*) .....screen-oriented (visual) display editor based on ex  
**volcopy(1M)** .....copy file systems with label checking  
**vprintf(3S)** .....print formatted output of a varargs argument list  
**vprintf(3X)** .....print formatted output of a varargs argument list  
**vsprintf** (*See vprintf(3S)*) .....print formatted output of a varargs argument list  
**vsprintf** (*See vprintf(3X)*) .....print formatted output of a varargs argument list  
**wait(2)** .....wait for child process to stop or terminate  
**wait3(2N)** .....wait for child process to stop or terminate  
**wall(1M)** .....write to all users  
**wc(1)** .....word count  
**what(1)** .....identify SCCS files  
**whereis(1)** .....locate source, binary, and/or manual for program  
**who(1)** .....who is on the system  
**whoami(1)** .....print effective current user id  
**whodo(1M)** .....who is doing what  
**worm(6)** .....play the growing worm game  
**worms(6)** .....animate worms on a display terminal  
**write(1)** .....write to another user

## COMMANDS

**write(3)** ..... write on a file  
**writeln(3N)** ..... write on a file  
**wtmp** (*See utmp(4)*) ..... wtmp entry format  
**wtmpfix** (*See fwtmp(1M)*) ..... manipulate connect accounting records  
**wump(6)** ..... the game of hunt-the-wumpus  
**xargs(1)** ..... construct argument list(s) and execute command  
**xor** (*See bool(3F)*) ..... Fortran bitwise boolean function  
**xstr(1)** ..... extract strings from C programs to implement shared strings  
**y0** (*See bessel(3M)*) ..... Bessel function  
**y1** (*See bessel(3M)*) ..... Bessel function  
**yacc(1)** ..... yet another compiler-compiler  
**yn** (*See bessel(3M)*) ..... Bessel function  
**zabs** (*See abs(3F)*) ..... Fortran absolute value



## PERMUTED INDEX

functions of DASI 300 and/ 300, 300s: handle special . . . . . 300(1)  
 /special functions of DASI 300 and 300s terminals. . . . . 300(1)  
 of DASI 300 and 300s/ 300, 300s: handle special functions . . . . . 300(1)  
 functions of DASI 300 and 300s terminals. /special . . . . . 300(1)  
 l3tol, lto13: convert between 3-byte integers and long/ . . . . . l3tol(3C)  
 comparison. diff3: 3-way differential file . . . . . diff3(1)  
 Tektronix 4014 terminal. 4014: paginator for the . . . . . 4014(1)  
 paginator for the Tektronix 4014 terminal. 4014: . . . . . 4014(1)  
 of the DASI 450 terminal. 450: handle special functions . . . . . 450(1)  
 special functions of the DASI 450 terminal. 450: handle . . . . . 450(1)  
 f77: Fortran 77 compiler. . . . . f77(1)  
 long integer and base-64/ a64l, l64a: convert between . . . . . a64l(3C)  
 program. abort: generate an IOT fault. . . . . abort(3C)  
 Fortran absolute value. abort: terminate Fortran . . . . . abort(3F)  
 value. abs, iabs, dabs, cabs, zabs: . . . . . abs(3F)  
 abs: return integer absolute . . . . . abs(3C)  
 value. abs: return integer absolute . . . . . abs(3C)  
 dabs, cabs, zabs: Fortran absolute value. abs, iabs, . . . . . abs(3F)  
 /floor, ceiling, remainder. absolute value functions. . . . . floor(3M)  
 socket. accept: accept a connection on a . . . . . accept(2N)  
 a socket. accept: accept a connection on . . . . . accept(2N)  
 LP requisit. accept, reject: allow/prevent . . . . . accept(1M)  
 of a file. touch: update access and modification times . . . . . touch(1)  
 utime: set file access and modification times. . . . . utime(2)  
 accessibility of a file. access: determine . . . . . access(2)  
 machine/ spuil, sgetl: access long integer data in a . . . . . spuil(3X)  
 phys: allow a process to access physical addresses. . . . . phys(2)  
 ldfcn: common object file access routines. . . . . ldfcn(4)  
 copy file systems for optimal access time. dcopy: . . . . . dcopy(1M)  
 /setutent, endutent, utmpname: access utmp file entry. . . . . getut(3C)  
 access: determine accessibility of a file. . . . . access(2)  
 enable or disable process accounting. acct: . . . . . acct(2)  
 acctcon2: connect-time accounting. acctcon1, . . . . . acctcon(1M)  
 acctprc1, acctprc2: process accounting. . . . . acctprc(1M)  
 turnacct: shell procedures for acctsh(1M)  
 /accton, acctwtmp: overview of accounting and miscellaneous/ acct(1M)  
 accounting commands. /of acct(1M)  
 diskusg: generate disk accounting data by user ID. . . . . diskusg(1M)  
 acct: per-process accounting file format. . . . . acct(4)  
 search and print process accounting file(s). acctcom: . . . . . acctcom(1)  
 acctmerg: merge or add total accounting files. . . . . acctmerg(1M)  
 mclck: return Fortran time accounting. . . . . mclck(3F)  
 summary from per-process accounting records. /command acctcms(1M)  
 wtmpfix: manipulate connect accounting records. fwtmp, . . . . . fwtmp(1M)  
 runacct: run daily accounting. . . . . runacct(1M)  
 process accounting. acct: enable or disable . . . . . acct(2)  
 file format. acct: per-process accounting . . . . . acct(4)  
 per-process accounting/ acctcms: command summary from . . . . . acctcms(1M)  
 process accounting file(s). acctcom: search and print . . . . . acctcom(1)  
 connect-time accounting. acctcon1, acctcon2: . . . . . acctcon(1M)  
 accounting. acctcon1, acctcon2: connect-time . . . . . acctcon(1M)  
 acctwtmp: overview of acctdisk, acctdusg, accton, . . . . . acct(1M)  
 overview of/ acctdisk, acctdusg, accton, acctwtmp: . . . . . acct(1M)  
 accounting files. acctmerg: merge or add total . . . . . acctmerg(1M)  
 acctdisk, acctdusg, accton, acctwtmp: overview of/ . . . . . acct(1M)  
 accounting. acctprc1, acctprc2: process . . . . . acctprc(1M)  
 acctprc1, acctprc2: process accounting. . . . . acctprc(1M)  
 acctdisk, acctdusg, accton, acctwtmp: overview of/ . . . . . acct(1M)  
 sin, cos, tan, asin, acos, atan, atan2:/ . . . . . trig(3M)  
 intrinsic function. acos, dacos: Fortran arccosine . . . . . acos(3F)  
 killall: kill all active processes. . . . . killall(1M)





ar5.0:	archive (library) file format. . . . .	ar5.0(4)
ip: manipulate tape	archive. . . . .	ip(1)
tar: tape file	archiver. . . . .	tar(1)
maintainer for portable	archives. /archive and library . . . . .	ar(1)
cpio: copy file	archives in and out. . . . .	cpio(1)
asin, dasin: Fortran	arcsine intrinsic function. . . . .	asin(3F)
atan2, datan2: Fortran	arctangent intrinsic function. . . . .	atan2(3F)
atan, datan: Fortran	arctangent intrinsic function. . . . .	atan(3F)
imaginary part of complex	argument. /dimag: Fortran . . . . .	aimag(3F)
return Fortran command-line	argument. getarg. . . . .	getarg(3F)
varargs: handle variable	argument list. . . . .	varargs(5)
formatted output of a varargs	argument list. /print . . . . .	vprintf(3S)
formatted output of a varargs	argument list. /print . . . . .	vprintf(3X)
command. xargs: construct	argument list(s) and execute . . . . .	xargs(1)
getopt: get option letter from	argument vector. . . . .	getopt(3C)
expr: evaluate	arguments as an expression. . . . .	expr(1)
echo: echo	arguments. . . . .	echo(1)
bc: arbitrary-precision	arithmetic language. . . . .	bc(1)
number facts.	arithmetic: provide drill in . . . . .	arithmetic(6)
Protocol.	arp: Address Resolution . . . . .	arp(5P)
expr: evaluate arguments	as an expression. . . . .	expr(1)
	as: common assembler. . . . .	as(1)
	as5.0: assembler. . . . .	as5.0(1)
characters. asa: interpret	ASA carriage control . . . . .	asa(1)
control characters.	asa: interpret ASA carriage . . . . .	asa(1)
ascii: map of	ASCII character set. . . . .	ascii(5)
set.	ascii: map of ASCII character . . . . .	ascii(5)
long integer and base-64	ASCII string. /convert between . . . . .	a64i(3C)
number. atof: convert	ASCII string to floating-point . . . . .	atof(3C)
and/ ctime, localtime, gmtime,	asetime, tzset: convert date . . . . .	ctime(3C)
trigonometric/ sin, cos, tan,	asin, acos, atan, atan2: . . . . .	trig(3M)
intrinsic function.	asin, dasin: Fortran arcsine . . . . .	asin(3F)
help:	ask for help in using SCCS. . . . .	help(1)
output. a.out: common	assembler and link editor . . . . .	a.out(4)
output. a.out5.0:	assembler and link editor . . . . .	a.out5.0(4)
as: common	assembler. . . . .	as(1)
as5.0:	assembler. . . . .	as5.0(1)
assertion.	assert: verify program . . . . .	assert(3X)
assert: verify program	assertion. . . . .	assert(3X)
setbuf, setvbuf:	assign buffering to a stream. . . . .	setbuf(3S)
a later time.	at, batch: execute commands at . . . . .	at(1)
sin, cos, tan, asin, acos,	atan, atan2: trigonometric/ . . . . .	trig(3M)
arctangent intrinsic/	atan, datan: Fortran . . . . .	atan(3F)
arctangent intrinsic/	atan2, datan2: Fortran . . . . .	atan2(3F)
cos, tan, asin, acos, atan,	atan2: trigonometric/ sin, . . . . .	trig(3M)
floating-point number.	atof: convert ASCII string to . . . . .	atof(3C)
double-precision/ strtod,	atof: convert string to . . . . .	strtod(3C)
integer. strtol, atol,	atoi: convert string to . . . . .	strtol(3C)
integer. strtol,	atol, atoi: convert string to . . . . .	strtol(3C)
aliens: The alien invaders	attack the earth. . . . .	aliens(6)
autorobots: Escape from the	automatic robots. . . . .	autorobots(6)
automatic robots.	autorobots: Escape from the . . . . .	autorobots(6)
lav: print load	average statistics. . . . .	lav(1)
processing language.	awk: pattern scanning and . . . . .	awk(1)
ungetc: push character	back into input stream. . . . .	ungetc(3S)
	back: the game of backgammon. . . . .	back(6)
	backgammon. . . . .	back(6)
back: the game of	backup. filesave, tapesave: . . . . .	filesave(1M)
daily/weekly UNIX file system	backup. . . . .	finc(1M)
finc: fast incremental	backup tape. . . . .	frec(1M)
frec: recover files from a	bad block handling. /alternate . . . . .	altblk(4)
block information for	bad block information. . . . .	badblk(1M)
/program to set or update	badblk: program to set or . . . . .	badblk(1M)
update bad block information.	banner: make posters. . . . .	banner(1)

Permuted Index

banner7: print large printer.	banner on printer.	banner7(1)
hosts: host name data	banner7: print large banner on	banner7(1)
networks: network name data	base.	hosts(4N)
port. ttytype: data	base.	networks(4N)
protocols: protocol name data	base of terminal types by	ttytype(4)
services: service name data	base.	protocols(4N)
terminal capability data	base. termcap:	termcap(5)
terminal capability data	base. terminfo:	terminfo(4)
between long integer and (visual) display editor	base-64 ASCII string. /convert	a64l(3C)
portions of path names.	based on ex. /screen-oriented	vi(1)
later time. at, arithmetic language.	basename, dirname: deliver	basename(1)
	batch: execute commands at a	at(1)
	bc: arbitrary-precision	bc(1)
	bcd: convert to antique media.	bcd(6)
system initialization/ brc, string operations. bcopy, and byte string operations.	bcheckrc, rc, powerfail:	brc(1M)
	bcmp, bzero, ffs: bit and byte	bstring(3N)
	bcopy, bcmp, bzero, ffs: bit	bstring(3N)
	bcopy: interactive block copy.	bcopy(1M)
	bdiff: big diff.	bdiff(1)
	beautifier.	cb(1)
	Bessel functions.	bessel(3M)
	bfs: big file scanner.	bfs(1)
	binary, and/or manual for/	whereis(1)
	binary directories.	cpset(1M)
	binary file. /the printable	strings(1)
	binary input/output.	freed(3S)
	binary search a sorted table.	bsearch(3C)
	binary search trees. tsearch,	tsearch(3C)
	bind a name to a socket.	bind(2N)
	bind: bind a name to a socket.	bind(2N)
	bit and byte string/	bstring(3N)
	bits. strip5.0:	strip5.0(1)
	bits to a sensible state.	tset(1)
	bitwise boolean functions.	bool(3F)
	bj: the game of black jack.	bj(6)
	black jack.	bj(6)
	block copy.	bcopy(1M)
	block count of a file.	sum(1)
	block handling. /alternate	altblk(4)
	block information. badblk:	badblk(1M)
	block information for bad	altblk(4)
	block.	sync(1)
	block transfer data.	blt(3C)
	blocks.	df(1M)
	blocks in a file.	sum7(1)
	blt, blt512: block transfer	blt(3C)
	blt512: block transfer data.	blt(3C)
	B-NET.	netmailer(8N)
	B-NET network mail system.	netmail(8N)
	boolean functions. /lshift,	bool(3F)
	boot: startup procedures.	boot(8)
	brc, bcheckrc, rc, powerfail:	brc(1M)
	brk, sbrk: change data segment	brk(2)
	bs: a compiler/interpreter for	bs(1)
	bsearch: binary search a	bsearch(3C)
	buffered input/output package.	stdio(3S)
	buffering to a stream.	setbuf(3S)
	build special file.	mknod(1M)
	byte order. /convert values	byteorder(3N)
	byte string operations.	bstring(3N)
	bytes.	swab(3C)
	bzero, ffs: bit and byte	bstring(3N)
	C compiler.	cc(1)
	cb: C program	
	j0, j1, jn, y0, y1, yn:	
	whereis: locate source,	
	cpset: install object files in	
	strings in an object, or other	
	freed, fwrite:	
	bsearch:	
	tfind, tdelete, twalk: manage	
	bind:	
	bcopy, bcmp, bzero, ffs:	
	remove symbols and relocation	
	/set or reset the teletype	
	/not, lshift, rshift: Fortran	
	bj: the game of	
	bcopy: interactive	
	sum: print checksum and	
	block information for bad	
	program to set or update bad	
	block/ altblk: alternate	
	sync: update the super	
	bit, blt512:	
	df: report number of free disk	
	sum7: sum and count	
	data.	
	blt,	
	netmailer: deliver mail to	
	netmail: the	
	rshift: Fortran bitwise	
	system initialization shell/	
	space allocation.	
	modest-sized programs.	
	sorted table.	
	stdio: standard	
	setbuf, setvbuf: assign	
	mknod:	
	between host and network	
	/bcmp, bzero, ffs: bit and	
	swab: swap	
	string/ bcopy, bcmp,	
	cc:	

cc5.0:	C compiler. . . . .	cc5.0(1)
mc68cc:	C compiler. . . . .	mc68cc(1)
cflow: generate	C flowgraph. . . . .	cflow(1)
cpp: the	C language preprocessor. . . . .	cpp(1)
cpp: the	C language preprocessor. . . . .	cpp5.0(1)
cb:	C program beautifier. . . . .	cb(1)
lint: a	C program checker. . . . .	lint(1)
cxref: generate	C program cross-reference. . . . .	cxref(1)
maintain a tags file for a	C program. ctags: . . . . .	ctags(1)
ctrace:	C program debugger. . . . .	ctrace(1)
xstr: extract strings from	C programs to implement shared/ . . . . .	xstr(1)
message file by massaging	C source. /create an error . . . . .	mkstr(1)
value. abs, iabs, dabs,	cabs, zabs: Fortran absolute . . . . .	abs(3F)
	cal: print calendar. . . . .	cal(1)
dc: desk	calculator. . . . .	dc(1)
cal: print	calendar. . . . .	cal(1)
	calendar: reminder service. . . . .	calendar(1)
cu:	call another UNIX system. . . . .	cu(1C)
data returned by stat system	call. stat: . . . . .	stat(5)
malloc, free, realloc,	calloc: main memory allocator. . . . .	malloc(3C)
fast/ malloc, free, realloc,	calloc, mallopt, mallinfo: . . . . .	malloc(3X)
intro: introduction to system	calls and error numbers. . . . .	intro(2)
link and unlink system	calls. link, unlink: exercise . . . . .	link(1M)
to an LP line printer. lp,	cancel: send/cancel requests . . . . .	lp(1)
termcap: terminal	capability data base. . . . .	termcap(5)
terminfo: terminal	capability data base. . . . .	terminfo(4)
cribbage: the	card game cribbage. . . . .	cribbage(6)
pnch: file format for	card images. . . . .	pnch(4)
asa: interpret ASA	carriage control characters. . . . .	asa(1)
files.	cat: concatenate and print . . . . .	cat(1)
	cb: C program beautifier. . . . .	cb(1)
	cc: C compiler. . . . .	cc(1)
function. cos, dcos,	cc5.0: C compiler. . . . .	cc5.0(1)
	ccos: Fortran cosine intrinsic . . . . .	cos(3F)
	cd: change working directory. . . . .	cd(1)
commentary of an SCCS delta.	cdc: change the delta . . . . .	cdc(1)
ceiling, remainder,/ floor,	ceil, fmod, fabs: floor, . . . . .	floor(3M)
/ceil, fmod, fabs: floor,	ceiling, remainder, absolute/ . . . . .	floor(3M)
intrinsic/ exp, dexp,	cexp: Fortran exponential . . . . .	exp(3F)
	cflow: generate C flowgraph. . . . .	cflow(1)
delta: make a delta	(change) to an SCCS file. . . . .	delta(1)
pipe: create an interprocess	channel. . . . .	pipe(2)
/dble, cmplx, dcmplx, ichar,	char: explicit Fortran type/ . . . . .	ftype(3F)
stream. ungetc: push	character back into input . . . . .	ungetc(3S)
and neqn. eqnchar: special	character definitions for eqn . . . . .	eqnchar(5)
file. freq: report on	character frequencies in a . . . . .	freq(1)
user. cuserid: get	character login name of the . . . . .	cuserid(3S)
/getchar, fgetc, getw: get	character or word from a/ . . . . .	getc(3S)
/putchar, fputc, putw: put	character or word on a stream. . . . .	putc(3S)
ascii: map of ASCII	character set. . . . .	ascii(5)
interpret ASA carriage control	characters. asa: . . . . .	asa(1)
_tolower, toascii: translate	characters. /_toupper, . . . . .	conv(3C)
iscntrl, isascii: classify	characters. /isprint, isgraph, . . . . .	ctype(3C)
given/ sumdir: sum and count	characters in the files in the . . . . .	sumdir(1)
tr: translate	characters. . . . .	tr(1)
lastlogin, monacct, nulladm,/	chargefee, ckpacct, dodisk, . . . . .	acctsh(1M)
killer robots.	chase: Try to escape the . . . . .	chase(6)
directory.	chdir: change working . . . . .	chdir(2)
/dfscck: file system consistency	check and interactive repair. . . . .	fsck(1M)
checking procedure.	checkall: faster file system . . . . .	checkall(1M)
constant-width text for/ cw,	checkcw: prepare . . . . .	cw(1)
text for nroff or/ eqn, neqn,	checkeq: format mathematical . . . . .	eqn(1)
lint: a C program	checker. . . . .	lint(1)
grpck: password/group file	checkers. pwck, . . . . .	pwck(1M)

Permuted Index

checkall: faster file system	checking procedure. . . . .	checkall(1M)
copy file systems with label	checking. volcopy, labelit: . . . . .	volcopy(1M)
systems processed by fsck.	checklist: list of file . . . . .	checklist(4)
formatted with the/ mm, osdd,	checkmm: print/check documents . . . . .	mm(1)
file. sum: print	checksum and block count of a . . . . .	sum(1)
vchk: version	checkup. . . . .	vchk(1M)
system nodename.	chgnod: change current UNIX . . . . .	chgnod(1M)
chown,	chgrp: change owner or group. . . . .	chown(1)
times: get process and	child process times. . . . .	times(2)
terminate. wait: wait for	child process to stop or . . . . .	wait(2)
terminate. wait3: wait for	child process to stop or . . . . .	wait3(2N)
	chmod: change mode. . . . .	chmod(1)
	chmod: change mode of file. . . . .	chmod(2)
of a file.	chown: change owner and group . . . . .	chown(2)
group.	chown, chgrp: change owner or . . . . .	chown(1)
	chroot: change root directory. . . . .	chroot(2)
for a command.	chroot: change root directory . . . . .	chroot(1M)
monacct, nulladm,/ chargefee,	ckpact, dodisk, lastlogin. . . . .	acctsh(1M)
isgraph, iscntrl, isascii:	classify characters. /isprint,	ctype(3C)
uuclean: uucp spool directory	clean-up. . . . .	uuclean(1M)
	clear: clear terminal screen. . . . .	clear(1)
	clear i-node. . . . .	cli(1M)
	clear terminal screen. . . . .	clear(1)
	clearerr, fileno: stream . . . . .	errorr(3S)
status/ ferror, feof,	C-like syntax. csh: a shell . . . . .	csh(1)
(command interpreter) with	clock. . . . .	alarm(2)
alarm: set a process's alarm	clock daemon. . . . .	cron(1M)
cron:	clock: report CPU time used. . . . .	clock(3C)
	clock: Fortran natural . . . . .	log(3F)
logarithm/ log, alog, dlog,	close a common object file. . . . .	ldclose(3X)
ldclose, ldaclose:	close a file descriptor. . . . .	close(2)
close:	close: close a file . . . . .	close(2)
descriptor.	close or flush a stream. . . . .	fclose(3S)
fclose, fflush:	closedir: flexible length/ . . . . .	directory(3X)
/telldir, seekdir, rewinddir,	cli: clear i-node. . . . .	cli(1M)
	cmp: compare two files. . . . .	cmp(1)
	cmplx, dcmplx, ichtar, char:/ . . . . .	ctype(3F)
/real, float, sngl, dbie,	col: filter reverse . . . . .	col(1)
line-feeds.	comb: combine SCCS deltas. . . . .	comb(1)
	combine SCCS deltas. . . . .	comb(1)
comb:	comm: select or reject lines . . . . .	comm(1)
common to two sorted files.	command at low priority. . . . .	nice(1)
nice: run a	command. chroot: . . . . .	chroot(1M)
change root directory for	command execution. . . . .	env(1)
env: set environment for	command execution. . . . .	uux(1C)
uux: UNIX-to-UNIX system	command from Fortran. . . . .	system(3F)
system: issue a shell	command immune to hangups (sh . . . . .	nohup(1)
only). nohup: run a	(command interpreter) with . . . . .	csh(1)
C-like syntax. csh: a shell	command options. . . . .	getopt(1)
getopt: parse	command programming language. . . . .	sh(1)
/shell, the standard/restricted	command. /routines for . . . . .	rcmd(3N)
returning a stream to a remote	command; report process data . . . . .	times(1)
and system/ timex: time a	command. rexec: . . . . .	rexec(3N)
return stream to a remote	command summary from . . . . .	acctcms(1M)
per-process/ acctcms:	command. . . . .	system(3S)
system: issue a shell	command. . . . .	test(1)
test: condition evaluation	command. . . . .	time(1)
time: time a	command. xargs: construct . . . . .	xargs(1)
argument list(s) and execute	command-line argument. . . . .	getarg(3F)
getarg: return Fortran	commands. /of accounting . . . . .	acct(1M)
and miscellaneous accounting	commands and application/ . . . . .	intro(1)
intro: introduction to	commands and application/ . . . . .	intro(1M)
/to system maintenance	commands at a later time. . . . .	at(1)
at, batch: execute	commands. . . . .	install(1M)
install: install		

cdc: change the delta	commentary of an SCCS delta. . . .	cdc(1)
ar: common archive file format. . . .		ar(4)
editor output. a.out: common assembler and link . . . .		a.out(4)
as: common assembler. . . . .		as(1)
log10, alog10, dlog10: Fortran common logarithm intrinsic/		log10(3F)
routines. ldfcn: common object file access . . . . .		ldfcn(4)
ldopen, ldaopen: open a common object file for/ . . . . .		ldopen(3X)
/line number entries of a common object file function. . . . .		ldlread(3X)
ldclose, ldclose: close a common object file. . . . .		ldclose(3X)
read the file header of a common object file. ldhread: . . . . .		ldhread(3X)
entries of a section of a common object file. /number . . . . .		ldlseek(3X)
the optional file header of a common object file. /seek to . . . . .		ldohseek(3X)
/entries of a section of a common object file. . . . .		ldrseek(3X)
/section header of a common object file. . . . .		ldshread(3X)
an indexed/named section of a common object file. /seek to . . . . .		ldsseek(3X)
of a symbol table entry of a common object file. /the index . . . . .		ldtindex(3X)
symbol table entry of a common object file. /indexed . . . . .		ldtbread(3X)
seek to the symbol table of a common object file. ldtbseek: . . . . .		ldtbseek(3X)
line number entries in a common object file. linenum: . . . . .		linenum(4)
nm: print name list of common object file. . . . .		nm(1)
relocation information for a common object file. reloc: . . . . .		reloc(4)
scnhdr: section header for a common object file. . . . .		scnhdr(4)
table format. syms: common object file symbol . . . . .		syms(4)
aouthdr.h - a.out header for common object files. . . . .		aouthdr(4)
filehdr: file header for common object files. . . . .		filehdr(4)
ld: link editor for common object files. . . . .		ld(1)
size: print section sizes of common object files. . . . .		size(1)
comm: select or reject lines common to two sorted files. . . . .		comm(1)
ipcs: report inter-process communication facilities/ . . . . .		ipcs(1)
ftok: standard interprocess communication package. . . . .		stdipc(3C)
socket: create an endpoint for communication. . . . .		socket(2N)
diff: differential file comparator. . . . .		diff(1)
cmp: compare two files. . . . .		cmp(1)
SCCS file. sccsdiff: compare two versions of an . . . . .		sccsdiff(1)
lge, lgt, lle, llt: string comparison intrinsic/ . . . . .		strcmp(3F)
diff3: 3-way differential file comparison. . . . .		diff3(1)
dircmp: directory comparison. . . . .		dircmp(1)
expression. regcmp, regex: compile and execute a regular . . . . .		regcmp(3X)
regex: regular expression compile and match routines. . . . .		regex(5)
regcmp: regular expression compile. . . . .		regcmp(1)
term: format of compiled term file.. . . . .		term(4)
cc: C compiler. . . . .		cc(1)
cc5.0: C compiler. . . . .		cc5.0(1)
f77: Fortran 77 compiler. . . . .		f77(1)
mc68cc: C compiler. . . . .		mc68cc(1)
tic: terminfo compiler. . . . .		tic(1M)
yacc: yet another compiler-compiler. . . . .		yacc(1)
modest-sized programs. bs: a compiler/interpreter for . . . . .		bs(1)
erf, erf: error function and complementary error function. . . . .		erf(3M)
Fortran imaginary part of complex argument. /dimag: . . . . .		aimag(3F)
conjg, dconjg: Fortran complex conjugate intrinsic/ . . . . .		conjg(3F)
pack, pcat, unpack: compress and expand files. . . . .		pack(1)
table entry of a/ ldtbindex: compute the index of a symbol . . . . .		ldtindex(3X)
cat: concatenate and print files. . . . .		cat(1)
test: condition evaluation command. . . . .		test(1)
config: configure system. . . . .		config(1M)
uvar: returns system-specific configuration information. . . . .		uvar(2)
parameters. ifconfig: configure network interface . . . . .		ifconfig(8N)
config: configure system. . . . .		config(1M)
system. lpadmin: configure the LP spooling . . . . .		lpadmin(1M)
conjugate intrinsic function. conjg, dconjg: Fortran complex . . . . .		conjg(3F)
conjg, dconjg: Fortran complex conjugate intrinsic function. . . . .		conjg(3F)
fwtmp, wtmpfix: manipulate connect accounting records. . . . .		fwtmp(1M)
on a socket. connect: initiate a connection . . . . .		connect(2N)

Permuted Index

getpeername: get name of an out-going terminal line	connected peer.	getpeername(2N)
accept: accept a connection	connection. dial: establish connection on a socket.	dial(3C) accept(2N)
connect: initiate a down part of a full-duplex	connection on a socket.	connect(2N)
listen: listen for acctcon1, acctcon2:	connection. shutdown: shut connections on a socket.	shutdown(2N) listen(2N)
fsock, dfsock: file system	connect-time accounting.	acctcon(1M)
math: math functions and cw, checkcw: prepare	consistency check and/ constants.	fsock(1M) math(5)
mkfslb: construct a file system.	constant-width text for troff.	mkfslb(1M)
mkfs: construct a file system.	construct a file system.	mkfs(1M)
execute command. xargs: construct argument list(s) and	constructs. deroff: remove contents of directory.	xargs(1) deroff(1) ls(1)
nroff/troff, tbl, and eqn ls: list	context split.	csplit(1)
cpplit: interpret ASA carriage	control characters.	asa(1)
asa: interpret ASA carriage	control device.	ioctl(2)
ioctl: control.	control.	fcntl(2)
fcntl: file control initialization.	control operations.	init(1M) msgctl(2)
init, telinit: process	control operations.	semctl(2)
msgctl: message	control operations.	shmctl(2)
semctl: semaphore	control options.	fcntl(5)
shmctl: shared memory	Control Protocol.	tcp(5P)
fcntl: file	control. uustat:	uustat(1C)
tcp: Internet Transmission	control.	vc(1)
uucp status inquiry and job	controlling terminal	tty(7)
vc: version	conv: object file converter.	conv(1)
interface. tty:	conventional names for	term(5)
terminals. term:	conversion. /dcmplx, ichar,	fltype(3F)
char: explicit Fortran type	conversion program.	units(1)
units:	convert and copy a file.	dd(1)
dd:	convert Arabic numerals to	number(6)
English. number:	convert ASCII string to	atof(3C)
floating-point number. atof:	convert between 3-byte	l3tol(3C)
integers and/ l3tol, l3l3:	convert between long integer	a64l(3C)
and base-64 ASCII/ a64l, l64a:	convert date and time to/	ctime(3C)
/gmtime, asctime, tzset:	convert files between M68000	fs cv(1M)
and VAX-11/780/ fscv:	convert floating-point number	ecvt(3C)
to string. ecvt, fcvt, gcv:	convert formatted input.	scanf(3S)
scanf, fscanf, sscanf:	convert string to/	strtod(3C)
strtod, atof:	convert string to integer.	strtol(3C)
strtod, atol, atoi:	convert to antique media.	bcd(6)
bcd:	convert values between host/	byteorder(3N)
htonl, htons, ntohl, ntohs:	converter.	conv(1)
conv: object file	copy a file.	dd(1)
dd: convert and	copy.	bcopy(1M)
bcopy: interactive block	copy file archives in and out.	cpio(1)
cpio:	copy file systems for optimal	dcopy(1M)
access time. dcopy:	copy file systems with label	volcopy(1M)
checking. volcopy, labelit:	copy, link or move files.	cp(1)
cp, ln, mv:	copy.	rcp(1N)
rcp: remote file	copy. uucp, uulog, uuname:	uucp(1C)
UNIX system to UNIX system	copy. uuto, uupick: public	uuto(1C)
UNIX-to-UNIX system file	core: format of core image	core(4)
file.	core image file.	core(4)
core: format of	core memory.	mem(7)
mem, kmem:	cos, dcos, ccos: Fortran	cos(3F)
cosine intrinsic function.	cos, tan, asin, acos, atan,	trig(3M)
atan2: trigonometric/ sin,	cosh, dcosh: Fortran	cosh(3F)
hyperbolic cosine intrinsic/	cosh, tanh: hyperbolic	sinh(3M)
functions. sinh,	cosine intrinsic function.	cos(3F)
cos, dcos, ccos: Fortran	cosine intrinsic function.	cosh(3F)
/dcosh: Fortran hyperbolic		

sum7: sum and	count blocks in a file. . . . .	sum7(1)
in the given/ sumdir: sum and	count characters in the files . . . . .	sumdir(1)
sum: print checksum and block	count of a file. . . . .	sum(1)
	count. . . . .	wc(1)
	cp, ln, mv: copy, link or move . . . . .	cp(1)
cpio: format of	cpio archive. . . . .	cpio(4)
and out.	cpio: copy file archives in . . . . .	cpio(1)
	cpio: format of cpio archive. . . . .	cpio(4)
preprocessor.	cpp: the C language . . . . .	cpp(1)
preprocessor.	cpp: the C language . . . . .	cpp5.0(1)
binary directories.	cpset: install object files in . . . . .	cpset(1M)
clock: report	CPU time used. . . . .	clock(3C)
craps: the game of	craps. . . . .	craps(6)
	craps: the game of craps. . . . .	craps(6)
system crashes.	crash: what to do when the . . . . .	crash(8)
what to do when the system	crashes. crash: . . . . .	crash(8)
rewrite an existing one.	creat: create a new file or . . . . .	creat(2)
file. tmpnam, tmpnam:	create a name for a temporary . . . . .	tmpnam(3S)
an existing one. creat:	create a new file or rewrite . . . . .	creat(2)
	fork: create a new process. . . . .	fork(2)
	tmpfile: create a temporary file. . . . .	tmpfile(3S)
communication. socket:	create an endpoint for . . . . .	socket(2N)
by massaging C source. mkstr:	create an error message file . . . . .	mkstr(1)
	channel. pipe: create an interprocess . . . . .	pipe(2)
files. admin:	create and administer SCCS . . . . .	admin(1)
umask: set and get file	creation mask. . . . .	umask(2)
cribbage: the card game	cribbage. . . . .	cribbage(6)
cribbage.	cribbage: the card game . . . . .	cribbage(6)
	cron: clock daemon. . . . .	cron(1M)
crontab: user	crontab file. . . . .	crontab(1)
	crontab: user crontab file. . . . .	crontab(1)
cxref: generate C program	cross-reference. . . . .	cxref(1)
optimization package. curses:	CRT screen handling and . . . . .	curses(3X)
more: file perusal filter for	crt viewing. . . . .	more(1)
	crypt: encode/decode. . . . .	crypt(1)
generate DES encryption.	crypt, setkey, encrypt: . . . . .	crypt(3C)
interpreter) with C-like/	csh: a shell (command . . . . .	csh(1)
function. sin, dsin,	csin: Fortran sine intrinsic . . . . .	sin(3F)
	csplit: context split. . . . .	csplit(1)
intrinsic/ sqrt, dsqrt,	csqrt: Fortran square root . . . . .	sqrt(3F)
terminal.	ct: spawn getty to a remote . . . . .	ct(1C)
for a C program.	ctags: maintain a tags file . . . . .	ctags(1)
terminal.	ctermid: generate filename for . . . . .	ctermid(3S)
asctime, tzset: convert date/	ctime, localtime, gmtime. . . . .	ctime(3C)
	ctrace: C program debugger. . . . .	ctrace(1)
	cu: call another UNIX system. . . . .	cu(1C)
ttl,	cubic: tic-tac-toe. . . . .	ttt(6)
get/set unique identifier of	current host. /sethostid: . . . . .	gethostid(2N)
sethostname: get/set name of	current host. gethostname. . . . .	gethostname(2N)
set or print identifier of	current host system. hostid: . . . . .	hostid(1N)
hostname: set or print name of	current host system. . . . .	hostname(1N)
activity. sact: print	current SCCS file editing . . . . .	sact(1)
chgnod: change	current UNIX system nodename. . . . .	chgnod(1M)
uname: print name of	current UNIX system. . . . .	uname(1)
uname: get name of	current UNIX system. . . . .	uname(2)
whoami: print effective	current user id. . . . .	whoami(1)
slot in the utmp file of the	current user. /find the . . . . .	ttyslot(3C)
getcwd: get pathname of	current working directory. . . . .	getcwd(3C)
and optimization package.	curses: CRT screen handling . . . . .	curses(3X)
spline: interpolate smooth	curve. . . . .	spline(1G)
name of the user.	cuserid: get character login . . . . .	cuserid(3S)
of each line of a file.	cut: cut out selected fields . . . . .	cut(1)
each line of a file. cut:	cut out selected fields of . . . . .	cut(1)
constant-width text for/	cw, checkcw: prepare . . . . .	cw(1)

cross-reference.	cxref: generate C program . . . . .	cxref(1)
absolute value. abs, labs,	dabs, cabs, zabs: Fortran . . . . .	abs(3F)
intrinsic function. acos,	dacos: Fortran arccosine . . . . .	acos(3F)
cron: clock	daemon. . . . .	cron(1M)
errdemon: error-logging	daemon. . . . .	errdemon(1M)
terminate the error-logging	daemon. errstop: . . . . .	errstop(1M)
routed: network routing	daemon. . . . .	routed(8N)
runacct: run	daily accounting. . . . .	runacct(1M)
backup. filesave, tapesave:	daily/weekly UNIX file system . . . . .	filesave(1M)
Protocol server. ftpd:	DARPA Internet File Transfer . . . . .	ftpd(8N)
telnetd:	DARPA TELNET protocol server. . . . .	telnetd(8N)
Protocol server. iftpd:	DARPA Trivial File Transfer . . . . .	iftpd(8N)
/handle special functions of	DASI 300 and 300s terminals. . . . .	300(1)
special functions of the	DASI 450 terminal. /handle . . . . .	450(1)
intrinsic function. asin,	dasin: Fortran arcsine . . . . .	asin(3F)
/time a command; report process	data and system activity. . . . .	time(x)1
hosts: host name	data base. . . . .	hosts(4N)
networks: network name	data base. . . . .	networks(4N)
port. ttytype:	data base of terminal types by . . . . .	ttytype(4)
protocols: protocol name	data base. . . . .	protocols(4N)
services: service name	data base. . . . .	services(4N)
termcap: terminal capability	data base. . . . .	termcap(5)
terminfo: terminal capability	data base. . . . .	terminfo(4)
bit, bit512: block transfer	data. . . . .	bit(3C)
generate disk accounting	data by user ID. diskusg: . . . . .	diskusg(1M)
/sgctl: access long integer	data in a machine independent/ . . . . .	sputl(3X)
plock: lock process, text, or	data in memory. . . . .	plock(2)
prof: display profile	data. . . . .	prof(1)
call. stat:	data returned by stat system . . . . .	stat(5)
brk, sbrk: change	data segment space allocation. . . . .	brk(2)
types: primitive system	data types. . . . .	types(5)
join: relational	database operator. . . . .	join(1)
tput: query terminfo	database. . . . .	tput(1)
udp: Internet User	Datagram Protocol. . . . .	udp(5P)
intrinsic function. atan,	datan: Fortran arctangent . . . . .	atan(3F)
intrinsic function. atan2,	datan2: Fortran arctangent . . . . .	atan2(3F)
/asctime, tzset: convert	date and time to string. . . . .	ctime(3C)
date: print and set the	date. . . . .	date(1)
/date: print and set the date.	date: print and set the date. . . . .	date(1)
/ldint, real, float, snl,	dbl, cmplx, dcmplx, ichar,/ . . . . .	fyte(3F)
/float, snl, dbl, cmplx,	dc: desk calculator. . . . .	dc(1)
conjugate intrinsic/ conjg,	dcmplx, ichar, char: explicit/ . . . . .	fyte(3F)
optimal access time.	dconjg: Fortran complex . . . . .	conjg(3F)
intrinsic function. cos,	dcopy: copy file systems for . . . . .	dcopy(1M)
cosine intrinsic/ cosh,	dcos, ccos: Fortran cosine . . . . .	cos(3F)
difference intrinsic/ dim,	dcosh: Fortran hyperbolic . . . . .	cosh(3F)
ctrace: C program	dd: convert and copy a file. . . . .	dd(1)
fsdb: file system	ddim, idim: positive . . . . .	dim(3F)
sdb: symbolic	debugger. . . . .	ctrace(1)
sysdef: system	debugger. . . . .	fsdb(1M)
eqnchar: special character	debugger. . . . .	sdb(1)
people. delivermail:	definition. . . . .	sysdef(1M)
netmailer:	definitions for eqn and neqn. . . . .	eqnchar(5)
names. basename, dirname:	deliver mail to arbitrary . . . . .	delivermail(8N)
file. tail:	deliver mail to B-NET. . . . .	netmailer(8N)
aliases: aliases file for	deliver portions of path . . . . .	basename(1)
arbitrary people.	deliver the last part of a . . . . .	tail(1)
delta commentary of an SCCS	delivermail. . . . .	aliases(7N)
file. delta: make a	delivermail: deliver mail to . . . . .	delivermail(8N)
delta. cdc: change the	delta. cdc: change the . . . . .	cdc(1)
delta commentary of an SCCS	delta (change) to an SCCS . . . . .	delta(1)
delta. cdc: change the	delta commentary of an SCCS . . . . .	cdc(1)
delta from an SCCS file.	delta from an SCCS file. . . . .	rmdel(1)
delta: make a delta (change)	delta: make a delta (change) . . . . .	delta(1)



comb: combine SCCS	deltas.	comb(1)
mesg: permit or	deny messages.	mesg(1)
tbl, and eqn constructs.	deroff: remove nroff/troff,	deroff(1)
setkey, encrypt: generate	DES encryption. crypt,	crypt(3C)
close: close a file	descriptor.	close(2)
dup2: duplicate a	descriptor.	dup2(3N)
dup: duplicate a	descriptor.	dup(3)
getdtablesize: get	descriptor table size.	getdtablesize(3N)
dc:	desk calculator.	dc(1)
file. access:	determine accessibility of a	access(2)
file:	determine file type.	file(1)
errors in the specified	device. /on/off the extended	exterr(1)
master: master	device information table.	master(4)
ioctl: control	device.	ioctl(2)
devnm:	device name.	devnm(1M)
	devnm: device name.	devnm(1M)
exponential intrinsic/ exp,	dexp, cexp: Fortran	exp(3F)
blocks.	df: report number of free disk	df(1M)
check and interactive/ fsck,	dfscck: file system consistency	fsck(1M)
terminal line connection.	dial: establish an out-going	dial(3C)
ratfor: rational Fortran	dialect.	ratfor(1)
bdiff: big	diff.	bdiff(1)
comparator.	diff: differential file	diff(1)
diffdir:	diff directories.	diffdir(1)
comparison.	diff3: 3-way differential file	diff3(1)
	diffdir: diff directories.	diffdir(1)
dim, ddim, idim: positive	difference intrinsic/	dim(3F)
sdiff: side-by-side	difference program.	sdiff(1)
diffmk: mark	differences between files.	diffmk(1)
diff:	differential file comparator.	diff(1)
diff3: 3-way	differential file comparison.	diff3(1)
between files.	diffmk: mark differences	diffmk(1)
difference intrinsic/	dim, ddim, idim: positive	dim(3F)
of complex argument. aimag,	dimag: Fortran imaginary part	aimag(3F)
intrinsic function. aint,	dint: Fortran integer part	aint(3F)
	dir: format of directories.	dir(4)
install object files in binary	dircmp: directory comparison.	dircmp(1)
diffdir: diff	directories. cpset:	cpset(1M)
dir: format of	directories.	diffdir(1)
rm, rmdir: remove files or	directories.	dir(4)
in the files in the given	directories.	rm(1)
cd: change working	directories. /count characters	sumdir(1)
chdir: change working	directory.	cd(1)
chroot: change root	directory.	chdir(2)
uuclean: uucp spool	directory.	chroot(2)
dircmp:	directory clean-up.	uuclean(1M)
unlink: remove	directory comparison.	dircmp(1)
chroot: change root	directory entry.	unlink(2)
/make a lost+found	directory for a command.	chroot(1M)
pathname of current working	directory for fsck.	mklost+found(1M)
ls: list contents of	directory. getcwd: get	getcwd(3C)
mkdir: make a	directory.	ls(1)
mkdir: move a	directory.	mkdir(1)
pwd: working	directory.	mvdir(1M)
/closedir: flexible length	directory name.	pwd(1)
ordinary file. mknod: make a	directory operations.	directory(3X)
path names. basename,	directory, or a special or	mknod(2)
	dirname: deliver portions of	basename(1)
printers. enable,	dis: disassembler.	dis(1)
acct: enable or	disable: enable/disable LP	enable(1)
dis:	disable process accounting.	acct(2)
type, modes, speed, and line	disassembler.	dis(1)
ID. diskusg: generate	discipline. /set terminal	getty(1M)
	disk accounting data by user	diskusg(1M)

df: report number of free disk blocks.	df(1M)
diskformat: format a disk.	diskformat(1M)
disktune: tune floppy disk settling time parameters.	disktune(1M)
du: summarize disk usage.	du(1)
diskformat: format a disk.	diskformat(1M)
disktune: tune floppy disk settling time parameters.	disktune(1M)
accounting data by user ID.	diskusg(1M)
mount, umount: mount and unmount.	mount(1M)
vi: screen-oriented (visual) editor.	vi(1)
prof: display profile data.	prof(1)
rain: animated raindrops.	rain(6)
worms: animate worms on a display terminal.	worms(6)
hypot: Euclidean distance function.	hypot(3M)
/lcong48: generate uniformly distributed pseudo-random/	drand48(3C)
logarithm/ log, alog, dlog, clog: Fortran natural logarithm/	log(3F)
log10, alog10, dlog10: Fortran common logarithm/	log10(3F)
max, max0, amax0, max1, amax1, dmax1: Fortran maximum-value/	max(3F)
min, min0, amin0, min1, amin1, dmin1: Fortran minimum-value/	min(3F)
intrinsic/ mod, amod, dmod: Fortran remaindering	mod(3F)
nearest integer/ anint, dnint, nint, adnint: Fortran	round(3F)
mm, osdd, checkmm: print/check documents formatted with the	mm(1)
macro package for formatting documents. mm: the MM	mm(5)
macro package for formatting documents. /the OSDD adapter	mosd(5)
slides. mmt, mvt: typeset documents, view graphs, and	mmt(1)
nulladm,/ chargefee, ckpacct, dodisk, lastlogin, monacct,	acctsh(1M)
whodo: who is doing what.	whodo(1M)
double precision product	dprod(3F)
double-precision number.	strtod(3C)
downloading into a file.	rcvhex(1)
dprod: double precision	dprod(3F)
drand48, erand48, lrand48,	drand48(3C)
draw a graph.	graph(1G)
drill in number facts.	arithmetic(6)
driver.	pty(5)
driver.	sxt(7)
dsign: Fortran	sign(3F)
dsin, csin: Fortran sine	sin(3F)
dsinh: Fortran hyperbolic sine	sinh(3F)
dsqrt, csqrt: Fortran square	sqrt(3F)
dtan: Fortran tangent	tan(3F)
dtanh: Fortran hyperbolic	tanh(3F)
du: summarize disk usage.	du(1)
dump: dump selected parts of	dump(1)
dump. errdead:	errdead(1M)
dump.	od(1)
dump selected parts of an	dump(1)
dup: duplicate a descriptor.	dup(3)
dup2: duplicate a descriptor.	dup2(3N)
duplicate a descriptor.	dup2(3N)
dup: duplicate a descriptor.	dup(3)
The alien invaders attack the	aliens(6)
echo: echo arguments.	echo(1)
echo: echo arguments.	echo(1)
ecvt, fcvt, gcvt: convert	ecvt(3C)
ed, red: text editor.	ed(1)
edata: last locations in	end(3C)
edit: text editor.	ex(1)
editing activity.	sact(1)
editor based on ex.	vi(1)
editor.	ed(1)
editor.	ex(1)
editor for common object	ld(1)
editor.	ld5.0(1)
editor output. a.out:	a.out(4)
df: report number of free disk blocks.	df(1M)
diskformat: format a disk.	diskformat(1M)
disktune: tune floppy disk settling time parameters.	disktune(1M)
du: summarize disk usage.	du(1)
diskformat: format a disk.	diskformat(1M)
disktune: tune floppy disk settling time parameters.	disktune(1M)
accounting data by user ID.	diskusg(1M)
mount, umount: mount and unmount.	mount(1M)
vi: screen-oriented (visual) editor.	vi(1)
prof: display profile data.	prof(1)
rain: animated raindrops.	rain(6)
worms: animate worms on a display terminal.	worms(6)
hypot: Euclidean distance function.	hypot(3M)
/lcong48: generate uniformly distributed pseudo-random/	drand48(3C)
logarithm/ log, alog, dlog, clog: Fortran natural logarithm/	log(3F)
log10, alog10, dlog10: Fortran common logarithm/	log10(3F)
max, max0, amax0, max1, amax1, dmax1: Fortran maximum-value/	max(3F)
min, min0, amin0, min1, amin1, dmin1: Fortran minimum-value/	min(3F)
intrinsic/ mod, amod, dmod: Fortran remaindering	mod(3F)
nearest integer/ anint, dnint, nint, adnint: Fortran	round(3F)
mm, osdd, checkmm: print/check documents formatted with the	mm(1)
macro package for formatting documents. mm: the MM	mm(5)
macro package for formatting documents. /the OSDD adapter	mosd(5)
slides. mmt, mvt: typeset documents, view graphs, and	mmt(1)
nulladm,/ chargefee, ckpacct, dodisk, lastlogin, monacct,	acctsh(1M)
whodo: who is doing what.	whodo(1M)
double precision product	dprod(3F)
double-precision number.	strtod(3C)
downloading into a file.	rcvhex(1)
dprod: double precision	dprod(3F)
drand48, erand48, lrand48,	drand48(3C)
draw a graph.	graph(1G)
drill in number facts.	arithmetic(6)
driver.	pty(5)
driver.	sxt(7)
dsign: Fortran	sign(3F)
dsin, csin: Fortran sine	sin(3F)
dsinh: Fortran hyperbolic sine	sinh(3F)
dsqrt, csqrt: Fortran square	sqrt(3F)
dtan: Fortran tangent	tan(3F)
dtanh: Fortran hyperbolic	tanh(3F)
du: summarize disk usage.	du(1)
dump: dump selected parts of	dump(1)
dump. errdead:	errdead(1M)
dump.	od(1)
dump selected parts of an	dump(1)
dup: duplicate a descriptor.	dup(3)
dup2: duplicate a descriptor.	dup2(3N)
duplicate a descriptor.	dup2(3N)
dup: duplicate a descriptor.	dup(3)
The alien invaders attack the	aliens(6)
echo: echo arguments.	echo(1)
echo: echo arguments.	echo(1)
ecvt, fcvt, gcvt: convert	ecvt(3C)
ed, red: text editor.	ed(1)
edata: last locations in	end(3C)
edit: text editor.	ex(1)
editing activity.	sact(1)
editor based on ex.	vi(1)
editor.	ed(1)
editor.	ex(1)
editor for common object	ld(1)
editor.	ld5.0(1)
editor output. a.out:	a.out(4)

a.out5.0: assembler and link  
 sed: stream  
 whoami: print  
 setregid: set real and  
 /user, real group, and  
 setreuid: set real and  
 and/ /getegid: get real user,  
 Language.  
 fsplit: split /77, ratfor, or  
 for a pattern. grep.  
 insque, remque: insert/remove  
 enable/disable LP printers.  
 accounting. acct:  
 enable, disable:  
 crypt:  
 encryption. crypt, setkey,  
 setkey, encrypt: generate DES  
 makekey: generate  
 locations in program.  
 /getgrgid, getgrnam, setgrent,  
 /gethostbyname, sethostent,  
 /getnetbyname, setnetent,  
 socket: create an  
 /getprotobyname, setprotoent,  
 /getpwuid, getpwnam, setpwent,  
 /getservbyname, setservent,  
 utmp/ /pututline, setutent,  
 convert Arabic numerals to  
 nlist: get  
 file. linenum: line number  
 man: print  
 man: macros for formatting  
 file/ /manipulate line number  
 /ldnseek: seek to line number  
 /ldnrseek: seek to relocation  
 utmp, wtmp: utmp and wtmp  
 /fgetgrent: obtain group file  
 endhostent: get network host  
 endnetent: get network  
 endprotoent: get protocol  
 fgetpwent: get password file  
 endservent: get service  
 utmpname: access utmp file  
 /the index of a symbol table  
 /read an indexed symbol table  
 putpwent: write password file  
 unlink: remove directory  
 command execution.  
 profile: setting up an  
 environ: user  
 execution. env: set  
 getenv: return value for  
 printenv: print out the  
 putenv: change or add value to  
 getenv: return Fortran  
 character definitions for  
 remove nroff/troff, tbl, and  
 mathematical text for nroff/  
 definitions for eqn and neqn.  
 mrand48, jrand48, / drand48,  
 complementary error function.  
 complementary error/ erf,  
 from dumpt.  
 editor output.  
 editor.  
 effective current user id.  
 effective group ID.  
 effective group IDs.  
 effective user ID's.  
 effective user, real group,  
 eft: Extended Fortran  
 efi files.  
 egrep, fgrep: search a file  
 element from a queue.  
 enable, disable:  
 enable or disable process  
 enable/disable LP printers.  
 encode/decode.  
 encrypt: generate DES  
 encryption. crypt,  
 encryption key.  
 end, e1ext, edata: last  
 endgrent, fgetgrent: obtain/  
 endhostent: get network host/  
 endnetent: get network entry.  
 endpoint for communication.  
 endprotoent: get protocol/  
 endpwent, fgetpwent: get/  
 endservent: get service entry.  
 endutent, utmpname: access  
 English. number:  
 entries from name list.  
 entries in a common object  
 entries in this manual.  
 entries in this manual.  
 entries of a common object  
 entries of a section of a/  
 entries of a section of a/  
 entry formats.  
 entry from a group file.  
 entry. /sethostent,  
 entry. /setnetent,  
 entry. /setprotoent,  
 entry. /setpwent, endpwent,  
 entry. /setservent,  
 entry. /setutent, endutent,  
 entry of a common object file.  
 entry of a common object file.  
 entry.  
 env: set environment for  
 environ: user environment.  
 environment at login time.  
 environment.  
 environment for command  
 environment name.  
 environment.  
 environment.  
 environment variable.  
 eqn and neqn. /special  
 eqn constructs. deroff:  
 eqn, neqn, checkeq: format  
 eqnchar: special character  
 erand48, lrand48, nrand48,  
 erf, erfc: error function and  
 erfc: error function and  
 errdead: extract error records  
 a.out5.0(4)  
 sed(1)  
 whoami(1)  
 setregid(2)  
 getuid(2)  
 setreuid(2)  
 getuid(2)  
 efi(1)  
 fsplit(1)  
 grep(1)  
 insque(3N)  
 enable(1)  
 acct(2)  
 enable(1)  
 crypt(1)  
 crypt(3C)  
 crypt(3C)  
 makekey(1)  
 end(3C)  
 getgrent(3C)  
 gethostent(3N)  
 getnetent(3N)  
 socket(2N)  
 getprotoent(3N)  
 getpwent(3C)  
 getservent(3N)  
 getut(3C)  
 number(6)  
 nlist(3C)  
 linenum(4)  
 man(1)  
 man(5)  
 ldread(3X)  
 ldseek(3X)  
 ldrseek(3X)  
 utmp(4)  
 getgrent(3C)  
 gethostent(3N)  
 getnetent(3N)  
 getprotoent(3N)  
 getpwent(3C)  
 getservent(3N)  
 getut(3C)  
 ldtbindex(3X)  
 ldtbread(3X)  
 putpwent(3C)  
 unlink(2)  
 env(1)  
 environ(5)  
 profile(4)  
 environ(5)  
 env(1)  
 getenv(3C)  
 printenv(1)  
 putenv(3C)  
 getenv(3F)  
 eqnchar(5)  
 deroff(1)  
 eqn(1)  
 eqnchar(5)  
 drand48(3C)  
 erf(3M)  
 erf(3M)  
 errdead(1M)

daemon	errdemon: error-logging	errdemon(1M)
format	errfile: error-log file	errfile(4)
system error/ perror,	errno, sys_errlist, sys_nerr:	perror(3C)
interface.	error: error-logging	error(7)
complementary/ erf, erfc:	error function and	erf(3M)
function and complementary	error function. /erfc: error	erf(3M)
massaging C/ mkstr: create an	error message file by	mkstr(1)
sys_errlist, sys_nerr: system	error messages. /errno,	perror(3C)
to system calls and	error numbers. /introduction	intro(2)
errdead: extract	error records from dump.	errdead(1M)
matherr:	error-handling function.	matherr(3M)
errfile:	error-log file format.	errfile(4)
errdemon:	error-logging daemon.	errdemon(1M)
errstop: terminate the	error-logging daemon.	errstop(1M)
error:	error-logging interface.	error(7)
process a report of logged	errors. errpt:	errpt(1M)
/turn on/off the extended	errors in the specified/	exterr(1)
hashcheck: find spelling	errors. /hashmake, spellin,	spell(1)
logged errors.	errpt: process a report of	errpt(1M)
error-logging daemon.	errstop: terminate the	errstop(1M)
robots. autorobots:	Escape from the automatic	autorobots(6)
robots:	Escape from the robots.	robots(6)
chase: Try to	escape the killer robots.	chase(6)
terminal line/ dial:	establish an out-going	dial(3C)
setmnt:	establish mount table.	setmnt(1M)
in program. end,	etext, edata: last locations	end(3C)
hypot:	Euclidean distance function.	hypot(3M)
expression. expr:	evaluate arguments as an	expr(1)
test: condition	evaluation command.	test(1)
	ex, edit: text editor.	ex(1)
display editor based on	ex. /screen-oriented (visual)	vi(1)
reading or/ locking: provide	exclusive file regions for	locking(2)
execlp, execvp: execute a/	execl, execv, execl, execlp,	exec(2)
execvp: execute/ execl, execv,	execle, execve, execvp,	exec(2)
execl, execv, execl, execlp,	execlp, execvp: execute a/	exec(2)
execve, execlp, execvp:	execute a file. /execle,	exec(2)
regcmp, regex: compile and	execute a regular expression.	regcmp(3X)
construct argument list(s) and	execute command. xargs:	xargs(1)
time. at, batch:	execute commands at a later	at(1)
set environment for command	execution. env:	env(1)
sleep: suspend	execution for an interval.	sleep(1)
sleep: suspend	execution for interval.	sleep(3C)
monitor: prepare	execution profile.	monitor(3C)
rexecd: remote	execution server.	rexecd(8N)
profil:	execution time profile.	profil(2)
UNIX-to-UNIX system command	execution. uux:	uux(1C)
execvp: execute a/ execl,	execv, execl, execve, execlp,	exec(2)
execle/ execl, execv, execl,	execve, execlp, execvp:	exec(2)
/execv, execl, execve, execlp,	execvp: execute a file.	exec(2)
system calls. link, unlink:	exercise link and unlink	link(1M)
a new file or rewrite an	existing one. creat: create	creat(2)
process.	exit, _exit: terminate	exit(2)
exit,	_exit: terminate process.	exit(2)
exponential intrinsic/	exp, dexp, cexp: Fortran	exp(3F)
exponential, logarithm,/	exp, log, log10, pow, sqrt:	exp(3M)
pcat, unpack: compress and	expand files. pack,	pack(1)
cmplx, dcmplx, ichtar, char:	explicit Fortran type/ /dble,	fctype(3F)
adventure: an	exploration game.	adventure(6)
exp, dexp, cexp: Fortran	exponential intrinsic/	exp(3F)
exp, log, log10, pow, sqrt:	exponential, logarithm, power,/	exp(3M)
expression.	expr: evaluate arguments as an	expr(1)
routines. regexp: regular	expression compile and match	regexp(5)
regcmp: regular	expression compile.	regcmp(1)
expr: evaluate arguments as an	expression.	expr(1)

compile and execute a regular	expression. regcmp, regex:	regcmp(3X)
exterr: turn on/off the	extended errors in the/	exterr(1)
efl:	Extended Fortran Language.	efl(1)
greek: graphics for the	extended TTY-37 type-box.	greek(5)
extended errors in the/	exterr: turn on/off the	exterr(1)
dump. errdead:	extract error records from	errdead(1M)
programs to implement/ xstr:	extract strings from C	xstr(1)
	f77: Fortran 77 compiler.	f77(1)
	f77, ratfor, or efl files.	fsplit(1)
fsplit: split	fabs: floor, ceiling,	floor(3M)
remainder,/ floor, ceil, fmod,	factor a number.	factor(1)
factor:	factor: factor a number.	factor(1)
	false: provide truth values.	true(1)
true,	fashion. /access long integer	sput(3X)
data in a machine independent	fast incremental backup.	finc(1M)
finc:	fast main memory allocator.	malloc(3X)
/calloc, malloc, mallinfo:	faster file system checking	checkall(1M)
procedure. checkall:	fault.	abort(3C)
abort: generate an IOT	fclose, fflush: close or flush	fclose(3S)
a stream.	fcntl: file control.	fcntl(2)
	fcntl: file control options.	fcntl(5)
floating-point number/ ecvt,	fcvt, gcvt: convert	ecvt(3C)
fopen, freopen,	fdopen: open a stream.	fopen(3S)
status inquiries. ferror,	feof, clearerr, fileno: stream	ferror(3S)
fileno: stream status/	ferror, feof, clearerr,	ferror(3S)
statistics for a file system.	ff: list file names and	ff(1M)
stream. fclose,	flush: close or flush a	fclose(3S)
bcopy, bcopy, bzero,	ffs: bit and byte string/	bstring(3N)
word from a/ getch, getchar,	fgetc, getw: get character or	getc(3S)
/getgrnam, setgrent, endgrent,	fgetgrent: obtain group file/	getgrent(3C)
/getpwnam, setpwent, endpwent,	fgetpwent: get password file/	getpwent(3C)
stream. gets,	fgets: get a string from a	gets(3S)
pattern. grep, egrep,	fgrep: search a file for a	grep(1)
times. utime: set	file access and modification	utime(2)
ldfcn: common object	file access routines.	ldfcn(4)
determine accessibility of a	file. access:	access(2)
tar: tape	file archiver.	tar(1)
cpio: copy	file archives in and out.	cpio(1)
mkstr: create an error message	file by massaging C source.	mkstr(1)
pwck, grpck: password/group	file checkers.	pwck(1M)
chmod: change mode of	file.	chmod(2)
change owner and group of a	file. chown:	chown(2)
diff: differential	file comparator.	diff(1)
diff3: 3-way differential	file comparison.	diff3(1)
fcntl:	file control.	fcntl(2)
fcntl:	file control options.	fcntl(5)
conv: object	file converter.	conv(1)
rcp: remote	file copy.	rcp(1N)
public UNIX-to-UNIX system	file copy. uuto, uupick:	uuto(1C)
core: format of core image	file.	core(4)
umask: set and get	file creation mask.	umask(2)
crontab: user crontab	file.	crontab(1)
fields of each line of a	file. cut: cut out selected	cut(1)
dd: convert and copy a	file.	dd(1)
a delta (change) to an SCCS	file. delta: make	delta(1)
close: close a	file descriptor.	close(2)
	file: determine file type.	file(1)
selected parts of an object	file. dump: dump	dump(1)
sact: print current SCCS	file editing activity.	sact(1)
/fgetgrent: obtain group	file entry from a group file.	getgrent(3C)
fgetpwent: get password	file entry. /endpwent,	getpwent(3C)
utmpname: access utmp	file entry. /endutent,	getut(3C)
putpwent: write password	file entry.	putpwent(3C)
execlp, execlvp: execute a	file. /execv, execl, execve,	exec(2)

ctags: maintain a tags	file for a C program. . . . .	ctags(1)
grep, egrep, fgrep: search a	file for a pattern. . . . .	grep(1)
aliases: aliases	file for delivermail. . . . .	aliases(7N)
ldopen: open a common object	file for reading. ldopen. . . . .	ldopen(3X)
acct: per-process accounting	file format. . . . .	acct(4)
ar: common archive	file format. . . . .	ar(4)
ar5.0: archive (library)	file format. . . . .	ar5.0(4)
errfile: error-log	file format. . . . .	errfile(4)
pnch: . . . . .	file format for card images. . . . .	pnch(4)
intro: introduction to	file formats. . . . .	intro(4)
on character frequencies in a	file. freq: report . . . . .	freq(1)
take: takes a	file from a remote machine. . . . .	take(1C)
entries of a common object	file function. /line number . . . . .	ldread(3X)
get: get a version of an SCCS	file. . . . .	get(1)
group file entry from a group	file. /fgetgrent: obtain . . . . .	getgrent(3C)
group: group	file. . . . .	group(4)
files. filehdr:	file header for common object . . . . .	filehdr(4)
file. ldhread: read the	file header of a common object . . . . .	ldhread(3X)
ldhseek: seek to the optional	file header of a common object/	ldhseek(3X)
split: split a	file into pieces. . . . .	split(1)
issue: issue identification	file. . . . .	issue(4)
of a member of an archive	file. /read the archive header . . . . .	ldhread(3X)
close a common object	file. ldclose, ldclose: . . . . .	ldclose(3X)
file header of a common object	file. ldhread: read the . . . . .	ldhread(3X)
symbol name for object	file. ldgetname: retrieve . . . . .	ldgetname(3X)
a section of a common object	file. /line number entries of . . . . .	ldlseek(3X)
file header of a common object	file. /seek to the optional . . . . .	ldhseek(3X)
a section of a common object	file. /relocation entries of . . . . .	ldrseek(3X)
header of a common object	file. /indexed/named section . . . . .	ldshread(3X)
section of a common object	file. /to an indexed/named . . . . .	ldsseek(3X)
table entry of a common object	file. /the index of a symbol . . . . .	ldtbindex(3X)
table entry of a common object	file. /read an indexed symbol . . . . .	ldtbread(3X)
table of a common object	file. /seek to the symbol . . . . .	ldtbseek(3X)
entries in a common object	file. linenum: line number . . . . .	linenum(4)
link: link to a	file. . . . .	link(2)
mknod: build special	file. . . . .	mknod(1M)
or a special or ordinary	file. /make a directory, . . . . .	mknod(2)
a file system. ff: list	file names and statistics for . . . . .	ff(1M)
change the format of a text	file. newform: . . . . .	newform(1)
name list of common object	file. nm: print . . . . .	nm(1)
null: the null	file. . . . .	null(7)
/find the slot in the utmp	file of the current user. . . . .	ttyslot(3C)
put: puts a	file onto a remote machine. . . . .	put(1C)
/identify processes using a	file or file structure. . . . .	fuser(1M)
one. creat: create a new	file or rewrite an existing . . . . .	creat(2)
passwd: password	file. . . . .	passwd(4)
or subsequent lines of one	file. /lines of several files . . . . .	paste(1)
viewing. more:	file perusal filter for crt . . . . .	more(1)
soft-copy terminals. pg:	file perusal filter for . . . . .	pg(1)
/rewind, ftell: reposition a	file pointer in a stream. . . . .	fseek(3S)
lseek: move read/write	file pointer. . . . .	lseek(2)
prs: print an SCCS	file. . . . .	prs(1)
from downloading into a	file. /Motorola S-records . . . . .	rcvhex(1)
read: read from	file. . . . .	read(2)
readv: read from	file. . . . .	readv(3N)
locking: provide exclusive	file regions for reading or/	locking(2)
for a common object	file. /relocation information . . . . .	reloc(4)
remove a delta from an SCCS	file. rmdel: . . . . .	rmdel(1)
bfs: big	file scanner. . . . .	bfs(1)
two versions of an SCCS	file. secsdiff: compare . . . . .	secsdiff(1)
secsfile: format of SCCS	file. . . . .	secsfile(4)
header for a common object	file. scnhdr: section . . . . .	scnhdr(4)
size5.0: size of an object	file. . . . .	size5.0(1)
stat, fstat: get	file status. . . . .	stat(2)

in an object, or other binary information from an object	file. /the printable strings	strings(1)
processes using a file or checksum and block count of a sum and count blocks in a syms: common object	file. /symbol and line number	strip(1)
tapesave: daily/weekly UNIX procedure. checkall: faster and interactive/ fsck, dfsck: fsdb: names and statistics for a volume.	file structure. /identify	fuser(1M)
mkfs1b: construct a mkfs: construct a umount: mount and dismount mount: mount a ustat: get mnttab: mounted umount: unmount a access time. dcopy: copy fsck. checklist: list of volcopy, labelit: copy deliver the last part of a term: format of compiled term tmpfile: create a temporary create a name for a temporary and modification times of a ftp: ftpd: DARPA Internet tftpd: DARPA Trivial ftw: walk a file: determine undo a previous get of an SCCS report repeated lines in a val: validate SCCS write: write on a writev: write on a umask: set common object files. ctermid: generate mktemp: make a unique ferror, feof, clearerr, and print process accounting merge or add total accounting create and administer SCCS a.out header for common object VAX-11/780/ fscv: convert updater: update updater: update cat: concatenate and print cmp: compare two lines common to two sorted cp, ln, mv: copy, link or move mark differences between file header for common object find: find frec: recover format specification in text split f77, ratfor, or efl hex: translates object cpset: install object and count characters in the intro: introduction to special link editor for common object	file system backup. filesave, file system checking file system consistency check file system debugger. file system. ff: list file file system: format of system file system. file system. file system. mount, file system. file system statistics. file system table. file system. file systems for optimal file systems processed by file systems with label/ file. tail: file.. file. file. tmpnam, tempnam: file. touch: update access file transfer program. File Transfer Protocol server. File Transfer Protocol server. file tree. file type. file. unget: file. uniq: file. file. file. file-creation mode mask. filehdr: file header for filename for terminal. filename. fileno: stream status/ file(s). acctcom: search files. acctmerg: files. admin: files. aouthdr.h - files between M68000 and files between two machines. files between two machines. files. files. files. comm: select or reject files. files. diffmk: files. filehdr: files. files from a backup tape. files. fspec: files. split: files. files in binary directories. files in the given/ /sum files. files. ld:	sum(1) sum7(1) syms(4) filesave(1M) checkall(1M) fsck(1M) fsdb(1M) ff(1M) fs(4) mkfs1b(1M) mkfs(1M) mount(1M) mount(2) ustat(2) mnttab(4) umount(2) dcopy(1M) checklist(4) volcopy(1M) tail(1) term(4) tmpfile(3S) tmpnam(3S) touch(1) ftp(1N) ftpd(8N) tftpd(8N) ftw(3C) file(1) unget(1) uniq(1) val(1) write(3) writev(3N) umask(1) filehdr(4) ctermid(3S) mktemp(3C) ferror(3S) acctcom(1) acctmerg(1M) admin(1) aouthdr(4) fscv(1M) updater(1) updater(1M) cat(1) cmp(1) comm(1) cp(1) diffmk(1) filehdr(4) find(1) frec(1M) fspec(4) split(1) hex(1) cpset(1M) sumdir(1) intro(7) ld(1)

Permuted Index

lockf: record locking on	files.	lockf(3C)
rm, rmdir: remove	files or directories.	rm(1)
/merge same lines of several	files or subsequent lines of/	paste(1)
unpack: compress and expand	files. pack, pcat,	pack(1)
pr: print	files.	pr(1)
section sizes of common object	files. size: print	size(1)
sort: sort and/or merge	files.	sort(1)
reports version number of	files. version:	version(1)
what: identify SCCS	files.	what(1)
daily/weekly UNIX file system/	filesave, tapcsave:	filesave(1M)
more: file perusal	filter for crt viewing.	more(1)
terminals. pg: file perusal	filter for soft-copy	pg(1)
greek: select terminal	filter.	greek(1)
nl: line numbering	filter.	nl(1)
col:	filter reverse line-feeds.	col(1)
tplot: graphics	filters.	tplot(1G)
find:	find: fast incremental backup.	find(1M)
hyphen:	find files.	find(1)
ttyname, isatty:	find: find files.	find(1)
object library. lorder:	find hyphenated words.	hyphen(1)
object library. lorder5.0:	find name of a terminal.	ttyname(3C)
hashmake, spellin, hashcheck:	find ordering relation for an	lorder(1)
an object, or other/ strings:	find ordering relation for an	lorder5.0(1)
of the current user. ttypslot:	find spelling errors. spell,	spell(1)
fish: play "Go	find the printable strings in	strings(1)
tee: pipe	find the slot in the utmp file	ttypslot(3C)
/seekdir, rewinddir, closedir:	Fish".	fish(6)
int, ifix, idint, real,	fish: play "Go Fish".	fish(6)
atof: convert ASCII string to	fitting.	tee(1)
ecvt, fcvt, gvvt: convert	flexible length directory/	directory(3X)
/modf: manipulate parts of	float, snfl, dble, cmplx,/	stype(3F)
floor, ceiling, remainder,/	floating-point number.	atof(3C)
floor, ceil, fmod, fabs:	floating-point number to/	ecvt(3C)
parameters. disk tune: tune	floating-point numbers.	frexp(3C)
cflow: generate C	floor, ceil, fmod, fabs:	floor(3M)
fclose, fflush: close or	floor, ceiling, remainder,/	floor(3M)
remainder,/ floor, ceil,	floppy disk settling time	disktune(1M)
stream.	flowgraph.	cflow(1)
diskformat:	flush a stream.	fclose(3S)
per-process accounting file	fmod, fabs: floor, ceiling,	floor(3M)
ar: common archive file	fopen, freopen, fdopen: open a	fopen(3S)
ar5.0: archive (library) file	fork: create a new process.	fork(2)
errfile: error-log file	format a disk.	diskformat(1M)
punch: file	format. acct:	acct(4)
nroff or/ eqn, neqn, checkeq:	format.	ar(4)
newform: change the	format.	ar5.0(4)
inode:	format.	errfile(4)
term:	format for card images.	punch(4)
core:	format mathematical text for	eqn(1)
cpio:	format of a text file.	newform(1)
dir:	format of an inode.	inode(4)
sccsfile:	format of compiled term file..	term(4)
file system:	format of core image file.	core(4)
files. fspec:	format of cpio archive.	cpio(4)
object file symbol table	format of directories.	dir(4)
troff. tbl:	format of SCCS file.	sccsfile(4)
nroff:	format of system volume.	fs(4)
intro: introduction to file	format specification in text	fspec(4)
wtmp: utmp and wtmp entry	format. syms: common	syms(4)
scanf, fscanf, sscanf: convert	format tables for nroff or	tbl(1)
	format text.	nroff(1)
	formats.	intro(4)
	formats. utmp,	utmp(4)
	formatted input.	scanf(3S)



/vfprintf, vsprintf: print	formatted output of a varargs/	vprintf(3S)
/vfprintf, vsprintf: print	formatted output of a varargs/	vprintf(3X)
fprintf, sprintf: print	formatted output. printf,	printf(3S)
/checkmm: print/check documents	formatted with the MM macros.	mm(1)
mptx: the macro package for	formatting a permuted index.	mptx(5)
mm: the MM macro package for	formatting documents.	mm(5)
OSDD adapter macro package for	formatting documents. /the	mostd(5)
manual. man: macros for	formatting entries in this	man(5)
f77:	Fortran 77 compiler.	f77(1)
abs, iabs, dabs, cabs, zabs:	Fortran absolute value.	abs(3F)
system/ signal: specify	Fortran action on receipt of a	signal(3F)
function. acos, dacos:	Fortran arccosine intrinsic	acos(3F)
function. asin, dasin:	Fortran arcsine intrinsic	asin(3F)
function. atan2, datan2:	Fortran arctangent intrinsic	atan2(3F)
function. atan, datan:	Fortran arctangent intrinsic	atan(3F)
or, xor, not, lshift, rshift:	Fortran bitwise boolean/ and,	bool(3F)
getarg: return	Fortran command-line argument.	getarg(3F)
log10, alog10, dlog10:	Fortran common logarithm/	log10(3F)
intrinsic/ conjg, dconjg:	Fortran complex conjugate	conjg(3F)
function. cos, dcos, ccos:	Fortran cosine intrinsic	cos(3F)
ratfor: rational	Fortran dialect.	ratfor(1)
getenv: return	Fortran environment variable.	getenv(3F)
function. exp, dexp, cexp:	Fortran exponential intrinsic	exp(3F)
intrinsic/ cosh, dcosh:	Fortran hyperbolic cosine	cosh(3F)
intrinsic/ sinh, dsinh:	Fortran hyperbolic sine	sinh(3F)
intrinsic/ tanh, dtanh:	Fortran hyperbolic tangent	tanh(3F)
complex/ aimag, dimag:	Fortran imaginary part of	aimag(3F)
function. aint, dint:	Fortran integer part intrinsic	aint(3F)
efl: Extended	Fortran Language.	efl(1)
amax0, max1, amax1, dmax1:	Fortran maximum-value/ /max0,	max(3F)
amin0, min1, amin1, dmin1:	Fortran minimum-value/ /min0,	min(3F)
log, alog, dlog, clog:	Fortran natural logarithm/	log(3F)
anint, dnint, nint, idnint:	Fortran nearest integer/	round(3F)
abort: terminate	Fortran program.	abort(3F)
functions. mod, amod, dmod:	Fortran remaindering intrinsic	mod(3F)
function. sin, dsin, csin:	Fortran sine intrinsic	sin(3F)
function. sqrt, dsqrt, csqrt:	Fortran square root intrinsic	sqrt(3F)
len: return length of	Fortran string.	len(3F)
index: return location of	Fortran substring.	index(3F)
issue a shell command from	Fortran. system:	system(3F)
function. tan, dtan:	Fortran tangent intrinsic	tan(3F)
mclock: return	Fortran time accounting.	mclock(3F)
intrinsic/ sign, isign, dsign:	Fortran transfer-of-sign	sign(3F)
/dcmplx, ichar, char: explicit	Fortran type conversion.	ftype(3F)
irand, srand, rand:	Fortran uniform random-number/	rand(3F)
hopefully interesting. adage.	fortune: print a random,	fortune(6)
formatted output. printf,	fprintf, sprintf: print	printf(3S)
word on a/putc, putchar,	fputc, putw: put character or	putc(3S)
stream. puts,	fputs: put a string on a	puts(3S)
input/output.	fread, fwrite: binary	fread(3S)
backup tape.	frec: recover files from a	frec(1M)
df: report number of	free disk blocks.	df(1M)
memory allocator. malloc,	free, realloc, calloc: main	malloc(3C)
mallopt, mallinfo:/ malloc,	free, realloc, calloc,	malloc(3X)
stream. fopen,	freopen, fdopen: open a	fopen(3S)
frequencies in a file.	freq: report on character	freq(1)
freq: report on character	frequencies in a file.	freq(1)
parts of floating-point/	frexp, ldexp, modf: manipulate	frexp(3C)
freq: recover files	from a backup tape.	frec(1M)
obtain group file entry	from a group file. /fgetgrent:	getgrent(3C)
remque: insert/remove element	from a queue. insque,	insque(3N)
take: takes a file	from a remote machine.	take(1C)
recvmsg: receive a message	from a socket. /recvfrom,	recv(2N)
sendmsg: send a message	from a socket. send, sendto,	send(2N)

getw: get character or word	from a stream. /fgetc, . . . . .	getc(3S)
gets, fgets: get a string	from a stream. . . . .	gets(3S)
and line number information	from an object file. /symbol . . . . .	strip(1)
rmdei: remove a delta	from an SCCS file. . . . .	rmdei(1)
getopt: get option letter	from argument vector. . . . .	getopt(3C)
shared/ xstr: extract strings	from C programs to implement	xstr(1)
/translates Motorola S-records	from downloading into a file. . . . .	rvhex(1)
errdead: extract error records	from dump. . . . .	errdead(1M)
read: read	from file. . . . .	read(2)
readv: read	from file. . . . .	readv(3N)
system: issue a shell command	from Fortran. . . . .	system(3F)
ncheck: generate names	from i-numbers. . . . .	ncheck(1M)
nlist: get entries	from name list. . . . .	nlist(3C)
acctcms: command summary	from per-process accounting/	acctcms(1M)
autorobots: Escape	from the automatic robots. . . . .	autorobots(6)
robots: Escape	from the robots. . . . .	robots(6)
getpw: get name	from UID. . . . .	getpw(3C)
formatted input. scanf,	fscanf, sscanf: convert	scanf(3S)
of file systems processed by	fsck. checklist: list . . . . .	checklist(4)
consistency check and/	fsck, dfck: file system . . . . .	fsck(1M)
a lost + found directory for	fsck, mklost + found: make	mklost + fnd(1M)
M68000 and VAX-11/780/	fsck: convert files between	fsck(1M)
	fsdb: file system debugger. . . . .	fsdb(1M)
reposition a file pointer in/	fseek, rewind, ftell: . . . . .	fseek(3S)
text files.	fspec: format specification in	fspec(4)
ell files.	fsplit: split f77, ratfor, or	fsplit(1)
stat.	fstat: get file status. . . . .	stat(2)
pointer in a/ fseek, rewind,	ftell: reposition a file . . . . .	fseek(3S)
communication package.	ftok: standard interprocess	stdipc(3C)
	ftp: file transfer program. . . . .	ftp(1N)
Transfer Protocol server.	ftpd: DARPA Internet File	ftpd(8N)
	ftw: walk a file tree. . . . .	ftw(3C)
shutdown: shut down part of a	full-duplex connection. . . . .	shutdown(2N)
Fortran arccosine intrinsic	function. acos, dacos: . . . . .	acos(3F)
Fortran integer part intrinsic	function. aint, dint: . . . . .	aint(3F)
error/ erf, erfc: error	function and complementary	erf(3M)
Fortran arcsine intrinsic	function. asin, dasin: . . . . .	asin(3F)
Fortran arctangent intrinsic	function. atan2, datan2: . . . . .	atan2(3F)
Fortran arctangent intrinsic	function. atan, datan: . . . . .	atan(3F)
complex conjugate intrinsic	function. /dconjg: Fortran . . . . .	conjg(3F)
ccos: Fortran cosine intrinsic	function. cos, dcos, . . . . .	cos(3F)
hyperbolic cosine intrinsic	function. /dcosh: Fortran . . . . .	cosh(3F)
precision product intrinsic	function. dprod: double . . . . .	dprod(3F)
and complementary error	function. /error function . . . . .	erf(3M)
Fortran exponential intrinsic	function. exp, dexp, cexp: . . . . .	exp(3F)
gamma: log gamma	function. . . . .	gamma(3M)
hypot: Euclidean distance	function. . . . .	hypot(3M)
of a common object file	function. /line number entries . . . . .	ldread(3X)
common logarithm intrinsic	function. /dlog10: Fortran . . . . .	log10(3F)
natural logarithm intrinsic	function. /dlog, clog: Fortran . . . . .	log(3F)
matherr: error-handling	function. . . . .	matherr(3M)
prof: profile within a	function. . . . .	prof(5)
transfer-of-sign intrinsic	function. /dsign: Fortran . . . . .	sign(3F)
csin: Fortran sine intrinsic	function. sin, dsin, . . . . .	sin(3F)
hyperbolic sine intrinsic	function. /dsinh: Fortran . . . . .	sinh(3F)
Fortran square root intrinsic	function. sqrt, dsqrt, csqrt: . . . . .	sqrt(3F)
Fortran tangent intrinsic	function. tan, dtan: . . . . .	tan(3F)
hyperbolic tangent intrinsic	function. /dtanh: Fortran . . . . .	tanh(3F)
math: math	functions and constants. . . . .	math(5)
j0, j1, jn, y0, y1, yn: Bessel	functions. . . . .	bessel(3M)
Fortran bitwise boolean	functions. /lshift, rshift: . . . . .	bool(3F)
positive difference intrinsic	functions. dim, ddim, idim: . . . . .	dim(3M)
logarithm, power, square root	functions. /sqrt: exponential, . . . . .	exp(3M)
remainder, absolute value	functions. /floor, ceiling, . . . . .	floor(3M)

dmaxl: Fortran maximum-value	functions. /maxl, amaxl, . . . . .	max(3F)
dminl: Fortran minimum-value	functions. /minl, aminl, . . . . .	min(3F)
Fortran remaindering intrinsic	functions. mod, amod, dmod: . . . . .	mod(3F)
300, 300s: handle special	functions of DASI 300 and 300s/ . . . . .	300(1)
terminal. 450: handle special	functions of the DASI 450 . . . . .	450(1)
Fortran nearest integer	functions. /nint, idnint: . . . . .	round(3F)
sinh, cosh, tanh: hyperbolic	functions. . . . .	sinh(3M)
string comparison intrinsic	functions. /lgt, lle, llt: . . . . .	strcmp(3F)
atan, atan2: trigonometric	functions. /tan, asin, acos, . . . . .	trig(3M)
using a file or file/	fuser: identify processes . . . . .	fuser(1M)
fread,	fwrite: binary input/output . . . . .	fread(3S)
connect accounting records.	fwtmp, wtmpfix: manipulate . . . . .	fwtmp(1M)
adventure: an exploration	game. . . . .	adventure(6)
cribbage: the card	game cribbage. . . . .	cribbage(6)
moo: guessing	game. . . . .	moo(6)
back: the	game of backgammon. . . . .	back(6)
bj: the	game of black jack. . . . .	bj(6)
craps: the	game of craps. . . . .	craps(6)
wump: the	game of hunt-the-wumpus. . . . .	wump(6)
life: play the	game of life. . . . .	life(6)
trek: trekkie	game. . . . .	trek(6)
worm: Play the growing worm	game. . . . .	worm(6)
intro: introduction to	games. . . . .	intro(6)
gamma: log	gamma function. . . . .	gamma(3M)
number to string. ecvt, fcvt,	gamma: log gamma function. . . . .	gamma(3M)
maze:	gcvt: convert floating-point . . . . .	ecvt(3C)
abort:	generate a maze. . . . .	maze(6)
cflow:	generate an IOT fault. . . . .	abort(3C)
cross-reference. cxref:	generate C flowgraph. . . . .	cflow(1)
crypt, setkey, encrypt:	generate C program . . . . .	cxref(1)
by user ID. diskusg:	generate DES encryption. . . . .	crypt(3C)
makekey:	generate disk accounting data . . . . .	diskusg(1M)
terminal. ctermid:	generate encryption key. . . . .	makekey(1)
ncheck:	generate filename for . . . . .	ctermid(3S)
lexical tasks. lex:	generate names from i-numbers. . . . .	ncheck(1M)
/srand48, seed48, lcong48:	generate programs for simple . . . . .	lex(1)
srand: simple random-number	generate uniformly distributed/ . . . . .	drand48(3C)
Fortran uniform random-number	generator. rand, . . . . .	rand(3C)
gets, fgets:	generator. /srand, rand: . . . . .	rand(3F)
get:	get a string from a stream. . . . .	gets(3S)
getsockopt, setsockopt:	get a version of an SCCS file. . . . .	get(1)
ulimit:	get and set options on/ . . . . .	getsockopt(2N)
the user. cuserid:	get and set user limits. . . . .	ulimit(2)
getc, getchar, fgets, getw:	get character login name of . . . . .	cuserid(3S)
getdtablesize:	get character or word from a/ . . . . .	getc(3S)
nlist:	get descriptor table size. . . . .	getdtablesize(3N)
umask: set and	get entries from name list. . . . .	nlist(3C)
stat, fsstat:	get file creation mask. . . . .	umask(2)
ustat:	get file status. . . . .	stat(2)
file:	get file system statistics. . . . .	ustat(2)
getlogin:	get: get a version of an SCCS . . . . .	get(1)
logname:	get login name. . . . .	getlogin(3C)
msgget:	get login name. . . . .	logname(1)
getpw:	get message queue. . . . .	msgget(2)
getpeername:	get name from UID. . . . .	getpw(3C)
system. uname:	get name of connected peer. . . . .	getpeername(2N)
/setnetent, endnetent:	get name of current UNIX . . . . .	uname(2)
/sethostent, endhostent:	get network entry. . . . .	getnetent(3N)
unset:	get network host entry. . . . .	gethostent(3N)
argument vector. getopt:	get of an SCCS file. . . . .	unset(1)
/setsptent, endsptent, fgetpwent:	get option letter from . . . . .	getopt(3C)
working directory. getcwd:	get password file entry. . . . .	getpwent(3C)
times. times:	get pathname of current . . . . .	getcwd(3C)
	get process and child process . . . . .	times(2)

Permuted Index

and/ getpid, getpgrp, getppid:	get process, process group, . . . . .	getpid(2)
/setprotoent, endprotoent:	get protocol entry. . . . .	getprotoent(3N)
/geteuid, getgid, getegid:	get real user, effective user,/ . . . . .	getuid(2)
/setservernt, endservernt:	get service entry. . . . .	getservernt(3N)
semget:	get set of semaphores. . . . .	semget(2)
shmget:	get shared memory segment. . . . .	shmget(2)
getsockname:	get socket name. . . . .	getsockname(2N)
tty:	get the terminal's name. . . . .	tty(1)
time:	get time. . . . .	time(2)
command-line argument.	getarg: return Fortran . . . . .	getarg(3F)
get character or word from a/	getc, getchar, fgetc, getw:	getc(3S)
character or word from/ getc,	getchar, fgetc, getw: get . . . . .	getc(3S)
current working directory.	getcwd: get pathname of . . . . .	getcwd(3C)
table size.	getdtablesize: get descriptor . . . . .	getdtablesize(3N)
getuid, geteuid, getgid,	getegid: get real user,/ . . . . .	getuid(2)
environment variable.	getenv: return Fortran . . . . .	getenv(3F)
environment name.	getenv: return value for . . . . .	getenv(3C)
real user, effective/ getuid,	geteuid, getgid, getegid: get . . . . .	getuid(2)
user,/ getuid, geteuid,	getgid, getegid: get real . . . . .	getuid(2)
setgrent, endgrent,/	getgrent, getgrgid, getgrnam, . . . . .	getgrent(3C)
endgrent,/ getgrent,	getgrgid, getgrnam, setgrent, . . . . .	getgrent(3C)
getgrent, getgrgid,	getgrnam, setgrent, endgrent,/ . . . . .	getgrent(3C)
sethostent,/ gethostent,	gethostbyaddr, gethostbyname, . . . . .	gethostent(3N)
gethostent, gethostbyaddr,	gethostbyname, sethostent,/ . . . . .	gethostent(3N)
gethostbyname, sethostent,/	gethostent, gethostbyaddr, . . . . .	gethostent(3N)
unique identifier of current/	gethostid, sethostid: get/set . . . . .	gethostid(2N)
get/set name of current host.	gethostname, sethostname: . . . . .	gethostname(2N)
setnetent,/ getnetent,	getlogin: get login name. . . . .	getlogin(3C)
getnetent, getnetbyaddr,	getnetbyaddr, getnetbyname, . . . . .	getnetent(3N)
getnetbyname, setnetent,/	getnetbyname, setnetent,/ . . . . .	getnetent(3N)
argument vector.	getnetent, getnetbyaddr, . . . . .	getnetent(3N)
connected peer.	getopt: get option letter from . . . . .	getopt(3C)
process group, and/ getpid,	getopt: parse command options. . . . .	getopt(1)
process, process group, and/	getpass: read a password. . . . .	getpass(3C)
group, and/ getpid, getpgrp,	getpeername: get name of . . . . .	getpeername(2N)
getprotoent, getprotobynumber,	getpgrp, getppid: get process, . . . . .	getpid(2)
getprotobyname,/ getprotoent,	getpid, getpgrp, getppid: get . . . . .	getpid(2)
getprotobyname, setprotoent,/	getppid: get process, process . . . . .	getpid(2)
getprotobyname, setprotoent,/	getprotoent, setprotoent,/ . . . . .	getprotoent(3N)
setpwent, endpwent,/	getprotobynumber, . . . . .	getprotoent(3N)
getpwent, getpwuid,	getprotoent, getprotobynumber, . . . . .	getprotoent(3N)
endpwent,/ getpwent,	getpw: get name from UID. . . . .	getpw(3C)
a stream.	getpwent, getpwuid, getpwnam, . . . . .	getpwent(3C)
getservernt, getservbyport,	getpwnam, setpwent, endpwent,/ . . . . .	getpwent(3C)
setservnt,/ getservnt,	getpwuid, getpwnam, setpwent, . . . . .	getpwent(3C)
getservbyname, setservnt,/	gets, fgets: get a string from . . . . .	gets(3S)
getservbyname, setservnt,/	getservbyname, setservnt,/ . . . . .	getservnt(3N)
gethostname, sethostname:	getservbyport, getservbyname, . . . . .	getservnt(3N)
current/ gethostid, sethostid:	getservernt, getservbyport, . . . . .	getservnt(3N)
and set options on sockets.	get/set name of current host. . . . .	gethostname(2N)
and terminal settings used by	get/set unique identifier of . . . . .	gethostid(2N)
modes, speed, and line/	getsockname: get socket name. . . . .	getsockname(2N)
ct: spawn	getsockopt, setsockopt: get . . . . .	getsockopt(2N)
settings used by getty.	getty. gettydefs: speed . . . . .	gettydefs(4)
getegid: get real user,/	getty: set terminal type, . . . . .	getty(1M)
pututline, setutent,/	getty to a remote terminal. . . . .	ct(1C)
setutent, endutent,/ getutent,	gettydefs: speed and terminal . . . . .	gettydefs(4)
setutent,/ getutent, getutid,	getuid, geteuid, getgid, . . . . .	getuid(2)
from a/ getc, getchar, fgetc,	getutent, getutid, getutline, . . . . .	getut(3C)
convert/ ctime, localtime,	getutid, getutline, pututline, . . . . .	getut(3C)
fish: play	getutline, pututline, . . . . .	getut(3C)
	getw: get character or word . . . . .	getc(3S)
	gmtime, asctime, tzset: . . . . .	ctime(3C)
	“Go Fish”. . . . .	fish(6)

setjmp, longjmp: non-local	goto.	setjmp(3C)
graph: draw a	graph: draw a graph.	graph(1G)
sag: system activity	graph.	graph(1G)
tplot:	graphics filters.	sag(1G)
TTY-37 type-box. greek:	graphics for the extended	tplot(1G)
plot:	graphics interface.	greek(5)
subroutines. plot:	graphics interface	plot(4)
mvt: typeset documents, view	graphs, and slides. mmt.	plot(3X)
package for typesetting view	graphs and slides. /macro	mnt(1)
extended TTY-37 type-box.	greek: graphics for the	mv(5)
	greek: select terminal filter.	greek(5)
file for a pattern.	grep, egrep, fgrep: search a	greek(1)
/user, effective user, real	group, and effective group/	grep(1)
/getppid: get process, process	group, and parent process IDs.	getuid(2)
chown, chgrp: change owner or	group.	getpid(2)
/endgrent, fgetgrent: obtain	group file entry from a group/	chown(1)
obtain group file entry from a	group file. /fgetgrent:	getgrent(3C)
group:	group file.	getgrent(3C)
	group: group file.	group(4)
setpgrp: set process	group ID.	group(4)
set real and effective	group ID. setregid:	setpgrp(2)
id: print user and	group IDs and names.	setregid(2)
real group, and effective	group IDs. /effective user,	id(1)
setuid, setgid: set user and	group IDs.	getuid(2)
send signal to a process	group. killpg:	setuid(2)
newgrp: log in to a new	group.	killpg(3N)
chown: change owner and	group of a file.	newgrp(1)
a signal to a process or a	group of processes. /send	chown(2)
update, and regenerate	groups of programs. /maintain,	kill(2)
worm: Play the	growing worm game.	make(1)
checkers. pwck,	grpck: password/group file	worm(6)
ssignal,	gsignal: software signals.	pwck(1M)
hangman:	guess the word.	ssignal(3C)
moo:	guessing game.	hangman(6)
DASI 300 and 300s/ 300, 300s:	handle special functions of	moo(6)
the DASI 450 terminal. 450:	handle special functions of	300(1)
varargs:	handle variable argument list.	450(1)
information for bad block	handling. /alternate block	varargs(5)
package. curses: CRT screen	handling and optimization	altblk(4)
	hangman: guess the word.	curses(3X)
	hangups (sh only).	hangman(6)
nohup: run a command immune to	hash search tables. hsearch,	nohup(1)
hcreate, hdestroy: manage	hashcheck: find spelling/	hsearch(3C)
spell, hashmake, spellin,	hashmake, spellin, hashcheck:	spell(1)
find spelling errors. spell,	hcreate, hdestroy: manage hash	spell(1)
search tables. hsearch,	hdestroy: manage hash search	hsearch(3C)
tables. hsearch, hcreate,	header for a common object	hsearch(3C)
file. scnhdr: section	header for a common object	scnhdr(4)
files. aouthdr.h - a.out	header for common object	aouthdr(4)
files. filehdr: file	header of a common object	filehdr(4)
file. ldhread: read the file	header of a common object/	ldhread(3X)
/seek to the optional file	header of a common object/	ldhseek(3X)
/read an indexed/named section	header of a member of an/	ldhread(3X)
ldahread: read the archive	SCCS.	ldahread(3X)
SCCS.	help: ask for help in using	help(1)
help: ask for	help in using SCCS.	help(1)
	hex: translates object files.	hex(1)
fortune: print a random,	hopefully interesting, adage.	fortune(6)
/ntohs: convert values between	host and network byte order.	byteorder(3N)
endhostent: get network	host entry. /sethostent,	gethostent(3N)
unique identifier of current	hostl. /sethostid: get/set	gethostid(2N)
get/set name of current	host. /sethostname:	gethostname(2N)
hosts:	host name data base.	hosts(4N)
ruptime: show	host status of local machines.	ruptime(1N)

*Permuted Index*

or print identifier of current host system. `hostid`: set  
 set or print name of current host system. `hostname`:  
 identifier of current host/ `hostid`: set or print  
 current host system. `hostname`: set or print name of  
 hosts: host name data base.  
 manage hash search tables. `hsearch`, `hcreate`, `hdestroy`:  
 convert values between host/ `htonl`, `htons`, `ntohl`, `ntohs`:  
 values between host/ `htonl`, `htons`, `ntohl`, `ntohs`: convert  
     `wump`: the game of `wump`:  
     `cosh`, `dcosh`: Fortran hyperbolic cosine intrinsic/  
     `sinh`, `cosh`, `tanh`: hyperbolic functions.  
     `sinh`, `dsinh`: Fortran hyperbolic sine intrinsic/  
     `tanh`, `dtanh`: Fortran hyperbolic tangent intrinsic/  
     `hyphen`: find hyphenated words.  
     `hyphenated words`:  
     `hypot`: Euclidean distance  
 Fortran absolute value. `abs`, `iabs`, `dabs`, `cabs`, `zabs`:  
     `iargc`:  
     `ichar`, `char`: explicit Fortran/  
     `ID`. `diskusg`: generate  
     `id`. /remove a message queue.  
     `id`: print user and group IDs  
     `ID`.  
     `ID`. `setregid`:  
     `id`. `whoami`:  
     identification file.  
     identifier of current host.  
     identifier of current host.  
     identify processes using a  
     identify SCCS files.  
     intrinsic/ `dim`, `ddim`,  
     `dim`: positive difference  
     `idint`, `real`, `float`, `sngl`,  
     integer/ `anint`, `dnint`, `nint`,  
     `idnint`: Fortran nearest  
     IDs and names.  
     IDs. /get process, process  
     IDs. /effective user. `real`  
     ID's. `setreuid`:  
     IDs. `setuid`,  
     interface parameters.  
     `sngl`, `dbl`, `cmplx`, / `int`,  
     `iflx`, `idint`, `real`, `float`,  
     `core`: format of core  
     `pnc`: file format for card  
     `aimag`, `dimag`: Fortran  
     `nohup`: run a command  
     /strings from C programs to  
     implement shared strings.  
     incremental backup.  
     independent fashion. /access  
     independent operation/  
     index. /the macro package  
     index of a symbol table entry  
     index.  
     index: return location of  
     indexed symbol table entry of  
     indexed/named section header/  
     indexed/named section of a/  
     indicate last logins of users  
     `inet`: Internet protocol  
     `inet_addr`, `inet_network`,  
     `inet_1naof`, `inet_netof`:  
     `inet_makeaddr`, `inet_1naof`/  
     `inet_netof`: Internet address/  
     `inet_network`, `inet_ntoa`,  
     `inet_addr`, `inet_network`,  
     `inet_ntoa`, `inet_makeaddr`/  
     `initab`: script for the

initialization. . . . . init, telinit: process control . . . . . init(1M)  
 init, telinit: process control . . . . . init(1M)  
   /rc, powerfail: system . . . . . brc(1M)  
   socket. connect: . . . . . connect(2N)  
   process. popen, pclose: . . . . . popen(3S)  
   process. . . . . inittab(4)  
   cli: clear . . . . . cli(1M)  
   inode: format of an . . . . . inode(4)  
   inode. . . . . inode(4)  
   sscanf: convert formatted . . . . . scanf(3S)  
   push character back into . . . . . scanf(3S)  
   fread, fwrite: binary . . . . . ungetc(3S)  
   stdio: standard buffered . . . . . fread(3S)  
   fileno: stream status . . . . . stdio(3S)  
   ustat: uucp status . . . . . ferror(3S)  
   queue. insque, remque: . . . . . ustat(1C)  
   element from a queue. . . . . insque(3N)  
   install: . . . . . insque(3N)  
   install commands. . . . . install(1M)  
   install: install commands. . . . . install(1M)  
   install object files in binary . . . . . cpset(1M)  
   int, ifix, idint, real, float, . . . . . ftype(3F)  
   integer absolute value. . . . . abs(3C)  
   integer and base-64 ASCII/ . . . . . a64(3C)  
   integer data in a machine/ . . . . . sputl(3X)  
   integer functions. /dnint, . . . . . round(3F)  
   integer part intrinsic . . . . . aint(3F)  
   integer. strtol, . . . . . strtol(3C)  
   integers and long integers. . . . . l3tol(3C)  
   integers. /convert between . . . . . l3tol(3C)  
   interactive block copy. . . . . bcopy(1M)  
   interactive message processing . . . . . mailx(1)  
   interactive repair. /file . . . . . fsck(1M)  
   interesting, adage. fortune: . . . . . fortune(6)  
   interface. . . . . error(7)  
   interface. . . . . lo(5)  
   interface parameters. . . . . ifconfig(8N)  
   interface. . . . . plot(4)  
   interface subroutines. . . . . plot(3X)  
   interface. . . . . termio(7)  
   interface to the TELNET . . . . . telnet(1N)  
   interface. . . . . tty(7)  
   Internet address manipulation/ . . . . . inet(3N)  
   Internet File Transfer . . . . . ftpd(8N)  
   Internet protocol family. . . . . inet(5F)  
   Internet Protocol. . . . . ip(5P)  
   Internet Transmission Control . . . . . tcp(5P)  
   Internet User Datagram . . . . . udp(5P)  
   interpolate smooth curve. . . . . spline(1G)  
   interpret ASA carriage control . . . . . asa(1)  
   interpreter. . . . . sno(1)  
   interpreter) with C-like . . . . . csh(1)  
   interprocess channel. . . . . pipe(2)  
   inter-process communication . . . . . ipc(1)  
   interprocess communication . . . . . stdipc(3C)  
   interval. sleep: . . . . . sleep(1)  
   interval. . . . . sleep(3C)  
   intrinsic function. . . . . acos(3F)  
   intrinsic function. aint, . . . . . aint(3F)  
   intrinsic function. . . . . asin(3F)  
   intrinsic function. atan2, . . . . . atan2(3F)  
   intrinsic function. atan, . . . . . atan(3F)  
   intrinsic function. /dconjg: . . . . . conjg(3F)  
   intrinsic function. cos, . . . . . cos(3F)  
   intrinsic function. /dcosh: . . . . . cosh(3F)

double precision product	intrinsic function. dprod: . . . . .	dprod(3F)
cexp: Fortran exponential	intrinsic function. /dexp: . . . . .	exp(3F)
Fortran common logarithm	intrinsic function. /dlog10: . . . . .	log10(3F)
Fortran natural logarithm	intrinsic function. /clog: . . . . .	log(3F)
Fortran transfer-of-sign	intrinsic function. /dsign: . . . . .	sign(3F)
sin, dsin, csin: Fortran sine	intrinsic function. . . . .	sin(3F)
dsinh: Fortran hyperbolic sine	intrinsic function. sinh, . . . . .	sinh(3F)
csqrt: Fortran square root	intrinsic function. /dsqrt, . . . . .	sqrt(3F)
tan, dtan: Fortran tangent	intrinsic function. . . . .	tan(3F)
Fortran hyperbolic tangent	intrinsic function. /dtanh: . . . . .	tanh(3F)
idim: positive difference	intrinsic functions. /ddim, . . . . .	dim(3F)
dmod: Fortran remaindering	intrinsic functions. /amod, . . . . .	mod(3F)
lle, llt: string comparison	intrinsic functions. /lgt, . . . . .	strcmp(3F)
commands and application/ formats.	intro: introduction to . . . . .	intro(1)
	intro: introduction to file . . . . .	intro(4)
	intro: introduction to games. . . . .	intro(6)
miscellany.	intro: introduction to . . . . .	intro(5)
files.	intro: introduction to special . . . . .	intro(7)
subroutines and libraries.	intro: introduction to . . . . .	intro(3)
calls and error numbers.	intro: introduction to system . . . . .	intro(2)
maintenance commands and/ maintenance procedures.	intro: introduction to system . . . . .	intro(1M)
application programs. intro:	intro: introduction to commands and . . . . .	intro(1)
	intro: introduction to file formats. . . . .	intro(4)
	intro: introduction to games. . . . .	intro(6)
	intro: introduction to miscellany. . . . .	intro(5)
facilities. networking:	intro: introduction to networking . . . . .	intro(5N)
	intro: introduction to special files. . . . .	intro(7)
and libraries. intro:	intro: introduction to subroutines . . . . .	intro(3)
and error numbers. intro:	intro: introduction to system calls . . . . .	intro(2)
maintenance commands/ intro:	intro: introduction to system . . . . .	intro(1M)
maintenance/ intro:	intro: introduction to system . . . . .	intro(8)
ncheck: generate names from	i-numbers. . . . .	ncheck(1M)
aliens: The alien	invaders attack the earth. . . . .	aliens(6)
select: synchronous	i/o multiplexing. . . . .	select(2N)
	ioctl: control device. . . . .	ioctl(2)
abort: generate an	IOT fault. . . . .	abort(3C)
	ip: Internet Protocol. . . . .	ip(5P)
semaphore set or shared/ communication facilities/ uniform random-number/ /islower, isdigit, isxdigit, isidigit, isxdigit, isalnum,/ /isprint, isgraph, iscntrl, terminal. ttyname, /ispunct, isprint, isgraph, isalpha, isupper, islower, /isspace, ispunct, isprint, transfer-of-sign/ sign, isalnum, / isalpha, isupper, /isalnum, isspace, ispunct, /isxdigit, isalnum, isspace, /isidigit, isxdigit, isalnum, Fortran. system: system: issue: file. isxdigit, isalnum, / isalpha, /isupper, islower, isdigit, news: print news functions. functions. j0, bj: the game of black functions. j0, j1,	ipcrm: remove a message queue, . . . . .	ipcrm(1)
	ipcs: report inter-process . . . . .	ipcs(1)
	irand, srand, rand: Fortran . . . . .	rand(3F)
	isalnum, isspace, ispunct, / . . . . .	ctype(3C)
	isalpha, isupper, islower, . . . . .	ctype(3C)
	isascii: classify characters. . . . .	ctype(3C)
	isatty: find name of a . . . . .	ttyname(3C)
	iscntrl, isascii: classify / . . . . .	ctype(3C)
	isidigit, isxdigit, isalnum, / . . . . .	ctype(3C)
	isgraph, iscntrl, isascii: / . . . . .	ctype(3C)
	isign, dsign: Fortran . . . . .	sign(3F)
	islower, isdigit, isxdigit, . . . . .	ctype(3C)
	isprint, isgraph, iscntrl, / . . . . .	ctype(3C)
	ispunct, isprint, isgraph, / . . . . .	ctype(3C)
	isspace, ispunct, isprint, / . . . . .	ctype(3C)
	issue a shell command from . . . . .	system(3F)
	issue a shell command. . . . .	system(3S)
	issue identification file. . . . .	issue(4)
	issue: issue identification . . . . .	issue(4)
	isupper, islower, isdigit, . . . . .	ctype(3C)
	isxdigit, isalnum, isspace, / . . . . .	ctype(3C)
	items. . . . .	news(1)
	j0, j1, jn, y0, y1, yn: Bessel . . . . .	bessel(3M)
	j1, jn, y0, y1, yn: Bessel . . . . .	bessel(3M)
	jack. . . . .	bj(6)
	jn, y0, y1, yn: Bessel . . . . .	bessel(3M)



operator. join: relational database . . . . . join(1)  
 /rand48, nrand48, mrand48, jrand48, srand48, seed48./ . . . . . drand48(3C)  
 makekey: generate encryption key. . . . . makekey(1)  
 killall: kill all active processes. . . . . killall(1M)  
 process or a group of/ kill: send a signal to a . . . . . kill(2)  
 kill: terminate a process. . . . . kill(1)  
 processes. killall: kill all active . . . . . killall(1M)  
 chase: Try to escape the killer robots. . . . . chase(6)  
 process group. killpg: send signal to a . . . . . killpg(3N)  
 mem. kmem: core memory. . . . . mem(7)  
 quiz: test your knowledge. . . . . quiz(6)  
 3-byte integers and long/ l3tol, ltol3: convert between . . . . . l3tol(3C)  
 integer and base-64/ a64l, l64a: convert between long . . . . . a64l(3C)  
 copy file systems with label checking. /labelit: . . . . . volcopy(1M)  
 with label checking. volcopy, labelit: copy file systems . . . . . volcopy(1M)  
 scanning and processing language. awk: pattern . . . . . awk(1)  
 arbitrary-precision arithmetic language. bc: . . . . . bc(1)  
 efl: Extended Fortran Language. . . . . efl(1)  
 cpp: the C language preprocessor. . . . . cpp(1)  
 cpp: the C language preprocessor. . . . . cpp5.0(1)  
 command programming language. /standard/restricted . . . . . sh(1)  
 chargefee, ckpact, dodisk, lastlogin, monacct, nulladm,/ . . . . . acctsh(1M)  
 statistics. lav: print load average . . . . . lav(1)  
 shi: shell layer manager. . . . . shi(1)  
 /rand48, srand48, seed48, kcong48: generate uniformly/ . . . . . drand48(3C)  
 object files. ld: link editor for common . . . . . ld(1)  
 ld5.0: link editor. . . . . ld5.0(1)  
 object file. ldclose: close a common . . . . . ldclose(3X)  
 header of a member of an/ ldahread: read the archive . . . . . ldahread(3X)  
 file for reading. ldopen: open a common object . . . . . ldopen(3X)  
 common object file. ldclose, ldaclose: close a . . . . . ldclose(3X)  
 of floating-point/ frexp, ldexp, modf: manipulate parts . . . . . frexp(3C)  
 access routines. ldfcn: common object file . . . . . ldfcn(4)  
 of a common object file. ldhread: read the file header . . . . . ldhread(3X)  
 name for object file. ldgetname: retrieve symbol . . . . . ldgetname(3X)  
 line number entries/ ldhread, ldlimit, ldliitem: manipulate . . . . . ldhread(3X)  
 number/ ldread, ldlimit, ldliitem: manipulate line number/ . . . . . ldhread(3X)  
 manipulate line number/ ldliitem: . . . . . ldhread(3X)  
 line number entries of a/ ldliitem: . . . . . ldliitem(3X)  
 entries of a section/ ldliitem: seek to . . . . . ldliitem(3X)  
 entries of a section/ ldliitem: seek to line number . . . . . ldliitem(3X)  
 indexed/named/ ldshread, ldnsseek: seek to relocation . . . . . ldshread(3X)  
 indexed/named/ ldshread, ldnsseek: read an . . . . . ldshread(3X)  
 file header of a common/ ldnsseek: seek to an . . . . . ldnsseek(3X)  
 object file for reading. ldohseek: seek to the optional . . . . . ldohseek(3X)  
 relocation entries of a/ ldopen, ldaopen: open a common . . . . . ldopen(3X)  
 indexed/named section header/ ldread, ldnsseek: seek to . . . . . ldread(3X)  
 indexed/named section of a/ ldshread, ldnsseek: seek to an . . . . . ldshread(3X)  
 of a symbol table entry of a/ ldnsseek, ldnsseek: seek to the index . . . . . ldnsseek(3X)  
 symbol table entry of a/ ldnsseek: compute the index . . . . . ldnsseek(3X)  
 table of a common object/ ldnsseek: read an indexed . . . . . ldnsseek(3X)  
 string. ldnsseek: seek to the symbol . . . . . ldnsseek(3X)  
 /rewinddir, closedir: len: return length of Fortran . . . . . len(3F)  
 len: return length directory operations. . . . . directory(3X)  
 len: return length of Fortran string. . . . . len(3F)  
 getopt: get option letter from argument vector. . . . . getopt(3C)  
 simple lexical tasks. lex: generate programs for . . . . . lex(1)  
 generate programs for simple lexical tasks. lex: . . . . . lex(1)  
 update. lsearch, lfind: linear search and . . . . . lsearch(3C)  
 comparison intrinsic/ lge, lgt, lle, llt: string . . . . . strcmp(3F)  
 comparison intrinsic/ lge, lgt, lle, llt: string . . . . . strcmp(3F)  
 to subroutines and libraries. /introduction . . . . . intro(3)  
 ar5.0: archive (library) file format. . . . . ar5.0(4)  
 relation for an object library. /find ordering . . . . . lorder(1)  
 relation for an object library. /find ordering . . . . . lorder5.0(1)

*Permuted Index*

ar5.0: archive and  
 portable/ ar: archive and  
 ulimit: get and set user  
 an out-going terminal  
 type, modes, speed, and  
 line: read one  
 common object file. linenum:  
 /ldlinit, ldliem: manipulate  
 ldlseek, ldnlseek: seek to  
 an/ strip: strip symbol and  
 nl:  
 out selected fields of each  
 send/cancel requests to an LP  
  
 lsearch, lfind:  
 col: filter reverse  
 in a common object file.  
 files. comm: select or reject  
 head: give first few  
 uniq: report repeated  
 of several files or subsequent  
 subsequent/ paste: merge same  
 link, unlink: exercise  
 files. ld:  
 ld5.0:  
 a.out: common assembler and  
 a.out5.0: assembler and  
  
 cp, ln, mv: copy,  
 link:  
 and unlink system calls.  
  
 ls:  
 for a file system. ff:  
 nlist: get entries from name  
 nm5.0: print name  
 nm: print name  
 by fsck. checklist:  
 handle variable argument  
 output of a varargs argument  
 output of a varargs argument  
 socket. listen:  
 on a socket.  
 xargs: construct argument  
 intrinsic/ lge, lgt,  
 intrinsic/ lge, lgt, lle,  
 files. cp,  
 interface.  
 lav: print  
 tzset: convert date/ ctime,  
 manual for program. whereis:  
 index: return  
 end, etext, edata: last  
 memory. plock:  
 files.  
 lockf: record  
 file regions for reading or/  
 natural logarithm intrinsic/  
 gamma:  
 newgrp:  
 exponential, logarithm./ exp,  
 common logarithm intrinsic/  
 logarithm, power./ exp, log,  
 /alog10, dlog10: Fortran common  
 library maintainer. . . . . ar5.0(1)  
 library maintainer for . . . . . ar(1)  
 limits. . . . . ulimit(2)  
 line connection. /establish . . . . . dial(3C)  
 line discipline. /set terminal . . . . . getty(1M)  
 line. . . . . line(1)  
 line number entries in a . . . . . linenum(4)  
 line number entries of a/ . . . . . ldread(3X)  
 line number entries of a/ . . . . . ldlseek(3X)  
 line number information from . . . . . strip(1)  
 line numbering filter. . . . . nl(1)  
 line of a file. cut: cut . . . . . cut(1)  
 line printer. lp, cancel: . . . . . lp(1)  
 line: read one line. . . . . line(1)  
 linear search and update. . . . . lsearch(3C)  
 line-feeds. . . . . col(1)  
 linenum: line number entries . . . . . linenum(4)  
 lines common to two sorted . . . . . comm(1)  
 lines. . . . . head(1)  
 lines in a file. . . . . uniq(1)  
 lines of one file. /same lines . . . . . paste(1)  
 lines of several files or . . . . . paste(1)  
 link and unlink system calls. . . . . link(1M)  
 link editor for common object . . . . . ld(1)  
 link editor. . . . . ld5.0(1)  
 link editor output. . . . . a.out(4)  
 link editor output. . . . . a.out5.0(4)  
 link: link to a file. . . . . link(2)  
 link or move files. . . . . cp(1)  
 link to a file. . . . . link(2)  
 link, unlink: exercise link . . . . . link(1M)  
 lint: a C program checker. . . . . lint(1)  
 list contents of directory. . . . . ls(1)  
 list file names and statistics . . . . . ff(1M)  
 list. . . . . nlist(3C)  
 list. . . . . nm5.0(1)  
 list of common object file. . . . . nm(1)  
 list of file systems processed . . . . . checklist(4)  
 list. varargs: . . . . . varargs(5)  
 list. /print formatted . . . . . vprintf(3S)  
 list. /print formatted . . . . . vprintf(3X)  
 listen for connections on a . . . . . listen(2N)  
 listen: listen for connections . . . . . listen(2N)  
 list(s) and execute command. . . . . xargs(1)  
 lle, lli: string comparison . . . . . strcmp(3F)  
 llt: string comparison . . . . . strcmp(3F)  
 ln, mv: copy, link or move . . . . . cp(1)  
 lo: software loopback network . . . . . lo(5)  
 load average statistics. . . . . lav(1)  
 localtime, gmtime, asctime, . . . . . ctime(3C)  
 locate source, binary, and/or . . . . . whereis(1)  
 location of Fortran substring. . . . . index(3F)  
 locations in program. . . . . end(3C)  
 lock process, text, or data in . . . . . plock(2)  
 lockf: record locking on . . . . . lockf(3C)  
 locking on files. . . . . lockf(3C)  
 locking: provide exclusive . . . . . locking(2)  
 log, alog, dlog, clog: Fortran . . . . . log(3F)  
 log gamma function. . . . . gamma(3M)  
 log in to a new group. . . . . newgrp(1)  
 log, log10, pow, sqrt: . . . . . exp(3M)  
 log10, alog10, dlog10: Fortran . . . . . log10(3F)  
 log10, pow, sqrt: exponential, . . . . . exp(3M)  
 logarithm intrinsic function. . . . . log10(3F)

/dlog, clog: Fortran natural	logarithm intrinsic function. . . . .	log(3F)
/log10, pow, sqrt: exponential,	logarithm, power, square root/	exp(3M)
errpt: process a report of	logged errors. . . . .	errpt(1M)
rwho: who's	logged in on local machines. . . . .	rwho(1N)
getlogin: get	login name. . . . .	getlogin(3C)
logname: get	login name. . . . .	logname(1)
userid: get character	login name of the user. . . . .	userid(3S)
logname: return	login name of user. . . . .	logname(3X)
passwd: change	login password. . . . .	passwd(1)
rlogin: remote	login. . . . .	rlogin(1N)
rlogind: remote	login server. . . . .	rlogind(8N)
setting up an environment at	login: sign on. . . . .	login(1)
last: indicate last	login time. profile: . . . . .	profile(4)
user.	logins of users and teletypes. . . . .	last(1)
a64l, l64a: convert between	logname: get login name. . . . .	logname(1)
sputl, sgetl: access	logname: return login name of . . . . .	logname(3X)
between 3-byte integers and	long integer and base-64 ASCII/ . . . . .	a64l(3C)
setjmp.	long integer data in a machine/ . . . . .	sputl(3X)
lo: software	long integers. /lto13: convert . . . . .	l3tol(3C)
for an object library.	longjmp: non-local goto. . . . .	setjmp(3C)
relation for an object/	loopback network interface. . . . .	lo(5)
mklost+found: make a	lorder: find ordering relation . . . . .	lorder(1)
nice: run a command at	lorder5.0: find ordering . . . . .	lorder5.0(1)
requests to an LP line/	lost+found directory for fsck. . . . .	mklost+found(1M)
send/cancel requests to an	lost priority. . . . .	nice(1)
disable: enable/disable	lp, cancel: send/cancel . . . . .	lp(1)
/lpshut, lpmove: start/stop the	LP line printer. lp, cancel: . . . . .	lp(1)
accept, reject: allow/prevent	LP printers. enable, . . . . .	enable(1)
lpadmin: configure the	LP request scheduler and move/ . . . . .	lpsched(1M)
lpstat: print	LP requests. . . . .	accept(1M)
spooling system.	LP spooling system. . . . .	lpadmin(1M)
request/ lpsched, lpshut,	LP status information. . . . .	lpstat(1)
start/stop the LP request/	lpadmin: configure the LP . . . . .	lpadmin(1M)
LP request scheduler/ lpsched,	lpmove: start/stop the LP . . . . .	lpsched(1M)
information.	lpsched, lpshut, lpmove: . . . . .	lpsched(1M)
jrand48, drand48, erand48,	lpshut, lpmove: start/stop the . . . . .	lpsched(1M)
directory.	lpstat: print LP status . . . . .	lpstat(1)
and update.	lrand48, rrand48, mrand48, . . . . .	drand48(3C)
pointer.	ls: list contents of . . . . .	ls(1)
bitwise/ and, or, xor, not,	lsearch, lfind: linear search . . . . .	lsearch(3C)
integers and long/ l3tol,	lseek: move read/write file . . . . .	lseek(2)
facv: convert files between	lshift, rshift: Fortran . . . . .	bool(3F)
provide truth value about/	lto13: convert between 3-byte . . . . .	l3tol(3C)
/access long integer data in a	m4: macro processor. . . . .	m4(1)
put: puts a file onto a remote	M68000 and VAX-11/780/ . . . . .	fsvc(1M)
takes a file from a remote	m68k, pdp11, u3b, u3b5, vax: . . . . .	machid(1)
values:	machine independent fashion. . . . .	sputl(3X)
show host status of local	machine. . . . .	put(1C)
rwho: who's logged in on local	machine. take: . . . . .	take(1C)
update files between two	machine-dependent values. . . . .	values(5)
update files between two	machines. runtime: . . . . .	runtime(1N)
permuted index. mptx: the	machines. . . . .	rwho(1N)
documents. mm: the MM	machines. updater: . . . . .	updater(1)
mosd: the OSD adapter	machines. updater: . . . . .	updater(1M)
view graphs and/ mv: a troff	macro package for formatting a . . . . .	mptx(5)
m4:	macro package for formatting . . . . .	mm(5)
in this manual. man:	macro package for formatting/ . . . . .	mosd(5)
formatted with the MM	macro package for typesetting . . . . .	mv(5)
send mail to users or read	macro processor. . . . .	m4(1)
users or read mail.	macros for formatting entries . . . . .	man(5)
netmail: the B-NET network	macros. /print/check documents . . . . .	mm(1)
	mail. mail, rmail: . . . . .	mail(1)
	mail, rmail: send mail to . . . . .	mail(1)
	mail system. . . . .	netmail(8N)

Permuted Index

delivermail: deliver mail to arbitrary people. . . . . delivermail(8N)  
 netmailer: deliver mail to B-NET. . . . . netmailer(8N)  
 mail, rmail: send mail to users or read mail. . . . . mail(1)  
 processing system. . . . . mailx(1)  
 malloc, free, realloc, calloc: main memory allocator. . . . . malloc(3C)  
 /malloc, mallinfo: fast main memory allocator. . . . . malloc(3X)  
 program. ctags: maintain a tags file for a C . . . . . ctags(1)  
 regenerate groups off/ make: maintain, update, and make(1)  
 ar5.0: archive and library maintainer. . . . . ar5.0(1)  
 ar: archive and library maintainer for portable/ . . . . . ar(1)  
 intro: introduction to system maintenance commands and/ intro(1M)  
 intro: introduction to system maintenance procedures. . . . . intro(8)  
 SCCS file. delta: make a delta (change) to an . . . . . delta(1)  
 mkdir: make a directory. . . . . mkdir(1)  
 or ordinary file. mknod: make a directory, or a special . . . . . mknod(2)  
 for fsck. mklost+found: make a lost+found directory . . . . . mklost+fnd(1M)  
 mktemp: make a unique filename. . . . . mktemp(3C)  
 regenerate groups off/ make: maintain, update, and make(1)  
 ssp: make output single spaced. . . . . ssp(1)  
 banner: make posters. . . . . banner(1)  
 session. script: make typescript of terminal . . . . . script(1)  
 key. makekey: generate encryption . . . . . makekey(1)  
 /realloc, calloc, malloc, mallinfo: fast main memory/ malloc(3X)  
 main memory allocator. . . . . malloc(3C)  
 malloc, free, realloc, calloc: malloc, free, realloc, calloc. . . . . malloc(3X)  
 malloc, free, realloc, calloc: malloc, free, realloc, calloc. . . . . malloc(3X)  
 malloc, free, realloc, calloc, malloc, free, realloc, calloc, . . . . . malloc(3X)  
 entries in this manual. man: macros for formatting . . . . . man(5)  
 manual. man: print entries in this . . . . . man(1)  
 /tfind, tdelete, twalk: manage binary search trees. . . . . tsearch(3C)  
 hsearch, hcreate, hdestroy: manage hash search tables. . . . . hsearch(3C)  
 sh: shell layer sh(1)  
 records. fwtmp, wtmpfix: manipulate connect accounting . . . . . fwtmp(1M)  
 of/ ldread, ldlinet, lditem: manipulate line number entries . . . . . ldread(3X)  
 frexp, ldexp, modf: manipulate parts off/ . . . . . frexp(3C)  
 tp: manipulate tape archive. . . . . tp(1)  
 route: manually manipulate the routing tables. . . . . route(8N)  
 /inet\_netof: Internet address manipulation routines. . . . . inet(3N)  
 locate source, binary, and/or manual for program. whereis: . . . . . whereis(1)  
 man: print entries in this manual. . . . . man(1)  
 for formatting entries in this manual. man: macros . . . . . man(5)  
 routing tables. route: manually manipulate the . . . . . route(8N)  
 asci: map of ASCII character set. . . . . ascii(5)  
 files. diffmk: mark differences between . . . . . diffmk(1)  
 umask: set file-creation mode mask. . . . . umask(1)  
 set and get file creation mask. umask: . . . . . umask(2)  
 an error message file by massaging C source. /create . . . . . mkstr(1)  
 table. master: master device information . . . . . master(4)  
 information table. master: master device . . . . . master(4)  
 regular expression compile and match routines. regexp: . . . . . regexp(5)  
 math: math functions and constants. . . . . math(5)  
 constants. math: math functions and . . . . . math(5)  
 eqn, neqn, checkeq: format mathematical text for nroff or/ . . . . . eqn(1)  
 function. matherr: error-handling . . . . . matherr(3M)  
 dmax1: Fortran maximum-value/ max, max0, amax0, max1, amax1, . . . . . max(3F)  
 dmax1: Fortran/ max, max0, amax0, max1, amax1, . . . . . max(3F)  
 max, max0, amax0, max1, amax1, dmax1: Fortran/ . . . . . max(3F)  
 /max1, amax1, dmax1: Fortran maximum-value functions. . . . . max(3F)  
 maze: generate a maze. . . . . maze(6)  
 maze: generate a maze. . . . . maze(6)  
 mc68cc: C compiler. . . . . mc68cc(1)  
 mclock: return Fortran time . . . . . mclock(3F)  
 media. . . . . bcd(6)  
 mem, kmem: core memory. . . . . mem(7)  
 memcopy, memset: memory/ memcopy, memchr, memcmp, . . . . . memory(3C)

memset: memory/ memccpy, memchr, memcmp, memcpy, . . . . . memory(3C)  
 operations. memccpy, memchr, memcmp, memcpy, memset: memory memory(3C)  
 memccpy, memchr, memcmp, memcpy, memset: memory/ . . . . . memory(3C)  
 free, realloc, calloc: main memory allocator. malloc, . . . . . malloc(3C)  
 malloc, mallinfo: fast main memory allocator. /calloc, . . . . . malloc(3X)  
 shmctl: shared memory control operations. . . . . shmctl(2)  
 queue, semaphore set or shared memory id. /remove a message . . . . . ipcrm(1)  
 mem. kmem: core memory. . . . . mem(7)  
 memcmp, memcpy, memset: memory operations. /memchr, . . . . . memory(3C)  
 shmop: shared memory operations. . . . . shmop(2)  
 lock process, text, or data in memory. plock. . . . . plock(2)  
 shmget: get shared memory segment. . . . . shmget(2)  
 /memchr, memcmp, memcpy, memset: memory operations. . . . . memory(3C)  
 sort: sort and/or merge files. . . . . sort(1)  
 files. acctmrg: merge or add total accounting . . . . . acctmrg(1M)  
 files or subsequent/ paste: merge same lines of several . . . . . paste(1)  
 msgctl: message control operations. . . . . msgctl(2)  
 mkstr: create an error message file by massaging C/ . . . . . mkstr(1)  
 recvfrom, recvmsg: receive a message from a socket. recv, . . . . . recv(2N)  
 send, sendto, sendmsg: send a message from a socket. . . . . send(2N)  
 msgop: message operations. . . . . msgop(2)  
 mailx: interactive message processing system. . . . . mailx(1)  
 msgget: get message queue. . . . . msgget(2)  
 or shared/ ipcrm: remove a message queue, semaphore set . . . . . ipcrm(1)  
 msg: permit or deny messages. . . . . msg(1)  
 sys\_err: system error messages. /errno, sys\_errlist, . . . . . perror(3C)  
 dminl: Fortran minimum-value/ min, min0, amin0, minl, aminl, . . . . . min(3F)  
 dminl: Fortran/ min, min0, amin0, minl, aminl, . . . . . min(3F)  
 /minl, aminl, dminl: Fortran/ minimum-value functions. . . . . min(3F)  
 mkdir: make a directory. . . . . mkdir(1)  
 mkfs: construct a file system. . . . . mkfs(1M)  
 mkfslb: construct a file system. . . . . mkfslb(1M)  
 mklost+found: make a message queue. . . . . mklost+found(1M)  
 mknod: build special file. . . . . mknod(1M)  
 mknod: make a directory, or a message queue. . . . . mknod(2)  
 mkstr: create an error message . . . . . mkstr(1)  
 mktemp: make a unique . . . . . mktemp(3C)  
 MM macro package for . . . . . mm(5)  
 MM macros. /print/check . . . . . mm(1)  
 mm, osdd, checkmm: print/check . . . . . mm(1)  
 mm: the MM macro package for . . . . . mm(5)  
 mmt, mvt: typeset documents, . . . . . mmt(1)  
 mnttab: mounted file system . . . . . mnttab(4)  
 mod, amod, dmod: Fortran . . . . . mod(3F)  
 mode. . . . . chmod(1)  
 mode mask. . . . . umask(1)  
 mode of file. . . . . chmod(2)  
 modes, speed, and line/ . . . . . getty(1M)  
 modest-sized programs. . . . . bs(1)  
 modf: manipulate parts of . . . . . frexp(3C)  
 modification times of a file. . . . . touch(1)  
 modification times. . . . . utime(2)  
 monacct, nulladm, pretmp,/ . . . . . acctsh(1M)  
 monitor: prepare execution . . . . . monitor(3C)  
 monitor uucp network. . . . . uucp(1M)  
 moo: guessing game. . . . . moo(6)  
 mosd: the OSDD adapter macro . . . . . mosd(5)  
 Motorola S-records from/ . . . . . rcvhex(1)  
 mount a file system. . . . . mount(2)  
 mount and dismount file . . . . . mount(1M)  
 mount: mount a file system. . . . . mount(2)  
 mount table. . . . . setmnt(1M)  
 setmnt: establish

*Permuted Index*

dismount file system.	mount, umount: mount and	mount(1M)
mnttab:	mounted file system table.	mnttab(4)
mvdir:	move a directory.	mvdir(1M)
cp, ln, mv: copy, link or	move files.	cp(1)
lseek:	move read/write file pointer.	lseek(2)
the LP request scheduler and	move requests. /start/stop	lpsched(1M)
formatting a permuted index.	mptx: the macro package for	mptx(5)
/erand48, lrand48, nrand48,	mrand48, jrand48, srand48,/	drand48(3C)
operations.	msgctl: message control	msgctl(2)
	msgget: get message queue.	msgget(2)
	msgop: message operations.	msgop(2)
	multiplexing.	select(2N)
select: synchronous i/o	mv: a troff macro package for	mv(5)
typesetting view graphs and/	mv: copy, link or move files.	cp(1)
cp, ln,	mvdir: move a directory.	mvdir(1M)
	mvt: typeset documents, view	mmt(1)
graphs, and slides. mmt,	natural logarithm intrinsic/	log(3F)
log, alog, dlog, clog: Fortran	ncheck: generate names from	ncheck(1M)
i-numbers.	nearest integer functions.	round(3F)
/dnint, nint, idnint: Fortran	neqn, checkeq: format	eqn(1)
mathematical text for/ eqn,	neqn. /special character	eqnchar(5)
definitions for eqn and	netmail: the B-NET network	netmail(8N)
mail system.	netmailer: deliver mail to	netmailer(8N)
B-NET.	netstat: show network status.	netstat(1N)
	network byte order. /convert	byteorder(3N)
values between host and	network entry. /getnetbyname.	getnetent(3N)
setnetent, endnetent: get	network host entry.	gethostent(3N)
/sethostent, endhostent: get	network interface.	lo(5)
lo: software loopback	network interface parameters.	ifconfig(8N)
ifconfig: configure	network mail system.	netmail(8N)
netmail: the B-NET	network name data base.	networks(4N)
networks:	network routing daemon.	routed(8N)
routed:	network status.	netstat(1N)
netstat: show	network.	uusb(1M)
uusb: monitor uucp	networking facilities.	intro(5N)
networking: introduction to	networking: introduction to	intro(5N)
networking facilities.	networks: network name data	networks(4N)
base.	newform: change the format of	newform(1)
a text file.	newgrp: log in to a new group.	newgrp(1)
	news items.	news(1)
news: print	news: print news items.	news(1)
process.	nice: change priority of a	nice(2)
priority.	nice: run a command at low	nice(1)
integer/ anint, dnint,	nint, idnint: Fortran nearest	round(3F)
	nl: line numbering filter.	nl(1)
list.	nlist: get entries from name	nlist(3C)
object file.	nm: print name list of common	nm(1)
	nm5.0: print name list.	nm5.0(1)
change current UNIX system	nodename. chgnod:	chgnod(1M)
hangups (sh only).	nohup: run a command immune to	nohup(1)
setjmp, longjmp:	non-local goto.	setjmp(3C)
bitwise boolean/ and, or, xor,	not, lshift, rshift: Fortran	bool(3F)
drand48, erand48, lrand48,	nrand48, mrand48, jrand48,/	drand48(3C)
	nroff: format text.	nroff(1)
format mathematical text for	nroff or troff. /checkeq:	eqn(1)
tbl: format tables for	nroff or troff.	tbl(1)
constructs. deroff: remove	nroff/troff, tbl, and eqn	deroff(1)
between host/ htonl, htons,	ntohl, ntohs: convert values	byteorder(3N)
host and/ htonl, htons, ntohl,	ntohs: convert values between	byteorder(3N)
null: the	null file.	null(7)
	null: the null file.	null(7)
/dodisk, lastlogin, monacct,	nulladm, pretmp, prdaily,/	acctsh(1M)
nl: line	numbering filter.	nl(1)
number: convert Arabic	numerals to English.	number(6)

ldfcn: common	object file access routines. . . . .	ldfcn(4)
conv:	object file converter. . . . .	conv(1)
dump selected parts of an	object file. dump: . . . . .	dump(1)
ldopen, ldaopen: open a common	object file for reading. . . . .	ldopen(3X)
number entries of a common	object file function. /line . . . . .	ldfread(3X)
ldaclose: close a common	object file. ldclose. . . . .	ldclose(3X)
the file header of a common	object file. ldhread: read . . . . .	ldhread(3X)
retrieve symbol name for	object file. ldgetname: . . . . .	ldgetname(3X)
of a section of a common	object file. /number entries . . . . .	ldseek(3X)
file header of a common	object file. /to the optional . . . . .	ldohseek(3X)
of a section of a common	object file. /entries . . . . .	ldrseek(3X)
section header of a common	object file. /an indexed/named . . . . .	ldhread(3X)
section of a common	object file. /an indexed/named . . . . .	ldsseek(3X)
symbol table entry of a common	object file. /the index of a . . . . .	ldtbindx(3X)
symbol table entry of a common	object file. /read an indexed . . . . .	ldtbread(3X)
the symbol table of a common	object file. /seek to . . . . .	ldtseek(3X)
number entries in a common	object file. linenum: line . . . . .	linenum(4)
nm: print name list of common	object file. . . . .	nm(1)
information for a common	object file. /relocation . . . . .	reloc(4)
section header for a common	object file. scnhdr: . . . . .	scnhdr(4)
size5.0: size of an	object file. . . . .	size5.0(1)
number information from an	object file. /symbol and line . . . . .	strip(1)
format. syms: common	object file symbol table . . . . .	syms(4)
- a.out header for common	object files. aouthdr.h . . . . .	aouthdr(4)
file header for common	object files. filehdr: . . . . .	filehdr(4)
hex: translates	object files. . . . .	hex(1)
directories. cpset: install	object files in binary . . . . .	cpset(1M)
ld: link editor for common	object files. . . . .	ld(1)
print section sizes of common	object files. size: . . . . .	size(1)
find ordering relation for an	object library. lorder: . . . . .	lorder(1)
find ordering relation for an	object library. lorder5.0: . . . . .	lorder5.0(1)
/the printable strings in an	object, or other binary file. . . . .	strings(1)
/setgrent, endgrent, fgetgrent:	obtain group file entry from a/	getgrent(3C)
od:	octal dump. . . . .	od(1)
od: octal dump.	only). nohup: run a . . . . .	od(1)
command immune to hangups (sh	on/off the extended errors in . . . . .	nohup(1)
the specified/ exterr: turn	onto a remote machine. . . . .	exterr(1)
put: puts a file	open a common object file for . . . . .	put(1C)
reading. ldopen, ldaopen:	open a stream. . . . .	ldopen(3X)
fopen, freopen, fdopen:	open for reading or writing. . . . .	fopen(3S)
open:	open: open for reading or . . . . .	open(2)
writing.	opendir, readdir, telldir, . . . . .	open(2)
seekdir, rewinddir, closedir:/	operation routines. /tgoto, . . . . .	directory(3X)
tputs: terminal independent	operations. /bcmp, bzero, . . . . .	termcap(3X)
ffs: bit and byte string	operations. /closedir: . . . . .	bstring(3N)
flexible length directory	operations. memccpy, memchr, . . . . .	directory(3X)
memcmp, memcpy, memset: memory	operations. . . . .	memory(3C)
msgctl: message control	operations. . . . .	msgctl(2)
msgop: message	operations. . . . .	msgop(2)
semctl: semaphore control	operations. . . . .	semctl(2)
semop: semaphore	operations. . . . .	semop(2)
shmctl: shared memory control	operations. . . . .	shmctl(2)
shmop: shared memory	operations. . . . .	shmop(2)
strncpy, strtok: string	operations. /strpbrk, strspn, . . . . .	string(3C)
join: relational database	operator. . . . .	join(1)
dcopy: copy file systems for	optimal access time. . . . .	dcopy(1M)
CRT screen handling and	optimization package. curses: . . . . .	curses(3X)
vector. getopt: get	option letter from argument . . . . .	getopt(3C)
common/ ldohseek: seek to the	optional file header of a . . . . .	ldohseek(3X)
fcntl: file control	options. . . . .	fcntl(5)
stty: set the	options for a terminal. . . . .	stty(1)
getopt: parse command	options. . . . .	getopt(1)
/setsockopt: get and set	options on sockets. . . . .	getsockopt(2N)
Fortran bitwise boolean/ and,	or, xor, not, lshift, rshift: . . . . .	bool(3F)

## Permuted Index

object library.	lorder: find	ordering relation for an . . . . .	lorder(1)
object/ lorder5.0: find	a directory, or a special or	ordering relation for an . . . . .	lorder5.0(1)
formatting/ mosd: the	documents formatted with/ mm,	ordinary file. mknod: make . . . . .	mknod(2)
dial: establish an	assembler and link editor	OSDD adapter macro package for . . . . .	mosd(5)
assembler and link editor	/vsprintf: print formatted	osdd, checkmm: print/check . . . . .	mm(1)
/vsprintf: print formatted	/vsprintf: print formatted	out-going terminal line/ . . . . .	dial(3C)
sprintf: print formatted	ssp: make	output. a.out: common . . . . .	a.out(4)
/acctdusg, accton, acctwtmp:	chown: change	output. a.out5.0: . . . . .	a.out5.0(4)
chown, chgrp: change	and expand files.	output of a varargs argument/ . . . . .	vprintf(3S)
handling and optimization	permuted/ mptx: the macro	output of a varargs argument/ . . . . .	vprintf(3X)
documents. mm: the MM macro	mosd: the OSDD adapter macro	output. printf, fprintf, . . . . .	printf(3S)
graphs and/ mv: a troff macro	sadc: system activity report	output single spaced. . . . .	ssp(1)
standard buffered input/output	interprocess communication	overview of accounting and/ . . . . .	acct(1M)
4014 terminal. 4014:	tune floppy disk settling time	owner and group of a file. . . . .	chown(2)
configure network interface	process, process group, and	owner or group. . . . .	chown(1)
getopt:		pack, pcat, unpack: compress . . . . .	pack(1)
		package. curses: CRT screen . . . . .	curses(3X)
/endpwent, fgetpwent: get	putpwent: write	package for formatting a . . . . .	mptx(5)
password:	getpass: read a	package for formatting/ . . . . .	mm(5)
password:	password: change login	package for typesetting view . . . . .	mosd(5)
password:	password:	package. sa1, sa2, . . . . .	mv(5)
password:	password:	package. sidio: . . . . .	sar(1M)
password:	password:	package. flock: standard . . . . .	stdio(3S)
password:	password:	paginator for the Tektronix . . . . .	stdipc(3C)
password:	password:	parameters. disktime: . . . . .	4014(1)
password:	password:	parameters. ifconfig: . . . . .	disktime(1M)
password:	password:	parent process IDs. /get . . . . .	ifconfig(8N)
password:	password:	parse command options. . . . .	getpid(2)
password:	password:	passwd: change login password. . . . .	getopt(1)
password:	password:	passwd: password file. . . . .	passwd(1)
password:	password:	passwd file entry. . . . .	passwd(4)
password:	password:	passwd file entry. . . . .	getpwent(3C)
password:	password:	passwd file. . . . .	putpwent(3C)
password:	password:	password. . . . .	passwd(4)
password:	password:	password. . . . .	getpass(3C)
password:	password:	password/group file checkers. . . . .	passwd(1)
password:	password:	paste: merge same lines of . . . . .	pwck(1M)
password:	password:	path names. basename, . . . . .	paste(1)
password:	password:	pathname of current working . . . . .	basename(1)
password:	password:	pattern. grep, egrep, . . . . .	getcwd(3C)
password:	password:	pattern scanning and . . . . .	grep(1)
password:	password:	pause: suspend process until . . . . .	awk(1)
password:	password:	pcat, unpack: compress and . . . . .	pause(2)
password:	password:	pclose: initiate pipe to/from . . . . .	pack(1)
password:	password:	pdp11, u3b, u3b5, vax: provide . . . . .	popen(3S)
password:	password:	peer. getpeername: . . . . .	machid(1)
password:	password:	permit or deny messages. . . . .	getpeername(2N)
password:	password:	permuted index. mptx: the . . . . .	mesg(1)
password:	password:	permuted index. . . . .	mptx(5)
password:	password:	per-process accounting file . . . . .	ptx(1)
password:	password:	per-process accounting/ . . . . .	acct(4)
password:	password:	perusr, errno, sys_errlist, . . . . .	acctcms(1M)
password:	password:	perusal filter for crt . . . . .	perror(3C)
password:	password:	perusal filter for soft-copy . . . . .	more(1)
password:	password:	pg: file perusal filter for . . . . .	pg(1)
password:	password:	phototypesetter simulator. . . . .	pg(1)
password:	password:	phys: allow a process to . . . . .	tc(1)
password:	password:	physical addresses. phys: . . . . .	phys(2)
password:	password:	pieces. . . . .	phys(2)
password:	password:	pipe: create an interprocess . . . . .	split(1)
password:	password:	pipe fitting. . . . .	pipe(2)
password:	password:	pipe to/from a process. . . . .	tee(1)
password:	password:	play "Go Fish". . . . .	popen(3S)
password:	password:		fish(6)



life:	play the game of life. . . . .	life(6)
worm:	Play the growing worm game. . . . .	worm(6)
data in memory.	plock: lock process, text, or . . . . .	plock(2)
	plot: graphics interface. . . . .	plot(4)
subroutines.	plot: graphics interface . . . . .	plot(3X)
images.	pnch: file format for card . . . . .	pnch(4)
fstell: reposition a file	pointer in a stream. /rewind, . . . . .	fseek(3S)
lseek: move read/write file	pointer. . . . .	lseek(2)
to/from a process.	popen, pclose: initiate pipe . . . . .	popen(3S)
data base of terminal types by	port. ttytype: . . . . .	ttytype(4)
and library maintainer for	portable archives. /archive . . . . .	ar(1)
basename, dirname: deliver	portions of path names. . . . .	basename(1)
functions. dim, ddim, idim:	positive difference intrinsic . . . . .	dim(3F)
banner: make	posters. . . . .	banner(1)
logarithm,/ exp, log, log10,	pow, sqrt: exponential, . . . . .	exp(3M)
/sqrt: exponential, logarithm,	power, square root functions. . . . .	exp(3M)
brc, bcheckrc, rc,	powerfail: system/ . . . . .	brc(1M)
/lastlogin, monacct, nulladm,	pr: print files. . . . .	pr(1)
/monacct, nulladm, prctmp,	prctmp, prdaily, prtacct,/ . . . . .	acctsh(1M)
function. dprod: double	prdaily, prtacct, runacct,/ . . . . .	acctsh(1M)
for troff. cw, checkcw:	precision product intrinsic . . . . .	dprod(3F)
monitor:	prepare constant-width text . . . . .	cw(1)
cpp: the C language	prepare execution profile. . . . .	monitor(3C)
cpp: the C language	preprocessor. . . . .	cpp(1)
unset: undo a	preprocessor. . . . .	cpp5.0(1)
types:	previous get of an SCCS file. . . . .	unset(1)
interesting, adage. fortune:	primitive system data types. . . . .	types(5)
prs:	print a random, hopefully . . . . .	fortune(6)
date:	print an SCCS file. . . . .	prs(1)
cal:	date: print and set the date. . . . .	date(1)
of a file. sum:	print calendar. . . . .	cal(1)
editing activity. sact:	print checksum and block count . . . . .	sum(1)
id. whoami:	print current SCCS file . . . . .	sact(1)
man:	print effective current user . . . . .	whoami(1)
cat: concatenate and	print entries in this manual. . . . .	man(1)
pr:	print files. . . . .	cat(1)
vprintf, vfprintf, vsprintf:	print files. . . . .	pr(1)
vprintf, vfprintf, vsprintf:	print formatted output of a/ . . . . .	vprintf(3S)
printf, fprintf, sprintf:	print formatted output of a/ . . . . .	vprintf(3X)
host system. hostid: set or	print formatted output. . . . .	printf(3S)
banner7:	print identifier of current . . . . .	hostid(1N)
lav:	print large banner on printer. . . . .	banner7(1)
lpstat:	print load average statistics. . . . .	lav(1)
nm5.0:	print LP status information. . . . .	lpstat(1)
object file. nm:	print name list. . . . .	nm5.0(1)
system. hostname: set or	print name list of common . . . . .	nm(1)
system. uname:	print name of current host . . . . .	hostname(1N)
news:	print name of current UNIX . . . . .	uname(1)
file(s). acctcom: search and	print news items. . . . .	news(1)
object files. size:	print out the environment. . . . .	printenv(1)
pstat:	print process accounting . . . . .	acctcom(1)
names. id:	print section sizes of common . . . . .	size(1)
object, or/ strings: find the	print system facts. . . . .	pstat(1M)
formatted/ mm, osdd, checkmm:	print user and group IDs and . . . . .	id(1)
environment.	printable strings in an . . . . .	strings(1)
banner7: print large banner on	print/check documents . . . . .	nm(1)
requests to an LP line	printenv: print out the . . . . .	printenv(1)
disable: enable/disable LP	printer. . . . .	banner7(1)
print formatted output.	printer. /cancel: send/cancel . . . . .	lp(1)
nice: run a command at low	printers. enable, . . . . .	enable(1)
nice: change	printf, fprintf, sprintf: . . . . .	printf(3S)
errors. errpt:	priority. . . . .	nice(1)
	priority of a process. . . . .	nice(2)
	process a report of logged . . . . .	errpt(1M)

Permuted Index

acct: enable or disable	process accounting.	acct(2)
acctprc1, acctprc2:	process accounting.	acctprc(1M)
acctcom: search and print	process accounting file(s).	acctcom(1)
times. times: get	process and child process	times(2)
init, telinit:	process control/	init(1M)
timex: time a command; report	process data and system/	timex(1)
exit, _exit: terminate	process.	exit(2)
fork: create a new	process.	fork(2)
/getgrp, getppid: get process,	process group, and parent/	getpid(2)
setppgrp: set	process group ID.	setppgrp(2)
killpg: send signal to a	process group.	killpg(3N)
process group, and parent	process IDs. /get process,	getpid(2)
inittab: script for the init	process.	inittab(4)
kill: terminate a	process.	kill(1)
nice: change priority of a	process.	nice(2)
kill: send a signal to a	process or a group of/	kill(2)
initiate pipe to/from a	process. popen, pclose:	popen(3S)
getpid, getppgrp, getppid: get	process, process group, and/	getpid(2)
ps: report	process status.	ps(1)
memory. plock: lock	process, text, or data in	plock(2)
times: get process and child	process times.	times(2)
addresses. phys: allow a	process to access physical	phys(2)
wait: wait for child	process to stop or terminate.	wait(2)
wait3: wait for child	process to stop or terminate.	wait3(2N)
ptrace:	process trace.	ptrace(2)
pause: suspend	process until signal.	pause(2)
list of file systems	processed by fsck. checklist:	checklist(4)
to a process or a group of	processes. /send a signal	kill(2)
killall: kill all active	processes.	killall(1M)
structure. fuser: identify	processes using a file or file	fuser(1M)
awk: pattern scanning and	processing language.	awk(1)
shutdown: terminate all	processing.	shutdown(1M)
mailx: interactive message	processing system.	mailx(1)
m4: macro	processor.	m4(1)
provide truth value about your	processor type. /u3b5, vax:	machid(1)
between M68000 and VAX-11/780	processors. /convert files	fscv(1M)
alarm: set a	process's alarm clock.	alarm(2)
dprod: double precision	product intrinsic function.	dprod(3F)
function.	prof: display profile data.	prof(1)
profile.	prof: profile within a	prof(5)
prof: display	profil: execution time	profil(2)
monitor: prepare execution	profile data.	prof(1)
profil: execution time	profile.	monitor(3C)
environment at login time.	profile: setting up an	profil(2)
prof:	profile within a function.	profile(4)
standard/restricted command	programming language. /the	prof(5)
arp: Address Resolution	Protocol.	sh(1)
/setprotoent, endprotoent: get	protocol entry.	arp(5P)
inet: Internet	protocol family.	getprotoent(3N)
ip: Internet	Protocol.	inet(5F)
protocols:	protocol name data base.	ip(5P)
DARPA Internet File Transfer	Protocol server. ftpd:	protocols(4N)
telnetd: DARPA TELNET	protocol server.	ftpd(8N)
DARPA Trivial File Transfer	Protocol server. tftpd:	telnetd(8N)
Internet Transmission Control	Protocol. tcp:	tftpd(8N)
user interface to the TELNET	protocol. telnet:	tcp(5P)
trpt: transliterate	protocol trace.	telnet(1N)
udp: Internet User Datagram	Protocol.	trpt(8N)
base.	protocols: protocol name data	udp(5P)
arithmetic:	provide drill in number facts.	protocols(4N)
for reading or/ locking:	provide exclusive file regions	arithmetic(6)
m68k, pdp11, u3b, u3b5, vax:	provide truth value about your/	locking(2)
true, false:	provide truth values.	machid(1)
		true(1)

	prs: print an SCCS file. . . . .	prs(1)
/nulladm, prctmp, prdaily,	prtacct, runacct, shutacct,/	acctsh(1M)
	ps: report process status. . . . .	ps(1)
	pty: pseudo terminal driver. . . . .	pty(5)
	sxt: pseudo-device driver. . . . .	sxt(7)
/generate uniformly distributed	pseudo-random numbers. . . . .	drand48(3C)
	pstat: print system facts. . . . .	pstat(1M)
	ptrace: process trace. . . . .	ptrace(2)
	plx: permuted index. . . . .	plx(1)
	pty: pseudo terminal driver. . . . .	pty(5)
stream. ungetc:	push character back into input . . . . .	ungetc(3S)
put character or word on a/	putc, putchar, fputc, putw:	putc(3S)
character or word on a/	putchar, fputc, putw: put . . . . .	putc(3S)
environment.	putenv: change or add value to . . . . .	putenv(3C)
entry.	putpwent: write password file . . . . .	putpwent(3C)
machine. put:	puts a file onto a remote . . . . .	put(1C)
stream.	puts, fputs: put a string on a . . . . .	puts(3S)
getutent, getutid, getutline,	pututline, setutent, endutent,/	getut(3C)
a/ putc, putchar, fputc,	putw: put character or word on . . . . .	putc(3S)
file checkers.	pwck, grpck: password/group . . . . .	pwck(1M)
	pwd: working directory name. . . . .	pwd(1)
	qsort: quicker sort. . . . .	qsort(3C)
(tput:	query terminfo database. . . . .	tput(1)
insert/remove element from a	queue. insque, remque: . . . . .	insque(3N)
msgget: get message	queue. . . . .	msgget(2)
ipcrm: remove a message	queue, semaphore set or shared/	ipcrm(1)
qsort:	quicker sort. . . . .	qsort(3C)
	quiz: test your knowledge. . . . .	quiz(6)
display.	rain: animated raindrops . . . . .	rain(6)
rain: animated	raindrops display. . . . .	rain(6)
random-number/ irand, srand,	rand: Fortran uniform . . . . .	rand(3F)
random-number generator.	rand, srand: simple . . . . .	rand(3C)
adage. fortune: print a	random, hopefully interesting, . . . . .	fortune(6)
rand, srand: simple	random-number generator. . . . .	rand(3C)
/srand, rand: Fortran uniform	random-number generator. . . . .	rand(3F)
fsplit: split f77,	ratfor, or elf files. . . . .	fsplit(1)
dialect.	ratfor: rational Fortran . . . . .	ratfor(1)
ratfor:	rational Fortran dialect. . . . .	ratfor(1)
initialization/ brc, bcheckrc,	rc, powerfail: system . . . . .	brc(1M)
routines for returning a/	rcmd, rresvport, ruserok:	rcmd(3N)
	rcp: remote file copy. . . . .	rcp(1N)
S-records from downloading/	rcvhex: translates Motorola . . . . .	rcvhex(1)
getpass:	read a password. . . . .	getpass(3C)
entry of a common/ ldtbread:	read an indexed symbol table . . . . .	ldtbread(3X)
header/ ldshread, ldnsbread:	read an indexed/named section . . . . .	ldshread(3X)
read:	read from file. . . . .	read(2)
readv:	read from file. . . . .	readv(3N)
rmail: send mail to users or	read mail. mail, . . . . .	mail(1)
line:	read one line. . . . .	line(1)
	read: read from file. . . . .	read(2)
member of an/ ldahread:	read the archive header of a . . . . .	ldahread(3X)
common object file. ldhread:	read the file header of a . . . . .	ldhread(3X)
rewinddir, closedir:/ opendir,	readdir, telldir, seekdir, . . . . .	directory(3X)
open a common object file for	reading. ldopen, ldaopen: . . . . .	ldopen(3X)
exclusive file regions for	reading or writing. /provide . . . . .	locking(2)
open: open for	reading or writing. . . . .	open(2)
	readv: read from file. . . . .	readv(3N)
lseek: move	read/write file pointer. . . . .	lseek(2)
cmplx,/ int, ifx, idint,	real, float, srngl, dble, . . . . .	ftype(3F)
allocator. malloc, free,	realloc, calloc: main memory . . . . .	malloc(3C)
mailinfo: fast/ malloc, free,	realloc, calloc, mallopt, . . . . .	malloc(3X)
	reboot: reboot the system. . . . .	reboot(1M)
	reboot: reboot the system. . . . .	reboot(2)
reboot:	reboot the system. . . . .	reboot(1M)

Permuted Index

reboot:	reboot the system. . . . .	reboot(2)
specify what to do upon	receipt of a signal. signal: . . . . .	signal(2)
/specify Fortran action on	receipt of a system signal. . . . .	signal(3F)
recv, recvfrom, recvmsg:	receive a message from a/	recv(2N)
lockf:	record locking on files. . . . .	lockf(3C)
from per-process accounting	records. /command summary . . . . .	acctcms(1M)
errdead: extract error	records from dump. . . . .	errdead(1M)
manipulate connect accounting	records. fwtmp, wtmpfx: . . . . .	fwtmp(1M)
tape. frec:	recover files from a backup . . . . .	frec(1M)
receive a message from a/	recv, recvfrom, recvmsg: . . . . .	recv(2N)
message from a socket. recv,	recvfrom, recvmsg: receive a	recv(2N)
from a/ recv, recvfrom,	recvmsg: receive a message . . . . .	recv(2N)
ed,	red: text editor. . . . .	ed(1)
execute a regular expression.	regcmp, regex: compile and	regcmp(3X)
compile.	regcmp: regular expression . . . . .	regcmp(1)
make: maintain, update, and	regenerate groups of programs. . . . .	make(1)
regular expression. regcmp,	regex: compile and execute a	regcmp(3X)
compile and match routines.	regexp: regular expression . . . . .	regexp(5)
/provide exclusive file	regions for reading or/ . . . . .	locking(2)
match routines. regexp:	regular expression compile and	regexp(5)
regcmp:	regular expression compile. . . . .	regcmp(1)
regex: compile and execute a	regular expression. regcmp, . . . . .	regcmp(3X)
requests. accept,	reject: allow/prevent LP . . . . .	accept(1M)
sorted files. comm: select or	reject lines common to two . . . . .	comm(1)
lorder: find ordering	relation for an object/ . . . . .	lorder(1)
lorder5.0: find ordering	relation for an object/ . . . . .	lorder5.0(1)
join:	relational database operator. . . . .	join(1)
for a common object file.	reloc: relocation information . . . . .	reloc(4)
strip5.0: remove symbols and	relocation bits. . . . .	strip5.0(1)
ldrseek, ldrseek: seek to	relocation entries of a/ . . . . .	ldrseek(3X)
common object file. reloc:	relocation information for a . . . . .	reloc(4)
/fmod, fabs: floor, ceiling,	remainder, absolute value/ . . . . .	floor(3M)
mod. amod, dmod: Fortran	remaindering intrinsic/ . . . . .	mod(3F)
calendar:	reminder service. . . . .	calendar(1)
for returning a stream to a	remote command. /routines . . . . .	rcmd(3N)
rexec: return stream to a	remote command. . . . .	rexec(3N)
rexecd:	remote execution server. . . . .	rexecd(8N)
rcp:	remote file copy. . . . .	rcp(1N)
rlogin:	remote login. . . . .	rlogin(1N)
rlogind:	remote login server. . . . .	rlogind(8N)
put: puts a file onto a	remote machine. . . . .	put(1C)
take: takes a file from a	remote machine. . . . .	take(1C)
remsh:	remote shell. . . . .	remsh(1N)
remshd:	remote shell server. . . . .	remshd(8N)
ct: spawn getty to a	remote terminal. . . . .	ct(1C)
file. rmdel:	remove a delta from an SCCS . . . . .	rmdel(1)
semaphore set or/ ipcrm:	remove a message queue. . . . .	ipcrm(1)
unlink:	remove directory entry. . . . .	unlink(2)
rm, rmdir:	remove files or directories. . . . .	rm(1)
eqn constructs. deroff:	remove nroff/troff, tbl, and . . . . .	deroff(1)
bits. strip5.0:	remove symbols and relocation . . . . .	strip5.0(1)
from a queue. insque,	remque: insert/remove element . . . . .	insque(3N)
remsh:	remote shell. . . . .	remsh(1N)
remshd:	remote shell server. . . . .	remshd(8N)
check and interactive	repair. /system consistency . . . . .	fsck(1M)
uniq: report	repeated lines in a file. . . . .	uniq(1)
clock:	report CPU time used. . . . .	clock(3C)
communication/ ipcsc:	report inter-process . . . . .	ipcsc(1)
blocks. df:	report number of free disk . . . . .	df(1M)
errpt: process a	report of logged errors. . . . .	errpt(1M)
frequencies in a file. freq:	report on character . . . . .	freq(1)
sa2, sadc: system activity	report package. sal, . . . . .	sar(1M)
timex: time a command;	report process data and system/ . . . . .	timex(1)
ps:	report process status. . . . .	ps(1)

file. uniq:	report repeated lines in a . . . . .	uniq(1)
sar: system activity	reporter. . . . .	sar(1)
files. version:	reports version number of . . . . .	version(1)
stream. fseek, rewind, ftell:	reposition a file pointer in a . . . . .	fseek(3S)
/lpmove: start/stop the LP	request scheduler and move/ . . . . .	lpsched(1M)
reject: allow/prevent LP	requests. accept. . . . .	accept(1M)
LP request scheduler and move	requests. /start/stop the . . . . .	lpsched(1M)
lp, cancel: send/cancel	requests to an LP line/ . . . . .	lp(1)
teletype bits to a/ tset,	reset: set or reset the . . . . .	tset(1)
sensible/ tset, reset: set or	reset the teletype bits to a . . . . .	tset(1)
arp: Address	Resolution Protocol. . . . .	arp(5P)
object file. lddgetname:	retrieve symbol name for . . . . .	lddgetname(3X)
argument. getarg:	return Fortran command-line . . . . .	getarg(3F)
variable. getenv:	return Fortran environment . . . . .	getenv(3F)
accounting. mclock:	return Fortran time . . . . .	mclock(3F)
abs:	return integer absolute value. . . . .	abs(3C)
string. len:	return length of Fortran . . . . .	len(3F)
substring. index:	return location of Fortran . . . . .	index(3F)
logname:	return login name of user. . . . .	logname(3X)
command. rexec:	return stream to a remote . . . . .	rexec(3N)
name. getenv:	return value for environment . . . . .	getenv(3C)
stat: data	returned by stat system call. . . . .	stat(5)
/ruserok: routines for	returning a stream to a remote/ . . . . .	rcmd(3N)
configuration/ uvar:	returns system-specific . . . . .	uvar(2)
col: filter	reverse line-feeds. . . . .	col(1)
file pointer in a/ fseek,	rewind, ftell: reposition a . . . . .	fseek(3S)
/readdir, tekdir, seekdir,	rewinddir, closedir: flexible/ . . . . .	directory(3X)
creat: create a new file or	rewrite an existing one. . . . .	creat(2)
remote command.	rexec: return stream to a . . . . .	rexec(3N)
server.	rexecd: remote execution . . . . .	rexecd(8N)
	rlogin: remote login. . . . .	rlogin(1N)
	rlogind: remote login server. . . . .	rlogind(8N)
directories.	rm, rmdir: remove files or . . . . .	rm(1)
read mail. mail,	rmail: send mail to users or . . . . .	mail(1)
SCCS file.	rmdel: remove a delta from an . . . . .	rmdel(1)
directories. rm,	rmdir: remove files or . . . . .	rm(1)
Escape from the automatic	robots. autorobots: . . . . .	autorobots(6)
Try to escape the killer	robots. chase: . . . . .	chase(6)
robots.	robots: Escape from the . . . . .	robots(6)
robots: Escape from the	robots. . . . .	robots(6)
chroot: change	root directory. . . . .	chroot(2)
chroot: change	root directory for a command. . . . .	chroot(1M)
logarithm, power, square	root functions. /exponential, . . . . .	exp(3M)
/dsqrt, csqrt: Fortran square	root intrinsic function. . . . .	sqrt(3F)
routing tables.	route: manually manipulate the . . . . .	route(8N)
daemon.	routed: network routing . . . . .	routed(8N)
rcmd, rresvport, ruserok:	routines for returning a/ . . . . .	rcmd(3N)
Internet address manipulation	routines. /inet_netof: . . . . .	inet(3N)
common object file access	routines. ldfcn: . . . . .	ldfcn(4)
expression compile and match	routines. regexp: regular . . . . .	regexp(5)
terminal independent operation	routines. /tgoto, tputs: . . . . .	termcap(3X)
routed: network	routing daemon. . . . .	routed(8N)
route: manually manipulate the	routing tables. . . . .	route(8N)
for returning a stream/ rcmd,	rresvport, ruserok: routines . . . . .	rcmd(3N)
standard/restricted/ sh,	rsh: shell, the . . . . .	sh(1)
and, or, xor, not, lshift,	rshift: Fortran bitwise/ . . . . .	bool(3F)
nice:	run a command at low priority. . . . .	nice(1)
hangups (sh only). nohup:	run a command immune to . . . . .	nohup(1)
runacct:	run daily accounting. . . . .	runacct(1M)
	runacct: run daily accounting. . . . .	runacct(1M)
/prctmp, prdaily, pracct,	runacct, shutacct, startup/ . . . . .	acctsh(1M)
local machines.	ruptime: show host status of . . . . .	ruptime(1N)
returning a/ rcmd, rresvport,	ruserok: routines for . . . . .	rcmd(3N)
machines.	rwho: who's logged in on local . . . . .	rwho(1N)

Permuted Index

	rwhod: system status server. . . . .	rwhod(8N)
activity report package.	sa1, sa2, sadc: system . . . . .	sar(1M)
report package. sa1,	sa2, sadc: system activity . . . . .	sar(1M)
editing activity.	sact: print current SCCS file . . . . .	sact(1)
package. sa1, sa2,	sadc: system activity report . . . . .	sar(1M)
	sag: system activity graph. . . . .	sag(1G)
	sar: system activity reporter. . . . .	sar(1)
space allocation. brk,	sbrk: change data segment . . . . .	brk(2)
formatted input.	scanf, fscanf, sscanf: convert . . . . .	scanf(3S)
bfs: big file	scanner. . . . .	bfs(1)
language. awk: pattern	scanning and processing . . . . .	awk(1)
the delta commentary of an	SCCS delta. cdc: change . . . . .	cdc(1)
comb: combine	SCCS deltas. . . . .	comb(1)
make a delta (change) to an	SCCS file. delta: . . . . .	delta(1)
sact: print current	SCCS file editing activity. . . . .	sact(1)
get: get a version of an	SCCS file. . . . .	get(1)
prs: print an	SCCS file. . . . .	prs(1)
rmddl: remove a delta from an	SCCS file. . . . .	rmddl(1)
compare two versions of an	SCCS file. sccsdiff: . . . . .	sccsdiff(1)
sccsfile: format of	SCCS file. . . . .	sccsfile(4)
undo a previous get of an	SCCS file. unget: . . . . .	unget(1)
val: validate	SCCS file. . . . .	val(1)
admin: create and administer	SCCS files. . . . .	admin(1)
what: identify	SCCS files. . . . .	what(1)
help: ask for help in using	SCCS. . . . .	help(1)
of an SCCS file.	sccsdiff: compare two versions . . . . .	sccsdiff(1)
	sccsfile: format of SCCS file. . . . .	sccsfile(4)
/start/stop the LP request	scheduler and move requests. . . . .	lpsched(1M)
common object file.	scnhdr: section header for a . . . . .	scnhdr(4)
clear: clear terminal	screen. . . . .	clear(1)
optimization/ curses: CRT	screen handling and . . . . .	curses(3X)
twinkle: twinkle stars on the	screen. . . . .	twinkle(6)
display editor based on/ vi:	screen-oriented (visual) . . . . .	vi(1)
inittab:	script for the init process. . . . .	inittab(4)
terminal session.	script: make typescript of . . . . .	script(1)
system initialization shell	scripts. /rc, powerfail: . . . . .	brc(1M)
	sdb: symbolic debugger. . . . .	sdb(1)
	sdiff: side-by-side difference . . . . .	sdiff(1)
program.	search a file for a pattern. . . . .	grep(1)
grep, egrep, fgrep:	search a sorted table. . . . .	bsearch(3C)
bsearch: binary	search and print process . . . . .	acctcom(1)
accounting file(s). acctcom:	search and update. . . . .	lsearch(3C)
lsearch, lfind: linear	search tables. hsearch, . . . . .	hsearch(3C)
hcreate, hdestroy: manage hash	search trees. tsearch, tfind, . . . . .	tsearch(3C)
tdelete, twalk: manage binary	section header for a common . . . . .	scnhdr(4)
object file. scnhdr:	section header of a common . . . . .	ldshdr(3X)
object/ /read an indexed/named	section of a common object/ . . . . .	ldlseek(3X)
/to line number entries of a	section of a common object/ . . . . .	ldrseek(3X)
/to relocation entries of a	section of a common object/ . . . . .	ldsseek(3X)
/seek to an indexed/named	section sizes of common object . . . . .	size(1)
files. size: print	sed: stream editor. . . . .	sed(1)
	seed48, lcong48: generate/ . . . . .	drand48(3C)
/mrand48, jrand48, srand48,	seek to an indexed/named . . . . .	ldsseek(3X)
section of/ ldsseek, ldnseek:	seek to line number entries of . . . . .	ldlseek(3X)
a section/ ldsseek, ldnseek:	seek to relocation entries of . . . . .	ldrseek(3X)
a section/ ldrseek, ldnseek:	seek to the optional file . . . . .	ldohseek(3X)
header of a common/ ldohseek:	seek to the symbol table of a . . . . .	ldtbseek(3X)
common object file. ldtbseek:	seekdir, rewinddir, closedir:/ . . . . .	directory(3X)
opendir, readdir, telldir,	segment. . . . .	shmget(2)
shmget: get shared memory	segment space allocation. . . . .	brk(2)
brk, sbrk: change data	select or reject lines common . . . . .	comm(1)
to two sorted files. comm:	select: synchronous i/o . . . . .	select(2N)
multiplexing.	select terminal filter. . . . .	greek(1)
greek:	selected fields of each line . . . . .	cut(1)
of a file. cut: cut out		

file. dump: dump selected parts of an object . . . . . dump(1)  
 semctl: semaphore control operations. . . . . semctl(2)  
 semop: semaphore operations. . . . . semop(2)  
 ipcrm: remove a message queue, semaphore set or shared memory/  
 semget: get set of semaphores. . . . . ipcrm(1)  
 operations. . . . . semget(2)  
 semctl: semaphore control . . . . . semctl(2)  
 semget: get set of semaphores. . . . . semget(2)  
 semop: semaphore operations. . . . . semop(2)  
 send, sendto, sendmsg: send a message from a socket. . . . . send(2N)  
 a group of processes. kill: send a signal to a process or . . . . . kill(2)  
 mail. mail, rmail: send mail to users or read . . . . . mail(1)  
 message from a socket. send, sendto, sendmsg: send a . . . . . send(2N)  
 group. killpg: send signal to a process . . . . . killpg(3N)  
 line printer. lp, cancel: send/cancel requests to an LP . . . . . lp(1)  
 socket. send, sendto, sendmsg: send a message from a . . . . . send(2N)  
 message from a socket. sendto, sendmsg: send a . . . . . send(2N)  
 reset the teletype bits to a sensible state. /reset: set or . . . . . tset(1)  
 File Transfer Protocol server. ftpd: DARPA Internet . . . . . ftpd(8N)  
 remshd: remote shell server. . . . . remshd(8N)  
 rexecd: remote execution server. . . . . rexecd(8N)  
 rlogind: remote login server. . . . . rlogind(8N)  
 rwhod: system status server. . . . . rwhod(8N)  
 telnetd: DARPA TELNET protocol server. . . . . telnetd(8N)  
 Trivial File Transfer Protocol server. tftpd: DARPA . . . . . tftpd(8N)  
 make typescript of terminal session. script: . . . . . script(1)  
 buffering to a stream. setbuf, setvbuf: assign . . . . . setbuf(3S)  
 IDs. setuid, setgid: set user and group . . . . . setuid(2)  
 getgrent, getgrgid, getgrnam, setgrent, endgrent, fgetgrent:/ . . . . . getgrent(3C)  
 /gethostbyaddr, gethostbyname, sethostent, endhostent: get/ . . . . . gethostent(3N)  
 identifier of/ gethostid, sethostid: get/set unique . . . . . gethostid(2N)  
 current host. gethostname, sethostname: get/set name of . . . . . gethostname(2N)  
 goto. setjmp, longjmp: non-local . . . . . setjmp(3C)  
 encryption. crypt, setkey, encrypt: generate DES . . . . . crypt(3C)  
 /getnetbyaddr, getnetbyname, setmnt: establish mount table. . . . . setmnt(1M)  
 protocol/ /getprotobyname, setnetent, endnetent: get/ . . . . . setnetent(3N)  
 getpwnid, getpwuid, getpwnam, setpgrp: set process group ID. . . . . setpgrp(2)  
 effective group ID. setprotoent, endprotoent: get . . . . . getprotoent(3N)  
 effective user ID's. setpwent, endpwent, fgetpwent:/ . . . . . getpwent(3C)  
 /getservbyport, getservbyname, setregid: set real and . . . . . setregid(2)  
 options on/ getsockopt, setreuid: set real and . . . . . setreuid(2)  
 login time. profile: setservent, endservent: get/ . . . . . getservent(3N)  
 gettydefs: speed and terminal setsockopt: get and set . . . . . getsockopt(2N)  
 disk tune: tune floppy disk setting up an environment at . . . . . profile(4)  
 group IDs. settings used by getty. . . . . gettydefs(4)  
 /getuid, getutline, pututline, settling time parameters. . . . . disk tune(1M)  
 stream. setbuf, setuid, setgid: set user and . . . . . setuid(2)  
 data in a machine/ sputl, setutent, endutent, utmpname:/ . . . . . getut(3C)  
 a command immune to hangups, setvbuf: assign buffering to a . . . . . setbuf(3S)  
 standard/restricted command/ sgetl: access long integer . . . . . sputl(3X)  
 operations. shmctl: (sh only). nohup: run . . . . . nohup(1)  
 queue, semaphore set or sh, rsh: shell, the . . . . . sh(1)  
 shmop: shared memory control . . . . . shmctl(2)  
 shmget: get shared memory id. /a message . . . . . ipcrm(1)  
 from C programs to implement shared memory operations. . . . . shmop(2)  
 system: issue a shared memory segment. . . . . shmget(2)  
 with C-like syntax. csh: a shared strings. /strings . . . . . xstr(1)  
 system: issue a shell command from Fortran. . . . . system(3F)  
 shl: shell (command interpreter) . . . . . csh(1)  
 shutacct, startup, turnacct: shell command. . . . . system(3S)  
 remsh: remote shell layer manager. . . . . shl(1)  
 system initialization shell procedures for/ /runacct. . . . . acctsh(1M)  
 remshd: remote shell. . . . . remsh(1N)  
 shell. shell scripts. /rc, powerfail: . . . . . brc(1M)  
 shell server. . . . . remshd(8N)

Permuted Index

command programming/ sh, rsh:	shell, the standard/restricted . . . . .	sh(1)
	sh: shell layer manager. . . . .	sh(1)
operations.	shmctl: shared memory control . . . . .	shmctl(2)
segment.	shmget: get shared memory . . . . .	shmget(2)
operations.	shmop: shared memory . . . . .	shmop(2)
full-duplex/ shutdown:	shut down part of a . . . . .	shutdown(2N)
/prdaily, prtacct, runacct,	shutacct, startup, turnacct:/ . . . . .	acctsh(1M)
full-duplex connection.	shutdown: shut down part of a . . . . .	shutdown(2N)
processing.	shutdown: terminate all . . . . .	shutdown(1M)
program. sdiff:	side-by-side difference . . . . .	sdiff(1)
transfer-of-sign intrinsic/	sign, isign, dsign: Fortran . . . . .	sign(3F)
login:	sign on. . . . .	login(1)
pause: suspend process until	signal. . . . .	pause(2)
what to do upon receipt of a	signal. signal: specify . . . . .	signal(2)
action on receipt of a system	signal. /specify Fortran . . . . .	signal(3F)
on receipt of a system/	signal: specify Fortran action . . . . .	signal(3F)
upon receipt of a signal.	signal: specify what to do . . . . .	signal(2)
killpg: send	signal to a process group. . . . .	killpg(3N)
of processes. kill: send a	signal to a process or a group . . . . .	kill(2)
ssignal, gsignal: software	signals. . . . .	ssignal(3C)
lex: generate programs for	simple lexical tasks. . . . .	lex(1)
generator. rand, srand:	simple random-number . . . . .	rand(3C)
tc: phototypesetter	simulator. . . . .	tc(1)
atan, atan2: trigonometric/	sin, cos, tan, asin, acos, . . . . .	trig(3M)
intrinsic function.	sin, dsin, csin: Fortran sine . . . . .	sin(3F)
sin, dsin, csin: Fortran	sine intrinsic function. . . . .	sin(3F)
/dsinh: Fortran hyperbolic	sine intrinsic function. . . . .	sinh(3F)
ssp: make output	single spaced. . . . .	ssp(1)
functions.	sinh, cosh, tanh: hyperbolic . . . . .	sinh(3M)
hyperbolic sine intrinsic/	sinh, dsinh: Fortran . . . . .	sinh(3F)
get descriptor table	size. getdtablesize: . . . . .	getdtablesize(3N)
size5.0:	size of an object file. . . . .	size5.0(1)
common object files.	size: print section sizes of . . . . .	size(1)
file.	size5.0: size of an object . . . . .	size5.0(1)
size: print section	sizes of common object files. . . . .	size(1)
an interval.	sleep: suspend execution for . . . . .	sleep(1)
interval.	sleep: suspend execution for . . . . .	sleep(3C)
documents, view graphs, and	slides. mmt, mvt: typeset . . . . .	mmt(1)
typesetting view graphs and	slides. /macro package for . . . . .	mv(5)
current/ ttypeslot: find the	slot in the utmp file of the . . . . .	ttypeslot(3C)
spline: interpolate	smooth curve. . . . .	spline(1G)
int, ifix, idint, real, float,	sngl, dble, cmplx, dcmplx,/ . . . . .	ftype(3F)
sno:	sno: SNOBOL interpreter. . . . .	sno(1)
sno:	SNOBOL interpreter. . . . .	sno(1)
accept a connection on a	socket. accept: . . . . .	accept(2N)
bind: bind a name to a	socket. . . . .	bind(2N)
initiate a connection on a	socket. connect: . . . . .	connect(2N)
communication.	socket: create an endpoint for . . . . .	socket(2N)
listen for connections on a	socket. listen: . . . . .	listen(2N)
getsockname: get	socket name. . . . .	getsockname(2N)
receive a message from a	socket. /recvfrom, recvmsg: . . . . .	recv(2N)
sendmsg: send a message from a	socket. send, sendto, . . . . .	send(2N)
get and set options on	sockets. /setsockopt: . . . . .	getsockopt(2N)
pg: file perusal filter for	soft-copy terminals. . . . .	pg(1)
interface. lo:	software loopback network . . . . .	lo(5)
ssignal, gsignal:	software signals. . . . .	ssignal(3C)
sort:	sort and/or merge files. . . . .	sort(1)
qsort: quicker	sort. . . . .	qsort(3C)
tsort: topological	sort: sort and/or merge files. . . . .	sort(1)
or reject lines common to two	sort. . . . .	tsort(1)
bsearch: binary search a	sorted files. comm: select . . . . .	comm(1)
for program. whereis: locate	sorted table. . . . .	bsearch(3C)
message file by massaging C	source, binary, and/or manual . . . . .	whereis(1)
	source. /create an error . . . . .	mkstr(1)



brk, sbrk: change data segment  
 ssp: make output single  
 terminal. ct:  
   fspec: format  
 the extended errors in the  
 receipt of a system/ signal:  
 receipt of a signal. signal:  
 /set terminal type, modes,  
 used by getty. gettydefs:  
 hashcheck: find spelling/  
 spelling/ spell, hashmake,  
 spellin, hashcheck: find  
 spelling errors. /hashmake,  
 spline: interpolate smooth  
 split a file into pieces.  
 split.  
 split f77, ratfor, or ell  
 split: split a file into  
 spool directory clean-up.  
 spooling system.  
 sprintf: print formatted  
 spurl, sgetl: access long  
 sqrt, dsqrt, csqrt: Fortran  
 sqrt: exponential, logarithm,  
 square root functions. /sqrt:  
 square root intrinsic/  
 srand, rand: Fortran uniform  
 srand: simple random-number  
 srand48, mrand48, jrand48,  
 rcvhex: translates Motorola  
 input. scanf, fscanf,  
 signals.  
 spaced.  
 package. stdio:  
 communication package. ftok:  
   sh, rsh: shell, the  
   twinkle: twinkle  
 lpsched, lpsht, lpmove:  
   boot:  
 /prtacct, runacct, shutacct,  
   system call.  
 stat: data returned by  
 ff: list file names and  
 lav: print load average  
 ustat: get file system  
 lpstat: print LP  
 feof, clearerr, fileno: stream  
 control. uustat: uucp  
 communication facilities  
 netstat: show network  
 ruptime: show host  
 ps: report process  
 rwhod: system  
 stat, fstat: get file  
 input/output package.  
 wait for child process to  
 wait for child process to  
 strncmp, strcpy, strncpy,  
 /strcpy, strncpy, strlen,  
 strncpy, / strcat, strncat,  
 /strncat, strcmp, strncmp,  
 /strchr, strpbrk, strspn,  
 sed:  
 space allocation.  
 spaced.  
 spawn getty to a remote  
 specification in text files.  
 specified device. /turn on/off  
 specify Fortran action on  
 specify what to do upon  
 speed, and line discipline.  
 speed and terminal settings  
 spell, hashmake, spellin,  
 spellin, hashcheck: find  
 spelling errors. /hashmake,  
 spline: interpolate smooth  
 split a file into pieces.  
 split.  
 split f77, ratfor, or ell  
 split: split a file into  
 spool directory clean-up.  
 spooling system.  
 sprintf: print formatted  
 spurl, sgetl: access long  
 sqrt, dsqrt, csqrt: Fortran  
 sqrt: exponential, logarithm,  
 square root functions. /sqrt:  
 square root intrinsic/  
 srand, rand: Fortran uniform  
 srand: simple random-number  
 srand48, seed48, lcong48:/  
 S-records from downloading/  
 sscanf: convert formatted  
 ssignal, gsignal: software  
 ssp: make output single  
 standard buffered input/output  
 standard interprocess  
 standard/restricted command/  
 stars on the screen.  
 start/stop the LP request/  
 startup procedures.  
 startup, turnacct: shell/  
 stat: data returned by stat  
 stat, fstat: get file status.  
 stat system call.  
 statistics for a file system.  
 statistics.  
 statistics.  
 status information.  
 status inquiries. ferror,  
 status inquiry and job  
 status. /report inter-process  
 status.  
 status of local machines.  
 status.  
 status server.  
 status.  
 stdio: standard buffered  
 stime: set time.  
 stop or terminate. wait:  
 stop or terminate. wait3:  
 strcat, strncat, strcmp,  
 strchr, strchr, strpbrk,  
 strcmp, strncmp, strcpy,  
 strcpy, strncpy, strlen,  
 strcspn, strtok: string/  
 stream editor.  
 brk(2)  
 ssp(1)  
 ct(1C)  
 fspec(4)  
 exterr(1)  
 signal(3F)  
 signal(2) -  
 getty(1M)  
 gettydefs(4)  
 spell(1)  
 spell(1)  
 spell(1)  
 spline(1G)  
 split(1)  
 csplit(1)  
 fsplit(1)  
 split(1)  
 uuclean(1M)  
 lpadmin(1M)  
 printf(3S)  
 sputl(3X)  
 sqrt(3F)  
 exp(3M)  
 exp(3M)  
 sqrt(3F)  
 rand(3F)  
 rand(3C)  
 drand48(3C)  
 rcvhex(1)  
 scanf(3S)  
 signal(3C)  
 ssp(1)  
 stdio(3S)  
 stidpc(3C)  
 sh(1)  
 twinkle(6)  
 lpsched(1M)  
 boot(8)  
 acctish(1M)  
 stat(5)  
 stat(2)  
 stat(5)  
 ff(1M)  
 lav(1)  
 ustat(2)  
 lpstat(1)  
 ferror(3S)  
 uustat(1C)  
 ipcs(1)  
 netstat(1N)  
 ruptime(1N)  
 ps(1)  
 rwhod(8N)  
 stat(2)  
 stdio(3S)  
 stime(2)  
 wait(2)  
 wait3(2N)  
 string(3C)  
 string(3C)  
 string(3C)  
 string(3C)  
 string(3C)  
 sed(1)

fflush: close or flush a stream. fclose, . . . . . fclose(3S)  
 fopen, freopen, fdopen: open a stream. . . . . fopen(3S)  
 reposition a file pointer in a stream. fseek, rewind, ftell: . . . . . fseek(3S)  
 get character or word from a stream. /getchar, fgetc, getw: . . . . . getc(3S)  
 fgets: get a string from a stream. gets, . . . . . gets(3S)  
 put character or word on a stream. /putchar, fputc, putw: . . . . . putc(3S)  
 puts, fputs: put a string on a stream. . . . . puts(3S)  
 setvbuf: assign buffering to a stream. setbuf, . . . . . setbuf(3S)  
 /feof, clearerr, fcntl: stream status inquiries. . . . . ferrord(3S)  
 /routines for returning a stream to a remote command. . . . . rcmd(3N)  
 rexec: return stream to a remote command. . . . . rexec(3N)  
 push character back into input stream. ungetc: . . . . . ungetc(3S)  
 long integer and base-64 ASCII string. /l64a: convert between lge, lgt, lle, llt: string comparison intrinsic/ . . . . . strcmp(3F)  
 convert date and time to string. /asctime, tzset: . . . . . ctime(3C)  
 floating-point number to string. /fcvt, gcvt: convert . . . . . ecvt(3C)  
 gets, fgets: get a string from a stream. . . . . gets(3S)  
 len: return length of Fortran string. . . . . len(3F)  
 puts, fputs: put a string on a stream. . . . . puts(3S)  
 bcmp, bzero, fs: bit and byte string operations. bcopy. . . . . bstring(3N)  
 strspn, strcspn, strtok: string operations. /strpbrk, . . . . . string(3C)  
 number. strtod, atof: convert string to double-precision . . . . . strtod(3C)  
 number. atof: convert ASCII string to floating-point . . . . . atof(3C)  
 strtol, atol, atoi: convert string to integer. . . . . strtol(3C)  
 strings in an object, or strings: find the printable . . . . . strings(1)  
 implement/ xstr: extract strings from C programs to . . . . . xstr(1)  
 strings: find the printable strings in an object, or other/ . . . . . strings(1)  
 C programs to implement shared sirings. /extract strings from . . . . . xstr(1)  
 number information from an/ strip: strip symbol and line . . . . . strip(1)  
 information from an/ strip: strip symbol and line number . . . . . strip(1)  
 relocation bits. strip5.0: remove symbols and . . . . . strip5.0(1)  
 /strncpy, strcpy, strncpy, strlen, strchr, strchr, / . . . . . string(3C)  
 strcpy, strncpy, / strcat, strncat, strncat, strncpy, / . . . . . string(3C)  
 strcat, strncat, strcmp, strcmp, strcpy, / . . . . . string(3C)  
 /strcmp, strncmp, strcpy, strncpy, strlen, strchr, / . . . . . string(3C)  
 /strlen, strchr, strchr, strpbrk, strspn, strcspn, / . . . . . string(3C)  
 /strncpy, strlen, strchr, strchr, strpbrk, strspn, / . . . . . string(3C)  
 /strchr, strchr, strpbrk, strtod, atof: convert string . . . . . strtod(3C)  
 to double-precision number. strtok: string operations. . . . . strtok(3C)  
 /strpbrk, strspn, strcspn, strtol, atol, atoi: convert . . . . . strtol(3C)  
 string to integer. structure. fuser: identify . . . . . fuser(1M)  
 processes using a file or file sty: set the options for a . . . . . sty(1)  
 terminal. su: become super-user or . . . . . su(1)  
 another user. subroutines and libraries. . . . . intro(3)  
 intro: introduction to subroutines. . . . . plot(3X)  
 plot: graphics interface subsequent lines of one file. . . . . paste(1)  
 /same lines of several files or substrings. index: . . . . . index(3M)  
 return location of Fortran file. sum7: . . . . . sum7(1)  
 file. sum7: sum and count blocks in a . . . . . sum7(1)  
 the files in the/ sumdir: sum and count characters in . . . . . sumdir(1)  
 count of a file. sum: print checksum and block . . . . . sum(1)  
 a file. sum7: sum and count blocks in . . . . . sum7(1)  
 characters in the files in/ sumdir: sum and count . . . . . sumdir(1)  
 du: summarize disk usage. . . . . du(1)  
 accounting/ acctcms: command summary from per-process . . . . . acctcms(1M)  
 sync: update the super block. . . . . sync(1)  
 sync: update super-block. . . . . sync(2)  
 su: become super-user or another user. . . . . su(1)  
 interval. sleep: suspend execution for an . . . . . sleep(1)  
 interval. sleep: suspend execution for . . . . . sleep(3C)  
 pause: suspend process until signal. . . . . pause(2)  
 swab: swap bytes. . . . . swab(3C)  
 swab: swap bytes. . . . . swab(3C)  
 sxt: pseudo-device driver. . . . . sxt(7)

information from/ strip: strip  
 ldgetname: retrieve  
 object/ /compute the index of a  
 ldtbody: read an indexed  
 syms: common object file  
 object/ ldtbseek: seek to the  
     sdb:  
     strip5.0: remove  
     symbol table format.  
     select:  
 interpreter) with C-like  
     error/ perror, errno,  
     perror, errno, sys\_errlist,  
     information. uvar: returns  
     binary search a sorted  
 /compute the index of a symbol  
 file. /read an indexed symbol  
 common object file symbol  
 master device information  
 mnttab: mounted file system  
 ldtbseek: seek to the symbol  
 setmnt: establish mount  
 getdtablesize: get descriptor  
 tbl: format  
 hdestroy: manage hash search  
 manipulate the routing  
 tabs: set  
     ctags: maintain a  
     a file.  
     remote machine.  
     machine. take:  
     talk:  
     trigonometric/ sin, cos,  
     intrinsic function.  
     tan, dtan: Fortran  
     /dtanh: Fortran hyperbolic  
     hyperbolic tangent intrinsic/  
     sinh, cosh,  
     tp: manipulate  
     tar:  
     tape file archiver.  
     recover files from a backup  
     file system backup. filesave,  
     programs for simple lexical  
     deroff: remove nroff/troff,  
     or troff.  
 Control Protocol.  
 search trees. tsearch, tfind,  
 4014: paginator for the  
 tset, reset: set or reset the  
 last logins of users and  
 initialization. init,  
 closedir:/ opendir, readdir,  
 telnetd: DARPA  
 telnet: user interface to the  
 TELNET protocol.  
     server.  
     symbol and line number . . . . . strip(1)  
     symbol name for object file. . . . . ldgetname(3X)  
     symbol table entry of a common . . . . . ldtbindex(3X)  
     symbol table entry of a common/ . . . . . ldttbody(3X)  
     symbol table format. . . . . syms(4)  
     symbol table of a common . . . . . ldtbseek(3X)  
     symbolic debugger. . . . . sdb(1)  
     symbols and relocation bits. . . . . strip5.0(1)  
     syms: common object file . . . . . syms(4)  
     sync: update super-block. . . . . sync(2)  
     sync: update the super block. . . . . sync(1)  
     synchronous i/o multiplexing. . . . . select(2N)  
     syntax. csh: a shell (command . . . . . csh(1)  
     sysdef: system definition. . . . . sysdef(1M)  
     sys\_errlist, sys\_nerr: system . . . . . perror(3C)  
     sys\_nerr: system error/ . . . . . perror(3C)  
     system-specific configuration . . . . . uvar(2)  
     table. bsearch: . . . . . bsearch(3C)  
     table entry of a common object/ . . . . . ldtbindex(3X)  
     table entry of a common object . . . . . ldttbody(3X)  
     table format. syms: . . . . . syms(4)  
     table. master: . . . . . master(4)  
     table. . . . . mnttab(4)  
     table of a common object file. . . . . ldtbseek(3X)  
     table. . . . . setmnt(1M)  
     table size. . . . . getdtablesize(3N)  
     tables for nroff or troff. . . . . tbl(1)  
     tables. hsearch, hcreate, . . . . . hsearch(3C)  
     tables. route: manually . . . . . route(8N)  
     tabs on a terminal. . . . . tabs(1)  
     tabs: set tabs on a terminal. . . . . tabs(1)  
     tags file for a C program. . . . . ctags(1)  
     tail: deliver the last part of . . . . . tail(1)  
     take: takes a file from a . . . . . take(1C)  
     takes a file from a remote . . . . . take(1C)  
     talk: talk to another user. . . . . talk(1N)  
     talk to another user. . . . . talk(1N)  
     tan, asin, acos, atan, atan2: . . . . . trig(3M)  
     tan, dtan: Fortran tangent . . . . . tan(3F)  
     tangent intrinsic function. . . . . tan(3F)  
     tangent intrinsic function. . . . . tanh(3F)  
     tanh, dtanh: Fortran . . . . . tanh(3F)  
     tanh: hyperbolic functions. . . . . sinh(3M)  
     tape archive. . . . . tp(1)  
     tape file archiver. . . . . tar(1)  
     tape. freq: . . . . . freq(1M)  
     tapesave: daily/weekly UNIX . . . . . filesave(1M)  
     tar: tape file archiver. . . . . tar(1)  
     tasks. lex: generate . . . . . lex(1)  
     tbl, and eqn constructs. . . . . deroff(1)  
     tbl: format tables for nroff . . . . . tbl(1)  
     tc: phototypesetter simulator. . . . . tc(1)  
     tcp: Internet Transmission . . . . . tcp(5P)  
     tdelete, twalk: manage binary . . . . . tsearch(3C)  
     tee: pipe fitting. . . . . tee(1)  
     Tektronix 4014 terminal. . . . . 4014(1)  
     teletype bits to a sensible/ . . . . . tset(1)  
     teletypes. last: indicate . . . . . last(1)  
     telinit: process control . . . . . init(1M)  
     telldir, seekdir, rewinddir, . . . . . directory(3X)  
     TELNET protocol server. . . . . telnetd(8N)  
     TELNET protocol. . . . . telnet(1N)  
     telnet: user interface to the . . . . . telnet(1N)  
     telnetd: DARPA TELNET protocol . . . . . telnetd(8N)

temporary file. tmpnam,	temporary file. . . . .	tmpnam(3S)
tmpfile: create a	temporary file. . . . .	tmpfile(3S)
tmpnam: create a name for a	temporary file. tmpnam, . . . . .	tmpnam(3S)
terminals.	term: conventional names for . . . . .	term(5)
term: format of compiled	term file. . . . .	term(4)
file..	term: format of compiled term . . . . .	term(4)
data base.	termcap: terminal capability . . . . .	termcap(5)
for the Tektronix 4014	terminal. 4014: paginator . . . . .	4014(1)
functions of the DASI 450	terminal. 450: handle special . . . . .	450(1)
termcap:	terminal capability data base. . . . .	termcap(5)
terminfo:	terminal capability data base. . . . .	terminfo(4)
ct: spawn getty to a remote	terminal. . . . .	ct(1C)
ctermid: generate filename for	terminal. . . . .	ctermid(3S)
pty: pseudo	terminal driver. . . . .	pty(5)
greek: select	terminal filter. . . . .	greek(1)
/tgetstr, tgoto, tputs:	terminal independent operation/ . . . . .	termcap(3X)
termio: general	terminal interface. . . . .	termio(7)
tty: controlling	terminal interface. . . . .	tty(7)
dial: establish an out-going	terminal line connection. . . . .	dial(3C)
clear: clear	terminal screen. . . . .	clear(1)
script: make typescript of	terminal session. . . . .	script(1)
getty. gettydefs: speed and	terminal settings used by . . . . .	gettydefs(4)
stty: set the options for a	terminal. . . . .	stty(1)
tabs: set tabs on a	terminal. . . . .	tabs(1)
isatty: find name of a	terminal. ttyname, . . . . .	ttyname(3C)
and line/ getty: set	terminal type, modes, speed, . . . . .	getty(1M)
ttytype: data base of	terminal types by port. . . . .	ttytype(4)
animate worms on a display	terminal. worms: . . . . .	worms(6)
functions of DASI 300 and 300s	terminals. /handle special . . . . .	300(1)
tty: get the	terminal's name. . . . .	tty(1)
perusal filter for soft-copy	terminals. pg: file . . . . .	pg(1)
term: conventional names for	terminals. . . . .	term(5)
kill:	terminate a process. . . . .	kill(1)
shutdown:	terminate all processing. . . . .	shutdown(1M)
abort:	terminate Fortran program. . . . .	abort(3F)
exit, _exit:	terminate process. . . . .	exit(2)
daemon, errstop:	terminate the error-logging . . . . .	errstop(1M)
for child process to stop or	terminate. wait: wait . . . . .	wait(2)
for child process to stop or	terminate. wait3: wait . . . . .	wait3(2N)
tic:	terminfo compiler. . . . .	tic(1M)
tput: query	terminfo database. . . . .	tput(1)
data base.	terminfo: terminal capability . . . . .	terminfo(4)
interface.	termio: general terminal . . . . .	termio(7)
command.	test: condition evaluation . . . . .	test(1)
quiz:	test your knowledge. . . . .	quiz(6)
ed, red:	text editor. . . . .	ed(1)
ex, edit:	text editor. . . . .	ex(1)
change the format of a	text file. newform: . . . . .	newform(1)
fspec: format specification in	text files. . . . .	fspec(4)
/checkeq: format mathematical	text for nroff or troff. . . . .	eqn(1)
prepare constant-width	text for troff. cw, checkcw: . . . . .	cw(1)
nroff: format	text. . . . .	nroff(1)
plock: lock process,	text, or data in memory. . . . .	plock(2)
troff: typeset	text. . . . .	troff(1)
binary search trees. tsearch,	tfind, tdelete, twalk: manage . . . . .	tsearch(3C)
Transfer Protocol server.	tftpd: DARPA Trivial File . . . . .	tftpd(8N)
tgetstr, tgoto, tputs:/	tgetent, tgetnum, tgetflag, . . . . .	termcap(3X)
tputs:/ tgetent, tgetnum,	tgetflag, tgetstr, tgoto, . . . . .	termcap(3X)
tgoto, tputs:/ tgetent,	tgetnum, tgetflag, tgetstr, . . . . .	termcap(3X)
tgetent, tgetnum, tgetflag,	tgetstr, tgoto, tputs:/ . . . . .	termcap(3X)
/tgetnum, tgetflag, tgetstr,	tgoto, tputs: terminal/ . . . . .	termcap(3X)
ttt, cubic:	tic: terminfo compiler. . . . .	tic(1M)
data and system/ timex:	tic-tac-toe. . . . .	ttt(6)
	time a command; report process . . . . .	timex(1)

time:	time a command.	time(1)
mclock: return Fortran	time accounting.	mclock(3F)
execute commands at a later	time. at, batch:	at(1)
systems for optimal access	time. dcopy: copy file	dcopy(1M)
	time: get time.	time(2)
tune floppy disk settling	time parameters. disktime:	disktime(1M)
profit: execution	time profile.	profil(2)
up an environment at login	time. profile: setting	profile(4)
stime: set	time.	stime(2)
	time: time a command.	time(1)
time: get	time.	time(2)
tzset: convert date and	time to string. /asctime,	ctime(3C)
clock: report CPU	time used.	clock(3C)
process times.	times: get process and child	times(2)
update access and modification	times of a file. touch:	touch(1)
get process and child process	times. times:	times(2)
file access and modification	times. utime: set	utime(2)
process data and system/	timex: time a command; report	timex(1)
file.	tmpfile: create a temporary	tmpfile(3S)
for a temporary file.	tmpnam, tmpnam: create a name	tmpnam(3S)
/tolower, _toupper, _tolower,	toascii: translate characters.	conv(3C)
popen, pclose: initiate pipe	to/from a process.	popen(3S)
toupper, tolower, _toupper,	_tolower, toascii: translate/	conv(3C)
toascii: translate/ toupper,	tolower, _toupper, _tolower,	conv(3C)
tsort:	topological sort.	tsort(1)
acctmrg: merge or add	total accounting files.	acctmrg(1M)
modification times of a file.	touch: update access and	touch(1)
translate/ toupper, tolower,	_toupper, _tolower, toascii:	conv(3C)
_tolower, toascii: translate/	toupper, tolower, _toupper,	conv(3C)
	tp: manipulate tape archive.	tp(1)
/tgetflag, tgetstr, tgoto,	tplot: graphics filters.	tplot(1G)
	tput: query terminfo database.	tput(1)
	tputs: terminal independent/	termcap(3X)
	tr: translate characters.	tr(1)
ptrace: process	trace.	ptrace(2)
trpt: transliterate protocol	trace.	trpt(8N)
blt, blt512: block	transfer data.	blt(3C)
ftp: file	transfer program.	ftp(1N)
ftpd: DARPA Internet File	Transfer Protocol server.	ftpd(8N)
tftpd: DARPA Trivial File	Transfer Protocol server.	tftpd(8N)
sign, isign, dsign: Fortran	transfer-of-sign intrinsic/	sign(3F)
/_toupper, _tolower, toascii:	translate characters.	conv(3C)
	translate characters.	tr(1)
from downloading into/ rcvhex:	translates Motorola S-records	rcvhex(1)
hex:	translates object files.	hex(1)
trpt:	transliterate protocol trace.	trpt(8N)
tcp: Internet	Transmission Control Protocol.	tcp(5P)
ftw: walk a file	tree.	ftw(3C)
twalk: manage binary search	trees. /tfind, tdelete,	tsearch(3C)
	trek: trekkie game.	trek(6)
trek:	trekkie game.	trek(6)
tan, asin, acos, atan, atan2:	trigonometric functions. /cos,	trig(3M)
server. iftpd: DARPA	Trivial File Transfer Protocol	iftpd(8N)
constant-width text for	troff. cw, checkcw: prepare	cw(1)
mathematical text for nroff or	troff. /neqn, checkeq: format	eqn(1)
typesetting view graphs/ mv: a	troff macro package for	mv(5)
format tables for nroff or	troff. tbl:	tbl(1)
	troff: typeset text.	troff(1)
trace.	trpt: transliterate protocol	trpt(8N)
values.	true, false: provide truth	true(1)
pdpl1, u3b, u3b5, vax: provide	truth value about your/ m68k,	machid(1)
true, false: provide	truth values.	true(1)
robots. chase:	Try to escape the killer	chase(6)
twalk: manage binary search/	tsearch, tfind, tdelete,	tsearch(3C)

Permuted Index


teletype bits to a sensible/	tset, reset: set or reset the . . . . .	tset(1)
	tsort: topological sort. . . . .	tsort(1)
	ttt, cubic: tic-tac-toe. . . . .	ttt(6)
interface.	tty: controlling terminal . . . . .	tty(7)
	tty: get the terminal's name. . . . .	tty(1)
graphics for the extended	TTY-37 type-box. greek: . . . . .	greek(5)
a terminal.	typename, isatty: find name of . . . . .	typename(3C)
utmp file of the current/	ttyslot: find the slot in the . . . . .	ttyslot(3C)
types by port.	ttysize: data base of terminal . . . . .	ttysize(4)
parameters. disktime:	tune floppy disk settling time . . . . .	disktime(1M)
/runacct, shutacct, startup,	turnacct: shell procedures for/ . . . . .	acctsh(1M)
tsearch, tfind, tdelete,	twalk: manage binary search/ . . . . .	tsearch(3C)
twinkle:	twinkle stars on the screen. . . . .	twinkle(6)
screen.	twinkle: twinkle stars on the . . . . .	twinkle(6)
ichar, char: explicit Fortran	type conversion. /dcmplx, . . . . .	ttype(3F)
file: determine file	type. . . . .	file(1)
value about your processor	type. /vax: provide truth . . . . .	machid(1)
getty: set terminal	type, modes, speed, and line/ . . . . .	getty(1M)
for the extended TTY-37	type-box. greek: graphics . . . . .	greek(5)
ttytype: data base of terminal	types by port. . . . .	ttysize(4)
types.	types: primitive system data . . . . .	types(5)
types: primitive system data	types. . . . .	types(5)
session. script: make	typescript of terminal . . . . .	script(1)
graphs, and slides. mmt, mvt:	typeset documents, view . . . . .	mmt(1)
troff:	typeset text. . . . .	troff(1)
mv: a troff macro package for	typesetting view graphs and/ . . . . .	mv(5)
/localtime, gmtime, asctime,	tzset: convert date and time/ . . . . .	ctime(3C)
value about your/ m68k, pdp11,	u3b, u3b5, vax: provide truth . . . . .	machid(1)
about your/ m68k, pdp11, u3b,	u3b5, vax: provide truth value . . . . .	machid(1)
Protocol.	udp: Internet User Datagram . . . . .	udp(5P)
getpw: get name from	UID. . . . .	getpw(3C)
limits.	ul: do underlining. . . . .	ul(1)
creation mask.	ulimit: get and set user . . . . .	ulimit(2)
mask.	umask: set and get file . . . . .	umask(2)
file system. mount,	umask: set file-creation mode . . . . .	umask(1)
UNIX system.	umount: mount and dismount . . . . .	mount(1M)
UNIX system.	umount: unmount a file system. . . . .	umount(2)
ul: do	uname: get name of current . . . . .	uname(2)
file. unget:	uname: print name of current . . . . .	uname(1)
an SCCS file.	underlining. . . . .	ul(1)
into input stream.	undo a previous get of an SCCS . . . . .	unget(1)
irand, srand, rand: Fortran	unget: undo a previous get of . . . . .	unget(1)
/seed48, lcong48: generate	ungetc: push character back . . . . .	ungetc(3S)
a file.	uniform random-number/ . . . . .	rand(3F)
mktemp: make a	uniformly distributed/ . . . . .	drand48(3C)
gethostid, sethostid: get/set	uniq: report repeated lines in . . . . .	uniq(1)
execution. uux:	unique filename. . . . .	mktemp(3C)
uuto, uupick: public	unique identifier of current/ . . . . .	gethostid(2N)
unlink system calls. link,	units: conversion program. . . . .	units(1)
entry.	UNIX-to-UNIX system command . . . . .	uux(1C)
unlink: exercise link and	UNIX-to-UNIX system file copy. . . . .	uuto(1C)
unlink system calls. link,	unlink: exercise link and . . . . .	link(1M)
umount:	unlink: remove directory . . . . .	unlink(2)
files. pack, pcat,	unlink system calls. link, . . . . .	link(1M)
times of a file. touch:	unmount a file system. . . . .	umount(2)
of programs. make: maintain,	unpack: compress and expand . . . . .	pack(1)
badblk: program to set or	update access and modification . . . . .	touch(1)
machines. updater:	update, and regenerate groups . . . . .	make(1)
machines. updater:	update bad block information. . . . .	badblk(1M)
lfind: linear search and	update files between two . . . . .	updater(1)
sync:	update files between two . . . . .	updater(1M)
sync:	update. lsearch, . . . . .	lsearch(3C)
	update super-block. . . . .	sync(2)
	update the super block. . . . .	sync(1)

two machines.	updater: update files between	updater(1)
two machines.	updater: update files between	updater(1M)
du: summarize disk	usage.	du(1)
id: print	user and group IDs and names.	id(1)
setuid, setgid: set	user and group IDs.	setuid(2)
crontab:	user crontab file.	crontab(1)
character login name of the	user. cuserid: get	cuserid(3S)
udp: Internet	User Datagram Protocol.	udp(5P)
/getgid, getegid: get real	user, effective user, real/	getuid(2)
environ:	user environment.	environ(5)
disk accounting data by	user ID. diskusg: generate	diskusg(1M)
print effective current	user id. whoami:	whoami(1)
set real and effective	user ID's. setreuid:	setreuid(2)
protocol. telnet:	user interface to the TELNET	telnet(1N)
ulimit: get and set	user limits.	ulimit(2)
logname: return login name of	user.	logname(3X)
/get real user, effective	user, real group, and/	getuid(2)
become super-user or another	user. su:	su(1)
talk: talk to another	user.	talk(1N)
the utmp file of the current	user. /find the slot in	ttyslot(3C)
write: write to another	user.	write(1)
last: indicate last logins of	users and teletypes.	last(1)
mail, rmail: send mail to	users or read mail.	mail(1)
wall: write to all	users.	wall(1M)
fuser: identify processes	using a file or file/	fuser(1M)
help: ask for help in	using SCCS.	help(1)
statistics.	ustat: get file system	ustat(2)
modification times.	utime: set file access and	utime(2)
utmp, wtmp:	utmp and wtmp entry formats.	utmp(4)
endulent, utmpname: access	utmp file entry. /setutent,	getut(3C)
ttyslot: find the slot in the	utmp file of the current user.	ttyslot(3C)
entry formats.	utmp, wtmp: utmp and wtmp	utmp(4)
/pututline, setutent, endulent,	utmpname: access utmp file/	getut(3C)
clean-up.	uuclean: uucp spool directory	uuclean(1M)
uusub: monitor	uucp network.	uusub(1M)
uuclean:	uucp spool directory clean-up.	uuclean(1M)
control. uustat:	uucp status inquiry and job	uustat(1C)
system to UNIX system copy.	uucp, uulog, uuname: UNIX	uucp(1C)
UNIX system copy. uucp,	uulog, uuname: UNIX system to	uucp(1C)
system copy. uucp, uulog,	uuname: UNIX system to UNIX	uucp(1C)
system file copy. uuto,	uupick: public UNIX-to-UNIX	uuto(1C)
and job control.	uustat: uucp status inquiry	uustat(1C)
UNIX-to-UNIX system file/	uusub: monitor uucp network.	uusub(1M)
command execution.	uuto, uupick: public	uuto(1C)
configuration information.	uux: UNIX-to-UNIX system	uux(1C)
	uvar: returns system-specific	uvar(2)
	val: validate SCCS file.	val(1)
	validate SCCS file.	val(1)
	value about your processor/	machid(1)
/u3b, u3b5, vax: provide truth	value.	abs(3C)
abs: return integer absolute	value. abs, labs, dabs,	abs(3F)
cabs, zabs: Fortran absolute	value for environment name.	getenv(3C)
getenv: return	value functions. /fabs: floor,	floor(3M)
ceiling, remainder, absolute	value to environment.	putenv(3C)
putenv: change or add	values between host and/	byteorder(3N)
/htons, ntohl, ntohs: convert	values: machine-dependent	values(5)
values.	values.	true(1)
true, false: provide truth	values.	values(5)
values: machine-dependent	varargs argument list.	vprintf(3S)
/print formatted output of a	varargs argument list.	vprintf(3X)
/print formatted output of a	varargs: handle variable	varargs(5)
argument list.	variable argument list.	varargs(5)
varargs: handle	variable. getenv:	getenv(3F)
return Fortran environment	vax: provide truth value about	machid(1)
your/ m68k, pdp11, u3b, u3b5,		

/files between M68000 and	VAX-11/780 processors.	fscv(1M)
	vc: version control.	vc(1)
	vchk: version checkup.	vchk(1M)
option letter from argument	vector. getopt: get	getopt(3C)
assert:	verify program assertion.	assert(3X)
vchk:	version checkup.	vchk(1M)
vc:	version control.	vc(1)
version: reports	version number of files.	version(1)
get: get a	version of an SCCS file.	get(1)
number of files.	version: reports version	version(1)
scsdiff: compare two	versions of an SCCS file.	scsdiff(1)
formatted output of/ vprintf,	vfprintf, vsprintf: print	vprintf(3S)
formatted output of/ vprintf,	vfprintf, vsprintf: print	vprintf(3X)
display editor based on ex.	vi: screen-oriented (visual)	vi(1)
mmtl, mvt: typeset documents,	view graphs, and slides.	mmt(1)
macro package for typesetting	view graphs and slides. /troff	mv(5)
file perusal filter for crt	viewing. more:	more(1)
on ex. vi: screen-oriented	(visual) display editor based	vi(1)
systems with label checking.	volcopy, labelit: copy file	volcopy(1M)
file system: format of system	volume.	fs(4)
print formatted output of a/	vprintf, vfprintf, vsprintf:	vprintf(3S)
print formatted output of a/	vprintf, vfprintf, vsprintf:	vprintf(3X)
output of/ vprintf, vfprintf,	vsprintf: print formatted	vprintf(3S)
output of/ vprintf, vfprintf,	vsprintf: print formatted	vprintf(3X)
or terminate. wait:	wait for child process to stop	wait(2)
or terminate. wait3:	wait for child process to stop	wait3(2N)
to stop or terminate.	wait: wait for child process	wait(2)
to stop or terminate.	wait3: wait for child process	wait3(2N)
ftw:	walk a file tree.	ftw(3C)
	wall: write to all users.	wall(1M)
	wc: word count.	wc(1)
	what: identify SCCS files.	what(1)
signal. signal: specify	what to do upon receipt of a	signal(2)
crashes. crash:	what to do when the system	crash(8)
binary, and/or manual for/	whereis: locate source,	whereis(1)
who do:	who is doing what.	who(1M)
who:	who is on the system.	who(1)
	who: who is on the system.	who(1)
current user id.	whoami: print effective	whoami(1)
	whodo: who is doing what.	whodo(1M)
machines. rwho:	who's logged in on local	rwho(1N)
cd: change	working directory.	cd(1)
chdir: change	working directory.	chdir(2)
get pathname of current	working directory. getcwd:	getcwd(3C)
pwd:	working directory name.	pwd(1)
worm: Play the growing	worm game.	worm(6)
game.	worm: Play the growing worm	worm(6)
display terminal.	worms: animate worms on a	worms(6)
worms: animate	worms on a display terminal.	worms(6)
write:	write on a file.	write(3)
writev:	write on a file.	writev(3N)
putpwent:	write password file entry.	putpwent(3C)
wall:	write to all users.	wall(1M)
write:	write to another user.	write(1)
	write: write on a file.	write(3)
	write: write to another user.	write(1)
	writev: write on a file.	writev(3N)
file regions for reading or	writing. /provide exclusive	locking(2)
open: open for reading or	writing.	open(2)
utmp, wtmp: utmp and	wtmp entry formats.	utmp(4)
formats. utmp,	wtmp: utmp and wtmp entry	utmp(4)
accounting records. fwtmp,	wtmpfix: manipulate connect	fwtmp(1M)
hunt-the-wumpus.	wump: the game of	wump(6)
list(s) and execute command.	xargs: construct argument	xargs(1)



Fortran bitwise/ and, or,	xor, not, lshift, rshift: . . . . .	bool(3F)
programs to implement shared/	xstr: extract strings from C . . . . .	xstr(1)
j0, j1, jn,	y0, y1, yn: Bessel functions. . . . .	bessel(3M)
j0, j1, jn, y0,	y1, yn: Bessel functions. . . . .	bessel(3M)
compiler-compiler.	yacc: yet another . . . . .	yacc(1)
j0, j1, jn, y0, y1,	yn: Bessel functions. . . . .	bessel(3M)
abs, iabs, dabs, cabs,	zabs: Fortran absolute value. . . . .	abs(3F)



**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 returned value. This is almost always -1; the individual descriptions specify the details. An error number is also made available in the external variable *errno*. *Errno* is not cleared on successful calls, so it should be tested only after an error has been indicated.

There is a table of messages associated with each error, and a routine for printing the message; see *error* (3C). Each system call description attempts to list all possible error numbers. The following is a complete list of the error numbers and their names as defined in <errno.h>.

**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**

No process can be found corresponding to that specified by *pid* in *kill* or *ptrace*.

**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 has occurred. 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, a tape drive is not on-line or no disk pack is loaded on a drive.

7 E2BIG Arg list too long

An argument list longer than 5,120 bytes is presented to a member of the *exec* family.

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* (4)).

9 EBADF Bad file number

Either a file descriptor refers to no open file, or a read (respectively, write) request is made to a file which is open only for writing (respectively, reading).

10 ECHILD No child processes

A *wait* was executed by a process that had no existing or unwaited-for child processes.

11 EAGAIN No more processes

The system is out of a resource that may be available later. A *fork* failed because the system's process table is full or the user is not allowed to create any more processes. A system call which requires memory may also fail with this error if the system is out of memory or swap space but the request is less than the system-imposed per process limit.

12 ENOMEM Not enough space

During an *exec*, *brk*, or *sbrk*, a program asks for more space than the system is able to supply. This is not a temporary condition; the maximum space size is a system parameter. The error may also occur if the arrangement of text, data, and stack segments requires too many segmentation registers, or if there is not enough swap space during a *fork*.

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 use an argument of a system call.

15 ENOTBLK Block device required

A non-block file was mentioned where a block device was required, e.g.,

in *mount*.

- 16 **EBUSY** Device or resource busy  
An attempt was made to mount a device that was already mounted or an attempt was made to dismount a device on which there is an active file (open file, current directory, mounted-on file, active text segment). It will also occur if an attempt is made to enable accounting when it is already enabled. The device or resource is currently unavailable.
- 17 **EEXIST** File exists  
An existing file was mentioned in an inappropriate context, e.g., *link*.
- 18 **EXDEV** Cross-device link  
A 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 prefix or as an argument to *chdir* (2).
- 21 **EISDIR** Is a directory  
An attempt was made to write on a directory.
- 22 **EINVAL** Invalid argument  
Some invalid argument (e.g., dismounting a non-mounted device; mentioning an undefined signal in *signal*, or *kill*; reading or writing a file for which *lseek* has generated a negative pointer). Also set by the math functions described in the (3M) entries of this manual.
- 23 **ENFILE** File table overflow  
The system file table is full, and temporarily no more *opens* can be accepted.
- 24 **EMFILE** Too many open files  
No process may have more than 20 file descriptors open at a time. When a record lock is being created with *fcntl*, there are too many files with record locks on them.
- 26 **ETXTBSY** Text file busy  
An attempt was made to execute a pure-procedure program which is currently open for writing. 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 file size (1,082,201,088 bytes) or **ULIMIT**; see *ulimit*(2).
- 28 **ENOSPC** No space left on device  
During a *write* to an ordinary file, there is no free space left on the device. In *fcntl*, the setting or removing of record locks on a file cannot be accomplished because there are no more record entries left on the system
- 29 **ESPIPE** Illegal seek  
An *lseek* was issued to a pipe. This error should 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 the maximum number of links (1000) to a file.
- 32 **EPIPE** Broken pipe  
A write on a pipe 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 not representable within machine precision.
- 35 **ENOMSG** No message of desired type  
An attempt was made to receive a message of a type that does not exist on the specified message queue; see *msgop*(2).
- 36 **EIDRM** Identifier Removed  
This error is returned to processes that resume execution due to the removal of an identifier from the file system's name space (see *msgctl*(2), *semctl*(2), and *shmctl*(2)).
- 45 **EDEADLK** Deadlock  
A deadlock situation was detected and avoided.

- 55 **EWouldBlock** Operation would block  
An operation which would cause a process to block was attempted on an object in non-blocking mode (see *socket(2N)*).
- 56 **EINPROGRESS** Operation now in progress  
An operation which takes a long time to complete (such as a *connect(2N)*) was started on a non-blocking object (see *socket(2N)*).
- 57 **EALREADY** Operation already in progress  
An operation was attempted on a non-blocking object which already had an operation in progress.
- 58 **ENOTSOCK** Socket operation on non-socket  
Self-explanatory.
- 59 **EDESTADDRREQ** Destination address required  
A required address was omitted from an operation on a socket.
- 60 **EMSGSIZE** Message too long  
A message sent on a socket was larger than the internal message buffer.
- 61 **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 internet UDP protocol with type `SOCK_STREAM`.
- 62 **ENOPROTOOPT** Protocol not available  
In this incarnation of the system.
- 63 **EPROTONOSUPPORT** Protocol not supported  
In this incarnation of the system.
- 64 **ESOCKTNOSUPPORT** Socket type not supported  
In this incarnation of the system.
- 65 **EOPNOTSUPP** Operation not supported on socket  
For example, trying to *accept* a connection on a datagram socket.
- 66 **EPFNOSUPPORT** Protocol family not supported  
In this incarnation of the system.
- 67 **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.

- 68 EADDRINUSE Address already in use  
Only one usage of each address is normally permitted.
- 69 EADDRNOTAVAIL Can't assign requested address  
Normally results from an attempt to create a socket with an address not on this machine.
- 70 ENETDOWN Network is down  
A socket operation encountered a dead network.
- 71 ENETUNREACH Network is unreachable  
A socket operation was attempted to an unreachable network.
- 72 ENETRESET Network dropped connection on reset  
The host you were connected to crashed and rebooted.
- 73 ECONNABORTED Software caused connection abort  
A connection abort was caused internal to your host machine.
- 74 ECONNRESET Connection reset by peer
- 55 ENOBUFS No buffer space available  
For a socket or a pipe in the buffer pool.
- 76 EISCONN Socket is already connected
- 77 ENOTCONN Socket is not connected
- 78 ESHUTDOWN Can't send after socket shutdown
- 79 *unused*
- 80 ETIMEDOUT Connection timed out  
Due to failure to initiate properly or because keep-alives failed.
- 81 ECONNREFUSED Connection refused  
No connection could be made because the target machine actively refused it.
- 83 ENAMETOOLONG File name too long  
A component of a path name exceeded 14 characters, or an entire path name exceeded 1023 characters.
- 84 EHOSTDOWN Host is down  
A socket operation encountered a defunct host.
- 85 EHOSTUNREACH No route to host  
A socket operation was attempted to an unreachable host.



**100 EDEADLOCK Locking Deadlock**

Returned by *locking(2)* system call if deadlock would occur or when locktable overflows.

**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 1 to 30,000.

**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 ID is the process ID of the group leader. This grouping permits the signaling of related processes; see *kill(2)*.

**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 terminate a group of related processes upon termination of one of the processes in the group; see *exit(2)* and *signal(2)*.

**Real User ID and Real Group ID**

Each user allowed on the system is identified by a positive integer called a real user ID.

Each user is also a member of a group. The group is identified by a positive integer called the real group ID.

An active process has a real user ID and real group ID that are set to the real user ID and real group ID, respectively, of the user responsible for the creation of the process.

**Effective User ID and Effective Group ID**

An active process has an effective user ID and an effective group ID that are used to determine file access permissions (see below). The effective user ID and effective group ID are equal to the process's real user ID and real group ID respectively, unless the process or one of its ancestors evolved from a file that had the set-user-ID bit or set-group ID bit set; see *exec(2)*.

**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 of 0 and a process ID of 1 are special processes and are referred to as *proc0* and *proc1*.

*Proc0* is the scheduler. *Proc1* is the initialization process (*init*). *Proc1* is the ancestor of every other process in the system and is used to control the process structure.

### File Descriptor

A file descriptor is a small integer used to do I/O on a file. The value of a file descriptor is from 0 to 19. A process may have no more than 20 file descriptors (0-19) open simultaneously. A file descriptor is returned by system calls such as *open(2)*, or *pipe(2)*. The file descriptor is used as an argument by calls such as *read(2)*, *write(3)*, *ioctl(2)*, and *close(2)*.

### File Name.

Names consisting of 1 to 14 characters may be used to name an ordinary file, special file or directory.

These characters may be selected from the set of all character values excluding  $\backslash$  (null) and the ASCII code for / (slash).

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. See *sh(1)*. Although permitted, it is advisable to avoid the use of unprintable characters in file names.

### Path Name and Path Prefix

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.

More precisely, a path name is a null-terminated character string constructed as follows:

```
<path-name> ::= <file-name> | <path-prefix> <file-name> | /
<path-prefix> ::= <rtprefix> | / <rtprefix>
<rtprefix> ::= <dirname> / | <rtprefix> <dirname> /
```

where <file-name> is a string of 1 to 14 characters other than the ASCII slash and null, and <dirname> is a string of 1 to 14 characters (other than the ASCII slash and null) that names a directory.

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.

Unless specifically stated otherwise, the null path name is treated as if it named a non-existent file.

### Directory

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.

### 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. The root directory of a process need not be the root directory of the root file system.

### File Access Permissions

Read, write, and execute/search permissions on a file are granted to a process if one or more of the following is true:

- The effective user ID of the process is super-user.

- The effective user ID of the process matches the user ID of the owner of the file and the appropriate access bit of the "owner" portion (0700) of the file mode is set.

- The effective user ID of the process does not match the user ID of the owner of the file, and the effective group ID of the process matches the group of the file and the appropriate access bit of the "group" portion (070) of the file mode is set.

- The effective user ID of the process does not match the user ID of the owner of the file, and the effective group ID of the process does not match the group ID of the file, and the appropriate access bit of the "other" portion (07) of the file mode is set.

Otherwise, the corresponding permissions are denied.

### Message Queue Identifier

A message queue identifier (*msqid*) is a unique positive integer created by a *msgget(2)* system call. Each *msqid* has a message queue and a data structure associated with it. The data structure is referred to as *msqid\_ds* and contains the following members:

```

struct ipc_perm msg_perm; /* operation permission struct */
ushort msg_qnum;          /* number of msgs on q */
ushort msg_qbytes;       /* max number of bytes on q */
ushort msg_lspid;        /* pid of last msgsnd operation */
ushort msg_lrpid;        /* pid of last msgrcv operation */
time_t msg_stime;        /* last msgsnd time */
time_t msg_rtime;        /* last msgrcv time */
time_t msg_ctime;        /* last change time */
                          /* Times measured in secs since */
                          /* 00:00:00 GMT, Jan. 1, 1970 */

```

**Msg\_perm** is an **ipc\_perm** structure that specifies the message operation permission (see below). This structure includes the following members:

```

ushort cuid;             /* creator user id */
ushort cgid;             /* creator group id */
ushort uid;              /* user id */
ushort gid;              /* group id */
ushort mode;             /* r/w permission */

```

**Msg\_qnum** is the number of messages currently on the queue. **Msg\_qbytes** is the maximum number of bytes allowed on the queue. **Msg\_lspid** is the process id of the last process that performed a *msgsnd* operation. **Msg\_lrpid** is the process id of the last process that performed a *msgrcv* operation. **Msg\_stime** is the time of the last *msgsnd* operation, **msg\_rtime** is the time of the last *msgrcv* operation, and **msg\_ctime** is the time of the last *msgctl(2)* operation that changed a member of the above structure.

#### Message Operation Permissions

In the *msgop(2)* and *msgctl(2)* system call descriptions, the permission required for an operation is given as "{token}", where "token" is the type of permission needed interpreted as follows:

```

00400  Read by user
00200  Write by user
00060  Read, Write by group
00006  Read, Write by others

```

Read and Write permissions on a *msgid* are granted to a process if one or more of the following is true:

The effective user ID of the process is super-user.

The effective user ID of the process matches `msg_perm.[c]uid` in the data structure associated with `msgid` and the appropriate bit of the "user" portion (0600) of `msg_perm.mode` is set.

The effective user ID of the process does not match `msg_perm.[c]uid` and the process's effective group ID matches `msg_perm.[c]gid` and the appropriate bit of the "group" portion (060) of `msg_perm.mode` is set.

The effective user ID of the process does not match `msg_perm.[c]uid` and the effective group ID of the process does not match `msg_perm.[c]gid` and the appropriate bit of the "other" portion (06) of `msg_perm.mode` is set.

Otherwise, the corresponding permissions are denied.

#### Semaphore Identifier

A semaphore identifier (`semid`) is a unique positive integer created by a `semget(2)` system call. Each `semid` has a set of semaphores and a data structure associated with it. The data structure is referred to as `semid_ds` and contains the following members:

```
struct ipc_perm sem_perm; /* operation permission struct */
ushort sem_nsems;        /* number of sems in set */
time_t sem_otime;       /* last operation time */
time_t sem_ctime;       /* last change time */
                        /* Times measured in secs since */
                        /* 00:00:00 GMT, Jan. 1, 1970 */
```

`Sem_perm` is an `ipc_perm` structure that specifies the semaphore operation permission (see below). This structure includes the following members:

```
ushort cuid;           /* creator user id */
ushort cgid;           /* creator group id */
ushort uid;            /* user id */
ushort gid;            /* group id */
ushort mode;           /* r/a permission */
```

The value of `sem_nsems` is equal to the number of semaphores in the set. Each semaphore in the set is referenced by a positive integer referred to as a `sem_num`. `Sem_num` values run sequentially from 0 to the value of `sem_nsems` minus 1. `Sem_otime` is the time of the last `semop(2)` operation, and `sem_ctime` is the time of the last `semctl(2)` operation that changed a member of the above structure.

A semaphore is a data structure that contains the following members:

```

ushort semval; /* semaphore value */
short  sempid; /* pid of last operation */
ushort semncnt; /* # awaiting semval > cval */
ushort semzcnt; /* # awaiting semval = 0 */

```

**Semval** is a non-negative integer. **Sempid** is equal to the process ID of the last process that performed a semaphore operation on this semaphore. **Semncnt** is a count of the number of processes that are currently suspended awaiting this semaphore's **semval** to become greater than its current value. **Semzcnt** is a count of the number of processes that are currently suspended awaiting this semaphore's **semval** to become zero.

#### Semaphore Operation Permissions

In the *semop*(2) and *semctl*(2) system call descriptions, the permission required for an operation is given as "[token]", where "token" is the type of permission needed interpreted as follows:

```

00400  Read by user
00200  Alter by user
00060  Read, Alter by group
00006  Read, Alter by others

```

Read and Alter permissions on a **semid** are granted to a process if one or more of the following is true:

The effective user ID of the process is super-user.

The effective user ID of the process matches **sem\_perm.[c]uid** in the data structure associated with **semid** and the appropriate bit of the "user" portion (0600) of **sem\_perm.mode** is set.

The effective user ID of the process does not match **sem\_perm.[c]uid** and the effective group ID of the process matches **sem\_perm.[c]gid** and the appropriate bit of the "group" portion (060) of **sem\_perm.mode** is set.

The effective user ID of the process does not match **sem\_perm.[c]uid** and the effective group ID of the process does not match **sem\_perm.[c]gid** and the appropriate bit of the "other" portion (06) of **sem\_perm.mode** is set.

Otherwise, the corresponding permissions are denied.

#### Shared Memory Identifier

A shared memory identifier (**shmid**) is a unique positive integer created by a *shmget*(2) system call. Each **shmid** has a segment of memory (referred to as a

shared memory segment) and a data structure associated with it. The data structure is referred to as *shmid\_ds* and contains the following members:

```

struct ipc_perm shm_perm; /* operation permission struct */
int   shm_segsz;         /* size of segment */
ushort shm_cpid;         /* creator pid */
ushort shm_lpid;         /* pid of last operation */
short  shm_nattch;       /* number of current attaches */
time_t shm_atime;        /* last attach time */
time_t shm_dtime;        /* last detach time */
time_t shm_ctime;        /* last change time */
                                /* Times measured in secs since */
                                /* 00:00:00 GMT, Jan. 1, 1970 */

```

*Shm\_perm* is an *ipc\_perm* structure that specifies the shared memory operation permission (see below). This structure includes the following members:

```

ushort cuid;             /* creator user id */
ushort cgid;             /* creator group id */
ushort uid;              /* user id */
ushort gid;              /* group id */
ushort mode;             /* r/w permission */

```

*Shm\_segsz* specifies the size of the shared memory segment. *Shm\_cpid* is the process id of the process that created the shared memory identifier. *Shm\_lpid* is the process id of the last process that performed a *shmop*(2) operation. *Shm\_nattch* is the number of processes that currently have this segment attached. *Shm\_atime* is the time of the last *shmat* operation, *shm\_dtime* is the time of the last *shmdt* operation, and *shm\_ctime* is the time of the last *shmctl*(2) operation that changed one of the members of the above structure.

#### Shared Memory Operation Permissions

In the *shmop*(2) and *shmctl*(2) system call descriptions, the permission required for an operation is given as "{token}", where "token" is the type of permission needed interpreted as follows:

```

00400  Read by user
00200  Write by user
00060  Read, Write by group
00006  Read, Write by others

```

Read and Write permissions on a *shmid* are granted to a process if one or more of the following is true:

The effective user ID of the process is super-user.

The effective user ID of the process matches `shm_perm.[c]uid` in the data structure associated with `shmid` and the appropriate bit of the "user" portion (0600) of `shm_perm.mode` is set.

The effective user ID of the process does not match `shm_perm.[c]uid` and the effective group ID of the process matches `shm_perm.[c]gid` and the appropriate bit of the "group" portion (060) of `shm_perm.mode` is set.

The effective user ID of the process does not match `shm_perm.[c]uid` and the effective group ID of the process does not match `shm_perm.[c]gid` and the appropriate bit of the "other" portion (06) of `shm_perm.mode` is set.

Otherwise, the corresponding permissions are denied.

**SEE ALSO**

`close(2)`, `ioctl(2)`, `open(2)`, `pipe(2)`, `read(2)`, `write(3)`, `intro(3)`.



**\_EXIT (2)**

**SEE EXIT**

**\_EXIT (2)**

**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;
cc ... -lnet
```

**DESCRIPTION**

The argument *s* is a socket which has been created with `socket(2N)`, bound to an address with `bind(2N)`, and is listening for connections after a `listen(2N)`. `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(2N)` 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 <code>SOCK_STREAM</code> .                  |
| [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. |

**LINKING**

This library is accessed by specifying `-lnet` as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**SEE ALSO**

`bind(2N)`, `connect(2N)`, `listen(2N)`, `select(2N)`, `socket(2N)`

**NAME**

access — determine accessibility of a file

**SYNOPSIS**

```
int access (path, amode)
char *path;
int amode;
```

**DESCRIPTION**

*Path* points to a path name naming a file. *Access* checks the named file for accessibility according to the bit pattern contained in *amode*, using the real user ID in place of the effective user ID and the real group ID in place of the effective group ID. The bit pattern contained in *amode* is constructed as follows:

```
04    read
02    write
01    execute (search)
00    check existence of file
```

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]	Read, write, or execute (search) permission is requested for a null path name.
[ENOENT]	The named file does not exist.
[EACCES]	Search permission is denied on a component of the path prefix.
[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.
[EACCESS]	Permission bits of the file mode do not permit the requested access.
[EFAULT]	<i>Path</i> points outside the allocated address space for the process.

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 permissions checked with respect to the "group" mode bits, and all others have permissions checked with respect to the "other" mode bits.

The super-user is always granted execute permission even though 1. execute permission is meaningful only for directories and regular files, and 2. *exec* requires that at least one execute mode bit be set for regular file to be executable.

Notice that it is only access bits that are checked. A directory may be announced as writable by *access*, but an attempt to open it for writing will fail because it is not allowed to write into the directory structure itself, although files may be created there. A file may look executable, but *exec* will fail unless it is in proper format.

**RETURN VALUE**

If the requested access is permitted, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

**ACCESS (2)**

**ACCESS (2)**

**SEE ALSO**

**chmod(2), stat(2).**

**NAME**

acct — enable or disable process accounting

**SYNOPSIS**

```
int acct (path)
char *path;
```

**DESCRIPTION**

*Acct* is used to enable or disable the system process accounting routine. If the routine is enabled, an accounting record will be written on an accounting file for each process that terminates. Termination can be caused by one of two things: an *exit* call or a signal; see *exit(2)* and *signal(2)*. The effective user ID of the calling process must be super-user to use this call.

*Path* points to a path name naming the accounting file. The accounting file format is given in *acct(4)*.

The accounting routine is enabled if *path* is non-zero and no errors occur during the system call. It is disabled if *path* is zero and no errors occur during the system call.

*Acct* will fail if one or more of the following are true:

[EPERM]	The effective user of the calling process is not super-user.
[EBUSY]	An attempt is being made to enable accounting when it is already enabled.
[ENOTDIR]	A component of the path prefix is not a directory.
[ENOENT]	One or more components of the accounting file path name do not exist.
[EACCES]	A component of the path prefix denies search permission.
[EACCES]	The file named by <i>path</i> is not an ordinary file.
[EACCES]	<i>Mode</i> permission is denied for the named accounting file.
[EISDIR]	The named file is a directory.
[EROFS]	The named file resides on a read-only file system.
[EFAULT]	<i>Path</i> 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**

*exit(2)*, *signal(2)*, *acct(4)*.

**NAME**

alarm — set a process's alarm clock

**SYNOPSIS**

**unsigned alarm (sec)**  
**unsigned sec;**

**DESCRIPTION**

*Alarm* instructs the calling process's alarm clock to send the signal SIGALRM to the calling process after the number of real time seconds specified by *sec* have elapsed; see *signal(2)*.

Alarm requests are not stacked; successive calls reset the calling process's alarm clock. If the argument is 0, any alarm request is canceled. Because the clock has a 1-second resolution, the signal may occur up to one second early; because of scheduling delays, resumption of execution of when the signal is caught may be delayed an arbitrary amount. The longest specifiable delay time is 4,294,967,295 ( $2^{32}-1$ ) seconds, or 136 years.

**RETURN VALUE**

*Alarm* returns the amount of time previously remaining in the calling process's alarm clock.

**SEE ALSO**

pause(2), signal(2).

**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;

cc ... -lnet
```

**DESCRIPTION**

*Bind* assigns a name to an unnamed socket. When a socket is created with *socket(2N)* it exists in a name space (address family) but has no name assigned. *Bind* requests the *name*, be assigned to the socket.

**NOTES**

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.

**LINKING**

This library is accessed by specifying *-lnet* as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**SEE ALSO**

connect(2N), listen(2N), socket(2N), getsockname(2N)



## NAME

*brk*, *sbrk* – change data segment space allocation

## SYNOPSIS

```
int brk (endds)
char *endds;

char *sbrk (incr)
int incr;
```

## DESCRIPTION

*Brk* and *sbrk* are used to change dynamically the amount of space allocated for the calling process's data segment; see *exec(2)*. The change is made by resetting the process's break value and allocating the appropriate amount of space. The break value is the address of the first location beyond the end of the data segment. The amount of allocated space increases as the break value increases. The newly allocated space is set to zero.

*Brk* sets the break value to *endds* and changes the allocated space accordingly.

*Sbrk* adds *incr* bytes to the break value and changes the allocated space accordingly. *Incr* can be negative, in which case the amount of allocated space is decreased.

## ERRORS

*Brk* and *sbrk* will fail without making any change in the allocated space if one or more of the following are true:

Such a change would result in more space being allocated than is allowed by a system-imposed maximum (see *ulimit(2)*). Two types of this condition with associated error messages may be encountered:

## [ENOMEM]

Not enough space. Program asks for more space than the system is able to supply.

## [EAGAIN]

The system has temporarily exhausted its available memory or swap space.

Such a change would result in the break value being greater than or equal to the start address of any attached shared memory segment (see *shmop(2)*).

## RETURN VALUE

Upon successful completion, *brk* returns a value of 0 and *sbrk* returns the old break value. Otherwise, a value of -1 is returned and *errno* is set to indicate

**BRK(2)**

**BRK(2)**

the error.

**SEE ALSO**

**exec(2), shmop(2), ulimit(2).**

**NAME**

chdir - change working directory

**SYNOPSIS**

```
int chdir (path)
char *path;
```

**DESCRIPTION**

*Path* points to the path name of a directory. *Chdir* causes the named directory to become the current working directory, the starting point for path searches for path names not beginning with */*.

*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 path name is not a directory.
- [ENOENT]      The named directory does not exist.
- [EACCES]      Search permission is denied for any component of the path name.
- [EFAULT]      *Path* points outside the allocated address space of the process.

**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**

chroot(2).

## NAME

chmod — change mode of file

## SYNOPSIS

```
int chmod (path, mode)
char *path;
int mode;
```

## DESCRIPTION

*Path* points to a path name naming a file. *Chmod* sets the access permission portion of the named file's mode according to the bit pattern contained in *mode*.

Access permission bits are interpreted as follows:

```
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 if a directory) by owner.
00070  Read, write, execute (search) by group.
00007  Read, write, execute (search) by others.
```

The effective user ID of the process must match the owner of the file or be super-user to change the mode of a file.

If the effective user ID of the process is not super-user, mode bit 01000 (save text image on execution) is cleared.

If the effective user ID of the process is not super-user and the effective group ID of the process does not match the group ID of the file, mode bit 02000 (set group ID on execution) is cleared.

If an executable file is prepared for sharing (see the *cc -n* option), then mode bit 01000 prevents the system from abandoning the swap-space image of the program-text portion of the file when its last user terminates. Thus, when the next user of the file executes it, the text need not be read from the file system but can simply be swapped in, saving time.

Changing the owner of a file turns off the set-user-id bit, unless the superuser does it. This makes the system somewhat more secure at the expense of a degree of compatibility. *Chmod* will fail and the file mode will be unchanged if one or more of the following are true:

```
[ENOTDIR]  A component of the path prefix is not a directory.
[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 super-user.
[EROFS]   The named file resides on a read-only file system.
[EFAULT]  Path points outside the allocated address space of the pro-
           cess.
```

## RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value

**CHMOD(2)**

**CHMOD(2)**

of `-1` is returned and *errno* is set to indicate the error.

**SEE ALSO**

`chown(2)`, `mknod(2)`.

**NAME**

`chown` — change owner and group of a file

**SYNOPSIS**

```
int chown (path, owner, group)
char *path;
int owner, group;
```

**DESCRIPTION**

*Path* points to a path name naming a file. The owner ID and group ID of the named file are set to the numeric values contained in *owner* and *group* respectively.

Only processes with effective user ID equal to the file owner or super-user may change the ownership of a file.

If *chown* is invoked by other than the super-user, the set-user-ID and set-group-ID bits of the file mode, 04000 and 02000 respectively, will be cleared.

*Chown* will fail and the owner and group of the named file will remain unchanged if one or more of the following are true:

- [ENOTDIR] A component of the path prefix is not a directory.
- [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 super-user.
- [EROFS] The named file resides on a read-only file system.
- [EFAULT] *Path* points outside the allocated address space of the process.

**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**

`chown(1)`, `chmod(2)`.

**NAME**

chroot — change root directory

**SYNOPSIS**

```
int chroot (path)
char *path;
```

**DESCRIPTION**

*Path* points to a path name naming a directory. *Chroot* causes the named directory to become the root directory, the starting point for path searches for path names beginning with /. The user's working directory is unaffected by the *chroot* system call.

The effective user ID of the process must be super-user to change the root directory.

The .. entry in the root directory is interpreted to mean the root directory itself. Thus, .. cannot be used to access files outside the subtree rooted at the root directory.

*Chroot* will fail and the root directory will remain unchanged if one or more of the following are true:

- [ENOTDIR] Any component of the path name is not a directory.
- [ENOENT] The named directory does not exist.
- [EPERM] The effective user ID is not super-user.
- [EFAULT] *Path* points outside the allocated address space of the process.

**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**

chdir(2).

**NAME**

close — close a file descriptor

**SYNOPSIS**

```
int close (fildes)
int fildes;
```

**DESCRIPTION**

*Fildes* is a file descriptor obtained from a *creat*, *open*, *dup*, *fcntl*, *pipe*, or *socket* system call. *Close* closes the file descriptor indicated by *fildes*. A close of all files is automatic on *exit*, but since there is a 20 open file limit on the number of open files per process, *close* is necessary for programs which deal with many files. All outstanding record locks owned by the process (on the file indicated by *fildes*) are removed.

[EBADF] *Close* will fail if *fildes* is not a valid open file descriptor.

**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**

*creat*(2), *dup*(3), *exec*(2), *fcntl*(2), *open*(2), *pipe*(2), *socket*(2N).



**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;
cc ... -lnet
```

**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] *S* is not a valid descriptor.
- [ENOTSOCK] *S* 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 *name* parameter specifies an area outside the process address space.
- [EWOULDBLOCK] The socket is non-blocking and the connection cannot be completed immediately. It is possible to *select(2N)* the socket while it is connecting by selecting it for writing.

**LINKING**

This library is accessed by specifying `-lnet` as the last argument to the compile line, e.g.:

`cc -o prog prog.c -lnet`

**SEE ALSO**

`accept(2N)`, `select(2N)`, `socket(2N)`, `getsockname(2N)`

## NAME

*creat* — create a new file or rewrite an existing one

## SYNOPSIS

```
int creat (path, mode)
char *path;
int mode;
```

## DESCRIPTION

*Creat* creates a new ordinary file or prepares to rewrite an existing file named by the path name pointed to by *path*.

If the file exists, the length is truncated to 0 and the mode and owner are unchanged. Otherwise, the file's owner ID is set to the effective user ID, of the process the group ID of the process is set to the effective group ID, of the process and the low-order 12 bits of the file mode are set to the value of *mode* modified as follows:

All bits set in the process's file mode creation mask are cleared.  
See *umask(2)*.

The "save text image after execution bit" of the mode is cleared.  
See *chmod(2)*.

Upon successful completion, the file descriptor is returned and the file is open for writing, even if the mode does not permit writing. The file pointer is set to the beginning of the file. The file descriptor is set to remain open across *exec* system calls. See *fcntl(2)*. No process may have more than 20 files open simultaneously.

The *mode* given is arbitrary; it need not allow writing. This feature is used by programs which deal with temporary files of fixed names. The creation is done with a mode that forbids writing. Then, if a second instance of the program attempts a *creat*, an error is returned and the program knows that the name is unusable for the moment.

The system-scheduling algorithm does not make this a true uninterruptible operation, and a race condition may develop if *creat* is done at precisely the same time by two different processes.

*Creat* will fail if one or more of the following are true:

- |           |   |
|-----------|---|
| [ENOTDIR] | A component of the path prefix is not a directory.  |
| [ENOENT]  | A component of the path prefix does not exist.  |
| [EACCES]  | Search permission is denied on a component of the path prefix.  |
| [ENOENT]  | The path name is null.  |
| [EACCES]  | The file does not exist and the directory in which the file is to be created does not permit writing. |
| [EROFS]   | The named file resides or would reside on a read-only file system.                                    |
| [ETXTBSY] | The file is a pure procedure (shared text) file that is being executed.                               |

- [EACCES]       The file exists and write permission is denied.
- [EISDIR]       The named file is an existing directory.
- [EMFILE]       Twenty (20) file descriptors are currently open.
- [EFAULT]       *Path* points outside the allocated address space of the process.
- [ENFILE]       The system file table is full.

**RETURN VALUE**

Upon successful completion, a non-negative integer, namely the file descriptor, is returned. Otherwise, a value of  $-1$  is returned and *errno* is set to indicate the error.

**SEE ALSO**

`chmod(2)`, `close(2)`, `dup(3)`, `fcntl(2)`, `lseek(2)`, `open(2)`, `read(2)`, `umask(2)`, `write(3)`.

**NAME**

*dup* - duplicate a descriptor

**SYNOPSIS**

```
news = dup(oldd)
```

```
int newd, oldd;
```

**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(3N)*. 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.

**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* fails if:

[EBADF]

*Oldd* or *newd* is not a valid active descriptor

[EMFILE]

Too many descriptors are active.

**SEE ALSO**

*accept(2N)*, *open(2)*, *close(2)*, *pipe(2)*, *socket(2N)*, *getdtablesize(3N)*.

**NAME**

dup2 – duplicate a descriptor

**SYNOPSIS**

```
dup2(oldd, newd)
int oldd, newd;
```

**DESCRIPTION**

*Dup2* causes *newd* to become a duplicate of *oldd*. If *newd* is already in use, the descriptor is first deallocated as if a *close(2)* call had been done first.

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.

**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**

*Dup2* fails if:

[EBADF]

*Oldd* or *newd* is not a valid active descriptor

[EMFILE]

Too many descriptors are active.

**SEE ALSO**

*accept(2N)*, *open(2)*, *close(2)*, *pipe(2)*, *socket(2N)*, *getdtablesize(3N)*.

## NAME

`execl`, `execv`, `execle`, `execve`, `execlp`, `execvp` – execute a file

## SYNOPSIS

```
int execl (path, arg0, arg1, ..., argn, 0)
char *path, *arg0, *arg1, ..., *argn;

int execv (path, argv)
char *path, *argv[ ];

int execle (path, arg0, arg1, ..., argn, 0, envp)
char *path, *arg0, *arg1, ..., *argn, *envp[ ];

int execve (path, argv, envp)
char *path, *argv[ ], *envp[ ];

int execlp (file, arg0, arg1, ..., argn, 0)
char *file, *arg0, *arg1, ..., *argn;

int execvp (file, argv)
char *file, *argv[ ];
```

## DESCRIPTION

*Exec* in all its forms transforms the calling process into a new process. The new process is constructed from an ordinary, executable file called the *new process file*. This file consists of a header (see *a.out* (4)), a text segment, and a data segment. The data segment contains an initialized portion and an uninitialized portion (bss). There can be no return from a successful *exec* because the calling process is overlaid by the new process.

*Path* points to a path name that identifies the new process file.

*File* points to the new process file. The path prefix for this file is obtained by a search of the directories passed as the *environment* line "PATH =" (see *environ* (5)). The environment is supplied by the shell (see *sh* (1)). The shell is invoked if a command file is found by *execlp* or *execvp*.

*Arg0*, *arg1*, ..., *argn* are pointers to null-terminated character strings. These strings constitute the argument list available to the new process. By convention, at least *arg0* must be present and point to a string that is the same as *path* (or its last component).

*Argv* is an array of character pointers to null-terminated strings. These strings constitute the argument list available to the new process. By convention, *argv* must have at least one member, and it must point to a string that is the same as *path* (or its last component). *Argv* is terminated by a null pointer and is directly

usable in another *execv* because *argv[argc]* is 0.

*Envp* is an array of character pointers to null-terminated strings. These strings constitute the environment for the new process. *Envp* is terminated by a null pointer. For *exec1* and *execv*, the C run-time start-off routine places a pointer to the environment of the calling process in the global cell:

```
extern char **environ;
```

and it is used to pass the environment of the calling process to the new process.

File descriptors open in the calling process remain open in the new process, except for those whose close-on-exec flag is set; see *fcntl(2)*. For those file descriptors that remain open, the file pointer is unchanged.

Signals set to terminate the calling process will be set to terminate the new process. Signals set to be ignored by the calling process will be set to be ignored by the new process. Signals set to be caught by the calling process will be set to terminate new process; see *signal(2)*.

If the set-user-ID mode bit of the new process file is set (see *chmod(2)*), *exec* sets the effective user ID of the new process to the owner ID of the new process file. Similarly, if the set-group-ID mode bit of the new process file is set, the effective group ID of the new process is set to the group ID of the new process file. The real user ID and real group ID of the new process remain the same as those of the calling process.

The shared memory segments attached to the calling process will not be attached to the new process (see *shmop(2)*).

Profiling is disabled for the new process; see *profil(2)*.

The new process also inherits the following attributes from the calling process:

- nice value (see *nice(2)*)
- process ID
- parent process ID
- process group ID
- semadj values (see *semop(2)*)
- tty group ID (see *exit(2)* and *signal(2)*)
- trace flag (see *ptrace(2)* request 0)
- time left until an alarm clock signal (see *alarm(2)*)
- current working directory
- root directory
- file mode creation mask (see *umask(2)*)



file size limit (see *ulimit* (2))  
*utime*, *stime*, *cutime*, and *cstime* (see *times* (2))

From C, two interfaces are available. *execl* is useful when a known file with known arguments is being called; the arguments to *execl* are the character strings constituting the file and the arguments; the first argument is conventionally the same as the file name (or its last component). A 0 argument must end the argument list.

When a C program is executed, it is called as follows:

```
main(argc, argv, envp)
int argc;
char **argv, **envp;
```

where *argc* is the argument count and *argv* is an array of character pointers to the arguments themselves. As indicated, *argc* is conventionally at least one and the first member of the array points to a string containing the name of the file.

*Envp* is a pointer to an array of strings that constitute the *environment* of the process. 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*(5) for some conventionally used names. The C run-time start-off routine places a copy of *envp* in the global cell *environ*, which is used by *execv* and *execl* to pass the environment to any subprograms executed by the current program. The *exec* routines use lower-level routines as follows to pass an environment explicitly:

```
execve(file, argv, environ);
execle(file, arg0, arg1, . . . , argn, 0, environ);
```

*Execvp* and *execvp* are called with the same arguments as *execl* and *execv*, but duplicate the shell's actions in searching for an executable file in a list of directories. The directory list is obtained from the environment.

*Exec* 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's path prefix 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 execution permission.
- [EAGAIN] The system has temporarily exhausted its available memory or swap space.
- [ENOEXEC] The *exec* is not an *execip* or *execvp*, and the new process file has the appropriate access permission but 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 by some process.
- [ENOMEM] The new process requires more memory than is allowed by the system-imposed maximum *MAXMEM*.
- [E2BIG] The number of bytes in the new process's argument list is greater than the system-imposed limit of 5120 bytes.
- [EFAULT] The new process file is not as long as indicated by the size values in its header.
- [EFAULT] *Path*, *argv*, or *envp* point to an illegal address.

**RETURN VALUE**

If *exec* returns to the calling process an error has occurred; the return value will be *-1* and *errno* will be set to indicate the error.

**SEE ALSO**

*sh(1)*, *alarm(2)*, *exit(2)*, *fork(2)*, *nice(2)*, *ptrace(2)*, *semop(2)*, *signal(2)*, *times(2)*.

**NAME**

`exit`, `_exit` — terminate process

**SYNOPSIS**

```
void exit (status)
int status; void _exit (status)
int status;
```

**DESCRIPTION**

*Exit* terminates the calling process with the following consequences:

All of the file descriptors open in the calling process are closed.

If the parent process of the calling process is executing a *wait*, it is notified of the calling process's termination and the low order eight bits (i.e., bits 0377) of *status* are made available to it; see *wait(2)*.

If the parent process of the calling process is not executing a *wait*, the calling process is transformed into a zombie process. A *zombie process* is a process that only occupies a slot in the process table. It has no other space allocated either in user or kernel space. The process table slot that it occupies is partially overlaid with time accounting information (see `<sys/proc.h>`) to be used by *times*.

The parent process ID of all of the calling process's existing child processes and zombie processes is set to 1. This means the initialization process (see *intro(2)*) inherits each of these processes.

Each attached shared memory segment is detached and the value of `shm_nattach` in the data structure associated with its shared memory identifier is decremented by 1.

For each semaphore for which the calling process has set a `semadj` value (see *semop(2)*), that `semadj` value is added to the `semval` of the specified semaphore.

If the process has a process, text, or data lock, an *unlock* is performed (see *plock(2)*).

An accounting record is written on the accounting file if the system's accounting routine is enabled; see *acct(2)*.

If the process ID, tty group ID, and process group ID of the calling process are equal, the `SIGHUP` signal is sent to each process that has a process group ID equal to that of the calling process.

The C function *exit* may cause cleanup actions before the process exits. The function *\_exit* circumvents all cleanup.

**SEE ALSO**

*acct(2)*, *intro(2)*, *plock(2)*, *semop(2)*, *signal(2)*, *wait(2)*.

**WARNING**

See *WARNING* in *signal(2)*.

## NAME

`fcntl` — file control

## SYNOPSIS

```
#include <fcntl.h>
int fcntl (fildes, cmd, arg)
int fildes, cmd, arg;
```

## DESCRIPTION

*Fcntl* provides for control over open files. *Fildes* is an open file descriptor obtained from a *creat*, *open*, *dup*, *fcntl*, or *pipe* system call.

The *commands* available are:

- F\_DUPFD** Return a new file descriptor as follows:  
 Lowest numbered available file descriptor greater than or equal to *arg*.  
 Same open file (or pipe) as the original file.  
 Same file pointer as the original file (i.e., both file descriptors share one file pointer).  
 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 *exec(2)* system calls.
- F\_GETFD** Get the close-on-exec flag associated with the file descriptor *fildes*. 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 *fildes* to the low-order bit of *arg* (0 or 1 as above).
- F\_GETFL** Get *file* status flags.
- F\_SETFL** Set *file* status flags to *arg*. Only certain flags can be set; see *fcntl(5)*.
- 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.

<b>F_SETLKW</b>	This <i>cmd</i> is the same as <b>F_SETLK</b> 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.
<b>F_GETOWN</b>	Get the process ID or process group currently receiving SIGIO and SIGURG signals; process groups are returned as negative values.
<b>F_SETOWN</b>	Set the process or process group to receive SIGIO and SIGURG signals; process groups are specified by supplying <i>arg</i> as negative, otherwise <i>arg</i> is interpreted as a process ID.

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.

*Fcntl* will fail if one or more of the following are true:

<b>[EBADF]</b>	<i>Fildes</i> is not a valid open file descriptor.
<b>[EMFILE]</b>	<i>Cmd</i> is <b>F_DUPFD</b> and 20 file descriptors are currently open.
<b>[EINFILE]</b>	<i>Cmd</i> is <b>F_DUPFD</b> and <i>arg</i> is negative or greater than 20.
<b>[EINVAL]</b>	<i>Cmd</i> is <b>F_GETLK</b> , <b>F_SETLK</b> , or <b>SETLKW</b> and <i>arg</i> or the data it points to is not valid.
<b>[EACCESS]</b>	<i>Cmd</i> is <b>F_SETLK</b> the type of lock ( <i>l_type</i> ) is a read ( <b>F_RDLCK</b> ) or write ( <b>F_WRLCK</b> ) 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.
- [ENOTSOCK] *Cmd* is F\_GETOWN or F\_SETOWN and *fildev* is not a file descriptor for a socket.

**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_SETFD	Value other than -1.
F_GETFL	Value of file flags.
F_SETFL	Value other than -1.
F_GETLK	Value other than -1.
F_SETLK	Value other than -1.
F_SETLKW	Value other than -1.
F_GETOWN	Value other than -1.
F_SETOWN	Value other than -1.

Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

**SEE ALSO**

close(2), exec(2), open(2), fcntl(5).

**NAME**

fork – create a new process

**SYNOPSIS**

int fork ()

**DESCRIPTION**

*Fork* causes creation of a new process. The new process (child process) is an exact copy of the calling process (parent process). This means the child process inherits the following attributes from the parent process:

- environment
- close-on-exec flag (see *exec(2)*)
- signal handling settings (i.e., *SIG\_DFL*, *SIG\_IGN*, function address)
- set-user-ID mode bit
- set-group-ID mode bit
- profiling on/off status
- nice value (see *nice(2)*)
- all attached shared memory segments (see *shmop(2)*)
- process group ID
- tty group ID (see *exit(2)* and *signal(2)*)
- trace flag (see *ptrace(2)* request 0)
- time left until an alarm clock signal (see *alarm(2)*)
- current working directory
- root directory
- file mode creation mask (see *umask(2)*)
- file size limit (see *ulimit(2)*)

The child process differs from the parent process in the following ways:

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 file descriptors. Each of the child's file descriptors shares a common file pointer with the corresponding file descriptor of the parent.

All *semadj* values are cleared (see *semop(2)*).

Process locks, text locks and data locks are not inherited by the child (see *plock(2)*).

The child process's *utime*, *stime*, *cutime*, and *cstime* are set to 0 (see *times(2)*). The time left until an alarm clock signal is reset to 0.

*Fork* will fail and no child process will be created if one or more of the following are true:

- [EAGAIN] The system-imposed limit on the total number of processes under execution would be exceeded.
- [EAGAIN] The system-imposed limit on the total number of processes under execution by a single user would be exceeded.
- [EAGAIN] The system has temporarily exhausted its available memory or swap space.

#### 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 *errno* is set to indicate the error.

#### SEE ALSO

*exec(2)*, *nice(2)*, *plock(2)*, *ptrace(2)*, *semop(2)*, *shmop(2)*, *signal(2)*, *times(2)*.



**FSTAT (2)**

**SEE STAT**

**FSTAT (2)**

**GETEGID (2)**

**SEE GETUID**

**GETEGID (2)**

**GETEUID (2)**

**SEE GETUID**

**GETEUID (2)**

**GETGID (2)**

**SEE GETUID**

**GETGID (2)**

**NAME**

gethostid, sethostid — get/set unique identifier of current host

**SYNOPSIS**

```
hostid = gethostid()
int hostid;

sethostid(hostid)
int hostid;

cc ... -lnet
```

**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.

**LINKING**

This library is accessed by specifying `-lnet` as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**SEE ALSO**

hostid(1N), gethostname(2N)

**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;

cc ... -lnet

**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 *name* or *namelen* parameter gave an invalid address.

[EPERM] The caller was not the super-user.

**LINKING**

This library is accessed by specifying *-lnet* as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**SEE ALSO**

gethostid(2N)

**BUGS**

Host names are limited to 255 characters.

**NAME**

getpeername — get name of connected peer

**SYNOPSIS**

```
getpeername(s, name, namelen)
int s;
struct sockaddr *name;
int *namelen;
cc ... -lnet
```

**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.

**LINKING**

This library is accessed by specifying `-lnet` as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**SEE ALSO**

bind(2N), socket(2N), getsockname(2N)

**GETPGRP (2)**

**SEE *GETPID***

**GETPGRP (2)**

**NAME**

`getpid`, `getpgrp`, `getppid` — get process, process group, and parent process IDs

**SYNOPSIS**

`int getpid ( )`

`int getpgrp ( )`

`int getppid ( )`

**DESCRIPTION**

*Getpid* returns the process ID of the calling process.

*Getpgrp* returns the process group ID of the calling process.

*Getppid* returns the parent process ID of the calling process.

These system calls are useful for generating uniquely-named temporary files.

**SEE ALSO**

`exec(2)`, `fork(2)`, `intro(2)`, `setpgrp(2)`, `signal(2)`.

**GETPPID (2)**

**SEE *GETPID***

**GETPPID (2)**

**NAME**

getsockname — get socket name

**SYNOPSIS**

```
getsockname(s, name, namelen)
int s;
struct sockaddr *name;
int *namelen;
cc ... -lnet
```

**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 *s* is not a valid descriptor.
- [ENOTSOCK]    The argument *s* is a file, not a socket.
- [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.

**LINKING**

This library is accessed by specifying `-lnet` as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**SEE ALSO**

bind(2N), socket(2N)



**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;

cc ... -lnet
```

**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 of the requested options(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*(2N). Options at other protocol levels vary in format and name, consult the appropriate entries in (5P).

**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.

**LINKING**

This library is accessed by specifying `-lnet` as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**SEE ALSO**

`socket(2N)`, `getprotoent(3N)`.

**NAME**

*getuid*, *geteuid*, *getgid*, *getegid* — get real user, effective user, real group, and effective group IDs

**SYNOPSIS**

**unsigned short** *getuid* ()

**unsigned short** *geteuid* ()

**unsigned short** *getgid* ()

**unsigned short** *getegid* ()

**DESCRIPTION**

*Getuid* returns the real user ID of the calling process.

*Geteuid* returns the effective user ID of the calling process.

*Getgid* returns the real group ID of the calling process.

*Getegid* returns the effective group ID of the calling process.

**SEE ALSO**

*intro(2)*, *setuid(2)*.

**NAME**

ioctl — control device

**SYNOPSIS**

```
ioctl (fildes, request, arg)
int fildes, request;
```

**DESCRIPTION**

*ioctl* performs a variety of functions on character special files (devices). The write-ups of various devices in Section 7 of the *UniPlus<sup>+</sup> System Administrator Reference Manual* discuss how *ioctl* applies to them.

*ioctl* will fail if one or more of the following are true:

- [EBADF] *Fildes* is not a valid open file descriptor.
- [ENOTTY] *Fildes* is not associated with a character special device.
- [EINVAL] *Request* or *arg* is not valid. See Section 7 of the *UniPlus<sup>+</sup> System Administrator Reference Manual*.
- [EINTR] A signal was caught during the *ioctl* system call.

**RETURN VALUE**

If an error has occurred, a value of `-1` is returned and *errno* is set to indicate the error.

**SEE ALSO**

*termio(7)* in the *UniPlus<sup>+</sup> System Administrator Reference Manual*.

## NAME

kill — send a signal to a process or a group of processes

## SYNOPSIS

```
int kill (pid, sig)
int pld, sig;
```

## DESCRIPTION

*Kill* sends a signal to a process or a group of processes. The process or group of processes to which the signal is to be sent is specified by *pid*. The signal that is to be sent is specified by *sig* and is either one from the list given in *signal(2)*, or 0. If *sig* is 0 (the null signal), error checking is performed but no signal is actually sent. This can be used to check the validity of *pid*.

The real or effective user ID of the sending process must match the real or effective user ID of the receiving process, unless the effective user ID of the sending process is super-user, or the process is sending to itself.

The processes with a process ID of 0 and a process ID of 1 are special processes (see *intro(2)*) and will be referred to below as *proc0* and *proc1* respectively.

If *pid* is greater than zero, *sig* will be sent to the process whose process ID is equal to *pid*. *Pid* may equal 1.

If *pid* is 0, *sig* will be sent to all processes excluding *proc0* and *proc1* whose process group ID is equal to the process group ID of the sender.

If *pid* is -1 and the effective user ID of the sender is not super-user, *sig* will be sent to all processes excluding *proc0* and *proc1* whose real user ID is equal to the effective user ID of the sender.

If *pid* is -1 and the effective user ID of the sender is super-user, *sig* will be sent to all processes excluding *proc0* and *proc1*.

If *pid* is negative but not -1, *sig* will be sent to all processes whose process group ID is equal to the absolute value of *pid*.

*Kill* will fail and no signal will be sent if one or more of the following are true:

[EINVAL]	<i>Sig</i> is not a valid signal number.
[EINVAL]	<i>Sig</i> is SIGKILL and <i>pid</i> is 1 ( <i>proc1</i> ).
[ESRCH]	No process can be found corresponding to that specified by <i>pid</i> .
[EPERM]	The sending process is not sending to itself, its effective user ID is not super-user, and its real or effective user ID does not match the real or effective user ID of the receiving process. [EPERM]

## 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

kill(1), getpid(2), setpggrp(2), signal(2).

## NAME

link - link to a file

## SYNOPSIS

```
int link (path1, path2)
char *path1, *path2;
```

## DESCRIPTION

*Path1* points to a path name naming an existing file. *Path2* points to a path name naming the new directory entry to be created. *Link* creates a new link (directory entry) for the existing file.

*Link* will fail and no link will be created if one or more 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>path1</i> does not exist.   |
| [EEXIST]  | The link named by <i>path2</i> exists.   |
| [EPERM]   | The file named by <i>path1</i> is a directory and the effective user ID is not super-user.                         |
| [EXDEV]   | The link named by <i>path2</i> and the file named by <i>path1</i> are on different logical devices (file systems). |
| [ENOENT]  | <i>Path2</i> points to a null path name.   |
| [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 allocated address space of the process.   |
| [EMLINK]  | The maximum number of links to a file would be exceeded.   |

## 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

unlink(2).

**NAME**

`listen` — listen for connections on a socket

**SYNOPSIS**

```
listen(s, backlog)
int s, backlog;
cc ... -lnet
```

**DESCRIPTION**

To accept connections, a socket is first created with `socket(2N)`, a backlog for incoming connections is specified with `listen(2N)` and then the connections are accepted with `accept(2N)`. The `listen` call applies only to sockets of type `SOCK_STREAM` or `SOCK_PKTSTREAM`.

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 <code>s</code> is not a valid descriptor.
[ENOTSOCK]	The argument <code>s</code> is not a socket.
[EOPNOTSUPP]	The socket is not of a type that supports the operation <code>listen</code> .

**LINKING**

This library is accessed by specifying `-lnet` as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**SEE ALSO**

`accept(2N)`, `connect(2N)`, `socket(2N)`

**BUGS**

The `backlog` is currently limited (silently) to 5.

**NAME**

locking — provide exclusive file regions for reading or writing

**SYNOPSIS**

```
locking(filides, mode, size)
int filides;
int mode;
int size;
```

**DESCRIPTION**

*Locking* will allow a specified number of bytes to be accessed only by the locking process. Other processes which attempt to lock, read, or write the locked area will sleep until the area becomes unlocked.

*Filides* is the word returned from a successful *open*, *creat*, *dup*, or *pipe* system call.

*Mode* is zero to unlock the area. *Mode* is one or two for making the area locked. If the *mode* is one and the area has some other lock on it, then the process will sleep until the entire area is available. If the *mode* is two and the area is locked, an error will be returned.

*Size* is the number of contiguous bytes to be locked or unlocked. The area to be locked starts at the current offset in the file. If *size* is zero, the area to the end of file is locked.

The potential for a deadlock occurs when a process controlling a locked area is put to sleep by accessing another process's locked area. Thus calls to *locking*, *read*, or *write* scan for a deadlock prior to sleeping on a locked area. An error return is made if sleeping on the locked area would cause a deadlock.

Lock requests may, in whole or part, contain or be contained by a previously locked area for the same process. When this or adjacent areas occur, the areas are combined into a single area. If the request requires a new lock element with the lock table full, an error is returned, and the area is not locked.

Unlock requests may, in whole or part, release one or more locked regions controlled by the process. When regions are not fully released, the remaining areas are still locked by the process. Release of the center section of a locked area requires an additional lock element to hold the cut off section. If the lock table is full, an error is returned, and the requested area is not released.

While locks may be applied to special files or pipes, read/write operations will not be blocked. Locks may not be applied to a directory.

Note that *close*(2) automatically removes any locks that were associated with the closed file descriptor.

**SEE ALSO**

*close*(2), *creat*(2), *dup*(3), *open*(2), *read*(2), *write*(3).

**DIAGNOSTICS**

The value -1 is returned if the file does not exist, or if a deadlock using file locks would occur. EACCES will be returned for lock requests in which the area is already locked by another process. EDEADLOCK will be returned



by: read, write, or locking if a deadlock would occur. EDEADLOCK will also be returned when the locktable overflows.

**NAME**

`lseek` — move read/write file pointer

**SYNOPSIS**

```
long lseek (fildes, offset, whence)
int fildes;
long offset;
int whence;
```

**DESCRIPTION**

*Fildes* is a file descriptor returned from a *creat*, *open*, *dup*, or *fcntl* system call. *Lseek* sets the file pointer associated with *fildes* as follows:

If *whence* is 0, the pointer is set to *offset* bytes.

If *whence* is 1, the pointer is set to its current location plus *offset*.

If *whence* is 2, the pointer is set to the size of the file plus *offset*.

Upon successful completion, the resulting pointer location, as measured in bytes from the beginning of the file, is returned.

*Lseek* will fail and the file pointer will remain unchanged if one or more of the following are true:

[EBADF] *Fildes* is not an open file descriptor.

[ESPIPE] *Fildes* is associated with a pipe or fifo.

[EINVAL and SIGSYS signal]  
*Whence* is not 0, 1, or 2.

[EINVAL] The resulting file pointer would be negative.

Some devices are incapable of seeking. The value of the file pointer associated with such a device is undefined.

**RETURN VALUE**

Upon successful completion, a non-negative integer indicating the file pointer value is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

**SEE ALSO**

`creat(2)`, `dup(3)`, `fcntl(2)`, `open(2)`.

**NAME**

`mknod` - make a directory, or a special or ordinary file

**SYNOPSIS**

```
int mknod (path, mode, dev)
char *path;
int mode, dev;
```

**DESCRIPTION**

*Mknod* creates a new file named by the path name pointed to by *path*. The mode of the new file is initialized from *mode*. Where the value of *mode* is interpreted as follows:

```
0170000 file type; one of the following:
    0010000 fifo special
    0020000 character special
    0040000 directory
    0060000 block special
    0100000 or 0000000 ordinary file
0004000 set user ID on execution
0002000 set group ID on execution
0001000 save text image after execution
0000777 access permissions; constructed from the following
    0000400 read by owner
    0000200 write by owner
    0000100 execute (search on directory) by owner
    0000070 read, write, execute (search) by group
    0000007 read, write, execute (search) by others
```

The owner ID of the file is set to the effective user ID of the process. The group ID of the file is set to the effective group ID of the process.

Values of *mode* other than those above are undefined and should not be used. 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)*. 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 for file types other than FIFO special.

*Mknod* will fail and the new file will not be created if one or more of the following are true:

[EPERM]	The effective user ID of the process is not super-user.
[ENOTDIR]	A component of the path prefix is not a directory.
[ENOENT]	A component of the path prefix does not exist.
[EROFS]	The directory in which the file is to be created is located on a read-only file system.
[EEXIST]	The named file exists.
[EFAULT]	<i>Path</i> points outside the allocated address space of the process.

**RETURN VALUE**

Upon successful completion a value of 0 is returned. Otherwise, a value of

## MKNOD (2)

## MKNOD (2)

-1 is returned and *errno* is set to indicate the error.

### SEE ALSO

mkdir(1), chmod(2), exec(2), umask(2), fs(4).

**NAME**

mount — mount a file system

**SYNOPSIS**

```
int mount (spec, dir, rwflag)
char *spec, *dir;
int rwflag;
```

**DESCRIPTION**

*Mount* requests that a removable file system contained on the block special file identified by *spec* be mounted on the directory identified by *dir*. *Spec* and *dir* are pointers to path names.

Upon successful completion, references to the file *dir* will refer to the root directory on the mounted file system.

The low-order bit of *rwflag* is used to control write permission on the mounted file system; if 1, writing is forbidden, otherwise writing is permitted according to individual file accessibility. 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.

*Mount* may be invoked only by the super-user.

*Mount* will fail if one or more of the following are true:

- |           |   |
|-----------|---|
| [EPERM]   | The effective user ID is not super-user.  |
| [ENOENT]  | Any of the named files does not exist.  |
| [ENOTDIR] | A component of a path prefix is not a directory.  |
| [ENOTBLK] | <i>Spec</i> is not a block special device.  |
| [ENXIO]   | The device associated with <i>spec</i> does not exist.  |
| [ENOTDIR] | <i>Dir</i> is not a directory.  |
| [EFAULT]  | <i>Spec</i> or <i>dir</i> points outside the allocated address space of the process.              |
| [EBUSY]   | <i>Dir</i> is currently mounted on, is someone's current working directory, or is otherwise busy. |
| [EBUSY]   | The device associated with <i>spec</i> is currently mounted.                                      |
| [EBUSY]   | There are no more mount table entries.  |

**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**

umount(2).

## 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] | <i>Msqid</i> is not a valid message queue identifier.  |
| [EINVAL] | <i>Cmd</i> is not a valid command.   |
| [EACCES] | <i>Cmd</i> is equal to IPC_STAT and [READ] operation permission is denied to the calling process (see <i>intro(2)</i> ).   |
| [EPERM]  | <i>Cmd</i> 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 <i>msg_perm.uid</i> in the data structure associated with <i>msqid</i> . |
| [EPERM]  | <i>Cmd</i> is equal to IPC_SET, an attempt is being made to increase to the value of <i>msg_qbytes</i> , 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 *key*, but operation permission (see *intro(2)*) as specified by the low-order 9 bits of *msgflg* would not be granted.
- [ENOENT] A message queue identifier does not exist for *key* 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 *key* 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, msgsnd, msgrcv** – 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.

*Msgid* 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] *Msgid* is not a valid message queue identifier.
- [EACCESS] Operation permission is denied to the calling process (see *intro*(2)).
- [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 *msgid* (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 *msgid* 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] | <i>Msqid</i> is not a valid message queue identifier.   |
| [EACCES] | Operation permission is denied to the calling process.  |
| [EINVAL] | <i>Msgsz</i> is less than 0.  |
| [E2BIG]  | <i>Mtext</i> is greater than <i>msgsz</i> and ( <i>msgflg</i> & MSG_NOERROR) is "false".              |
| [ENOMSG] | The queue does not contain a message of the desired type and ( <i>msgtyp</i> & IPC_NOWAIT) is "true". |
| [EFAULT] | <i>Msgp</i> 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 *msgId* 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**

*nice* – change priority of a process

**SYNOPSIS**

```
int nice (incr)
```

```
int incr;
```

**DESCRIPTION**

*Nice* adds the value of *incr* to the nice value of the calling process. A process's *nice value* is a positive number for which a more positive value results in lower CPU priority.

A maximum nice value of 39 and a minimum nice value of 0 are imposed by the system. Requests for values above or below these limits result in the nice value being set to the corresponding limit.

[EPERM] *Nice* will fail and not change the nice value if *incr* is negative or greater than 40 and the effective user ID of the calling process is not super-user.

**RETURN VALUE**

Upon successful completion, *nice* returns the new nice value minus 20. Otherwise, a value of -1 is returned and *errno* is set to indicate the error. If a value of -1 is a valid return value on successful completion (i.e., if your new nice value is 19), *errno* is not changed.

**SEE ALSO**

*nice(1)*, *exec(2)*.

## NAME

open — open for reading or writing

## SYNOPSIS

```
#include <fcntl.h>
int open (path, oflag [ , mode ] )
char *path;
int oflag, mode;
```

## DESCRIPTION

*Path* points to a path name naming a file. *Open* opens a file descriptor for the named file and sets the file status flags according to the value of *oflag*. *Oflag* values are constructed by or-ing flags from the following list (only one of the first three flags below may be used):

- O\_RDONLY** Open for reading only.
- O\_WRONLY** Open for writing only.
- O\_RDWR** Open for reading and writing.
- O\_NDELAY** This flag may affect subsequent reads and writes. See *read(2)* and *write(3)*.

When opening a FIFO with **O\_RDONLY** or **O\_WRONLY** set:

If **O\_NDELAY** is set:

An *open* for reading-only will return without delay. An *open* for writing-only will return an error if no process currently has the file open for reading.

If **O\_NDELAY** is clear:

An *open* for reading-only will block until a process opens the file for writing. An *open* for writing-only will block until a process opens the file for reading.

When opening a file associated with a communication line:

If **O\_NDELAY** is set:

The open will return without waiting for carrier.

If **O\_NDELAY** is clear:

The open will block until carrier is present.

- O\_APPEND** If set, the file pointer will be set to the end of the file prior to each write.

- O\_CREAT** If the file exists, this flag has no effect. Otherwise, the owner ID of the file is set to the effective user ID of the process, the group ID of the file is set to the effective group ID of the process, and the low-order 12 bits of the file mode are set to the value of *mode* modified as follows (see *creat(2)*):

All bits set in the file mode creation mask of the process are cleared. See *umask(2)*.

The "save text image after execution bit" of the mode is cleared. See *chmod(2)*.

- O\_TRUNC** If the file exists, its length is truncated to 0 and the mode and owner are unchanged.
- O\_EXCL** If **O\_EXCL** and **O\_CREAT** are set, *open* will fail if the file exists.

The file pointer used to mark the current position within the file is set to the beginning of the file.

The new file descriptor is set to remain open across *exec* system calls. See *fcntl(2)*.

The named file is opened unless one or more of the following are true:

- [ENOTDIR] A component of the path prefix is not a directory.
- [ENOENT] **O\_CREAT** is not set and the named file does not exist.
- [EACCES] A component of the path prefix denies search permission.
- [EACCES] *Oflag* permission is denied for the named file.
- [EISDIR] The named file is a directory and *oflag* is write or read/write.
- [EROFS] The named file resides on a read-only file system and *oflag* is write or read/write.
- [EMFILE] Twenty (20) 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 *oflag* is write or read/write.
- [EFAULT] *Path* points outside the allocated address space of the process.
- [EEXIST] **O\_CREAT** and **O\_EXCL** are set, and the named file exists.
- [ENXIO] **O\_NDELAY** is set, the named file is a FIFO, **O\_WRONLY** is set, and no process has the file open for reading.
- [EINTR] A signal was caught during the *open* system call.
- [ENFILE] The system file table is full.

#### RETURN VALUE

Upon successful completion, the file descriptor is returned. Otherwise, a value of  $-1$  is returned and *errno* is set to indicate the error.

#### SEE ALSO

*chmod(2)*, *close(2)*, *creat(2)*, *fcntl(2)*, *lseek(2)*, *read(2)*, *umask(2)*, *write(3)*.

**NAME**

pause — suspend process until signal

**SYNOPSIS**

pause ()

**DESCRIPTION**

*Pause* suspends the calling process until it receives a signal. The signal must be one that is not currently set to be ignored by the calling process.

If the signal causes termination of the calling process, *pause* will not return.

If the signal is *caught* by the calling process and control is returned from the signal-catching function (see *signal(2)*), the calling process resumes execution from the point of suspension; with a return value of -1 from *pause* and *errno* set to EINTR.

**SEE ALSO**

alarm(2), kill(2), signal(2), wait(2).



**NAME**

*phys* — allow a process to access physical addresses

**SYNOPSIS**

```
phys(physnum, virtaddr, size, physaddr)  
int physnum  
char *virtaddr;  
long size;  
char *physaddr;
```

**DESCRIPTION**

The *phys*(2) call maps arbitrary physical memory into a process's virtual address space. The virtual address used by *phys* must not otherwise be used. *Physnum* is a number (0-3) that specifies which of 4 physical spaces to set up. Up to 4 *phys*(2) calls can be active at any one time. *Virtaddr* is the process's virtual address. *Size* is the number of bytes to map in. *Physaddr* is the physical address to map in.

Valid *virtaddr* and *physaddr* values are constrained by hardware and must be at an address multiple of the resolution of the CPU's memory management scheme. If *size* is non zero, *size* is rounded up to the next MMU resolution boundary. If *size* is zero, any previous *phys*(2) mapping for that *physnum* segment is nullified.

For example, the call:

```
phys(2, 0x100000, 32768, 0)
```

will allow a process to access physical locations 0 through 32767 by referencing virtual address 0x100000 through 0x100000+32767.

In actuality, the CPU MMU register is loaded with *physaddr* shifted to account for page resolution.

*Phys*(2) may only be executed by the super-user.

**DIAGNOSTICS**

The value zero is returned if the *phys* call was successful. The value -1 is returned if not super-user, if *virtaddr* or *physaddr* is not in the proper range, or if the specified *virtaddr* segment register is already in use.

**BUGS**

This system call is very machine dependent.

**NAME**

pipe — create an interprocess channel

**SYNOPSIS**

```
int pipe (fildes)
int fildes[2];
```

**DESCRIPTION**

*Pipe* creates an I/O mechanism called a pipe and returns two file descriptors, *fildes*[0] and *fildes*[1]. *Fildes*[0] is opened for reading and *fildes*[1] is opened for writing.

Up to 5120 bytes of data are buffered by the pipe before the writing process is blocked. A read only file descriptor *fildes*[0] accesses the data written to *fildes*[1] on a first-in-first-out (FIFO) basis.

[EMFILE] *Pipe* will fail if 19 or more file descriptors are currently open.

[ENFILE] The system file table is full.

**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**

sh(1), read(2), write(3).

## NAME

`plock` – lock process, text, or data in memory

## SYNOPSIS

```
#include <sys/lock.h>
```

```
int plock (op)
```

```
int op;
```

## DESCRIPTION

*Plock* allows the calling process to lock its text segment (text lock), its data segment (data lock), or both its text and data segments (process lock) into memory. Locked segments are immune to all routine swapping. *Plock* also allows these segments to be unlocked. The effective user ID of the calling process must be super-user to use this call. *Op* specifies the following:

**PROCLOCK** – lock text and data segments into memory (process lock)

**TXTLCK** – lock text segment into memory (text lock)

**DATLOCK** – lock data segment into memory (data lock)

**UNLOCK** – remove locks

*Plock* will fail and not perform the requested operation if one or more of the following are true:

- |          |  |
|----------|--|
| [EPERM]  | The effective user ID of the calling process is not super-user.  |
| [EAGAIN] | The system has temporarily exhausted its available memory or swap space.   |
| [EINVAL] | <i>Op</i> is equal to <b>PROCLOCK</b> and a process lock, a text lock, or a data lock already exists on the calling process. |
| [EINVAL] | <i>Op</i> is equal to <b>TXTLCK</b> and a text lock, or a process lock already exists on the calling process.                |
| [EINVAL] | <i>Op</i> is equal to <b>DATLOCK</b> and a data lock, or a process lock already exists on the calling process.               |
| [EINVAL] | <i>Op</i> is equal to <b>UNLOCK</b> and no type of lock exists on the calling process.                                       |

## RETURN VALUE

Upon successful completion, a value of 0 is returned to the calling process. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

**PLOCK(2)**

**PLOCK(2)**

**SEE ALSO**

**exec(2), exit(2), fork(2).**

**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; *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 (16 bit), fixed-point fraction with binary point at the left: FFFF (hex) gives a 1-1 mapping of *pc*'s to words in *buff*; FFFF (hex) maps each pair of instruction words together. 2(hex) 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 *exec* is executed, but remains on in child and parent both after a *fork*. Profiling will be turned off if an update in *buff* would cause a memory fault.

**RETURN VALUE**

Not defined.

**SEE ALSO**

prof(1), monitor(3C).

## NAME

ptrace — process trace

## SYNOPSIS

```
int 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. Its primary use is for the implementation of breakpoint debugging. The child process behaves normally until it encounters a signal (see *signal(2)* for the list), at which time it enters a stopped state and its parent is notified via *wait(2)*. When the child is in the stopped state, its parent can examine and modify its "core image" using *ptrace*. Also, the parent can cause the child either to terminate or continue, with the possibility of ignoring the signal that caused it to stop.

The *request* argument determines the precise action to be taken by *ptrace* and is one of the following:

- 0 This request must be issued by the child process if it is to be traced by its parent. It turns on the child's trace flag that stipulates that the child should be left in a stopped state upon receipt of a signal rather than the state specified by *func*; see *signal(2)*. The *pid*, *addr*, and *data* arguments are ignored, and a return value is not defined for this request. Peculiar results will ensue if the parent does not expect to trace the child.

The remainder of the requests can only be used by the parent process. For each, *pid* is the process ID of the child. The child must be in a stopped state before these requests are made.

- 1, 2 With these requests, the word at location *addr* in the address space of the child is returned to the parent process. Either request 1 or request 2 may be used with equal results. The *data* argument is ignored. These two requests will fail if *addr* is not the start address of a word, in which case a value of  $-1$  is returned to the parent process and the parent's *errno* is set to  $EIO$ .
- 3 With this request, the word at location *addr* in the child's USER area in the system's address space (see  $\langle \text{sys/user.h} \rangle$ ) is returned to the parent process. Addresses are system dependent. The *data* argument is ignored. This request will fail if *addr* is not the start address of a word or is outside the USER area, in which case a value of  $-1$  is returned to the parent process and the parent's *errno* is set to  $EIO$ .
- 4, 5 With these requests, the value given by the *data* argument is written into the address space of the child at location *addr*. Either request 4 or request 5 may be used with equal results. Upon successful completion, the value written into the address space of the child is returned to the parent. These two requests will fail if *addr* is a location in a pure procedure space and another process is executing in that space, or *addr* is not the start address of a word. Upon failure a value of  $-1$  is returned to the parent process and the parent's *errno* is set to  $EIO$ .

- 6 With this request, a few entries in the child's USER area can be written. *Data* gives the value that is to be written and *addr* is the location of the entry. The few entries that can be written are:
- the general registers
  - the condition codes
  - certain bits of the Processor Status Word
- 7 This request causes the child to resume execution. If the *data* argument is 0, all pending signals including the one that caused the child to stop are canceled before it resumes execution. If the *data* argument is a valid signal number, the child resumes execution as if it had incurred that signal, and any other pending signals are canceled. The *addr* argument must be equal to 1 for this request. Upon successful completion, the value of *data* is returned to the parent. This request will fail if *data* is not 0 or a valid signal number, in which case a value of -1 is returned to the parent process and the parent's *errno* is set to EIO.
- 8 This request causes the child to terminate with the same consequences as *exit(2)*.
- 9 This request sets the trace bit in the Processor Status Word of the child and then executes the same steps as listed above for request 7. The trace bit causes an interrupt upon completion of one machine instruction. This effectively allows single stepping of the child.  
Note: the trace bit remains set after an interrupt.
- 10 Read user register; *pid* = child process id; *addr* = register number; *data* is ignored; returns value of child's register.
- 11 Write user register; *pid* = child process id; *addr* = register number; *data* = integer value to be written into named register.  
NOTE: For both requests 10 and 11, the register numbers are as shown below for the 68000 family (these numbers are system dependent).

Register	Register #	Register	Register #
d0	0	a1	9
d1	1	a2	10
d2	2	a3	11
d3	3	a4	12
d4	4	a5	13
d5	5	a6	14
d6	6	SP	15
d7	7	PC	16
a0	8	PS	17

To forestall possible fraud, *ptrace* inhibits the set-user-id facility on subsequent *exec(2)* calls. If a traced process calls *exec*, it will stop before executing the first instruction of the new image showing signal SIGTRAP.

#### GENERAL ERRORS

*Ptrace* will in general fail if one or more of the following are true:

*Request* is an illegal number. [EIO]

**PTRACE(2)**

**PTRACE(2)**

*Pid* identifies a child that does not exist or has not executed a *ptrace* with request 0. [ESRCH]

**NOTE**

Request 11 completely supercedes request 6, and request 10 largely supercedes request 3 (request 3 can read any part of the child's user area while request 10 can only read register values of the child).

**SEE ALSO**

`exec(2)`, `signal(2)`, `wait(2)`.



## NAME

read – read from file

## SYNOPSIS

```
int read (fildes, buf, nbyte)
int fildes;
char *buf;
unsigned nbyte;
```

## DESCRIPTION

*Fildes* is a file descriptor obtained from a *creat*, *open*, *dup*, *fcntl*, *pipe*, or *socket* system call.

*Read* attempts to read *nbyte* bytes from the file associated with *fildes* into the buffer pointed to by *buf*.

On devices capable of seeking, the *read* starts at a position in the file given by the file pointer associated with *fildes*. Upon return from *read*, the file pointer is incremented by the number of bytes actually read.

Devices that are incapable of seeking always read from the current position. The value of a file pointer associated with such a file is undefined.

Upon successful completion, *read* returns the number of bytes actually read and placed in the buffer; this number may be less than *nbyte* if the file is associated with a communication line (see *ioctl(2)*, *socket(2N)*, and *termio(7)*), or if the number of bytes left in the file is less than *nbyte* bytes. A value of 0 is returned when an end-of-file has been reached.

When attempting to read from an empty pipe (or FIFO):

If *O\_NDELAY* is set, the read will return a 0.

If *O\_NDELAY* is clear, the read will block until data is written to the file or the file is no longer open for writing.

When attempting to read a file associated with a tty that has no data currently available:

If *O\_NDELAY* is set, the read will return a 0.

If *O\_NDELAY* is clear, the read will block until data becomes available.

*Read* will fail if one or more of the following are true:

[EIO]            A physical I/O error has occurred.

- [ENXIO]           The device associated with the file descriptor is a block-special or character-special file and the value of the file pointer is out of range.
- [EBADF]           *Fildes* is not a valid file descriptor open for reading.
- [EFAULT]           *Buf* points outside the allocated address space.
- [EINTR]           A signal was caught during the *read* system call.

**RETURN VALUE**

Upon successful completion a non-negative integer is returned indicating the number of bytes actually read. Otherwise, a -1 is returned and *errno* is set to indicate the error.

**SEE ALSO**

*creat(2)*, *fcntl(2)*, *ioctl(2)*, *open(2)*, *pipe(2)*, *socket(2N)*, *termio(7)* in the *Administrator Reference Manual*.

## NAME

readv – read from file

## SYNOPSIS

```
#include <sys/types.h>
#include <sys/uio.h>
```

```
cc = readv(d,iov,iocnt)
int cc, d;
struct iovec *iov;
int iocnt;
```

## DESCRIPTION

*Fildes* is a file descriptor obtained from a *creat*, *open*, *dup*, *fcntl*, *pipe*, or *socket* system call.

*Readv* attempts to read *nbyte* bytes from the file associated with *fildes* and scatters the input data into the *iocnt* buffers specified by the members of the *iovec* array: *iov*[0], *iov*[1], . . . , *iov*[*iocnt* - 1].

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 devices capable of seeking, the *readv* starts at a position in the file given by the file pointer associated with *fildes*. Upon return from *readv*, the file pointer is incremented by the number of bytes actually read.

Devices that are incapable of seeking always read from the current position. The value of a file pointer associated with such a file is undefined.

Upon successful completion, *readv* returns the number of bytes actually read and placed in the buffer; this number may be less than *nbyte* if the file is associated with a communication line (see *ioctl*(2), *socket*(2N), and *termio*(7)), or if the number of bytes left in the file is less than *nbyte* bytes. A value of 0 is returned when an end-of-file has been reached.

When attempting to read from an empty pipe (or FIFO):

If `O_NDELAY` is set, the read will return a 0.

If `O_NDELAY` is clear, the read will block until data is written to the file or the file is no longer open for writing.

When attempting to read a file associated with a tty that has no data currently available:

If `O_NDELAY` is set, the read will return a 0.

If `O_NDELAY` is clear, the read will block until data becomes available.

*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 signal was caught during the *read* system call.

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.

#### RETURN VALUE

Upon successful completion a non-negative integer is returned indicating the number of bytes actually read. Otherwise, a -1 is returned and *errno* is set to indicate the error.

#### SEE ALSO

*creat(2)*, *fcntl(2)*, *ioctl(2)*, *open(2)*, *pipe(2)*, *socket(2N)*, *termio(7)* in the *Administrator Reference Manual*.

**NAME**

reboot – reboot the system

**SYNOPSIS**

reboot ( )

**DESCRIPTION**

*Reboot* causes the kernel to execute the initial bootstrap code that was used to boot the operating system.

The *reboot(2)* command takes the place of a manual restart. Reboot does not work on all systems.

**SEE ALSO**

reboot(1m).

## 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;

cc ... -lnet
```

## 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(2N)*), 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(2N)*.

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 EWOULD-BLOCK.

The *select(2N)* 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,

```
#defineMSG_PEEK    0x1    /* peek at incoming message */
#defineMSG_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 iov *msg_iov;       /* scatter/gather array */
    int     msg_iovlen;        /* # elements in msg_iov */
    caddr_t msg_accrights;     /* access rights sent/received */
    int     msg_accrightslen;
};

```

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. 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. |

#### LINKING

This library is accessed by specifying `-lnet` as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

#### SEE ALSO

`read(2)`, `send(2N)`, `socket(2N)`

**RECVFROM (2N)**

**SEE *RECV***

**RECVFROM (2N)**

**RECVMSG (2N)**

**SEE *RECV***

**RECVMSG (2N)**

**SBRK (2)**

**SEE *BRK***

**SBRK (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;
cc ... -lnet
```

**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] | A signal was delivered before any of the selected for events occurred or the time limit expired. |

**LINKING**

This library is accessed by specifying *-lnet* as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**SEE ALSO**

accept(2N), connect(2N), readv(3N), writev(3N), recv(2N), send(2N)

**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 <i>cmd</i> 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 <i>cmd</i> 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> : <i>sem_perm.uid</i> <i>sem_perm.gid</i> <i>sem_perm.mode</i> /* 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 <i>semval</i> .
GETPID	The value of <i>sempid</i> .
GETNCNT	The value of <i>semncnt</i> .
GETZCNT	The value of <i>semzcnt</i> .
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.

[EXIST] 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] | <i>Semid</i> is not a valid semaphore identifier.   |
| [EFBIG]  | <i>Sem_num</i> is less than zero or greater than or equal to the number of semaphores in the set associated with <i>semid</i> . |
| [E2BIG]  | <i>Nsops</i> is greater than the system-imposed maximum.  |
| [EACCES] | Operation permission is denied to the calling process (see <i>intro(2)</i> ).   |
| [EAGAIN] | The operation would result in suspension of the calling process but ( <i>sem_flg</i> & 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 <i>semval</i> 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;

cc ... -lnet
```

**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*(2N) call may be used to determine when it is possible to send more data.

The *flags* parameter may be set to MSG\_OOB to send out-of-band data on sockets which support this notion (e.g. SOCK\_STREAM).

See *recv*(2N) 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 <i>s</i> 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.

**LINKING**

This library is accessed by specifying `-lnet` as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**SEE ALSO**

`recv(2N)`, `socket(2N)`

**SENDMSG (2N)**

**SEE SEND**

**SENDMSG (2N)**

**SENDTO (2N)**

**SEE SEND**

**SENDTO (2N)**

**SETGID (2)**

**SEE SETUID**

**SETGID (2)**

**SETHOSTID (2N)**

**SEE GETHOSTID**

**SETHOSTID (2N)**

**SETHOSTNAME (2N)**

**SEE GETHOSTNAME**

**SETHOSTNAME (2N)**

**SETPGRP(2)****SETPGRP(2)****NAME**

*setpgrp* — set process group ID

**SYNOPSIS**

**int** *setpgrp* ( )

**DESCRIPTION**

*Setpgrp* sets the process group ID of the calling process to the process ID of the calling process and returns the new process group ID.

**RETURN VALUE**

*Setpgrp* returns the value of the new process group ID.

**SEE ALSO**

*exec(2)*, *fork(2)*, *getpid(2)*, *intro(2)*, *kill(2)*, *signal(2)*.

**NAME**

setregid — set real and effective group ID

**SYNOPSIS**

```
setregid(rgid, egid)
int rgid, egid;
cc ... -lnet
```

**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.

**LINKING**

This library is accessed by specifying `-lnet` as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**SEE ALSO**

getgid(2), setreuid(2), setuid(2)

**NAME**

setreuid — set real and effective user ID's

**SYNOPSIS**

```
setreuid(ruid, euid)
int ruid, euid;
cc ... -lnet
```

**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.

**LINKING**

This library is accessed by specifying `-lnet` as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**SEE ALSO**

getuid(2), setregid(2), setuid(2)

**NAME**

setuid, setgid — set user and group IDs

**SYNOPSIS**

```
int setuid (uid)
int uid;

int setgid (gid)
int gid;
```

**DESCRIPTION**

*Setuid (setgid)* is used to set the real user (group) ID and effective user (group) ID of the calling process.

If the effective user ID of the calling process is super-user, the real user (group) ID and effective user (group) ID are set to *uid (gid)*.

If the effective user ID of the calling process is not super-user, but its real user (group) ID is equal to *uid (gid)*, the effective user (group) ID is set to *uid (gid)*.

If the effective user ID of the calling process is not super-user, but the saved set-user (group) ID from *exec(2)* is equal to *uid (gid)*, the effective user (group) ID is set to *uid (gid)*.

*Setuid (setgid)* will fail if the real user (group) ID of the calling process is not equal to *uid (gid)* and its effective user ID is not super-user. [EPERM]

The *uid* is out of range. [EINVAL]

**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**

getuid(2), intro(2).

**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 *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*.

*Shmctl* will fail if one or more of the following are true:

- [EINVAL] *Shmid* is not a valid shared memory 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)*).
- [EAGAIN] The system has temporarily exhausted its available memory or swap space.
- [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 `shm_perm.uid` in the data structure associated with *shmid*.
- [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)*, *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] | <i>Size</i> is less than the system-imposed minimum or greater than the system-imposed maximum.  |
| [EACCES] | A shared memory identifier exists for <i>key</i> but operation permission (see <i>intro(2)</i> ) as specified by the low-order 9 bits of <i>shmflg</i> would not be granted. |

- [EAGAIN] The system has temporarily exhausted its available memory or swap space.
- [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> ).          |
| [EAGAIN] | The system has temporarily exhausted its available memory or swap space.               |
| [ENOMEM] | The available data space is not large enough to accommodate the shared memory segment. |

- [EINVAL] *Shmaddr* is not equal to zero, and the value of (*shmaddr* - (*shmaddr* modulus SHMLBA)) is an illegal address.
- [EINVAL] *Shmaddr* is not equal to zero, (*shmflg* & SHM\_RND) is "false", and the value of *shmaddr* is an illegal address.
- [EMFILE] The number of shared memory segments attached to the calling process would exceed the system-imposed limit.
- [EINVAL] *Shmdt* detaches from the calling process's data segment the shared memory segment located at the address specified by *shmaddr*.
- [EINVAL] *Shmdt* will fail and not detach the shared memory segment if *shmaddr* 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;
cc ... -lnet
```

**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.

**LINKING**

This library is accessed by specifying `-lnet` as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**SEE ALSO**

connect(2N), socket(2N)

## NAME

signal — specify what to do upon receipt of a signal

## SYNOPSIS

```
#include <signal.h>
int (*signal (sig, func))()
int sig;
void (*func)();
```

## DESCRIPTION

*Signal* allows the calling process to choose one of three ways in which it is possible to handle the receipt of a specific signal. *Sig* specifies the signal and *func* specifies the choice.

*Sig* can be assigned any one of the following except SIGKILL:

SIGHUP	01	hangup
SIGINT	02	interrupt
SIGQUIT	03*	quit
SIGILL	04*	illegal instruction (not reset when caught)
SIGTRAP	05*	trace trap (not reset when caught)
SIGIOT	06*	IOT instruction
SIGEMT	07*	EMT instruction
SIGFPE	08*	floating point exception
SIGKILL	09	kill (cannot be caught 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
SIGUSR1	16	user defined signal 1
SIGUSR2	17	user defined signal 2
SIGCLD	18	death of a child (see <i>WARNING</i> below)
SIGPWR	19	power fail (see <i>WARNING</i> below)

See below for the significance of the asterisk ( \* ) in the above list.

*Func* is assigned one of three values: SIG\_DFL, SIG\_IGN, or a *function address*. The actions prescribed by these values are as follows:

**SIG\_DFL** — terminate process upon receipt of a signal

Upon receipt of the signal *sig*, the receiving process is to be terminated with the following consequences:

All of the receiving process's open file descriptors will be closed.

If the parent process of the receiving process is executing a *wait*, it will be notified of the termination of the receiving process and the terminating signal's number will be made available to the parent process; see *wait(2)*.

If the parent process of the receiving process is not executing a *wait*, the receiving process will be transformed into a zombie process (see *exit(2)* for definition of zombie process).

The parent process ID of each of the receiving process's existing child processes and zombie processes will be set to 1. This means the initialization process (see *intro(2)*) inherits each of these processes.

Each attached shared memory segment is detached and the value of `shm_nattach` in the data structure associated with its shared memory identifier is decremented by 1.

For each semaphore for which the receiving process has set a `semadj` value (see `semop(2)`), that `semadj` value is added to the `semval` of the specified semaphore.

If the process has a process, text, or data lock, an `unlock` is performed (see `plock(2)`).

An accounting record will be written on the accounting file if the system's accounting routine is enabled; see `acct(2)`.

If the receiving process's process ID, tty group ID, and process group ID are equal, the signal `SIGHUP` will be sent to all of the processes that have a process group ID equal to the process group ID of the receiving process.

A 'core image' will be made in the current working directory of the receiving process if `sig` is one for which an asterisk appears in the above list *and* the following conditions are met:

The effective user ID and the real user ID of the receiving process are equal.

An ordinary file named `core` exists and is writable or can be created. If the file must be created, it will have the following properties:

a mode of 0666 modified by the file creation mask (see `umask(2)`)

a file owner ID that is the same as the effective user ID of the receiving process

a file group ID that is the same as the effective group ID of the receiving process

**SIG\_IGN** — ignore signal

The signal `sig` is to be ignored.

Note: the signal `SIGKILL` cannot be ignored.

*function address* — catch signal

Upon receipt of the signal `sig`, the receiving process is to execute the signal-catching function pointed to by `func`. The signal number `sig` will be passed as the only argument to the signal-catching function. Additional arguments are passed to the signal-catching function for hardware-generated signals. Before entering the signal-catching function, the value of `func` for the caught signal will be set to `SIG_DFL` unless the signal is `SIGILL`, `SIGTRAP`, or `SIGPWR`.

Upon return from the signal-catching function, the receiving process will resume execution at the point it was interrupted.

When a signal that is to be caught occurs during a `read`, a `write`, an `open`, or an `ioctl` system call on a slow device (like a terminal; but not a file), during a `pause` system call, or during a `wait` system call that does not return immediately due to the existence of a previously stopped or zombie process, the signal-catching function will be executed and then the interrupted system call may return a `-1` to the



calling process with *errno* set to EINTR.

Note: The signal SIGKILL cannot be caught.

A call to *signal* cancels a pending signal *sig* except for a pending SIGKILL signal.

*Signal* will fail if *sig* is an illegal signal number, including SIGKILL. [EINVAL]

#### RETURN VALUE

Upon successful completion, *signal* returns the previous value of *func* for the specified signal *sig*. Otherwise, a value of  $-1$  is returned and *errno* is set to indicate the error.

#### SEE ALSO

kill(1), kill(2), pause(2), ptrace(2), wait(2), setjmp(3C).

#### WARNING

Two other signals that behave differently than the signals described above exist in this release of the system; they are:

SIGCLD	18	death of a child (reset when caught)
SIGPWR	19	power fail (not reset when caught)

There is no guarantee that, in future releases of the UNIX system, these signals will continue to behave as described below; they are included only for compatibility with other versions of the UNIX system. Their use in new programs is strongly discouraged.

For these signals, *func* is assigned one of three values: SIG\_DFL, SIG\_IGN, or a *function address*. The actions prescribed by these values of are as follows:

**SIG\_DFL** - ignore signal

The signal is to be ignored.

**SIG\_IGN** - ignore signal

The signal is to be ignored. Also, if *sig* is SIGCLD, the calling process's child processes will not create zombie processes when they terminate; see *exit(2)*.

*function address* - catch signal

If the signal is SIGPWR, the action to be taken is the same as that described above for *func* equal to *function address*. The same is true if the signal is SIGCLD except, that while the process is executing the signal-catching function, any received SIGCLD signals will be queued and the signal-catching function will be continually reentered until the queue is empty.

The SIGCLD affects two other system calls (*wait(2)*, and *exit(2)*) in the following ways:

*wait* If the *func* value of SIGCLD is set to SIG\_IGN and a *wait* is executed, the *wait* will block until all of the calling process's child processes terminate; it will then return a value of  $-1$  with *errno* set to ECHILD.

*exit* If in the exiting process's parent process the *func* value of SIGCLD is set to SIG\_IGN, the exiting process will not create a zombie process.

When processing a pipeline, the shell makes the last process in the pipeline the parent of the proceeding processes. A process that may be piped into in

this manner (and thus become the parent of other processes) should take care not to set SIGCLD to be caught.

**BUGS**

If a repeated signal arrives before the last one can be reset, there is no chance to catch it.

The type specification of the routine and its *func* argument are problematical.

The symbols *signd* and *sigtrap* are globally defined symbols used by *signal(2)* and are reserved words.

**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;

cc ... -lnet
```

**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), and
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(4N)* and *protocols(4N)*.

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(2N)* call. Once connected, data may be transferred using *read(2)* and *write(3)* calls or some variant of the *send(2N)* and *recv(2N)* 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(2N)* and received as described in *recv(2N)*.

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 an 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(2N)` calls. It is also possible to receive datagrams at such a socket with `recv(2N)`.

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(2N)` are used to set and get options, respectively.

<code>SO_DEBUG</code>	turn on recording of debugging information
<code>SO_REUSEADDR</code>	allow local address reuse
<code>SO_KEEPAIVE</code>	keep connections alive
<code>SO_DONTROUTE</code>	do no apply routing on outgoing messages
<code>SO_LINGER</code>	linger on close if data present
<code>SO_DONTLINGER</code>	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(2N)` call should allow reuse of local addresses. `SO_KEEPAIVE` 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.

#### LINKING

This library is accessed by specifying `-lnet` as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

#### SEE ALSO

`accept(2N)`, `bind(2N)`, `connect(2N)`, `getsockname(2N)`, `getsockopt(2N)`, `ioctl(2)`, `listen(2N)`, `recv(2N)`, `select(2N)`, `send(2N)`, `shutdown(2N)`

#### BUGS

The use of `keepalives` is a questionable feature for this layer.

## NAME

stat, fstat — get file status

## SYNOPSIS

```
#include <sys/types.h>
#include <sys/stat.h>

int stat (path, buf)
char *path;
struct stat *buf;

int fstat (fildes, buf)
int fildes;
struct stat *buf;
```

## DESCRIPTION

*Path* points to a path name naming a file. 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 searchable. *Stat* obtains information about the named file.

Similarly, *fstat* obtains information about an open file known by the file descriptor *fildes*, obtained from a successful *open*, *creat*, *dup*, *fcntl*, or *pipe* system 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* include the following members:

```
ushort  st_mode;    /* File mode; see mknod(2) */
ino_t   st_ino;    /* Inode number */
dev_t   st_dev;    /* ID of device containing */
          /* a directory entry for this file */
dev_t   st_rdev;   /* ID of device */
          /* This entry is defined only for */
          /* character special or block special files */
short   st_nlink;  /* Number of links */
ushort  st_uid;    /* User ID of the file's owner */
ushort  st_gid;    /* Group ID of the file's group */
off_t   st_size;   /* File size in bytes */
time_t  st_atime;  /* Time of last access */
time_t  st_mtime;  /* Time of last data modification */
time_t  st_ctime;  /* Time of last file status change */
          /* Times measured in seconds since */
          /* 00:00:00 GMT, Jan. 1, 1970 */
```

**st\_atime** Time when file data was last accessed. Changed by the following system calls: *creat*(2), *mknod*(2), *pipe*(2), *utime*(2), and *read*(2).

**st\_mtime** Time when data was last modified. Changed by the following system calls: *creat*(2), *mknod*(2), *pipe*(2), *utime*(2), and *write*(3).

**st\_ctime** Time when file status was last changed. Changed by the following system calls: *chmod*(2), *chown*(2), *creat*(2), *link*(2), *mknod*(2), *pipe*(2), *unlink*(2), *utime*(2), and *write*(3).

*Stat* will fail if one or more of the following are true:

- [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.
- [EFAULT] *Buf* or *path* points to an invalid address.

*Fstat* will fail if one or more of the following are true:

- [EBADF] *Fildes* is not a valid open file descriptor.
- [EFAULT] *Buf* points to an invalid 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

*chmod(2)*, *chown(2)*, *creat(2)*, *link(2)*, *mknod(2)*, *pipe(2)*, *read(2)*, *time(2)*, *unlink(2)*, *utime(2)*, *write(3)*.

**STIME(2)****STIME(2)****NAME**

*stime* — set time

**SYNOPSIS**

```
int stime (tp)
long *tp;
```

**DESCRIPTION**

*Stime* sets the system's idea of the time and date. *Tp* points to the value of time as measured in seconds from 00:00:00 GMT January 1, 1970.

[EPERM] *Stime* will fail if the effective user ID of the calling process is not 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.

**SEE ALSO**

*time*(2).



**NAME**

sync - update super-block

**SYNOPSIS**

void sync ( )

**DESCRIPTION**

*Sync* causes all information in memory that should be on disk to be written out. This includes modified super blocks, modified i-nodes, and delayed block I/O.

It should be used by programs which examine a file system, for example *fsck*, *df*, etc. It is mandatory before a boot.

The writing, although scheduled, is not necessarily complete upon return from *sync*.

## TIME(2)

## TIME(2)

### NAME

*time* – get time

### SYNOPSIS

`long time ((long *) 0)`

`long time (tloc)`

`long *tloc;`

### DESCRIPTION

*Time* returns the value of time in seconds since 00:00:00 GMT, January 1, 1970.

If *tloc* (taken as an integer) is non-zero, the return value is also stored in the location to which *tloc* points.

[EFAULT] *Time* will fail if *tloc* points to an illegal address.

### RETURN VALUE

Upon successful completion, *time* returns the value of time. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

### SEE ALSO

*stime(2)*, *ctime(3)*.

**NAME**

*times* – get process and child process times

**SYNOPSIS**

```
#include <sys/types.h>
```

```
#include <sys/times.h>
```

```
long times (buffer)
```

```
struct tms *buffer;
```

**DESCRIPTION**

*Times* fills the structure pointed to by *buffer* with time-accounting information. The following are the contents of this structure:

```
struct tms {
    time_t tms_utime;
    time_t tms_stime;
    time_t tms_cutime;
    time_t tms_cstime;
};
```

This information comes from the calling process and each of its terminated child processes for which it has executed a *wait*. All times are in 60ths of a second.

*Tms\_utime* is the CPU time used while executing instructions in the user space of the calling process.

*Tms\_stime* is the CPU time used by the system on behalf of the calling process.

*Tms\_cutime* is the sum of the *tms\_utimes* and *tms\_cutimes* of the child processes.

*Tms\_cstime* is the sum of the *tms\_stimes* and *tms\_cstimes* of the child processes.

[EFAULT] *Times* will fail if *buffer* points to an illegal address.

**RETURN VALUE**

Upon successful completion, *times* returns the elapsed real time, in 60ths of a second, since an arbitrary point in the past (e.g., system start-up time). This point does not change from one invocation of *times* to another. If *times* fails, a *-1* is returned and *errno* is set to indicate the error.

**SEE ALSO**

*exec(2)*, *fork(2)*, *time(2)*, *wait(2)*.

**NAME**

ulimit — get and set user limits

**SYNOPSIS**

```
long ulimit (cmd, newlimit)
int cmd;
long newlimit;
```

**DESCRIPTION**

This function provides for control over process limits. The *cmd* values available are:

- 1 Get the file size limit of the process. The limit is in units of 512-byte blocks and is inherited by child processes. Files of any size can be read.
- 2 Set the file size limit of the process to the value of *newlimit*. Any process may decrease this limit, but only a process with an effective user ID of super-user may increase the limit. *Ulimit* will fail and the limit will be unchanged if a process with an effective user ID other than super-user attempts to increase its file size limit. [EPERM]
- 3 Get the maximum possible break value. See *brk(2)*.

**RETURN VALUE**

Upon successful completion, a non-negative value is returned. Otherwise, a value of *-1* is returned and *errno* is set to indicate the error.

**SEE ALSO**

*brk(2)*, *write(3)*.

**NAME**

umask -- set and get file creation mask

**SYNOPSIS**

```
int umask (cmask)
int cmask;
```

**DESCRIPTION**

*Umask* sets the process's file mode creation mask to *cmask* and returns the previous value of the mask. Only the low-order 9 bits of *cmask* and the file mode creation mask are used.

The file mode creation mask is used whenever a file is created by *creat(2)*, *mknod(2)* or *open(2)*. The actual mode (see *chmod(2)*) of the newly-created file is the difference between the given mode and *cmask*. In other words, *cmask* shows the bits to be turned off when a new file is created.

The previous value of *cmask* is returned by the call. The value is initially 022, which is an octal 'mask' number representing the complement of the desired mode. '022' here means that no permissions are withheld from the owner, but write permission is forbidden to group and to others. Its complement, the mode of the file, would be 755. The file mode creation mask is inherited by child processes.

**RETURN VALUE**

The previous value of the file mode creation mask is returned.

**SEE ALSO**

*mkdir(1)*, *sh(1)*, *chmod(2)*, *creat(2)*, *mknod(2)*, *open(2)*.

**NAME**

umount – unmount a file system

**SYNOPSIS**

```
int umount (spec)
char *spec;
```

**DESCRIPTION**

*Umount* requests that a previously mounted file system contained on the block special device identified by *spec* be unmounted. *Spec* is a pointer to a path name. After unmounting the file system, the directory upon which the file system was mounted reverts to its ordinary interpretation.

*Umount* may be invoked only by the super-user.

*Umount* will fail if one or more of the following are true:

- |           |  |
|-----------|--|
| [EPERM]   | The process's effective user ID is not super-user. |
| [ENXIO]   | <i>Spec</i> device does not exist.                 |
| [ENOTBLK] | <i>Spec</i> is not a block special device.         |
| [EINVAL]  | <i>Spec</i> is not mounted.                        |
| [EBUSY]   | A file on <i>spec</i> is busy.                     |
| [ENOENT]  | No such <i>spec</i> file or 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.

**SEE ALSO**

mount(2).

**NAME**

uname — get name of current UNIX system

**SYNOPSIS**

```
#include <sys/utsname.h>
int uname (name)
struct utsname *name;
```

**DESCRIPTION**

*Uname* stores information identifying the current UNIX system in the structure pointed to by *name*.

*Uname* uses the structure defined in <sys/utsname.h>:

```
struct utsname {
    char    sysname[9];
    char    nodename[9];
    char    release[9];
    char    version[9];
    char    machine[9];
};
extern struct utsname utsname;
```

*Uname* returns a null-terminated character string naming the current UNIX system in the character array *sysname*. Similarly, *nodename* contains the name that the system is known by on a communications network. *Release* and *version* further identify the operating system. *Machine* contains a standard name that identifies the hardware that the UNIX system is running on.

[DEFAULT] *Uname* will fail if *name* points to an invalid address.

**RETURN VALUE**

Upon successful completion, a non-negative value is returned. Otherwise, -1 is returned and *errno* is set to indicate the error.

**SEE ALSO**

uname(1).

**NAME**

unlink — remove directory entry

**SYNOPSIS**

```
int unlink (path)
char *path;
```

**DESCRIPTION**

*Unlink* removes the directory entry named by the path name pointed to be *path*.

The named file is unlinked unless one or more of the following are true:

- [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 super-user.
- [EBUSY] The entry to be unlinked is the mount point for a mounted file system.
- [ETXTBSY] The entry to be unlinked is the last link to a pure procedure (shared text) file that is being executed.
- [EROFS] The directory entry to be unlinked is part of a read-only file system.
- [EFAULT] *Path* points outside the process's allocated address space.

When all links to a file have been removed and no process has the file open, the space occupied by the file is freed and the file ceases to exist. If one or more processes have the file open when the last link is removed, the removal is postponed until all references to the file have been closed.

**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**

rm(1), close(2), link(2), open(2).



**NAME**

ustat — get file system statistics

**SYNOPSIS**

```
#include <sys/types.h>
#include <ustat.h>
```

```
int ustat (dev, buf)
int dev;
struct ustat *buf;
```

**DESCRIPTION**

*Ustat* returns information about a mounted file system. *Dev* is a device number identifying a device containing a mounted file system. *Buf* is a pointer to a *ustat* structure that includes the following elements:

```
daddr_t f_tfree;      /* Total free blocks */
ino_t   f_tinode;    /* Number of free inodes */
char    f_fname[6];  /* Filsys name */
char    f_fpack[6];  /* Filsys pack name */
```

*Ustat* will fail if one or more of the following are true:

- [EINVAL] *Dev* is not the device number of a device containing a mounted file system.
- [EFAULT] *Buf* points outside the process's allocated address space.

**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**

stat(2), fs(4).

## NAME

utime — set file access and modification times

## SYNOPSIS

```
#include <sys/types.h>
int utime (path, times)
char *path;
struct utimbuf *times;
```

## DESCRIPTION

*Path* points to a path name naming a file. *Utime* sets the access and modification times of the named file.

If *times* is NULL, the access and modification times of the file are set to the current time. A process must be the owner of the file or have write permission to use *utime* in this manner.

If *times* is not NULL, *times* is interpreted as a pointer to a *utimbuf* structure and the access and modification times are set to the values contained in the designated structure. Only the owner of the file or the super-user may use *utime* this way.

The times in the following structure are measured in seconds since 00:00:00 GMT, Jan. 1, 1970.

```
struct utimbuf {
    time_t actime;    /* access time */
    time_t modtime;  /* modification time */
};
```

*Utime* will fail if one or more of the following are true:

- [ENOENT] The named file does not exist.
- [ENOTDIR] A component of the path prefix is not a directory.
- [EACCES] Search permission is denied by a component of the path prefix.
- [EPERM] The effective user ID is not super-user and not the owner of the file and *times* is not NULL.
- [EACCES] The effective user ID is not super-user and not the owner of the file and *times* is NULL and write access is denied.
- [EROFS] The file system containing the file is mounted read-only.
- [EFAULT] *Times* is not NULL and points outside the process's allocated address space.
- [EFAULT] *Path* points outside the process's allocated address space.

## 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

stat(2).

**NAME**

uvar — returns system-specific configuration information

**SYNOPSIS**

```
#include <sys/var.h>
```

```
uvar(v)
struct var *v;
```

**DESCRIPTION**

*Uvar* returns system-specific configuration information contained in the kernel. The information returned contains table sizes, mask words, and other system-specific information for programs such as *ld(1)* and *ps(1)*.

Presently a maximum of 256 bytes of information is returned. This number is subject to change. *v* points to the *var* structure:

```
struct var {
    int     v_buf;           /* Number of system buffers */
    int     v_call;         /* Maximum number of simultaneous callouts */
    int     v_inode;        /* Maximum number of incore inodes */
    char *  ve_inode;       /* Pointer to last incore inode table */
    int     v_file;         /* Maximum number of open files */
    char *  ve_file;        /* Pointer to last open file table */
    int     v_mount;        /* Maximum number of file systems mountable */
    char *  ve_mount;       /* Pointer to last mounted file system table */
    int     v_proc;         /* Maximum number of processes */
    char *  ve_proc;        /* Pointer to last process table */
    int     v_text;         /* Maximum number of shared text segments */
    char *  ve_text;        /* Pointer to last shared text segment table */
    int     v_clist;        /* Maximum number of clists */
    int     v_sabuf;        /* Maximum number of system activity buffers */
    int     v_maxup;        /* Maximum number of user processes */
    int     v_cmap;         /* Size of core memory allocation map */
    int     v_smap;         /* Size of swap memory allocation map */
    int     v_hbuf;         /* Maximum number of buffer headers */
    int     v_hmask;        /* Maximum number of buffer headers - 1 */
    int     v_flock;        /* Maximum number of file locks */
    int     v_phys;         /* Maximum number of simultaneous phys calls */
    int     v_qlsize;       /* Click size */
    int     v_txrnd;        /* Number of clicks per segment */
    int     v_bsize;        /* Block size */
    int     v_cxmap;        /* Context map size */
    int     v_clktick;      /* Clock tick */
    int     v_hz;           /* Hz */
    int     v_usize;        /* Size of user structure */
    int     v_pageshift;    /* Page shift */
    int     v_pagemask;     /* Page mask */
    int     v_segshift;     /* Segment shift */
    int     v_segmask;      /* Segment mask */
    int     v_ustart;       /* Starting virtual address for user program */
    int     v_uend;         /* Ending virtual address for user program */
    char *  ve_call;        /* Pointer to last callout table */
    int     v_stkgap;       /* Obsolete */
};
```

UVAR (2)

UniSoft

UVAR (2)

```
int      v_cputype;      /* CPU type (1=68000) */
int      v_cpuver;      /* CPU version id (1=68000, 2=68010, 3=68020) */
int      v_mmutype;     /* MMU type (1=none, 2=SUN, 3=68451) */
int      v_doffset;     /* Data offset */
int      v_kvoffset;    /* Kernel virtual offset */
int      v_svttext;     /* Maximum number of text loitering segments */
char *   ve_svttext;    /* Pointer to last text loitering segment in table */
int      v_pbuf;        /* Maximum number of buffers for physio */
int      v_nscatload;   /* Maximum number of entries in scatter map */
int      v_udot;        /* Address of user structure */
int      v_fill[64-46]; /* Sized to make var 256 bytes long */
};
```

SEE ALSO

/usr/include/sys/space.h

**NAME**

`wait` - wait for child process to stop or terminate

**SYNOPSIS**

```
int wait (stat_loc)
int *stat_loc;
int wait ((int *)0)
```

**DESCRIPTION**

*Wait* suspends the calling process until one of the immediate children terminates or until a child that is being traced stops, because it has hit a break point. The *wait* system call will return prematurely if a signal is received and if a child process stopped or terminated prior to the call on *wait*, return is immediate.

If *stat\_loc* (taken as an integer) is non-zero, 16 bits of information called status are stored in the low order 16 bits of the location pointed to by *stat\_loc*. *Status* can be used to differentiate between stopped and terminated child processes and if the child process terminated, status identifies the cause of termination and passes useful information to the parent. This is accomplished in the following manner:

If the child process stopped, the high order 8 bits of status will contain the number of the signal that caused the process to stop and the low order 8 bits will be set equal to 0177.

If the child process terminated due to an *exit* call, the low order 8 bits of status will be zero and the high order 8 bits will contain the low order 8 bits of the argument that the child process passed to *exit*; see *exit(2)*.

If the child process terminated due to a signal, the high order 8 bits of status will be zero and the low order 8 bits will contain the number of the signal that caused the termination. In addition, if the low order seventh bit (i.e., bit 200) is set, a "core image" will have been produced; see *signal(2)*.

If a parent process terminates without waiting for its child processes to terminate, the parent process ID of each child process is set to 1. This means the initialization process inherits the child processes; see *intro(2)*.

*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] *Stat\_loc* points to an illegal address.

**RETURN VALUE**

If *wait* 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 *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.

**SEE ALSO**

*exec(2)*, *exit(2)*, *fork(2)*, *intro(2)*, *pause(2)*, *ptrace(2)*, *signal(2)*.

**WAIT(2)**

**WAIT(2)**

**WARNING**

See *WARNING* in *signal(2)*.

**NAME**

*wait3* – wait for child process to stop or terminate

**SYNOPSIS**

```
#include <sys/wait.h>

pid = wait3(status, options, 0)
int pid;
union wait *status;
int options;
```

**DESCRIPTION**

*Wait3* provides an 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 stats (WNOHANG).

When the WNOHANG option is specified and no processes wish to report status, *wait3* returns a *pid* of 0.

**RETURN VALUE**

*Wait3* 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.

**SEE ALSO**

*exit(2)*

## NAME

write – write on a file

## SYNOPSIS

```
int write (fildes, buf, nbyte)
int fildes;
char *buf;
unsigned nbyte;
```

## DESCRIPTION

*Fildes* is a file descriptor obtained from a *creat*, *open*, *dup*, *fcntl*, *pipe*, or *socket* system call.

*Write* attempts to write *nbyte* bytes from the buffer pointed to by *buf* to the file associated with the *fildes*.

On devices capable of seeking, the actual writing of data proceeds from the position in the file indicated by the file pointer. Upon return from *write*, the file pointer is incremented by the number of bytes actually written.

On devices incapable of seeking, writing always takes place starting at the current position. The value of a file pointer associated with such a device is undefined.

If the *O\_APPEND* flag of the file status flags is set, the file pointer will be set to the end of the file prior to each write.

*Write* will fail and the file pointer will remain unchanged if one or more of the following are true:

- [EIO]                   A physical I/O error has occurred.
- [ENXIO]                The device associated with the file descriptor is a block-special or character-special file and the value of the file pointer is out of range.
- [EBADF]                *Fildes* is not a valid file descriptor open for writing.
- [EPIPE and SIGPIPE signal]
  - 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 pipe that is not open for reading by any process.
- [EFBIG]                An attempt was made to write a file that exceeds the process's file size limit or the maximum file size. See



*ulimit(2)*.

- [EFAULT] Part of *iov* or data to be written to the file points outside the process's allocated address space.
- [EFAULT] *Buf* points outside the process's allocated address space.
- [EINTR] A signal was caught during the *write* system call.
- [ENOSPC] Not enough space is left on the device containing the file.

If a *write* requests that more bytes be written than there is room for (e.g., the *ulimit* (see *ulimit(2)*) or the physical end of a medium), only as many bytes as there is room for will be written. For example, suppose there is space for 20 bytes more in a file before reaching a limit. A write of 512 bytes will return 20. The next write of a non-zero number of bytes will give a failure return (except as noted below).

If the file being written is a pipe (or FIFO) and the `O_NDELAY` flag of the file flag word is set, then write to a full pipe (or FIFO) will return a count of 0. Otherwise (`O_NDELAY` clear), writes to a full pipe (or FIFO) will block until space becomes available.

#### RETURN VALUE

Upon successful completion the number of bytes actually written is returned. Otherwise, `-1` is returned and *errno* is set to indicate the error.

#### SEE ALSO

*creat(2)*, *lseek(2)*, *open(2)*, *pipe(2)*, *socket(2N)*, *ulimit(2)*.

**NAME**

*writev* – write on a file

**SYNOPSIS**

```
#include <sys/types.h>
#include <sys/uio.h>
```

```
writev(d, iov, iovectlen)
int d;
struct iovec *iov;
int iovectlen;
```

**DESCRIPTION**

*Fildes* is a file descriptor obtained from a *creat*, *open*, *dup*, *fcntl*, *pipe*, or *socket* system call.

*Writev* attempts to write *nbyte* bytes to the file associated with the *fildes* and gathers the output data from the *iovlen* buffers specified by the members of the *iovec* array: *iov*[0], *iov*[1], etc.

On devices capable of seeking, the actual writing of data proceeds from the position in the file indicated by the file pointer. Upon return from *writev*, the file pointer is incremented by the number of bytes actually written.

On devices incapable of seeking, writing always takes place starting at the current position. The value of a file pointer associated with such a device is undefined.

If the *O\_APPEND* flag of the file status flags is set, the file pointer will be set to the end of the file prior to each write.

*Writev* will fail and the file pointer will remain unchanged if one or more of the following are true:

[EBADF] *Fildes* is not a valid file descriptor open for writing.

[EPIPE and SIGPIPE signal]

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 pipe that is not open for reading by any process.

[EFBIG]

An attempt was made to write a file that exceeds the process's file size limit or the maximum file size. See *ulimit*(2).

[EFAULT] Part of *iov* or data to be written to the file points outside the process's allocated address space.

[EFAULT] *Buf* points outside the process's allocated address space.

[EINTR] A signal was caught during the *writv* system call.

If a *writv* requests that more bytes be written than there is room for (e.g., the *ulimit* (see *ulimit(2)*) or the physical end of a medium), only as many bytes as there is room for will be written. For example, suppose there is space for 20 bytes more in a file before reaching a limit. A write of 512 bytes will return 20. The next write of a non-zero number of bytes will give a failure return (except as noted below).

If the file being written is a pipe (or FIFO) and the O\_NDELAY flag of the file flag word is set, then write to a full pipe (or FIFO) will return a count of 0. Otherwise (O\_NDELAY clear), writes to a full pipe (or FIFO) will block until space becomes available.

#### RETURN VALUE

Upon successful completion the number of bytes actually written is returned. Otherwise, -1 is returned and *errno* is set to indicate the error.

#### SEE ALSO

*creat(2)*, *lseek(2)*, *open(2)*, *pipe(2)*, *socket(2N)*, *ulimit(2)*.



**NAME**

intro -- introduction to subroutines and libraries

**SYNOPSIS**

```
#include <stdio.h>
```

```
#include <math.h>
```

**DESCRIPTION**

This section describes functions found in various libraries, other than those functions that directly invoke system primitives, which are described in Section 2 of this volume. Certain major collections are identified by a letter after the section number:

- (3C) These functions, together with those of Section 2 and those marked (3S), constitute the Standard C Library, *libc*, which is automatically loaded by the C compiler, *cc(1)*. The link editor *ld(1)* searches this library under the *-lc* option. Some functions require declarations that can be included in the program being compiled by adding the line

```
#include <header filename>
```

The appropriate *#include* file is indicated in the SYNOPSIS part of a function description.

- (3F) These functions constitute the FORTRAN intrinsic function library, *libF77*. These functions are automatically available to the FORTRAN programmer and require no special invocation of the compiler.
- (3M) These functions constitute the Math Library, *libm*. They are automatically loaded as needed by the FORTRAN compiler *f77(1)*. They are not automatically loaded by the C compiler, *cc(1)*; however, the link editor searches this library under the *-lm* option. Declarations for these functions may be obtained from the *#include* file *<math.h>*.
- (3S) These functions constitute the "standard I/O package"; an introduction to this package is provided in *stdio(3S)*. The functions are in the library *libc*, already mentioned. Declarations should be obtained from the *#include* file *<stdio.h>*.
- (3X) Various specialized libraries. The files in which these libraries are found are given on the appropriate pages.

For descriptions and examples of *#include* files, refer to the "Libraries" section of the Programming Guide.

**DEFINITIONS**

A *character* is any bit pattern able to fit into a byte on the machine. The *null character* is a character with value 0, represented in the C language as *\0*. A *character array* is a sequence of characters. A *null-terminated character array* is a sequence of characters, the last of which is the *null character*. A *string* is a designation for a *null-terminated character array*. The *null string* is a character array containing only the null character. A *NULL pointer* is the value that is obtained by casting 0 into a pointer. The C language guarantees that this value will not match that of any legitimate pointer, so many functions that return pointers return it to indicate an error. *NULL* is defined as 0 in *<stdio.h>*; the user can include his own definition if he is not using *<stdio.h>*.

Many groups of FORTRAN intrinsic functions have *generic* function names that do not require explicit or implicit type declaration. The type of the function is determined by the type of its argument(s). For example, the generic function *max* returns an integer value if given integer arguments (*max0*), a real value if given real arguments (*amax1*), or a double-precision value if given double-precision arguments (*dmax1*).

**FILES**

/lib/libc.a  
/usr/lib/libF77.a  
/lib/libm.a

**SEE ALSO**

ar(1), cc(1), f77(1), ld(1), lint(1), nm(1), intro(2), stdio(3S), math(5).  
*Programming Guide.*

**DIAGNOSTICS**

Functions in the C and Math Libraries (3C and 3M) may return the conventional values 0 or  $\pm$ HUGE (the largest-magnitude single-precision floating-point numbers; HUGE is defined in the *<math.h>* header file) when the function is undefined for the given arguments or when the value is not representable. In these cases, the external variable *errno* (see *intro(2)*) is set to the value EDOM or ERANGE. Because many of the FORTRAN intrinsic functions use the routines found in the Math Library, the same conventions apply.

**WARNING**

Many of the functions in the libraries call and/or refer to other functions and external variables described in this section and in section 2 (*System Calls*). If a program inadvertently defines a function or external variable with the same name, the presumed library version of the function or external variable may not be loaded. The *lint(1)* program checker reports name conflicts of this kind as "multiple declarations" of the names in question. Definitions for sections 2, 3C, and 3S are checked automatically. Other definitions can be included by using the *-l* option (for example, *-lm* includes definitions for the Math Library, section 3M). Use of *lint* is highly recommended.

**NAME**

a64l, l64a — convert between long integer and base-64 ASCII string

**SYNOPSIS**

```
long a64l (s)
char *s;
char *l64a (l)
long l;
```

**DESCRIPTION**

These functions are used to maintain numbers stored in *base-64* ASCII characters. This is a notation by which long integers can be represented by up to 6 characters; each character represents a "digit" in a radix-64 notation.

The characters used to represent "digits" are . for 0, / for 1, 0 through 9 for 2-11, A through Z for 12-37, and a through z for 38-63.

*A64l* takes a pointer to a null-terminated base-64 representation and returns a corresponding long value. If the string pointed to by *s* contains more than 6 characters, *a64l* uses the first 6.

*L64a* takes a long argument and returns a pointer to the corresponding base-64 representation. If the argument is 0, *l64a* returns a pointer to a null string.

**BUGS**

The value returned by *l64a* is a pointer into a static buffer, the contents of which are overwritten by each call.

**NAME**

abort — generate an IOT fault

**SYNOPSIS**

int abort ( )

**DESCRIPTION**

*Abort* first closes all open files if possible, then causes an IOT signal to be sent to the process. This usually results in termination with a core dump.

It is possible for *abort* to return control if SIGIOT is caught or ignored, in which case the value returned is that of the *kill(2)* system call.

**SEE ALSO**

sdb(1), exit(2), kill(2), signal(2).

**DIAGNOSTICS**

If SIGIOT is neither caught nor ignored, and the current directory is writable, a core dump is produced and the message **abort - core dumped** is written by the shell.



**NAME**

abort — terminate Fortran program

**SYNOPSIS**

call abort ( )

**DESCRIPTION**

*Abort* terminates the program which calls it, closing all open files truncated to the current position of the file pointer.

**DIAGNOSTICS**

When invoked, *abort* prints Fortran abort routine called on the standard error output.

**SEE ALSO**

abort(3C).

**NAME**

*abs* — return integer absolute value

**SYNOPSIS**

```
int abs (i)
int i;
```

**DESCRIPTION**

*Abs* returns the absolute value of its integer operand.

**BUGS**

In two's-complement representation, the absolute value of the negative integer with largest magnitude is undefined. Some implementations trap this error, but others simply ignore it.

**SEE ALSO**

*floor*(3M).

**NAME**

*abs*, *iabs*, *dabs*, *cabs*, *zabs* — Fortran absolute value

**SYNOPSIS**

integer *i1*, *i2*  
real *r1*, *r2*  
double precision *dp1*, *dp2*  
complex *cx1*, *cx2*  
double complex *dx1*, *dx2*  
  
*r2* = *abs*(*r1*)  
*i2* = *iabs*(*i1*)  
*i2* = *abs*(*i1*)  
  
*dp2* = *dabs*(*dp1*)  
*dp2* = *abs*(*dp1*)  
  
*cx2* = *cabs*(*cx1*)  
*cx2* = *abs*(*cx1*)  
  
*dx2* = *zabs*(*dx1*)  
*dx2* = *abs*(*dx1*)

**DESCRIPTION**

*Abs* is the family of absolute value functions. *Iabs* returns the integer absolute value of its integer argument. *Dabs* returns the double-precision absolute value of its double-precision argument. *Cabs* returns the complex absolute value of its complex argument. *Zabs* returns the double-complex absolute value of its double-complex argument. The generic form *abs* returns the type of its argument.

**SEE ALSO**

*floor*(3M).

**NAME**

*acos*, *dacos* — Fortran arccosine intrinsic function

**SYNOPSIS**

```
real r1, r2
double precision dp1, dp2
r2 = acos(r1)
dp2 = dacos(dp1)
dp2 = acos(dp1)
```

**DESCRIPTION**

*Acos* returns the real arccosine of its real argument. *Dacos* returns the double-precision arccosine of its double-precision argument. The generic form *acos* may be used with impunity because its argument determines the type of the returned value.

**SEE ALSO**

*trig*(3M).

**NAME**

aimag, dimag — Fortran imaginary part of complex argument

**SYNOPSIS**

real r  
complex cxr  
double precision dp  
double complex cxd  
r = aimag(cxr)  
dp = dimag(cxd)

**DESCRIPTION**

*Aimag* returns the imaginary part of its single-precision complex argument.  
*Dimag* returns the double-precision imaginary part of its double-complex argument.

**NAME**

*aint*, *dint* — Fortran integer part intrinsic function

**SYNOPSIS**

```
real r1, r2
double precision dp1, dp2
r2 = aint(r1)
dp2 = dint(dp1)
dp2 = aint(dp1)
```

**DESCRIPTION**

*Aint* returns the truncated value of its real argument in a real. *Dint* returns the truncated value of its double-precision argument as a double-precision value. *Aint* may be used as a generic function name, returning either a real or double-precision value depending on the type of its argument.

**NAME**

asin, dasin — Fortran arcsine intrinsic function

**SYNOPSIS**

```
real r1, r2
double precision dp1, dp2
r2 = asin(r1)
dp2 = dasin(dp1)
dp2 = asin(dp1)
```

**DESCRIPTION**

*Asin* returns the real arcsine of its real argument. *Dasin* returns the double-precision arcsine of its double-precision argument. The generic form *asin* may be used with impunity as it derives its type from that of its argument.

**SEE ALSO**

trig(3M).

**NAME**

assert — verify program assertion

**SYNOPSIS**

```
#include <assert.h>
assert (expression)
int expression;
```

**DESCRIPTION**

This macro is useful for putting diagnostics into programs. If *expression* is false (zero) when *assert* is executed, *assert* prints

*Assertion failed: expression, file xyz, line nnn*

on the standard error output and aborts. In the error message, *xyz* is the name of the source file and *nnn* is the source line number of the *assert* statement.

Compiling with the preprocessor option `-DNDEBUG` (see *cpp(1)*), or with the preprocessor control statement `#define NDEBUG` ahead of the `#include <assert.h>` statement, stops assertions from being compiled into the program.

**SEE ALSO**

*cpp(1)*, *abort(3C)*.



**NAME**

atan, datan — Fortran arctangent intrinsic function

**SYNOPSIS**

```
real r1, r2
double precision dp1, dp2
r2 = atan(r1)
dp2 = datan(dp1)
dp2 = atan(dp1)
```

**DESCRIPTION**

*Atan* returns the real arctangent of its real argument. *Datan* returns the double-precision arctangent of its double-precision argument. The generic form *atan* may be used with a double-precision argument returning a double-precision value.

**SEE ALSO**

trig(3M).

**NAME**

atan2, datan2 -- Fortran arctangent intrinsic function

**SYNOPSIS**

```
real r1, r2, r3
double precision dp1, dp2, dp3
r3 = atan2(r1, r2)
dp3 = datan2(dp1, dp2)
dp3 = atan2(dp1, dp2)
```

**DESCRIPTION**

*Atan2* returns the arctangent of *arg1/arg2* as a real value. *Datan2* returns the double-precision arctangent of its double-precision arguments. The generic form *atan2* may be used with impunity with double-precision arguments.

**SEE ALSO**

trig(3M).

**NAME**

`atof` — convert ASCII string to floating-point number

**SYNOPSIS**

```
double atof (nptr)
char *nptr;
```

**DESCRIPTION**

*Atof* converts a character string pointed to by *nptr* to a double-precision floating-point number. The first unrecognized character ends the conversion. *Atof* recognizes an optional string of white-space characters (blanks or tabs), then an optional sign, then a string of digits optionally containing a decimal point, then an optional *e* or *E* followed by an optionally signed integer. If the string begins with an unrecognized character, *atof* returns the value zero.

**DIAGNOSTICS**

When the correct value would overflow, *atof* returns **HUGE**, and sets *errno* to **ERANGE**. Zero is returned on underflow.

**SEE ALSO**

`scanf(3S)`, `strtol(3C)`.

## NAME

$j_0$ ,  $j_1$ ,  $j_n$ ,  $y_0$ ,  $y_1$ ,  $y_n$  — Bessel functions

## SYNOPSIS

```
#include <math.h>

double j0 (x)
double x;

double j1 (x)
double x;

double jn (n, x)
int n;
double x;

double y0 (x)
double x;

double y1 (x)
double x;

double yn (n, x)
int n;
double x;
```

## DESCRIPTION

$j_0$  and  $j_1$  return Bessel functions of  $x$  of the first kind of orders 0 and 1 respectively.  $j_n$  returns the Bessel function of  $x$  of the first kind of order  $n$ .

$y_0$  and  $y_1$  return the Bessel functions of  $x$  of the second kind of orders 0 and 1 respectively.  $y_n$  returns the Bessel function of  $x$  of the second kind of order  $n$ . The value of  $x$  must be positive.

## DIAGNOSTICS

Non-positive arguments cause  $y_0$ ,  $y_1$ , and  $y_n$  to return the value -HUGE and to set *errno* to EDOM. In addition, a message indicating DOMAIN error is printed on the standard error output.

Arguments too large in magnitude cause  $j_0$ ,  $j_1$ ,  $y_0$  and  $y_1$  to return zero and to set *errno* to ERANGE. In addition, a message indicating TLOSS error is printed on the standard error output.

These error-handling procedures may be changed with the function *matherr*(3M).

## SEE ALSO

*matherr*(3M).

**NAME**

blt, blt512 — block transfer data

**SYNOPSIS**

```
int blt(to,from,count)
char *to;
char *from;
int count;
```

```
int blt512(to,from,count)
char *to;
char *from;
int count;
```

**DESCRIPTION**

*Blt* does a fast copy of *count* bytes of data starting at address *from* to address *to*.

*Blt512* does a fast copy of *count* number of consecutive 512 byte units starting at address *from* to address *to*.

**NAME**

*and*, *or*, *xor*, *not*, *lshift*, *rshift* — Fortran bitwise boolean functions

**SYNOPSIS**

```
integer i, j, k
real a, b, c
double precision dp1, dp2, dp3

k = and(i, j)
c = or(a, b)
j = xor(i, a)
j = not(i)
k = lshift(i, j)
k = rshift(i, j)
```

**DESCRIPTION**

The generic intrinsic boolean functions *and*, *or*, and *xor* return the value of the binary operations on their arguments. *Not* is a unary operator returning the one's complement of its argument. *Lshift* and *rshift* return the value of the first argument shifted left or right, respectively, the number of times specified by the second (integer) argument.

The boolean functions are generic, i.e., defined for all data types as arguments and return values. Where required, the compiler generates appropriate type conversions.

**NOTE**

Although defined for all data types, use of boolean functions on non-integer data is not productive.

**BUGS**

The implementation of the shift functions may cause large shift values to deliver unexpected results.

**NAME**

bsearch — binary search a sorted table

**SYNOPSIS**

```
#include <search.h>

char *bsearch ((char *) key, (char *) base, nel, width, compar)
unsigned nel; width,
int (*compar) ( );
```

**DESCRIPTION**

*Bsearch* is a binary search routine generalized from Knuth (6.2.1) Algorithm B. It returns a pointer into a table indicating where a datum may be found. The table must be previously sorted in increasing order according to a provided comparison function. *Key* points to a datum instance to be sought in the table. *Base* points to the element at the base of the table. *Nel* is the number of elements in the table. *Width* is the width of an element in bytes; *sizeof (\*key)* should be used. *Compar* is the name of the comparison function, which is called with two arguments that point to the elements being compared. The function must return an integer less than, equal to, or greater than zero as accordingly the first argument is to be considered less than, equal to, or greater than the second.

**EXAMPLE**

The example below searches a table containing pointers to nodes consisting of a string and its length. The table is ordered alphabetically on the string in the node pointed to by each entry.

This code fragment reads in strings and either finds the corresponding node and prints out the string and its length, or prints an error message.

```
#include <stdio.h>
#include <search.h>

#define TABSIZE 1000

struct node {
    char *string;
    int length;
};
struct node table[TABSIZE]; /* table to be searched */
.
.
.
{
    struct node *node_ptr, node;
    int node_compare( ); /* routine to compare 2 nodes */
    char str_space[20]; /* space to read string into */
    .
    .
    .
    node.string = str_space;
    while (scanf("%s", node.string) != EOF) {
        node_ptr = (struct node *)bsearch((char *)&node,
            (char *)table, TABSIZE,
            sizeof(struct node), node_compare);
        if (node_ptr != NULL) {
```

```

        (void)printf("string == %20s, length = %d\n",
                    node_ptr->string, node_ptr->length);
    } else {
        (void)printf("not found: %s\n", node.string);
    }
}
}
/*
   This routine compares two nodes based on an
   alphabetical ordering of the string field.
*/
int
node_compare(node1, node2)
struct node *node1, *node2;
{
    return strcmp(node1->string, node2->string);
}

```

**NOTES**

The pointers to the key and the element at the base of the table should be of type pointer-to-element, and cast to type pointer-to-character. The comparison function need not compare every byte, so arbitrary data may be contained in the elements in addition to the values being compared. Although declared as type pointer-to-character, the value returned should be cast into type pointer-to-element.

**SEE ALSO**

bsearch(3C), lsearch(3C), qsort(3C), tsearch(3C).

**DIAGNOSTICS**

A NULL pointer is returned if the key cannot be found in the table.



**NAME**

*bcopy*, *bcmp*, *bzero*, *ffs* — bit and byte string operations

**SYNOPSIS**

```
bcopy(b1, b2, length)  
char *b1, *b2;  
int length;  
  
bcmp(b1, b2, length)  
char *b1, *b2;  
int length;  
  
bzero(b, length)  
char *b;  
int length;  
  
ffs(i)  
int i;  
  
cc ... -lnet
```

**DESCRIPTION**

The functions *bcopy*, *bcmp*, and *bzero* operate on variable length strings of bytes. They do not check for null bytes as the routines in *string*(3C) do.

*Bcopy* copies *length* bytes from string *b1* to the string *b2*.

*Bcmp* compares byte string *b1* against byte string *b2*, returning zero if they are identical, non-zero otherwise. Both strings are assumed to be *length* bytes long.

*Bzero* places *length* 0 bytes in the string *b1*.

*Ffs* find the first bit set in the argument passed it and returns the index of that bit. Bits are numbered starting at 1. A return value of -1 indicates the value passed is zero.

**LINKING**

This library is accessed by specifying *-lnet* as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**BUGS**

The *bcmp* and *bcopy* routines take parameters backwards from *strcmp* and *strcpy*.

**NAME**

htonl, htons, ntohl, ntohs — convert values between host and network byte order

**SYNOPSIS**

```
#include <sys/types.h>
#include <netinet/in.h>
netlong = htonl(hostlong);
u_long netlong, hostlong;
netshort = htons(hostshort);
u_short netshort, hostshort;
hostlong = ntohl(netlong);
u_long hostlong, netlong;
hostshort = ntohs(netshort);
u_short hostshort, netshort;
cc ... -lnet
```

**DESCRIPTION**

These routines convert 16 and 32 bit quantities between network byte order and host byte order. On machines such as the SUN these routines are defined as null macros in the include file *<netinet/in.h>*.

These routines are most often used in conjunction with Internet addresses and ports as returned by *gethostent(3N)* and *getservent(3N)*.

**LINKING**

This library is accessed by specifying *-lnet* as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**SEE ALSO**

*gethostent(3N)*, *getservent(3N)*

**NAME**

*clock* — report CPU time used

**SYNOPSIS**

**long clock ( )**

**DESCRIPTION**

*Clock* returns the amount of CPU time (in microseconds) used since the first call to *clock*. The time reported is the sum of the user and system times of the calling process and its terminated child processes for which it has executed *wait(2)* or *system(3S)*.

**SEE ALSO**

*times(2)*, *wait(2)*, *system(3S)*.

**BUGS**

The value returned by *clock* is defined in microseconds for compatibility with systems that have CPU clocks with much higher resolution. Because of this, the value returned wraps around after accumulating only 2,147 seconds of CPU time (about 36 minutes).

**NAME**

*conjg*, *dconjg* — Fortran complex conjugate intrinsic function

**SYNOPSIS**

**complex** *cx1*, *cx2*  
**double complex** *dx1*, *dx2*  
*cx2* = **conjg**(*cx1*)  
*dx2* = **dconjg**(*dx1*)

**DESCRIPTION**

*Conjg* returns the complex conjugate of its complex argument. *Dconjg* returns the double-complex conjugate of its double-complex argument.

**NAME**

*toupper*, *tolower*, *\_toupper*, *\_tolower*, *toascii* — translate characters

**SYNOPSIS**

```
#include <ctype.h>

int toupper (c)
int c;

int tolower (c)
int c;

int _toupper (c)
int c;

int _tolower (c)
int c;

int toascii (c)
int c;
```

**DESCRIPTION**

*Toupper* and *tolower* have as domain the range of *getc*(3S): the integers from -1 through 255. If the argument of *toupper* represents a lower-case letter, the result is the corresponding upper-case letter. If the argument of *tolower* represents an upper-case letter, the result is the corresponding lower-case letter. All other arguments in the domain are returned unchanged.

The macros *\_toupper* and *\_tolower*, are macros that accomplish the same thing as *toupper* and *tolower* but have restricted domains and are faster. *\_toupper* requires a lower-case letter as its argument; its result is the corresponding upper-case letter. The macro *\_tolower* requires an upper-case letter as its argument; its result is the corresponding lower-case letter. Arguments outside the domain cause undefined results.

*Toascii* yields its argument with all bits turned off that are not part of a standard ASCII character; it is intended for compatibility with other systems.

**SEE ALSO**

*ctype*(3C), *getc*(3S).

**NAME**

cos, dcos, ccos — Fortran cosine intrinsic function

**SYNOPSIS**

```
real r1, r2
double precision dp1, dp2
r2 = cos(r1)
dp2 = dcos(dp1)
dp2 = cos(dp1)
cx2 = ccos(cx1)
cx2 = cos(cx1)
```

**DESCRIPTION**

*Cos* returns the real cosine of its real argument. *Dcos* returns the double-precision cosine of its double-precision argument. *Ccos* returns the complex cosine of its complex argument. The generic form *cos* may be used with impunity because its returned type is determined by that of its argument.

**SEE ALSO**

trig(3M).

**NAME**

*cosh*, *dcosh* — Fortran hyperbolic cosine intrinsic function

**SYNOPSIS**

```
real r1, r2
double precision dp1, dp2
r2 = cosh(r1)
dp2 = dcosh(dp1)
dp2 = cosh(dp1)
```

**DESCRIPTION**

*Cosh* returns the real hyperbolic cosine of its real argument. *Dcosh* returns the double-precision hyperbolic cosine of its double-precision argument. The generic form *cosh* may be used to return the hyperbolic cosine in the type of its argument.

**SEE ALSO**

*sinh*(3M).

**NAME**

`crypt`, `setkey`, `encrypt` — generate DES encryption

**SYNOPSIS**

```
char *crypt (key, salt)
char *key, *salt;

void setkey (key)
char *key;

void encrypt (block, edflag)
char *block;
int edflag;
```

**DESCRIPTION**

*Crypt* is the password encryption function. It is based on the NBS Data Encryption Standard (DES), with variations intended to frustrate use of hardware implementations of the DES for key search.

*Key* is a user's typed password. *Salt* is a 2-character string chosen from the set `la-zA-Z0-9./`; this string is used to perturb the DES algorithm in one of 4,096 different ways, after which the password is used as the key to encrypt repeatedly a constant string. The returned value points to the encrypted password. The first 2 characters are the salt itself.

The *setkey* and *encrypt* entries provide (rather primitive) access to the actual DES algorithm. The argument of *setkey* is a character array of length 64 containing only the characters with numerical value 0 and 1. If this string is divided into groups of 8, the low-order bit in each group is ignored; this gives a 56-bit key which is set into the machine. The 56-bit key is used with the above-mentioned algorithm to encrypt or decrypt the string *block* with the function *encrypt*.

The argument to the *encrypt* entry is a character array of length 64 containing only the characters with numerical value 0 and 1. The argument array is modified in place to a similar array representing the bits of the argument after having been subjected to the DES algorithm using the key set by *setkey*. If *edflag* is zero, the argument is encrypted; if non-zero, it is decrypted.

**SEE ALSO**

`login(1)`, `passwd(1)`, `getpass(3C)`, `passwd(4)`.

**BUGS**

The return value points to static data that is overwritten by each call.

**NOTE**

The international distribution of this family of subroutines has *setkey* removed and disallows decryption by the *encrypt* function.



**NAME**

*ctermid* -- generate filename for terminal

**SYNOPSIS**

```
#include <stdio.h>
char *ctermid(s)
char *s;
```

**DESCRIPTION**

*Ctermid* generates the pathname of the controlling terminal for the current process, and stores it in a string.

If *s* is a NULL pointer, the string is stored in an internal static area, the contents of which are overwritten at the next call to *ctermid*, and the address of which is returned. Otherwise, *s* is assumed to point to a character array of at least `L_ctermid` elements; the pathname is placed in this array and the value of *s* is returned. The constant `L_ctermid` is defined in the `<stdio.h>` header file.

**NOTES**

The difference between *ctermid* and *ttyname(3C)* is that *ttyname* must be handed a file descriptor and returns the actual name of the terminal associated with that file descriptor, while *ctermid* returns a string (`/dev/tty`) that refers to the terminal if used as a filename. For this reason, *ttyname* is useful only if the process already has at least one file open to a terminal.

**SEE ALSO**

*ttyname(3C)*.

## NAME

ctime, localtime, gmtime, asctime, tzset — convert date and time to string

## SYNOPSIS

```
#include <time.h>

char *ctime (clock)
long *clock;

struct tm *localtime (clock)
long *clock;

struct tm *gmtime (clock)
long *clock;

char *asctime (tm)
struct tm *tm;

extern long timezone;
extern int daylight;
extern char *tzname[2];
void tzset ( )
```

## DESCRIPTION

*Ctime* converts a long integer, pointed to by *clock*, representing the time in seconds since 00:00:00 GMT, January 1, 1970, and returns a pointer to a 26-character string in the following form. All the fields have constant width.

```
Sun Sep 16 01:03:52 1973\n\0
```

*Localtime* and *gmtime* return pointers to *tm* structures, described below. *Localtime* corrects for the time zone and possible Daylight Savings Time; *gmtime* converts directly to Greenwich Mean Time (GMT), which is the time the system uses.

*Asctime* converts a *tm* structure to a 26-character string, as shown in the above example, and returns a pointer to the string.

Declarations of all the functions and externals, and the *tm* structure, are in the *<time.h>* header file. The structure declaration is:

```
struct tm {
    int tm_sec; /* seconds (0 - 59) */
    int tm_min; /* minutes (0 - 59) */
    int tm_hour; /* hours (0 - 23) */
    int tm_mday; /* day of month (1 - 31) */
    int tm_mon; /* month of year (0 - 11) */
    int tm_year; /* year - 1900 */
    int tm_wday; /* day of week (Sunday = 0) */
    int tm_yday; /* day of year (0 - 365) */
    int tm_isdst;
};
```

*Tm\_isdst* is non-zero if Daylight Savings Time is in effect.

The external long variable *timezone* contains the difference, in seconds, between GMT and local standard time (in EST, *timezone* is 5\*60\*60); the external variable *daylight* is non-zero if, and only if, the standard U.S.A. Daylight Savings Time conversion should be applied. The program knows

about the peculiarities of this conversion in 1974 and 1975; if necessary, a table for these years can be extended.

If an environment variable named TZ is present, *asctime* uses the contents of the variable to override the default time zone. The value of TZ must be a 3-letter time zone name, followed by a number representing the difference between local time and Greenwich Mean Time in hours, followed by an optional 3-letter name for a daylight time zone. For example, the setting for New Jersey would be EST5EDT. The effects of setting TZ are thus to change the values of the external variables *timezone* and *daylight*; in addition, the time zone names contained in the external variable

```
char *tzname[2] = { "EST", "EDT" };
```

are set from the environment variable TZ. The function *tzset* sets these external variables from TZ; *tzset* is called by *asctime* and may also be called explicitly by the user.

Note that in most installations, TZ is set by default when the user logs on, to a value in the local */etc/profile* file (see *profile(4)*).

**SEE ALSO**

*time(2)*, *getenv(3C)*, *profile(4)*, *environ(5)*.

**BUGS**

The return values point to static data whose content is overwritten by each call.

**NAME**

*isalpha*, *isupper*, *islower*, *isdigit*, *isxdigit*, *isalnum*, *isspace*, *ispunct*, *isprint*, *isgraph*, *isctrl*, *isascii* – classify characters

**SYNOPSIS**

```
#include <ctype.h>
int isalpha (c)
int c;
. . .
```

**DESCRIPTION**

These macros classify character-coded integer values by table lookup. Each is a predicate returning nonzero for true, zero for false. *isascii* is defined on all integer values; the rest are defined only where *isascii* is true and on the single non-ASCII value EOF (-1); see *stdio*(3S).

<i>isalpha</i>	<i>c</i> is a letter.
<i>isupper</i>	<i>c</i> is an upper-case letter.
<i>islower</i>	<i>c</i> is a lower-case letter.
<i>isdigit</i>	<i>c</i> is a digit [0-9].
<i>isxdigit</i>	<i>c</i> is a hexadecimal digit [0-9], [A-F] or [a-f].
<i>isalnum</i>	<i>c</i> is an alphanumeric (letter or digit).
<i>isspace</i>	<i>c</i> is a space, tab, carriage return, new-line, vertical tab, or form-feed.
<i>ispunct</i>	<i>c</i> is a punctuation character (neither control nor alphanumeric).
<i>isprint</i>	<i>c</i> is a printing character, code 040 (space) through 0176 (tilde).
<i>isgraph</i>	<i>c</i> is a printing character, similar to <i>isprint</i> except false for space.
<i>isctrl</i>	<i>c</i> is a delete character (0177) or an ordinary control character (less than 040).
<i>isascii</i>	<i>c</i> is an ASCII character, code less than 0200.

**DIAGNOSTICS**

If the argument to any of these macros is not in the domain of the function, the result is undefined.

**SEE ALSO**

*stdio*(3S), *ascii*(5).

## NAME

curse - CRT screen handling and optimization package

## SYNOPSIS

```
#include <curse.h>
cc { flags } files -lcurse [ libraries ]
```

## DESCRIPTION

These routines give the user a method of updating screens with reasonable optimization. In order to initialize the routines, the routine *initscr()* must be called before any of the other routines that deal with windows and screens are used. The routine *endwin()* should be called before exiting. To get character-at-a-time input without echoing, (most interactive, screen oriented-programs want this) after calling *initscr()* you should call "*nonl(); cbreak(); noecho();*"

The full curses interface permits manipulation of data structures called *windows* which can be thought of as two dimensional arrays of characters representing all or part of a CRT screen. A default window called *stdscr* is supplied, and others can be created with *newwin*. Windows are referred to by variables declared "WINDOW \*\*", the type WINDOW is defined in *curse.h* to be a C structure. These data structures are manipulated with functions described below, among which the most basic are *move*, and *addch*. (More general versions of these functions are included with names beginning with 'w', allowing you to specify a window. The routines not beginning with 'w' affect *stdscr*.) Then *refresh()* is called, telling the routines to make the users CRT screen look like *stdscr*.

Mini-Curses is a subset of curses which does not allow manipulation of more than one window. To invoke this subset, use -DMINICURSES as a cc option. This level is smaller and faster than full curses.

If the environment variable TERMINFO is defined, any program using curses will check for a local terminal definition before checking in the standard place. For example, if the standard place is */usr/lib/terminfo*, and TERM is set to "vt100", then normally the compiled file is found in */usr/lib/terminfo/v/vt100*. (The "v" is copied from the first letter of "vt100" to avoid creation of huge directories.) However, if TERMINFO is set to */usr/mark/myterms*, curses will first check */usr/mark/myterms/v/vt100*, and if that fails, will then check */usr/lib/terminfo/v/vt100*. This is useful for developing experimental definitions or when write permission in */usr/lib/terminfo* is not available.

## SEE ALSO

*terminfo(4)*.

## FUNCTIONS

Routines listed here may be called when using the full curses. Those marked with an asterisk may be called when using Mini-Curses.

<i>addch(ch)*</i>	add a character to <i>stdscr</i> (like <i>putchar</i> ) (wraps to next line at end of line)
<i>addstr(str)*</i>	calls <i>addch</i> with each character in <i>str</i>
<i>attroff(attrs)*</i>	turn off attributes named
<i>attron(attrs)*</i>	turn on attributes named
<i>attrset(attrs)*</i>	set current attributes to <i>attrs</i>
<i>baudrate()*</i>	current terminal speed

## CURSES (3X)

## CURSES (3X)

beep()*	sound beep on terminal
box(win, vert, hor)	draw a box around edges of <i>win</i> <i>vert</i> and <i>hor</i> are chars to use for vert. and hor. edges of box
clear()	clear <i>stdscr</i>
clearok(win, bf)	clear screen before next redraw of <i>win</i>
clrrobot()	clear to bottom of <i>stdscr</i>
clrtoeol()	clear to end of line on <i>stdscr</i>
cbreak()*	set cbreak mode
delay_output(ms)*	insert <i>ms</i> millisecond pause in output
delch()	delete a character
deleteln()	delete a line
delwin(win)	delete <i>win</i>
doupdate()	update screen from all <i>wnooutrefresh</i>
echo()*	set echo mode
endwin()*	end window modes
erase()	erase <i>stdscr</i>
erasechar()	return user's erase character
fixterm()	restore tty to "in curses" state
flash()	flash screen or beep
flushinp()*	throw away any typesahead
getch()*	get a char from tty
getstr(str)	get a string through <i>stdscr</i>
getmode()	establish current tty modes
getyx(win, y, x)	get (y, x) co-ordinates
has_ic()	true if terminal can do insert character
has_il()	true if terminal can do insert line
idlok(win, bf)*	use terminal's insert/delete line if <i>bf</i> != 0
inch()	get char at current (y, x) co-ordinates
initscr()*	initialize screens
insch(c)	insert a char
insertln()	insert a line
intrflush(win, bf)	interrupts flush output if <i>bf</i> is TRUE
keypad(win, bf)	enable keypad input
killchar()	return current user's kill character
leaveok(win, flag)	OK to leave cursor anywhere after refresh if <i>flag</i> !=0 for <i>win</i> , otherwise cursor must be left at current position.
longname()	return verbose name of terminal
meta(win, flag)*	allow meta characters on input if <i>flag</i> != 0
move(y, x)*	move to (y, x) on <i>stdscr</i>
mvaddch(y, x, ch)	move(y, x) then addch(ch)
mvaddstr(y, x, str)	similar...
mvcur(olddrow, oldcol, newrow, newcol)	low level cursor motion
mvdelch(y, x)	like delch, but move(y, x) first
mvgetch(y, x)	etc.
mvgetstr(y, x)	
mvinch(y, x)	
mvinsch(y, x, c)	
mvprintw(y, x, fmt, args)	
mvscanw(y, x, fmt, args)	
mvwaddch(win, y, x, ch)	

mvwaddstr(win, y, x, str)  
 mvwdech(win, y, x)  
 mvwgetch(win, y, x)  
 mvwgetstr(win, y, x)  
 mvwin(win, by, bx)  
 mvwinch(win, y, x)  
 mvwinsch(win, y, x, c)  
 mvwprintw(win, y, x, fmt, args)  
 mvwscanw(win, y, x, fmt, args)  
 newpad(nlines, ncols)  
 newterm(type, fd)  
 newwin(lines, cols, begin\_y, begin\_x)

nl()\*  
 nocbreak()\*  
 nodelay(win, bf)  
 noecho()\*  
 nonl()\*  
 noraw()\*  
 overlay(win1, win2)  
 overwrite(win1, win2)  
 pnoutrefresh(pad, pminrow, pmincol, sminrow, smincol, smaxrow, smaxcol)

prefresh(pad, pminrow, pmincol, sminrow, smincol, smaxrow, smaxcol)

printw(fmt, arg1, arg2, ...)

raw()\*  
 refresh()\*  
 resetterm()\*  
 resetty()\*  
 saveterm()\*  
 savetty()\*  
 scanw(fmt, arg1, arg2, ...)

scroll(win)  
 scrollok(win, flag)  
 set\_term(new)  
 setscreg(t, b)  
 setterm(type)  
 setupterm(term, flenum, errret)  
 standend()\*  
 standout()\*  
 subwin(win, lines, cols, begin\_y, begin\_x)

touchwin(win)  
 traceoff()  
 traceon()  
 typeahead(fd)

create a new pad with given dimensions  
 set up new terminal of given type to output on fd

create a new window  
 set newline mapping  
 unset cbreak mode  
 enable nodelay input mode through getch  
 unset echo mode  
 unset newline mapping  
 unset raw mode  
 overlay win1 on win2  
 overwrite win1 on top of win2

like `prefresh` but with no output until `doupdate` called

refresh from pad starting with given upper left  
 corner of pad with output to given  
 portion of screen

print on *stdscr*  
 set raw mode  
 make current screen look like *stdscr*  
 set tty modes to "out of curses" state  
 reset tty flags to stored value  
 save current modes as "in curses" state  
 store current tty flags

scanf through *stdscr*  
 scroll *win* one line  
 allow terminal to scroll if flag != 0  
 now talk to terminal *new*  
 set user scrolling region to lines *t* through *b*  
 establish terminal with given type

clear standout mode attribute  
 set standout mode attribute

create a subwindow  
 change all of *win*  
 turn off debugging trace output  
 turn on debugging trace output  
 use file descriptor *fd* to check typeahead

<code>unctrl(ch)*</code>	printable version of <i>ch</i>
<code>waddch(win, ch)</code>	add char to <i>win</i>
<code>waddstr(win, str)</code>	add string to <i>win</i>
<code>wattroff(win, attrs)</code>	turn off <i>attrs</i> in <i>win</i>
<code>wattron(win, attrs)</code>	turn on <i>attrs</i> in <i>win</i>
<code>wattrset(win, attrs)</code>	set <i>attrs</i> in <i>win</i> to <i>attrs</i>
<code>wclear(win)</code>	clear <i>win</i>
<code>wclrtoeol(win)</code>	clear to bottom of <i>win</i>
<code>wcrtoeol(win)</code>	clear to end of line on <i>win</i>
<code>wdelch(win, c)</code>	delete char from <i>win</i>
<code>wdeleteln(win)</code>	delete line from <i>win</i>
<code>werase(win)</code>	erase <i>win</i>
<code>wgetch(win)</code>	get a char through <i>win</i>
<code>wgetstr(win, str)</code>	get a string through <i>win</i>
<code>winch(win)</code>	get char at current (y, x) in <i>win</i>
<code>winsch(win, c)</code>	insert char into <i>win</i>
<code>winsertln(win)</code>	insert line into <i>win</i>
<code>wmove(win, y, x)</code>	set current (y, x) co-ordinates on <i>win</i>
<code>wnoutrefresh(win)</code>	refresh but no screen output
<code>wprintw(win, fmt, arg1, arg2, ...)</code>	
	<code>printf</code> on <i>win</i>
<code>wrefresh(win)</code>	make screen look like <i>win</i>
<code>wscanw(win, fmt, arg1, arg2, ...)</code>	
	<code>scanf</code> through <i>win</i>
<code>wsetscrreg(win, t, b)</code>	set scrolling region of <i>win</i>
<code>wstandend(win)</code>	clear standout attribute in <i>win</i>
<code>wstandout(win)</code>	set standout attribute in <i>win</i>

### TERMINFO LEVEL ROUTINES

These routines should be called by programs wishing to deal directly with the terminfo database. Due to the low level of this interface, it is discouraged. Initially, *setupterm* should be called. This will define the set of terminal dependent variables defined in `terminfo(4)`. The include files `< curses.h >` and `< term.h >` should be included to get the definitions for these strings, numbers, and flags. Parameterized strings should be passed through *tparm* to instantiate them. All terminfo strings (including the output of *tparm*) should be printed with *tputs* or *putp*. Before exiting, *resetterm* should be called to restore the tty modes. (Programs desiring shell escapes or suspending with control Z can call *resetterm* before the shell is called and *fixterm* after returning from the shell.)

<code>fixterm()</code>	restore tty modes for terminfo use (called by <i>setupterm</i> )
<code>resetterm()</code>	reset tty modes to state before program entry
<code>setupterm(term, fd, rc)</code>	read in database. Terminal type is the character string <i>term</i> , all output is to UNIX System file descriptor <i>fd</i> . A status value is returned in the integer pointed to by <i>rc</i> : 1 is normal. The simplest call would be <i>setupterm(0, 1, 0)</i> which uses all defaults.
<code>tparm(str, p1, p2, ..., p9)</code>	instantiate string <i>str</i> with parms <i>p<sub>i</sub></i>
<code>tputs(str, affcnt, putc)</code>	apply padding info to string <i>str</i> . <i>affcnt</i> is the number of lines affected, or 1 if not applicable. <i>putc</i> is a



<p>putp(str)</p> <p>vidputs(attrs,putc)</p> <p>vidattr(attrs)</p>	<p>putchar-like function to which the characters are passed, one at a time.</p> <p>handy function that calls tputs (str, 1, putchar)</p> <p>output the string to put terminal in video attribute mode <i>attrs</i>, which is any combination of the attributes listed below. Chars are passed to putchar-like function <i>putc</i>.</p> <p>Like vidputs but outputs through putchar</p>
---	---

**TERMCAP COMPATIBILITY ROUTINES**

These routines were included as a conversion aid for programs that use *termcap*. Their parameters are the same as for *termcap*. They are emulated using the *terminfo* database. They may go away at a later date.

<p>tgetent(bp, name)</p> <p>tgetflag(id)</p> <p>tgetnum(id)</p> <p>tgetstr(id, area)</p> <p>tgoto(cap, col, row)</p> <p>tputs(cap, affcnt, fn)</p>	<p>look up <i>termcap</i> entry for name</p> <p>get boolean entry for id</p> <p>get numeric entry for id</p> <p>get string entry for id</p> <p>apply parms to given cap</p> <p>apply padding to cap calling fn as putchar</p>
--	---

**ATTRIBUTES**

The following video attributes can be passed to the functions *attron*, *attroff*, *attrset*.

<p>A_STANDOUT</p> <p>A_UNDERLINE</p> <p>A_REVERSE</p> <p>A_BLINK</p> <p>A_DIM</p> <p>A_BOLD</p> <p>A_BLANK</p> <p>A_PROTECT</p> <p>A_ALTCHARSET</p>	<p>Terminal's best highlighting mode</p> <p>Underlining</p> <p>Reverse video</p> <p>Blinking</p> <p>Half bright</p> <p>Extra bright or bold</p> <p>Blanking (invisible)</p> <p>Protected</p> <p>Alternate character set</p>
---	---

**FUNCTION KEYS**

The following function keys might be returned by *getch* if *keypad* has been enabled. Note that not all of these are currently supported, due to lack of definitions in *terminfo* or the terminal not transmitting a unique code when the key is pressed.

<i>Name</i>	<i>Value</i>	<i>Key name</i>
KEY_BREAK	0401	break key (unreliable)
KEY_DOWN	0402	The four arrow keys ...
KEY_UP	0403	
KEY_LEFT	0404	
KEY_RIGHT	0405	...
KEY_HOME	0406	Home key (upward+left arrow)
KEY_BACKSPACE	0407	backspace (unreliable)
KEY_F0	0410	Function keys. Space for 64 is reserved.
KEY_F(n)	(KEY_F0+(n))	Formula for fn.
KEY_DL	0510	Delete line
KEY_IL	0511	Insert line
KEY_DC	0512	Delete character
KEY_IC	0513	Insert char or enter insert mode

**CURSES (3X)****CURSES (3X)**

KEY_EIC	0514	Exit insert char mode
KEY_CLEAR	0515	Clear screen
KEY_EOS	0516	Clear to end of screen
KEY_EOL	0517	Clear to end of line
KEY_SF	0520	Scroll 1 line forward
KEY_SR	0521	Scroll 1 line backwards (reverse)
KEY_NPAGE	0522	Next page
KEY_PPAGE	0523	Previous page
KEY_STAB	0524	Set tab
KEY_CTAB	0525	Clear tab
KEY_CATAB	0526	Clear all tabs
KEY_ENTER	0527	Enter or send (unreliable)
KEY_SRESET	0530	soft (partial) reset (unreliable)
KEY_RESET	0531	reset or hard reset (unreliable)
KEY_PRINT	0532	print or copy
KEY_LL	0533	home down or bottom (lower left)

**WARNING**

The plotting library *plot(3X)* and the curses library *curses(3X)* both use the names *erase()* and *move()*. The curses versions are macros. If you need both libraries, put the *plot(3X)* code in a different source file than the *curses(3X)* code, and/or *#undef move()* and *erase()* in the *plot(3X)* code.

**NAME**

*cuserid* — get character login name of the user

**SYNOPSIS**

```
#include <stdio.h>
char *cuserid (s)
char *s;
```

**DESCRIPTION**

*Cuserid* generates a character-string representation of the login name that the owner of the current process is logged in under. If *s* is a NULL pointer, this representation is generated in an internal static area, the address of which is returned. Otherwise, *s* is assumed to point to an array of at least `L_cuserid` characters; the representation is left in this array. The constant `L_cuserid` is defined in the `<stdio.h>` header file.

**DIAGNOSTICS**

If the login name cannot be found, *cuserid* returns a NULL pointer; if *s* is not a NULL pointer, a null character (`\0`) is placed at *s[0]*.

**SEE ALSO**

*getlogin(3C)*, *getpwent(3C)*.

**BUGS**

*Cuserid* uses *getpwnam(3C)*; thus the results of a user's call to the latter will be obliterated by a subsequent call to the former.

The name *cuserid* is rather a misnomer.

## NAME

`dial` — establish an out-going terminal line connection

## SYNOPSIS

```
#include <dial.h>
```

```
int dial (call)
```

```
CALL call;
```

```
void undial (fd)
```

```
int fd;
```

## DESCRIPTION

*Dial* returns a file descriptor for a terminal line open for read/write. The argument to *dial* is a CALL structure (defined in the `<dial.h>` header file).

When finished with the terminal line, the calling program must invoke *undial* to release the semaphore that has been set during the allocation of the terminal device.

The CALL typedef in the `<dial.h>` header file is:

```
typedef struct {
    struct termio      *attr;      /* pointer to termio attribute struct */
    int                baud;       /* transmission data rate */
    int                speed;      /* 212A modem: low=300, high=1200 */
    char               *line;      /* device name for out-going line */
    char               *telno;     /* pointer to tel-no digits string */
    int                modem;      /* specify modem control for direct lines */
} CALL;

    char               *device;    /* Will hold the name of the device used
                                   to make a connection */
    int                dev_len     /* The length of the device used to
                                   make connection */
```

The CALL element *speed* is intended only for use with an outgoing dialed call, in which case its value should be either 300 or 1200 to identify the 113A modem, or the high-speed or low-speed setting on the 212A modem. Note that the 113A modem or the low-speed setting of the 212A modem will transmit at any rate between 0 and 300 bits per second. However, the high-speed setting of the 2121 modem transmits and receives at 1200 bits per second only. The CALL element *baud* is for the desired transmission baud rate. For example, one might set *baud* to 110 and *speed* to 300 (or 1200). However, if *speed* is set to 1200 *baud* must be set to high (1200).

If the desired terminal line is a direct line, a string pointer to its device name should be placed in the *line* element in the CALL structure. Legal values for such terminal device names are kept in the *L-devices* file. In this case, the value of the *baud* element need not be specified as it will be determined from the *L-devices* file.

The *telno* element is for a pointer to a character string representing the telephone number to be dialed. The termination symbol will be supplied by the *dial* function, and should not be included in the *telno* string passed to *dial* in the CALL structure.

The CALL element *modem* is used to specify modem control for direct lines. This element should be non-zero if modem control is required. The CALL element *attr* is a pointer to a *termio* structure, as defined in the `<termio.h>` header file. A NULL value for this pointer element may be passed to the *dial* function, but if such a structure is included, the elements specified in it will be set for the outgoing terminal line before the connection is established. This is important for attributes such as parity and baud rate.

The CALL element *device* is used to hold the device name (cul..) that establishes the connection.

The CALL element *dev\_len* is the length of the device name that is copied into the array *device*.

#### FILES

`/usr/lib/uucp/L-devices`  
`/usr/spool/uucp/LCK..tty-device`

#### SEE ALSO

`uucp(1C)`, `alarm(2)`, `read(2)`, `write(3)`,  
`termio(7)` in the *Administrator's Manual*.

#### DIAGNOSTICS

On failure, a negative value indicating the reason for the failure is returned. Mnemonics for these negative indices as listed here are defined in the `<dial.h>` header file.

INTRPT	-1	/* interrupt occurred */
D_HUNG	-2	/* dialer hung (no return from write) */
NO_ANS	-3	/* no answer within 10 seconds */
ILL_BD	-4	/* illegal baud-rate */
A_PROB	-5	/* acu problem (open() failure) */
L_PROB	-6	/* line problem (open() failure) */
NO_Ldv	-7	/* can't open LDEVS file */
DV_NT_A	-8	/* requested device not available */
DV_NT_K	-9	/* requested device not known */
NO_BD_A	-10	/* no device available at requested baud */
NO_BD_K	-11	/* no device known at requested baud */

#### WARNINGS

Including the `<dial.h>` header file automatically includes the `<termio.h>` header file.

Because the above routine uses `<stdio.h>`, the size of programs not otherwise using standard I/O is increased more than might be expected.

#### BUGS

An `alarm(2)` system call for 3,600 seconds is made (and caught) within the *dial* module for the purpose of "touching" the *LCK..* file and constitutes the device allocation semaphore for the terminal device. Otherwise, `uucp(1C)` may simply delete the *LCK..* entry on its 90-minute clean-up rounds. The alarm may go off while the user program is in a `read(2)` or `write(3)` system call, causing an apparent error return. If the user program is to run for an hour or more, error returns from `reads` should be checked for (`errno == EINTR`), and the `read` possibly reissued.

**DIM (3F)**

**DIM (3F)**

**NAME**

dim, ddim, idim — positive difference intrinsic functions

**SYNOPSIS**

integer a1,a2,a3  
a3 = idim(a1,a2)

real a1,a2,a3  
a3 = dim(a1,a2)

double precision a1,a2,a3  
a3 = ddim(a1,a2)

**DESCRIPTION**

These functions return:

a1-a2 if a1 > a2  
0 if a1 <= a2

## NAME

opendir, readdir, telldir, seekdir, rewinddir, closedir – flexible length directory operations

## SYNOPSIS

```
#include <sys/dir.h>

DIR *opendir(filename)
char *filename;

struct direct *readdir(dirp)
DIR *dirp;

long telldir(dirp)
DIR *dirp;

seekdir(dirp, loc)
DIR *dirp;
long loc;

rewinddir(dirp)
DIR *dirp;

closedir(dirp)
DIR *dirp;

cc ... -ldir
```

## DESCRIPTION

The purpose of this library is to simulate the new flexible length directory names of 4.2BSD UNIX on top of the old directory structure of 4.1BSD. It allows programs to be converted immediately to the new directory access interface, so that they need only be relinked when 4.2BSD becomes available.

*Opendir* opens the directory named by *filename* and associates a *directory stream* with it. *Opendir* returns a pointer to be used to identify the *directory stream* in subsequent operations. The pointer NULL is returned if *filename* cannot be accessed or is not a directory.

*Readdir* returns a pointer to the next directory entry. It returns NULL upon reaching the end of the directory or detecting an invalid *seekdir* operation.

*Telldir* returns the current location associated with the named *directory stream*.

*Seekdir* sets the position of the next *readdir* operation on the *directory stream*. The new position reverts to the one associated with the *directory stream* when the *telldir* operation was performed. Values returned by *telldir* are good only

for the lifetime of the DIR pointer from which they are derived. If the directory is closed and then reopened, the *telldir* value may be invalidated due to undetected directory compaction. It is safe to use a previous *telldir* value immediately after a call to *opendir* and before any calls to *readdir*.

*Rewinddir* resets the position of the named *directory stream* to the beginning of the directory.

*Closedir* causes the named *directory stream* to be closed, and the structure associated with the DIR pointer to be freed.

See */usr/include/dir.h* for a description of the fields available in a directory entry. The preferred way to search the current directory for entry "name" is:

```
len = strlen(name);
dirp = opendir(".");
for (dp = readdir(dirp); dp != NULL; dp = readdir(dirp))
    if (dp->d_namlen == len && !strcmp(dp->d_name, name)) {
        closedir(dirp);
        return FOUND;
    }
closedir(dirp);
return NOT_FOUND;
```

**LINKING**

This library is accessed by specifying "-lndir" as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lndir
```

**SEE ALSO**

*/usr/include/sys/ndir.h*, *open(2)*, *close(2)*, *read(2)*, *lseek(2)*

**AUTHOR**

Kirk McKusick



**NAME**

dprod — double precision product intrinsic function

**SYNOPSIS**

real a1,a2  
double precision a3  
a3 = dprod (a1,a2)

**DESCRIPTION**

Dprod returns the double precision product of its real arguments.

## NAME

*drand48*, *erand48*, *lrand48*, *nrand48*, *mrand48*, *jrand48*, *strand48*, *seed48*, *lcg48* — generate uniformly distributed pseudo-random numbers

## SYNOPSIS

```
double drand48 ( )
double erand48 (xsubi)
unsigned short xsubi[3];
long lrand48 ( )
long nrand48 (xsubi)
unsigned short xsubi[3];
long mrand48 ( )
long jrand48 (xsubi)
unsigned short xsubi[3];
void strand48 (seedval)
long seedval;
unsigned short *seed48 (seed16v)
unsigned short seed16v[3];
void lcg48 (param)
unsigned short param[7];
```

## DESCRIPTION

This family of functions generates pseudo-random numbers using the well-known linear congruential algorithm and 48-bit integer arithmetic.

Functions *drand48* and *erand48* return non-negative double-precision floating-point values uniformly distributed over the interval [0.0, 1.0).

Functions *lrand48* and *nrand48* return non-negative long integers uniformly distributed over the interval [0,  $2^{31}$ ).

Functions *mrand48* and *jrand48* return signed long integers uniformly distributed over the interval [ $-2^{31}$ ,  $2^{31}$ ).

Functions *strand48*, *seed48*, and *lcg48* are initialization entry points, one of which should be invoked before *drand48*, *lrand48*, or *mrand48* is called. (Although it is not recommended practice, constant default initializer values are supplied automatically if *drand48*, *lrand48*, or *mrand48* is called without a prior call to an initialization entry point.) Functions *erand48*, *nrand48*, and *jrand48* do not require an initialization entry point to be called first.

All the routines work by generating a sequence of 48-bit integer values,  $X_i$ , according to the linear congruential formula

$$X_{n+1} = (aX_n + c) \bmod m \quad n \geq 0.$$

The parameter  $m = 2^{48}$ ; hence 48-bit integer arithmetic is performed. Unless *lcg48* has been invoked, the multiplier value  $a$  and the addend value  $c$  are given by

$$\begin{aligned} a &= 5DEECE66D_{16} = 273673163155_8 \\ c &= B_{16} = 13_8. \end{aligned}$$

The value returned by any of the functions *drand48*, *erand48*, *lrand48*, *nrand48*, *mrand48*, or *jrand48* is computed by first generating the next 48-

bit  $X_i$  in the sequence. Then the appropriate number of bits, according to the type of data item to be returned, are copied from the high-order (left-most) bits of  $X_i$  and transformed into the returned value.

The functions *drand48*, *lrand48*, and *mrnd48* store the last 48-bit  $X_i$  generated in an internal buffer; that is why they must be initialized prior to being invoked. The functions *erand48*, *nrnd48*, and *jrnd48* require the calling program to provide storage for the successive  $X_i$  values in the array specified as an argument when the functions are invoked. That is why these routines do not have to be initialized; the calling program merely has to place the desired initial value of  $X_i$  into the array and pass it as an argument. By using different arguments, functions *erand48*, *nrnd48*, and *jrnd48* allow separate modules of a large program to generate several independent streams of pseudo-random numbers, i.e., the sequence of numbers in each stream does not depend upon how many times the routines have been called to generate numbers for the other streams.

The initializer function *srnd48* sets the high-order 32 bits of  $X_i$  to the 32 bits contained in its argument. The low-order 16 bits of  $X_i$  are set to the arbitrary value  $330E_{16}$ .

The initializer function *seed48* sets the value of  $X_i$  to the 48-bit value specified in the argument array. The previous value of  $X_i$  is copied into a 48-bit internal buffer, used only by *seed48*. A pointer to this buffer is the value returned by *seed48*. The returned pointer, which can be ignored if not needed, is useful if a program is to be restarted from a given point at some future time. Use the pointer to get and store the last  $X_i$  value; then use this value to reinitialize via *seed48* when the program is restarted.

The initialization function *lcng48* allows the user to specify the initial  $X_i$ , the multiplier value  $a$ , and the addend value  $c$ . Argument array elements *param*[0-2] specify  $X_i$ , elements *param*[3-5] specify the multiplier  $a$ , and *param*[6] specifies the 16-bit addend  $c$ . After *lcng48* has been called, a subsequent call to either *srnd48* or *seed48* will restore the "standard" multiplier and addend values,  $a$  and  $c$ , specified on the previous page.

#### NOTES

The routines are coded in portable C. The source code for the portable version can even be used on computers which do not have floating-point arithmetic. In such a situation, functions *drand48* and *erand48* do not exist; instead, they are replaced by the following two functions:

```
long lrand48 (m)
unsigned short m;

long krand48 (xsubi, m)
unsigned short xsubi[3], m;
```

Functions *lrnd48* and *krnd48* return non-negative long integers uniformly distributed over the interval  $[0, m-1]$ .

#### SEE ALSO

rand(3C).

**NAME**

*ecvt*, *fcvt*, *gcvt* -- convert floating-point number to string

**SYNOPSIS**

```
char *ecvt (value, ndigit, decpt, sign)
double value;
int ndigit, *decpt, *sign;

char *fcvt (value, ndigit, decpt, sign)
double value;
int ndigit, *decpt, *sign;

char *gcvt (value, ndigit, buf)
double value;
int ndigit;
char *buf;
```

**DESCRIPTION**

*Ecvt* converts *value* to a null-terminated string of *ndigit* digits and returns a pointer to this string. The high-order digit is non-zero, unless the value is zero. The low-order digit is rounded. The position of the decimal point relative to the beginning of the string is stored indirectly through *decpt* (negative means to the left of the returned digits). The decimal point is not included in the returned string. If the sign of the result is negative, the word pointed to by *sign* is non-zero; otherwise it is zero.

*Fcv* is identical to *ecvt*, except that the correct digit has been rounded for printf "%f" (Fortran F-format) output of the number of digits specified by *ndigit*.

*Gcv* converts the *value* to a null-terminated string in the array pointed to by *buf* and returns *buf*. It attempts to produce *ndigit* significant digits in Fortran F-format, ready for printing; E-format is produced when F-format is not possible. A minus sign, if there is one, or a decimal point is included as part of the returned string. Trailing zeros are suppressed.

**SEE ALSO**

printf(3S).

**BUGS**

The values returned by *ecvt* and *fcvt* point to a single static data array.

**NAME**

*end*, *etext*, *edata* — last locations in program

**SYNOPSIS**

```
extern end;
extern etext;
extern edata;
```

**DESCRIPTION**

These names refer neither to routines nor to locations with interesting contents. The address of *etext* is the first address above the program text, *edata* above the initialized data region, and *end* above the uninitialized data region.

When execution begins, the program break (the first location beyond the data) coincides with *end*, but the program break may be reset by the routines of *brk*(2), *malloc*(3C), standard input/output (*stdio*(3S)), the profile (*-p*) option of *cc*(1), and so on. Thus, the current value of the program break should be determined by *sbrk*(0) (see *brk*(2)).

These symbols are accessible from assembly language if it is remembered that they should be prefixed by *\_*.

**SEE ALSO**

*cc*(1), *brk*(2), *malloc*(3C), *stdio*(3S).

**NAME**

erf, erfc — error function and complementary error function

**SYNOPSIS**

```
#include <math.h>
```

```
double erf (x)
```

```
double x;
```

```
double erfc (x)
```

```
double x;
```

**DESCRIPTION**

*Erf* returns the error function of  $x$ , defined as  $\frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$ .

*Erfc*, which returns  $1.0 - erf(x)$ , is provided because of the extreme loss of relative accuracy if *erf*( $x$ ) is called for large  $x$  and the result subtracted from 1.0 (e.g. for  $x = 5$ , 12 places are lost).

**SEE ALSO**

exp(3M).

**NAME**

*exp*, *dexp*, *cexp* — Fortran exponential intrinsic function

**SYNOPSIS**

```
real r1, r2
double precision dp1, dp2
complex cx1, cx2
r2 = exp(r1)
dp2 = dexp(dp1)
dp2 = exp(dp1)
cx2 = clog(cx1)
cx2 = exp(cx1)
```

**DESCRIPTION**

*Exp* returns the real exponential function  $e^x$  of its real argument. *Dexp* returns the double-precision exponential function of its double-precision argument. *Cexp* returns the complex exponential function of its complex argument. The generic function *exp* becomes a call to *dexp* or *cexp*, as required, depending on the type of its argument.

**SEE ALSO**

*exp*(3M).

## NAME

*exp*, *log*, *log10*, *pow*, *sqrt* — exponential, logarithm, power, square root functions

## SYNOPSIS

```
#include <math.h>
double exp (x)
double x;

double log (x)
double x;

double log10 (x)
double x;

double pow (x, y)
double x, y;

double sqrt (x)
double x;
```

## DESCRIPTION

*Exp* returns  $e^x$ .

*Log* returns the natural logarithm of  $x$ . The value of  $x$  must be positive.

*Log10* returns the logarithm base ten of  $x$ . The value of  $x$  must be positive.

*Pow* returns  $x^y$ . If  $x$  is zero,  $y$  must be positive. If  $x$  is negative,  $y$  must be an integer.

*Sqrt* returns the non-negative square root of  $x$ . The value of  $x$  may not be negative.

## DIAGNOSTICS

*Exp* returns HUGE when the correct value would overflow, or 0 when the correct value would underflow, and sets *errno* to ERANGE.

*Log* and *log10* return -HUGE and set *errno* to EDOM when  $x$  is non-positive. A message indicating DOMAIN error (or SING error when  $x$  is 0) is printed on the standard error output.

*Pow* returns 0 and sets *errno* to EDOM when  $x$  is 0 and  $y$  is non-positive, or when  $x$  is negative and  $y$  is not an integer. In these cases a message indicating DOMAIN error is printed on the standard error output. When the correct value for *pow* would overflow or underflow, *pow* returns  $\pm$ HUGE or 0 respectively, and sets *errno* to ERANGE.

*Sqrt* returns 0 and sets *errno* to EDOM when  $x$  is negative. A message indicating DOMAIN error is printed on the standard error output.

These error-handling procedures may be changed with the function *matherr*(3M).

## SEE ALSO

*intro*(2), *hypot*(3M), *matherr*(3M), *sinh*(3M).



**NAME**

*fclose*, *fflush* -- close or flush a stream

**SYNOPSIS**

```
#include <stdio.h>
```

```
int fclose (stream)
```

```
FILE *stream;
```

```
int fflush (stream)
```

```
FILE *stream;
```

**DESCRIPTION**

*Fclose* causes any buffered data for the named *stream* to be written out and the *stream* to be closed.

*Fclose* is performed automatically for all open files upon calling *exit(2)*.

*Fflush* causes any buffered data for the named *stream* to be written to that file. The *stream* remains open.

**DIAGNOSTICS**

These functions return 0 for success, and EOF if any error (such as trying to write to a file that has not been opened for writing) was detected.

**SEE ALSO**

*close(2)*, *exit(2)*, *fopen(3S)*, *setbuf(3S)*.

**NAME**

`ferror`, `feof`, `clearerr`, `fileno` — stream status inquiries

**SYNOPSIS**

```
#include <stdio.h>
```

```
int feof (stream)
```

```
FILE *stream;
```

```
int ferror (stream)
```

```
FILE *stream;
```

```
void clearerr (stream)
```

```
FILE *stream;
```

```
int fileno (stream)
```

```
FILE *stream;
```

**DESCRIPTION**

*Feof* returns non-zero when EOF has previously been detected reading the named input *stream*; otherwise, it returns zero.

*Ferror* returns non-zero when an I/O error has previously occurred reading from or writing to the named *stream*; otherwise, it returns zero.

*Clearerr* resets the error indicator and EOF indicator to zero on the named *stream*.

*Fileno* returns the integer file descriptor associated with the named *stream*; see *open(2)*.

**NOTE**

All these functions are implemented as macros; they cannot be declared or redeclared.

**SEE ALSO**

*open(2)*, *fopen(3S)*.

**NAME**

floor, ceil, fmod, fabs — floor, ceiling, remainder, absolute value functions

**SYNOPSIS**

```
#include <math.h>

double floor (x)
double x;

double ceil (x)
double x;

double fmod (x, y)
double x, y;

double fabs (x)
double x;
```

**DESCRIPTION**

*Floor* returns the largest integer (as a double-precision number) not greater than  $x$ .

*Ceil* returns the smallest integer not less than  $x$ .

*Fmod* returns the floating-point remainder of the division of  $x$  by  $y$ : zero if  $y$  is zero or if  $x/y$  would overflow; otherwise the number  $f$  with the same sign as  $x$ , such that  $x = iy + f$  for some integer  $i$ , and  $|f| < |y|$ .

*Fabs* returns the absolute value of  $|x|$ .

**SEE ALSO**

abs(3C).

## NAME

`fopen`, `freopen`, `fdopen` – open a stream

## SYNOPSIS

```
#include <stdio.h>

FILE *fopen (filename, type )
char *filename, *type;

FILE *freopen (filename, type, stream)
char *filename, *type;
FILE *stream;

FILE *fdopen (fildes, type)
int fildes;
char *type;
```

## DESCRIPTION

*Fopen* opens the file named by *filename* and associates a *stream* with it. *Fopen* returns a pointer to the FILE structure associated with the *stream*.

*Filename* points to a character string that contains the name of the file to be opened.

*Type* is a character string having one of the following values:

r	open for reading
w	truncate or create for writing
a	append; open for writing at end of file, or create for writing
r+	open for update (reading and writing)
w+	truncate or create for update
a+	append; open or create for update at end-of-file

*Freopen* substitutes the named file in place of the open *stream*. The original *stream* is closed, regardless of whether the open ultimately succeeds. *Freopen* returns a pointer to the FILE structure associated with *stream*.

*Freopen* is typically used to attach the preopened *streams* associated with `stdin`, `stdout`, and `stderr` to other files.

*Fdopen* associates a *stream* with a file descriptor by formatting a file structure from the file descriptor. Thus, *fdopen* can be used to access the file descriptors returned by `open(2)`, `dup(3)`, `creat(2)`, or `pipe(2)`. (These calls open files but do not return pointers to a FILE structure.) The *type* of *stream* must agree with the mode of the open file.

When a file is opened for update, both input and output may be done on the resulting *stream*. However, output may not be directly followed by input without an intervening *fseek* or *rewind*, and input may not be directly followed by output without an intervening *fseek*, *rewind*, or an input operation which encounters end-of-file.

When a file is opened for append (i.e., when *type* is "a" or "a+"), it is impossible to overwrite information already in the file. *Fseek* may be used to reposition the file pointer to any position in the file, but when output is written to the file the current file pointer is disregarded. All output is

written at the end of the file and causes the file pointer to be repositioned at the end of the output. If two separate processes open the same file for append, each process may write freely to the file without fear of destroying output being written by the other. The output from the two processes will be intermixed in the file in the order in which it is written.

**SEE ALSO**

`creat(2)`, `dup(3)`, `open(2)`, `pipe(2)`, `fclose(3S)`, `fseek(3S)`.

**DIAGNOSTICS**

*Fopen* and *freopen* return a NULL pointer on failure.

**NAME**

*fread*, *fwrite* — binary input/output

**SYNOPSIS**

```
#include <stdio.h>
```

```
int fread (ptr, size, nitems, stream)
```

```
char *ptr;
```

```
int size, nitems;
```

```
FILE *stream;
```

```
int fwrite (ptr, size, nitems, stream)
```

```
char *ptr;
```

```
int size, nitems;
```

```
FILE *stream;
```

**DESCRIPTION**

*Fread* copies *nitems* items of data from the named input *stream* into an array beginning at *ptr*. An item of data is a sequence of bytes (not necessarily terminated by a null byte) of length *size*. *Fread* stops appending bytes if an end-of-file or error condition is encountered while reading *stream* or if *nitems* items have been read. *Fread* leaves the file pointer in *stream*, if defined, pointing to the byte following the last byte read if there is one. *Fread* does not change the contents of *stream*.

*Fwrite* appends at most *nitems* items of data from the the array pointed to by *ptr* to the named output *stream*. *Fwrite* stops appending when it has appended *nitems* items of data or if an error condition is encountered on *stream*. *Fwrite* does not change the contents of the array pointed to by *ptr*.

The variable *size* is typically *sizeof(\*ptr)* where the pseudo-function *sizeof* specifies the length of an item pointed to by *ptr*. If *ptr* points to a data type other than *char* it should be cast into a pointer to *char*.

**SEE ALSO**

*read*(2), *write*(3), *fopen*(3S), *getc*(3S), *gets*(3S), *printf*(3S), *putc*(3S), *puts*(3S), *scanf*(3S).

**DIAGNOSTICS**

*Fread* and *fwrite* return the number of items read or written. If *size* or *nitems* is non-positive, no characters are read or written and 0 is returned by both *fread* and *fwrite*.

**NAME**

*frexp*, *ldexp*, *modf* — manipulate parts of floating-point numbers

**SYNOPSIS**

```
double frexp (value, eptr)
double value;
int *eptr;

double ldexp (value, exp)
double value;
int exp ;

double modf (value, iptr)
double value, *iptr;
```

**DESCRIPTION**

Every non-zero number can be written uniquely as  $x * 2^n$ , where the “mantissa” (fraction)  $x$  is in the range  $0.5 \leq |x| < 1.0$ , and the “exponent”  $n$  is an integer. *Frexp* returns the mantissa of a double *value*, and stores the exponent indirectly in the location pointed to by *eptr*. If *value* is zero, both results returned by *frexp* are zero.

*Ldexp* returns the quantity  $value * 2^{exp}$ .

*Modf* returns the signed fractional part of *value* and stores the integral part indirectly in the location pointed to by *iptr*.

**DIAGNOSTICS**

If *ldexp* would cause overflow,  $\pm$ HUGE is returned (according to the sign of *value*), and *errno* is set to ERANGE.

If *ldexp* would cause underflow, zero is returned and *errno* is set to ERRANGE.

**NAME**

*fseek*, *rewind*, *ftell* — reposition a file pointer in a stream

**SYNOPSIS**

```
#include <stdio.h>

int fseek (stream, offset, ptrname)
FILE *stream;
long offset;
int ptrname;

void rewind (stream)
FILE *stream;

long ftell (stream)
FILE *stream;
```

**DESCRIPTION**

*Fseek* sets the position of the next input or output operation on the *stream*. The new position is at the signed distance *offset* bytes from the beginning, the current position, or the end of the file, when the value of *ptrname* is 0, 1, or 2, respectively.

*Rewind(stream)* is equivalent to *fseek(stream, 0L, 0)*, except that no value is returned.

*Fseek* and *rewind* undo any effects of *ungetc(3S)*.

After *fseek* or *rewind*, the next operation on a file opened for update may be either input or output.

*Ftell* returns the offset of the current byte relative to the beginning of the file associated with the named *stream*.

**SEE ALSO**

*lseek(2)*, *fopen(3S)*, *popen(3S)*, *ungetc(3S)*.

**DIAGNOSTICS**

*Fseek* returns non-zero for improper seeks; otherwise it returns zero. An improper seek can be, for example, an *fseek* done on a file that has not been opened via *fopen*; in particular, *fseek* may not be used on a terminal or on a file opened via *popen(3S)*.

**WARNING**

On an offset returned by *ftell* is measured in bytes, and it is permissible to seek to positions relative to that offset; however, portability to systems other than requires that an offset be used by *fseek* directly. Arithmetic may not meaningfully be performed on such an offset, which is not necessarily measured in bytes.



**NAME**

*ftok* – standard interprocess communication package

**SYNOPSIS**

```
#include <sys/types.h>
#include <sys/ipc.h>

key_t ftok(path, id)
char *path;
char id;
```

**DESCRIPTION**

All interprocess communication facilities require the user to supply a key to be used by the *msgget(2)*, *semget(2)*, and *shmget(2)* system calls to obtain interprocess communication identifiers. One method for forming a key is to use the *ftok* subroutine described below. Another way to compose keys is to include the project ID in the most significant byte and to use the remaining portion as a sequence number. There are many other ways to form keys, but it is necessary for each system to define standards for forming them. If a standard is not adhered to, unrelated processes may interfere with each other's operation. Therefore, it is strongly suggested that the most significant byte of a key in some sense refer to a project so that keys do not conflict across a given system.

*ftok* returns a key based on *path* and *id* that is usable in subsequent *msgget*, *semget*, and *shmget* system calls. *Path* must be the pathname of an existing file that is accessible to the process. *Id* is a character that uniquely identifies a project. *ftok* returns the same key for linked files when called with the same *id*; it returns different keys when called with the same filename but different *ids*.

**SEE ALSO**

*intro(2)*, *msgget(2)*, *semget(2)*, *shmget(2)*.

**DIAGNOSTICS**

*ftok* returns (key\_t) -1 if *path* does not exist or if it is not accessible to the process.

**WARNING**

If the file whose *path* is passed to *ftok* is removed when keys still refer to the file, future calls to *ftok* with the same *path* and *id* will return an error. If the same file is recreated, *ftok* is likely to return a different key than it did the original time it was called.

**NAME**

`ftw` - walk a file tree

**SYNOPSIS**

```
#include <ftw.h>

int ftw (path, fn, depth)
char *path;
int (*fn) ( );
int depth;
```

**DESCRIPTION**

*Ftw* recursively descends the directory hierarchy rooted in *path*. For each object in the hierarchy, *ftw* calls *fn*, passing it a pointer to a null-terminated character string containing the name of the object, a pointer to a *stat* structure (see *stat(2)*) containing information about the object, and an integer. Possible values of the integer, defined in the `<ftw.h>` header file, are `FTW_F` for a file, `FTW_D` for a directory, `FTW_DNR` for a directory that cannot be read, and `FTW_NS` for an object for which *stat* could not be executed successfully. If the integer is `FTW_DNR`, descendants of that directory will not be processed. If the integer is `FTW_NS`, the *stat* structure will contain garbage. An example of an object that would cause `FTW_NS` to be passed to *fn* is a file in a directory with read permission but not execute (search) permission.

*Ftw* visits a directory before visiting any of its descendants.

The tree traversal continues until the tree is exhausted, an invocation of *fn* returns a nonzero value, or an error is detected within *ftw* (such as an I/O error). If the tree is exhausted, *ftw* returns zero. If *fn* returns a nonzero value, *ftw* stops its tree traversal and returns whatever value was returned by *fn*. If *ftw* detects an error, it returns `-1`, and sets the error type in *errno*.

*Ftw* uses one file descriptor for each level in the tree. The *depth* argument limits the number of file descriptors so used. If *depth* is zero or negative, the effect is the same as if it were 1. *Depth* must not be greater than the number of file descriptors currently available for use. *Ftw* runs more quickly if *depth* is at least as large as the number of levels in the tree.

**SEE ALSO**

*stat(2)*, *malloc(3C)*.

**BUGS**

Because *ftw* is recursive, it is possible for it to terminate with a memory fault when applied to very deep file structures.

*Ftw* could be made to run faster and use less storage on deep structures at the cost of considerable complexity.

*Ftw* uses *malloc(3C)* to allocate dynamic storage during its operation. If *ftw* is forcibly terminated, such as by *longjmp* being executed by *fn* or an interrupt routine, *ftw* does not have a chance to free that storage, so it remains permanently allocated. A safe way to handle interrupts is to store the fact that an interrupt has occurred, and arrange to have *fn* return a nonzero value at its next invocation.

**NAME**

int, ifix, idint, real, float, sngl, dble, cmplx, dcmplx, ichar, char — explicit Fortran type conversion

**SYNOPSIS**

```

integer i, j
real r, s
double precision dp, dq
complex cx
double complex dcx
character *1 ch

i = int(r)
i = int(dp)
i = int(cx)
i = int(dcx)
i = ifix(r)
i = idint(dp)

r = real(i)
r = real(dp)
r = real(cx)
r = real(dcx)
r = float(i)
r = sngl(dp)

dp = dble(i)
dp = dble(r)
dp = dble(cx)
dp = dble(dcx)

cx = cmplx(i)
cx = cmplx(i, j)
cx = cmplx(r)
cx = cmplx(r, s)
cx = cmplx(dp)
cx = cmplx(dp, dq)
cx = cmplx(dcx)

dcx = dcmplx(i)
dcx = dcmplx(i, j)
dcx = dcmplx(r)
dcx = dcmplx(r, s)
dcx = dcmplx(dp)
dcx = dcmplx(dp, dq)
dcx = dcmplx(dcx)

i = ichar(ch)
ch = char(i)

```

**DESCRIPTION**

These functions perform conversion from one data type to another.

**Int** converts to *integer* form its *real*, *double precision*, *complex*, or *double complex* argument. If the argument is *real* or *double precision*, **int** returns the integer whose magnitude is the largest integer that does not exceed the magnitude of the argument and whose sign is the same as the sign of the argument (i.e., truncation). For complex types, the above rule is applied to

the real part. **Ifix** and **Idint** convert only *real* and *double precision* arguments respectively.

**Real** converts to *real* form an *integer*, *double precision*, *complex*, or *double complex* argument. If the argument is *double precision* or *double complex*, as much precision is kept as is possible. If the argument is one of the complex types, the real part is returned. **Float** and **sngl** convert only *integer* and *double precision* arguments, respectively.

**Dble** converts any *integer*, *real*, *complex*, or *double complex* argument to *double precision* form. If the argument is of a complex type, the real part is returned.

**Cmplx** converts its *integer*, *real*, *double precision*, or *double complex* argument(s) to *complex* form.

**Dcmplx** converts its *integer*, *real*, *double precision*, or *complex* argument(s) to *double complex* form.

Either one or two arguments may be supplied to **cmplx** and **dcmplx**. If there is only one argument, it is taken as the real part of the complex type and a imaginary part of zero is supplied. If two arguments are supplied, the first is taken as the real part and the second as the imaginary part.

**Ichar** converts from a character to an integer depending on the character's position in the collating sequence.

**Char** returns the character in the *i*th position in the processor collating sequence, where *i* is the supplied argument.

For a processor capable of representing *n* characters,

**ichar(char(i))** = *i* for  $0 < i < n$ , and

**char(ichar(ch))** = *ch* for any representable character *ch*.

**NAME**

gamma - log gamma function

**SYNOPSIS**

```
#include <math.h>
extern int signgam;
double gamma (x)
double x;
```

**DESCRIPTION**

*Gamma* returns the natural log of gamma as a function of the absolute value of a given value. *Gamma* returns  $\ln(|\Gamma(x)|)$ , where  $\Gamma(x)$  is defined as

$$\int_0^{\infty} e^{-t} t^{x-1} dt.$$

The sign of  $\Gamma(x)$  is returned in the external integer *signgam*. The argument *x* may not be a non-positive integer.

The following C program fragment might be used to calculate  $\Gamma$ :

```
if ((y = gamma(x)) > LN_MAXDOUBLE)
    error();
y = signgam * exp(y);
```

where LN\_MAXDOUBLE is the least value that causes *exp(3M)* to return a range error, and is defined in the *<values.h>* header file.

**DIAGNOSTICS**

For non-negative integer arguments HUGE is returned, and *errno* is set to EDOM. A message indicating SING error is printed on the standard error output.

If the correct value would overflow, *gamma* returns HUGE and sets *errno* to ERANGE.

These error-handling procedures may be changed with the function *matherr(3M)*.

**SEE ALSO**

*exp(3M)*, *matherr(3M)*, *values(5)*.

**NAME**

*getarg* — return Fortran command-line argument

**SYNOPSIS**

character\*N *c*

integer *i*

*getarg* (*i*, *c*)

**DESCRIPTION**

*Getarg* returns the *i*-th command-line argument of the current process. Thus, if a program were invoked via

*foo arg1 arg2 arg3*

*getarg*(2, *c*) would return the string *arg2* in the character variable *c*.

**SEE ALSO**

*getopt*(3C).

**NAME**

*getc*, *getchar*, *fgetc*, *getw* — get character or word from a stream

**SYNOPSIS**

```
#include <stdio.h>
```

```
int getc (stream)
```

```
FILE *stream;
```

```
int getchar ()
```

```
int fgetc (stream)
```

```
FILE *stream;
```

```
int getw (stream)
```

```
FILE *stream;
```

**DESCRIPTION**

*Getc* returns the next character (i.e., byte) from the named input *stream*, as an integer. It also moves the file pointer, if defined, ahead one character in *stream*. *Getchar* is defined as *getc(stdin)*. *Getc* and *getchar* are macros.

*Fgetc* behaves like *getc*, but is a function rather than a macro. *Fgetc* runs more slowly than *getc*, but takes less space per invocation and its name can be passed as an argument to a function.

*Getw* returns the next word (32-bit integer on a 68000) from the named input *stream*. *Getw* increments the associated file pointer, if defined, to point to the next word. *Getw* assumes no special alignment in the file.

**SEE ALSO**

*fclose(3S)*, *ferror(3S)*, *fopen(3S)*, *fread(3S)*, *gets(3S)*, *putc(3S)*, *scanf(3S)*, *ungetc(3S)*.

**DIAGNOSTICS**

These functions return the constant EOF at end-of-file or upon an error. Because EOF is a valid integer, *ferror(3S)* should be used to detect *getw* errors.

**WARNING**

If the integer value returned by *getc*, *getchar*, or *fgetc* is stored into a character variable and then compared against the integer constant EOF, the comparison may never succeed, because sign-extension of a character on widening to integer is machine-dependent.

**BUGS**

Because it is implemented as a macro, *getc* treats incorrectly a *stream* argument with side effects. In particular, *getc(\*f++)* does not work sensibly. *Fgetc* should be used instead.

Because of possible differences in word length and byte ordering, files written using *putw* are machine-dependent, and may not be read using *getw* on a different processor.

**NAME**

getcwd — get pathname of current working directory

**SYNOPSIS**

```
char *getcwd (buf, size)
char *buf;
int size;
```

**DESCRIPTION**

*Getcwd* returns a pointer to the current directory pathname. The value of *size* must be at least two greater than the length of the pathname to be returned.

If *buf* is a NULL pointer, *getcwd* obtains *size* bytes of space using *malloc*(3C). In this case, the pointer returned by *getcwd* may be used as the argument in a subsequent call to *free*.

The function is implemented by using *popen*(3S) to pipe the output of the *pwd*(1) command into the specified string space.

**EXAMPLE**

```
char *cwd, *getcwd();
:
:
:
if ((cwd = getcwd((char *)NULL, 64)) == NULL) {
    perror("pwd");
    exit(1);
}
printf("%s\n", cwd);
```

**SEE ALSO**

*pwd*(1), *malloc*(3C), *popen*(3S).

**DIAGNOSTICS**

Returns NULL with *errno* set if *size* is not large enough, or if an error occurs in a lower-level function.



**NAME**

getdtablesize — get descriptor table size

**SYNOPSIS**

```
nds = getdtablesize()
int nds;

cc ... -lnet
```

**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.

**LINKING**

This library is accessed by specifying *-lnet* as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**SEE ALSO**

close(2), dup(3), open(2)

**NAME**

getenv — return value for environment name

**SYNOPSIS**

```
char *getenv (name)
char *name;
```

**DESCRIPTION**

*Getenv* searches the environment list (see *environ(5)*) for a string of the form *name=value*, and returns a pointer to the *value* in the current environment if such a string is present; otherwise a NULL pointer is returned.

**SEE ALSO**

exec(2), putenv(3C), environ(5).

**NAME**

*getenv* — return Fortran environment variable

**SYNOPSIS**

character \*N c

*getenv*(TMPDIR, c)

**DESCRIPTION**

*Getenv* returns the character-string value of the environment variable represented by its first argument into the character variable of its second argument. If no such environment variable exists, all blanks are returned.

**SEE ALSO**

*getenv*(3C), *environ*(5).

## NAME

*getgrent*, *getgrgid*, *getgrnam*, *setgrent*, *endgrent*, *fgetgrent* — obtain group file entry from a group file

## SYNOPSIS

```
#include <grp.h>

struct group *getgrent ( )
struct group *getgrgid (gid)
int gid;
struct group *getgrnam (name)
char *name;
void setgrent ( )
struct group *fgetgrent (f)
FILE *f;
void endgrent ( )
```

## DESCRIPTION

*Getgrent*, *getgrgid*, and *getgrnam* each return pointers to an object with the following structure containing the broken-out fields of a line in the */etc/group* file. Each line contains a group structure, defined in the *<grp.h>* header file.

```
struct group {
    char    *gr_name; /* the name of the group */
    char    *gr_passwd; /* the encrypted group password */
    int     gr_gid; /* the numerical group ID */
    char    **gr_mem; /* vector of pointers to member names */
};
```

When first called, *getgrent* returns a pointer to the first group structure in the file; thereafter, it returns a pointer to the next group structure in the file; therefore, successive calls may be used to search the entire file. *Getgrgid* searches from the beginning of the file until a numerical group id matching *gid* is found; it returns a pointer to the particular structure in which the match was found. *Getgrnam* searches from the beginning of the file until a group name matching *name* is found; it returns a pointer to the particular structure in which the match was found. If an end-of-file or an error is encountered on reading, these functions return a NULL pointer.

A call to *setgrent* has the effect of rewinding the group file to allow repeated searches. *Endgrent* may be called to close the group file when processing is complete.

*Fgetgrent* returns a pointer to the next group structure in the stream *f*, which matches the format of */etc/group*.

## FILES

*/etc/group*

## SEE ALSO

*getlogin(3C)*, *getpwent(3C)*, *group(4)*.

## DIAGNOSTICS

A NULL pointer is returned on EOF or error.

## WARNING

The above routines use *<stdio.h>*. This causes them to increase the size

of programs not otherwise using standard I/O more than might be expected.

**BUGS**

All information is contained in a static area, so it must be copied if it is to be saved.

**NAME**

gethostent, gethostbyaddr, gethostbyname, sethostent, endhostent — get network host entry

**SYNOPSIS**

```
#include <netdb.h>

struct hostent *gethostent()
struct hostent *gethostbyname(name)
char *name;
struct hostent *gethostbyaddr(addr, len, type)
char *addr; int len, type;
sethostent(stayopen)
int stayopen
endhostent()
cc ... -lnet
```

**DESCRIPTION**

*Gethostent*, *gethostbyname*, and *gethostbyaddr* each return a pointer to an object with the following structure containing the broken-out fields of a line in the network host data base, */etc/hosts*.

```
struct hostent {
    char *h_name;      /* official name of host */
    char **h_aliases; /* alias list */
    int h_addrtype;   /* address type */
    int h_length;     /* length of address */
    char *h_addr;     /* address */
};
```

The members of this structure are:

**h\_name** Official name of the host.

**h\_aliases** A zero terminated array of alternate names for the host.

**h\_addrtype** The type of address being returned; currently always AF\_INET.

**h\_length** The length, in bytes, of the address.

**h\_addr** A pointer to the network address for the host. Host addresses are returned in network byte order.

*Gethostent* reads the next line of the file, opening the file if necessary.

*Sethostent* opens and rewinds the file. If the *stayopen* flag is non-zero, the host data base will not be closed after each call to *gethostent* (either directly, or indirectly through one of the other *gethost* calls).

*Endhostent* closes the file.

*Gethostbyname* and *gethostbyaddr* sequentially search from the beginning of the file until a matching host name or host address is found, or until EOF is encountered. Host addresses are supplied in network order.

**FILES**

*/etc/hosts*

**LINKING**

This library is accessed by specifying `-lnet` as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**SEE ALSO**

hosts(4N)

**DIAGNOSTICS**

Null pointer (0) returned on EOF or error.

**BUGS**

All information is contained in a static area so it must be copied if it is to be saved. Only the Internet address format is currently understood.

**NAME**

getlogin — get login name

**SYNOPSIS**

```
char *getlogin ( );
```

**DESCRIPTION**

*Getlogin* returns a pointer to the login name as found in */etc/utmp*. It may be used in conjunction with *getpwnam* to locate the correct password file entry when the same user ID is shared by several login names.

If *getlogin* is called within a process that is not attached to a terminal, it returns a NULL pointer. The correct procedure for determining the login name is to call *cuserid* or *getlogin*. If *getlogin* fails, call *getpwuid*.

**FILES**

*/etc/utmp*

**SEE ALSO**

*cuserid(3S)*, *getgrent(3C)*, *getpwent(3C)*, *utmp(4)*.

**DIAGNOSTICS**

*Getlogin* returns the NULL pointer if *name* is not found.

**BUGS**

The return values point to static data whose content is overwritten by each call.



**NAME**

setmntent, getmntent, addmntent, endmntent - get file system descriptor file entry

**SYNOPSIS**

```
#include <stdio.h>
#include <mntent.h>

FILE *setmntent(filep, type)
char *filep;
char *type;

struct mntent *getmntent(filep)
FILE *filep;

int addmntent(filep, mnt)
FILE *filep;
struct mntent *mnt;

int endmntent(filep)
FILE *filep;
```

**DESCRIPTION**

These routines access the file system description file */etc/fstab*, and the mounted file system description file */etc/mnttab*.

*Setmntent* opens a file system description file and returns a file pointer for use with *getmntent*, *addmntent*, or *endmntent*. The *type* argument is the same as in *fopen(3)*. *Getmntent* reads the next line from *filep* and returns a pointer to an object with the following structure containing broken-out fields of a line in the file system description file, *mnttab.h*. The fields have meanings described in as follows:

```
struct mntent {
    char *mnt_fstype; /* file system name */
    char *mnt_dir; /* file system path prefix */
};
```

*Addmntent* adds the *mntent* structure *mnt* to the end of the open file *filep*. Note that *filep* has to be opened for writing if this is to work. *Endmntent* closes the file.

**RETURN VALUE**

NULL pointer (0) returned on EOF or error.

**FILES**

*/etc/mnttab*

**SEE ALSO**

*mnttab(4)*

**BUGS**

The returned *mntent* structure points to static information that is overwritten in each call.

**NAME**

getnetent, getnetbyaddr, getnetbyname, setnetent, endnetent — get network entry

**SYNOPSIS**

```
#include <netdb.h>
struct netent *getnetent()
struct netent *getnetbyname(name)
char *name;
struct netent *getnetbyaddr(net)
long net;
setnetent(stayopen)
int stayopen
endnetent()
cc ... -lnet
```

**DESCRIPTION**

*Getnetent*, *getnetbyname*, and *getnetbyaddr* each return a pointer to an object with the following structure containing the broken-out fields of a line in the network data base, */etc/networks*.

```
struct netent {
    char *n_name; /* official name of net */
    char **n_aliases; /* alias list */
    int n_addrtype; /* net number type */
    long n_net; /* net number */
};
```

The members of this structure are:

- n\_name** The official name of the network.
- n\_aliases** A zero terminated list of alternate names for the network.
- n\_addrtype** The type of the network number returned; currently only AF\_INET.
- n\_net** The network number. Network numbers are returned in machine byte order.

*Getnetent* reads the next line of the file, opening the file if necessary.

*Setnetent* opens and rewinds the file. If the *stayopen* flag is non-zero, the net data base will not be closed after each call to *getnetent* (either directly, or indirectly through one of the other getnet calls).

*Endnetent* closes the file.

*Getnetbyname* and *getnetbyaddr* sequentially search from the beginning of the file until a matching net name or net address is found, or until EOF is encountered. Network numbers are supplied in host order.

**FILES**

*/etc/networks*

**LINKING**

This library is accessed by specifying `-lnet` as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**SEE ALSO**

networks(4N)

**DIAGNOSTICS**

Null pointer (0) returned on EOF or error.

**BUGS**

All information is contained in a static area so it must be copied if it is to be saved. Only Internet network numbers are currently understood. Expecting network numbers to fit in no more than 32 bits is probably naive.

**NAME**

getopt — get option letter from argument vector

**SYNOPSIS**

```
int getopt (argc, argv, optstring)
int argc;
char **argv ; *optstring ;
extern char *optarg;
extern int optind, opterr;
```

**DESCRIPTION**

*Getopt* returns the next option letter in *argv* that matches a letter in *optstring*. *Optstring* is a string of recognized option letters; if a letter is followed by a colon, the option is expected to have an argument that may or may not be separated from it by white space. *Optarg* is set to point to the start of the option argument on return from *getopt*.

*Getopt* places in *optind* the *argv* index of the next argument to be processed. Because *optind* is external, it is normally initialized to zero automatically before the first call to *getopt*.

When all options have been processed (i.e., up to the first non-option argument), *getopt* returns EOF. The special option -- may be used to delimit the end of the options; EOF will be returned, and -- will be skipped.

**DIAGNOSTICS**

*Getopt* prints an error message on *stderr* and returns a question mark (?) when it encounters an option letter not included in *optstring*. This error message may be disabled by setting *opterr* to 0.

**EXAMPLE**

The following code fragment shows how one might process the arguments for a command that can take the mutually exclusive options **a** and **b**, and the options **f** and **o**, both of which require arguments:

```
main (argc, argv)
int argc;
char **argv;
{
    int c;
    extern int optind;
    extern char *optarg;
    :
    while ((c = getopt (argc, argv, "abf:o:")) != EOF)
        switch (c) {
            case 'a':
                if (bflg)
                    errflg++;
                else
                    aflg++;
                break;
            case 'b':
                if (aflg)
                    errflg++;
```

```

        else
            bproc ( );
        break;
    case 'f':
        ifile = optarg;
        break;
    case 'o':
        ofile = optarg;
        break;
    case '?':
        errflg ++;
    }
    if (errflg) {
        fprintf (stderr, "usage: . . . ");
        exit (2);
    }
    for ( ; optind < argc; optind++) {
        if (access (argv[optind], 4)) {
            :
        }
    }
}

```

**SEE ALSO**  
 getopt(1).

**NAME**

getpass — read a password

**SYNOPSIS**

```
char *getpass (prompt)
char *prompt;
```

**DESCRIPTION**

*Getpass* reads up to a newline or EOF from the file */dev/tty*, after prompting on the standard error output with the null-terminated string *prompt* and disabling echo. A pointer is returned to a null-terminated string of at most 8 characters. If */dev/tty* cannot be opened, a NULL pointer is returned. An interrupt terminates input and sends an interrupt signal to the calling program before returning.

**FILES**

*/dev/tty*

**SEE ALSO**

crypt(3C).

**WARNING**

The above routine uses `<stdio.h>`. This causes the size of programs not otherwise using standard I/O to increase more than might be expected.

**BUGS**

The return value points to static data whose content is overwritten by each call.

**NAME**

getprotoent, getprotobynumber, getprotobynname, setprotoent, endprotoent  
 - get protocol entry

**SYNOPSIS**

```
#include <netdb.h>

struct protoent *getprotoent()
struct protoent *getprotobynname(name)
char *name;
struct protoent *getprotobynumber(proto)
int proto;

setprotoent(stayopen)
int stayopen
endprotoent()
cc ... -lnet
```

**DESCRIPTION**

*Getprotoent*, *getprotobynname*, and *getprotobynumber* each return a pointer to an object with the following structure containing the broken-out fields of a line in the network protocol data base, */etc/protocols*.

```
struct protoent {
    char *p_name; /* official name of protocol */
    char **p_aliases; /* alias list */
    long p_proto; /* protocol number */
};
```

The members of this structure are:

**p\_name** The official name of the protocol.

**p\_aliases** A zero terminated list of alternate names for the protocol.

**p\_proto** The protocol number.

*Getprotoent* reads the next line of the file, opening the file if necessary.

*Setprotoent* opens and rewinds the file. If the *stayopen* flag is non-zero, the net data base will not be closed after each call to *getprotoent* (either directly, or indirectly through one of the other *getproto* calls).

*Endprotoent* closes the file.

*Getprotobynname* and *getprotobynumber* sequentially search from the beginning of the file until a matching protocol name or protocol number is found, or until EOF is encountered.

**FILES**

/etc/protocols

**LINKING**

This library is accessed by specifying **-lnet** as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```



**SEE ALSO**

protocols(4N)

**DIAGNOSTICS**

Null pointer (0) returned on EOF or error.

**BUGS**

All information is contained in a static area so it must be copied if it is to be saved. Only the Internet protocols are currently understood.

**NAME**

getpw — get name from UID

**SYNOPSIS**

```
int getpw (uid, buf)
int uid;
char *buf;
```

**DESCRIPTION**

*Getpw* searches the password file for a user id number that equals *uid*, copies the line of the password file in which *uid* was found into the array pointed to by *buf*, and returns 0. The line is null terminated. *Getpw* returns non-zero if *uid* cannot be found.

This routine is included only for compatibility with prior systems and should not be used; see *getpwent(3C)* for routines to use instead.

**FILES**

/etc/passwd

**SEE ALSO**

getpwent(3C), passwd(4).

**DIAGNOSTICS**

*Getpw* returns non-zero on error.

**WARNING**

The above routine uses `<stdio.h>`. Therefore, the size of programs not otherwise using standard I/O is increased more than might be expected.

**NAME**

getpwent, getpwuid, getpwnam, setpwent, endpwent, fgetpwent — get password file entry

**SYNOPSIS**

```
#include <pwd.h>

struct passwd *getpwent ( )
struct passwd *getpwuid (uid)
int uid;
struct passwd *getpwnam (name)
char *name;
void setpwent ( )
void endpwent ( )
struct passwd *fgetpwent (f)
FILE *f;
```

**DESCRIPTION**

*Getpwent*, *getpwuid*, and *getpwnam* each return a pointer to an object with the following structure containing the broken-out fields of a line in the */etc/passwd* file. Each line in the file contains a *passwd* structure, declared in the *<pwd.h>* header file:

```
struct passwd {
    char *pw_name;
    char *pw_passwd;
    int pw_uid;
    int pw_gid;
    char *pw_age;
    char *pw_comment;
    char *pw_gecos;
    char *pw_dir;
    char *pw_shell;
};
```

Because this structure is declared in *<pwd.h>*, it is not necessary to redeclare it.

The *pw\_comment* field is unused; the others have meanings described in *passwd(4)*.

When first called, *getpwent* returns a pointer to the first *passwd* structure in the file; thereafter, it returns a pointer to the next *passwd* structure in the file; therefore, successive calls can be used to search the entire file. *Getpwuid* searches from the beginning of the file until a numerical user id matching *uid* is found; it returns a pointer to the particular structure in which the match was found. *Getpwnam* searches from the beginning of the file until a login name matching *name* is found; it returns a pointer to the particular structure in which the match was found. If an end-of-file or an error is encountered on reading, these functions return a NULL pointer.

A call to *setpwent* has the effect of rewinding the password file to allow repeated searches. *Endpwent* may be called to close the password file when processing is complete.

**GETPWENT(3C)****GETPWENT(3C)**

*Fgetpwent* returns a pointer to the next passwd structure in the stream *f*, which matches the format of */etc/passwd*.

**FILES**

*/etc/passwd*

**SEE ALSO**

*cuserid(3S)*, *getlogin(3C)*, *getgrent(3C)*, *passwd(4)*.

**DIAGNOSTICS**

A NULL pointer is returned on EOF or error.

**WARNING**

The above routines use *<stdio.h>*. Therefore the size of programs not otherwise using standard I/O is increased more than might be expected.

**BUGS**

All information is contained in a static area, so it must be copied if it is to be saved.

**NAME**

gets, fgets — get a string from a stream

**SYNOPSIS**

```
#include <stdio.h>
char *gets (s)
char *s;

char *fgets (s, n, stream)
char *s;
int n;
FILE *stream;
```

**DESCRIPTION**

*Gets* reads characters from the standard input stream, *stdin*, into the array pointed to by *s*, until a new-line character is read or an end-of-file condition is encountered. The new-line character is discarded and the string is terminated with a null character.

*Fgets* reads characters from the *stream* into the array pointed to by *s* until *n*-1 characters are read, or a new-line character is read and transferred to *s*, or an end-of-file condition is encountered. The string is then terminated with a null character.

**SEE ALSO**

ferror(3S), fopen(3S), fread(3S),getc(3S), scanf(3S).

**DIAGNOSTICS**

If end-of-file is encountered and no characters have been read, no characters are transferred to *s* and a NULL pointer is returned. If a read error (e.g., trying to use these functions on a file that has not been opened for reading) occurs, a NULL pointer is returned. Otherwise *s* is returned.

**NOTE**

*Gets* deletes the new-line ending its input, but *fgets* keeps it.

## NAME

getservent, getservbyport, getservbyname, setservent, endservent — get service entry

## SYNOPSIS

```
#include <netdb.h>
struct servent *getservent()
struct servent *getservbyname(name, proto)
char *name, *proto;
struct servent *getservbyport(port, proto)
int port; char *proto;
setservent(stayopen)
int stayopen
endservent()
cc ... -lnet
```

## DESCRIPTION

*Getservent*, *getservbyname*, and *getservbyport* each return a pointer to an object with the following structure containing the broken-out fields of a line in the network services data base, */etc/services*.

```
struct servent {
    char    *s_name;        /* official name of service */
    char    **s_aliases;    /* alias list */
    long    s_port;        /* port service resides at */
    char    *s_proto;       /* protocol to use */
};
```

The members of this structure are:

**s\_name** The official name of the service.

**s\_aliases** A zero terminated list of alternate names for the service.

**s\_port** The port number at which the service resides. Port numbers are returned in network byte order.

**s\_proto** The name of the protocol to use when contacting the service.

*Getservent* reads the next line of the file, opening the file if necessary.

*Setservent* opens and rewinds the file. If the *stayopen* flag is non-zero, the net data base will not be closed after each call to *getservent* (either directly, or indirectly through one of the other *getserv* calls).

*Endservent* closes the file.

*Getservbyname* and *getservbyport* sequentially search from the beginning of the file until a matching protocol name or port number is found, or until EOF is encountered. If a protocol name is also supplied (non-NULL), searches must also match the protocol.

## FILES

*/etc/services*

## LINKING

This library is accessed by specifying `-lnet` as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**SEE ALSO**

`getprotoent(3N)`, `services(4N)`

**DIAGNOSTICS**

Null pointer (0) returned on EOF or error.

**BUGS**

All information is contained in a static area so it must be copied if it is to be saved. Expecting port numbers to fit in a 32 bit quantity is probably naive.

## NAME

getutent, getutid, getutline, pututline, setutent, endutent, utmpname – access utmp file entry

## SYNOPSIS

```
#include <sys/types.h>
#include <utmp.h>

struct utmp *getutent ( )
struct utmp *getutid (id)
struct utmp *id;

struct utmp *getutline (line)
struct utmp *line;

void pututline (utmp)
struct utmp *utmp;

void setutent ( )
void endutent ( )

void utmpname (file )
char *file;
```

## DESCRIPTION

*Getutent*, *getutid*, and *getutline* each return a pointer to a structure of the following type:

```
struct utmp {
    char    ut_user[8];        /* User login name */
    char    ut_id[4];         /* /etc/inittab id (usually line #) */
    char    ut_line[12];      /* device name (console, lxxx) */
    short   ut_pid;           /* process id */
    short   ut_type;          /* type of entry */
    struct  exit_status {
        short   e_termination; /* Process termination status */
        short   e_exit;         /* Process exit status */
    } ut_exit;                /* The exit status of a process
                               /* marked as DEAD_PROCESS. */
    time_t  ut_time;          /* time entry was made */
};
```

*Getutent* reads in the next entry from a *utmp*-like file. If the file is not already open, it opens it. If it reaches the end of the file, it fails.



*Getutid* searches forward from the current point in the *utmp* file until it finds an entry with a *ut\_type* matching *id*->*ut\_type* if the type specified is RUN\_LVL, BOOT\_TIME, OLD\_TIME, or NEW\_TIME. If the type specified in *id* is INIT\_PROCESS, LOGIN\_PROCESS, USER\_PROCESS, or DEAD\_PROCESS, *getutid* will return a pointer to the first entry whose type is one of these four and whose *ut\_id* field matches *id*->*ut\_id*. *Getutid* fails if the end of file is reached without a match.

*Getutline* searches forward from the current point in the *utmp* file until it finds an entry of the type LOGIN\_PROCESS or USER\_PROCESS which also has a *ut\_line* string matching the *line*->*ut\_line* string. If the end of file is reached without a match, it fails.

*Pututline* writes out the supplied *utmp* structure into the *utmp* file. It uses *getutid* to search forward for the proper place if it finds that it is not already at the proper place. It is assumed that the user of *pututline* has searched for the proper entry using one of the *getut* routines. If this has been done, *pututline* will not search. If *pututline* does not find a matching slot for the new entry, it will add a new entry to the end of the file.

*Setutent* resets the input stream to the beginning of the file. This should be done before each search for a new entry if it is desired that the entire file be examined.

*Endutent* closes the currently open file.

*Utmpname* allows the user to change the name of the file examined from */etc/utmp* to any other filename. It is expected that most often this other file will be */etc/wtmp*. If the file doesn't exist, this will not be apparent until the first attempt to reference the file is made. *Utmpname* does not open the file. It just closes the old file, if it is currently open, and saves the new filename.

#### FILES

*/etc/utmp*

*/etc/wtmp*

#### SEE ALSO

*ttyslot(3C)*, *utmp(4)*.

#### DIAGNOSTICS

A NULL pointer is returned upon failure to read or write. Failure to read may be due to permissions or because end-of-file has been reached.

## COMMENTS

The most current entry is saved in a static structure. Multiple accesses require that it be copied before further accesses are made. Each call to either *getutid* or *getutline* sees the routine examine the static structure before performing more I/O. If the search of the static structure results in a match, no further search is performed. To use *getutline* to search for multiple occurrences, zero out the static structure after each success; otherwise *getutline* will just return the same pointer over and over again. There is one exception to the rule about removing the structure before further reads are done. If the implicit read done by *pututline* finds that it isn't already at the correct place in the file, the contents of the static structure returned by the *getutent*, *getutid*, or *getutline* routines are not harmed, if the user has just modified those contents and passed the pointer back to *pututline*.

These routines use buffered standard I/O for input, but *pututline* uses an unbuffered non-standard write to avoid race conditions between processes trying to modify the *utmp* and *wtmp* files.

**NAME**

`hsearch`, `hcreate`, `hdestroy` — manage hash search tables

**SYNOPSIS**

```
#include <search.h>

ENTRY *hsearch (item, action)
ENTRY item;
ACTION action;

int hcreate (nel)
unsigned nel;

void hdestroy ( )
```

**DESCRIPTION**

*Hsearch* is a hash-table search routine generalized from Knuth (6.4) Algorithm D. It returns a pointer into a hash table indicating the location at which an entry can be found. *Item* is a structure of type `ENTRY` (defined in the `<search.h>` header file) containing two pointers: *item.key* points to the comparison key, and *item.data* points to any other data to be associated with that key. (Pointers to types other than character should be cast to pointer-to-character.) *Action* is a member of an enumeration type `ACTION` indicating the disposition of the entry if it cannot be found in the table. `ENTER` indicates that the item should be inserted in the table at an appropriate point. `FIND` indicates that no entry should be made. Unsuccessful resolution is indicated by the return of a `NULL` pointer.

*Hcreate* allocates sufficient space for the table, and must be called before *hsearch* is used. *Nel* is an estimate of the maximum number of entries that the table will contain. This number may be adjusted upward by the algorithm in order to obtain certain mathematically favorable circumstances.

*Hdestroy* destroys the search table, and may be followed by another call to *hcreate*.

**NOTES**

*Hsearch* uses *open addressing* with a *multiplicative* hash function. However, its source code has many other options available which the user may select by compiling the *hsearch* source with the following symbols defined to the preprocessor:

- DIV** Use the *remainder modulo table size* as the hash function instead of the multiplicative algorithm.
- USCR** Use a User Supplied Comparison Routine for ascertaining table membership. The routine should be named *hcompare* and should behave in a manner similar to *strcmp* (see *string(3C)*).
- CHAINED** Use a linked list to resolve collisions. If this option is selected, the following other options become available.
  - START** Place new entries at the beginning of the linked list (default is at the end).
  - SORTUP** Keep the linked list sorted by key in ascending order.
  - SORTDOWN** Keep the linked list sorted by key in descending order.

Additionally, there are preprocessor flags for obtaining debugging printout (-DDEBUG) and for including a test driver in the calling routine (-DDRIVER). The source code should be consulted for further details.

**EXAMPLE**

The following example will read in strings followed by two numbers and store them in a hash table, discarding duplicates. It will then read in strings and find the matching entry in the hash table and print it out.

```
#include <stdio.h>
#include <search.h>

struct info {          /* this is the info stored in the table */
    int age, room;     /* other than the key. */
};
#define NUM_EMPL      5000    /* # of elements in search table */

main( )
{
    /* space to store strings */
    char string_space[NUM_EMPL*20];
    /* space to store employee info */
    struct info info_space[NUM_EMPL];
    /* next avail space in string_space */
    char *str_ptr = string_space;
    /* next avail space in info_space */
    struct info *info_ptr = info_space;
    ENTRY item, *found_item, *hsearch( );
    /* name to look for in table */
    char name_to_find[30];
    int i = 0;

    /* create table */
    (void) hcreate(NUM_EMPL);
    while (scanf("%s%d%d", str_ptr, &info_ptr->age,
                &info_ptr->room) != EOF && i++ < NUM_EMPL) {
        /* put info in structure, and structure in item */
        item.key = str_ptr;
        item.data = (char *)info_ptr;
        str_ptr += strlen(str_ptr) + 1;
        info_ptr++;
        /* put item into table */
        (void) hsearch(item, ENTER);
    }

    /* access table */
    item.key = name_to_find;
    while (scanf("%s", item.key) != EOF) {
        if ((found_item = hsearch(item, FIND)) != NULL) {
            /* if item is in the table */
            (void)printf("found %s, age = %d, room = %d\n",
                        found_item->key,
                        (struct info *)found_item->data->age,
                        ((struct info *)found_item->data->room);
        }
    }
}
```

```
        ) else {  
            (void)printf("no such employee %s\n",  
                name_to_find)  
        }  
    )  
}
```

**SEE ALSO**

bsearch(3C), lsearch(3C), malloc(3C), malloc(3X), string(3C),  
tsearch(3C).

**DIAGNOSTICS**

*Hsearch* returns a NULL pointer if either the action is **FIND** and the item could not be found or the action is **ENTER** and the table is full.

*Hcreate* returns zero if it cannot allocate sufficient space for the table.

**WARNING**

*Hsearch* and *hcreate* use *malloc*(3C) to allocate space.

**BUGS**

Only one hash search table may be active at any given time.

**NAME**

hypot — Euclidean distance function

**SYNOPSIS**

```
#include <math.h>
double hypot (x, y)
double x, y;
```

**DESCRIPTION**

*Hypot* returns the following, taking precautions against unwarranted overflows:

$$\text{sqrt}(x * x + y * y)$$
**DIAGNOSTICS**

When the correct value would overflow, *hypot* returns HUGE and sets *errno* to ERANGE.

These error-handling procedures may be changed with the function *matherr*(3M).

**SEE ALSO**

*matherr*(3M).

**NAME***iargc***SYNOPSIS****integer i**  
**i = *iargc*()****DESCRIPTION**

The *iargc* function returns the number of command line arguments passed to the program. Thus, if a program were invoked via

**foo arg1 arg2 arg3**

***iargc*()** would return "3".

**SEE ALSO**

***getarg*(3F).**

**NAME**

*index* — return location of Fortran substring

**SYNOPSIS**

**character** \*N1 *ch1*

**character** \*N2 *ch2*

**integer** *i*

*i* = **index**(*ch1*, *ch2*)

**DESCRIPTION**

*Index* returns the location of substring *ch2* in string *ch1*. The value returned is either the position at which substring *ch2* starts or 0 if *ch2* is not present in string *ch1*.



**NAME**

*inet\_addr*, *inet\_network*, *inet\_ntoa*, *inet\_makeaddr*, *inet\_lnaof*, *inet\_netof*  
 – Internet address manipulation routines

**SYNOPSIS**

```
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>

struct in_addr inet_addr(cp)
char *cp;

int inet_network(cp)
char *cp;

char *inet_ntoa(in)
struct in_addr in;

struct in_addr inet_makeaddr(net, lna)
int net, lna;

int inet_lnaof(in)
struct in_addr in;

int inet_netof(in)
struct in_addr in;

cc ... -lnet
```

**DESCRIPTION**

The routines *inet\_addr* and *inet\_network* each interpret character strings representing numbers expressed in the Internet standard “.” notation, returning numbers suitable for use as Internet addresses and Internet network numbers, respectively. The routine *inet\_ntoa* takes an Internet address and returns an ASCII string representing the address in “.” notation. The routine *inet\_makeaddr* takes an Internet network number and a local network address and constructs an Internet address from it. The routines *inet\_netof* and *inet\_lnaof* break apart Internet host addresses, returning the network number and local network address part, respectively.

All Internet address are returned in network order (bytes ordered from left to right). All network numbers and local address parts are returned as machine format integer values.

**INTERNET ADDRESSES**

Values specified using the “.” notation take one of the following forms:

```
a.b.c.d
a.b.c
a.b
a
```

When four parts are specified, each is interpreted as a byte of data and assigned, from left to right, to the four bytes of an Internet address.

When a three part address is specified, the last part is interpreted as a 16-bit quantity and placed in the right most two bytes of the network address. This makes the three part address format convenient for specifying Class B network addresses as “128.net.host”.

When a two part address is supplied, the last part is interpreted as a 24-bit quantity and placed in the right most three bytes of the network address. This makes the two part address format convenient for specifying Class A network addresses as "net.host".

When only one part is given, the value is stored directly in the network address without any byte rearrangement.

All numbers supplied as "parts" in a "." notation may be decimal, octal, or hexadecimal, as specified in the C language (i.e. a leading 0x or 0X implies hexadecimal; otherwise, a leading 0 implies octal; otherwise, the number is interpreted as decimal).

#### LINKING

This library is accessed by specifying `-lnet` as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

#### SEE ALSO

`gethostent(3N)`, `getnetent(3N)`, `hosts(4N)`, `networks(4N)`,

#### DIAGNOSTICS

The value `-1` is returned by `inet_addr` and `inet_network` for malformed requests.

#### BUGS

The problem of host byte ordering versus network byte ordering is confusing. A simple way to specify Class C network addresses in a manner similar to that for Class B and Class A is needed. The string returned by `inet_ntoa` resides in a static memory area.

**NAME**

*insque*, *remque* — insert/remove element from a queue

**SYNOPSIS**

```
struct qelem {
    struct qelem *q_forw;
    struct qelem *q_back;
    char    q_data[];
};

insque(elem, pred)
struct qelem *elem, *pred;

remque(elem)
struct qelem *elem;

cc ... -lnet
```

**DESCRIPTION**

*Insque* and *remque* manipulate queues built from doubly linked lists. Each element in the queue must in the form of "struct qelem". *Insque* inserts *elem* in a queue immediately after *pred*; *remque* removes an entry *elem* from a queue.

**LINKING**

This library is accessed by specifying *-lnet* as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**NAME**

killpg — send signal to a process group

**SYNOPSIS**

```
killpg(pgrp, sig)
int pgrp, sig;
cc ... -lnet
```

**DESCRIPTION**

*Killpg* sends the signal *sig* to the process group *pgrp*.

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] *Sig* is not a valid signal number.
- [ESRCH] No process can be found corresponding to that specified by *pid*.
- [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.

**LINKING**

This library is accessed by specifying *-lnet* as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**SEE ALSO**

kill(2), getpid(2)

**NAME**

*l3tol*, *ltoi3* — convert between 3-byte integers and long integers

**SYNOPSIS**

```
void l3tol (lp, cp, n)
long *lp;
char *cp;
int n;

void ltoi3 (cp, lp, n)
char *cp;
long *lp;
int n;
```

**DESCRIPTION**

*L3tol* converts a list of *n* 3-byte integers (packed into a character string pointed to by *cp*) into a list of long integers pointed to by *lp*.

*Ltoi3* performs the reverse conversion from long integers (*lp*) to 3-byte integers (*cp*).

These functions are useful for file system maintenance where the block numbers are 3 bytes long.

**SEE ALSO**

*fs(4)*.

**BUGS**

Because of possible differences in byte ordering, the numerical values of the long integers are machine-dependent.

**NAME**

`ldahread` — read the archive header of a member of an archive file

**SYNOPSIS**

```
#include <stdio.h>
#include <ar.h>
#include <filehdr.h>
#include <ldfcn.h>
```

```
int ldahread (ldptr, arhead)
LDFILE *ldptr;
ARCHDR *arhead;
```

**DESCRIPTION**

If `TYPE(ldptr)` is the archive file magic number, `ldahread` reads the archive header of the common object file currently associated with `ldptr` into the area of memory beginning at `arhead`.

`Ldahread` returns `SUCCESS` or `FAILURE`. `Ldahread` fails if `TYPE(ldptr)` does not represent an archive file or if it cannot read the archive header.

The program must be loaded with the object file access routine library `libld.a`.

**SEE ALSO**

`ldclose(3X)`, `ldopen(3X)`, `ar(4)`, `ldfcn(4)`.

**NAME**

*ldclose*, *ldaclose* — close a common object file

**SYNOPSIS**

```
#include <stdio.h>
#include <filehdr.h>
#include <ldfcn.h>
```

```
int ldclose (ldptr)
LDFILE *ldptr;
```

```
int ldaclose (ldptr)
LDFILE *ldptr;
```

**DESCRIPTION**

*Ldopen(3X)* and *ldclose* are designed to provide uniform access to both simple object files and object files that are members of archive files. Thus an archive of common object files can be processed as if it were a series of simple common object files.

If *TYPE(ldptr)* does not represent an archive file, *ldclose* closes the file and frees the memory allocated to the *LDFILE* structure associated with *ldptr*. If *TYPE(ldptr)* is the magic number of an archive file, and if there are any more files in the archive, *ldclose* reinitializes *OFFSET(ldptr)* to the file address of the next archive member and returns *FAILURE*. The *LDFILE* structure is prepared for a subsequent *ldopen(3X)*. In all other cases, *ldclose* returns *SUCCESS*.

*Ldaclose* closes the file and frees the memory allocated to the *LDFILE* structure associated with *ldptr* regardless of the value of *TYPE(ldptr)*. *Ldaclose* always returns *SUCCESS*. The function is often used in conjunction with *ldaopen*.

The program must be loaded with the object file access routine library *libld.a*.

**SEE ALSO**

*fclose(3S)*, *ldopen(3X)*, *ldfcn(4)*.

**NAME**

`ldfhread` — read the file header of a common object file

**SYNOPSIS**

```
#include <stdio.h>
#include <filehdr.h>
#include <ldfcn.h>
```

```
int ldfhread (ldptr, filehead)
LDFILE *ldptr;
FILHDR *filehead;
```

**DESCRIPTION**

*Ldfhread* reads the file header of the common object file currently associated with *ldptr* into the area of memory beginning at *filehead*.

*Ldfhread* returns SUCCESS or FAILURE. *Ldfhread* fails if it cannot read the file header.

In most cases the use of *ldfhread* can be avoided by using the macro `HEADER(ldptr)` defined in `<ldfcn.h>` (see `ldfcn(4)`). The information in any field, *fieldname*, of the file header may be accessed using `HEADER(ldptr).fieldname`.

The program must be loaded with the object file access routine library `libld.a`.

**SEE ALSO**

`ldclose(3X)`, `ldopen(3X)`, `ldfcn(4)`.



**NAME**

*ldgetname* -- retrieve symbol name for object file symbol table entry

**SYNOPSIS**

```
#include <stdio.h>
#include <filehdr.h>
#include <syms.h>
#include <ldfcn.h>
```

```
char ldgetname (ldptr, symbol)
LDFILE ldptr;
SYMENT symbol;
```

**DESCRIPTION**

*Ldgetname* returns a pointer to the name associated with *symbol* as a string. The string is contained in a static buffer local to *ldgetname*. Because the buffer is overwritten by each call to *ldgetname*, it must be copied by the caller if the name is to be saved.

The common object file format has been extended to handle arbitrary length symbol names with the addition of a "string table". *Ldgetname* returns the symbol name associated with a symbol table entry for either an object file or a pre-object file. Thus, *ldgetname* can be used to retrieve names from object files without any backward compatibility problems. *Ldgetname* returns NULL (defined in <stdio.h>) for an object file if the name cannot be retrieved. This occurs when:

- the string table cannot be found.
- not enough memory can be allocated for the string table.
- the string table appears not to be a string table (e.g., if an auxiliary entry is handed to *ldgetname* that looks like a reference to a name in a non-existent string table).
- the name's offset into the string table is beyond the end of the string table.

Typically, *ldgetname* is called immediately after a successful call to *ldtbread* to retrieve the name associated with the symbol table entry filled by *ldtbread*.

The program must be loaded with the object file access routine library *libld.a*.

**SEE ALSO**

*ldclose*(3X), *ldopen*(3X), *ldtbseek*(3X), *ldtbread*(3X), *ldfcn*(4).

## NAME

`ldlread`, `ldlinit`, `ldlitem` — manipulate line number entries of a common object file function

## SYNOPSIS

```
#include <stdio.h>
#include <filehdr.h>
#include <linenum.h>
#include <ldfcn.h>
```

```
int ldlread (ldptr, fcndx, linenum, linent)
```

```
LDFILE *ldptr;
```

```
long fcndx;
```

```
unsigned short linenum;
```

```
LINENO linent;
```

```
int ldlinit (ldptr, fcndx)
```

```
LDFILE *ldptr;
```

```
long fcndx;
```

```
int ldlitem (ldptr, linenum, linent)
```

```
LDFILE *ldptr;
```

```
unsigned short linenum;
```

```
LINENO linent;
```

## DESCRIPTION

*Ldlread* searches the line number entries of the common object file currently associated with *ldptr*. *Ldlread* begins its search with the line number entry for the beginning of a function and confines its search to the line numbers associated with a single function. The function is identified by *fcndx*, the index of its entry in the object file symbol table. *Ldlread* reads the entry with the smallest line number equal to or greater than *linenum* into *linent*.

*Ldlinit* and *ldlitem* together perform exactly the same function as *ldlread*. After an initial call to *ldlread* or *ldlinit*, *ldlitem* may be used to retrieve a series of line number entries associated with a single function. *Ldlinit* simply locates the line number entries for the function identified by *fcndx*. *Ldlitem* finds and reads the entry with the smallest line number equal to or greater than *linenum* into *linent*.

*Ldlread*, *ldlinit*, and *ldlitem* each return either SUCCESS or FAILURE. *Ldlread* fails if there are no line number entries in the object file, if *fcndx* does not index a function entry in the symbol table, or if it finds no line number equal to or greater than *linenum*. *Ldlinit* fails if there are no line number entries in the object file or if *fcndx* does not index a function entry in the symbol table. *Ldlitem* fails if it finds no line number equal to or greater than *linenum*.

The programs must be loaded with the object file access routine library `libld.a`.

## SEE ALSO

`ldclose(3X)`, `ldopen(3X)`, `ldtbindx(3X)`, `ldfcn(4)`.

**NAME**

*ldlseek*, *ldnlseek* — seek to line number entries of a section of a common object file

**SYNOPSIS**

```
#include <stdio.h>
#include <filehdr.h>
#include <ldfcn.h>

int ldlseek (ldptr, sectindx)
LDFILE *ldptr;
unsigned short sectindx;

int ldnlseek (ldptr, sectname)
LDFILE *ldptr;
char *sectname;
```

**DESCRIPTION**

*Ldlseek* seeks to the line number entries of the section specified by *sectindx* of the common object file currently associated with *ldptr*.

*Ldnlseek* seeks to the line number entries of the section specified by *sectname*.

*Ldlseek* and *ldnlseek* return **SUCCESS** or **FAILURE**. *Ldlseek* fails if *sectindx* is greater than the number of sections in the object file; *ldnlseek* fails if there is no section name corresponding to *\*sectname*. Either function fails if the specified section has no line number entries or if it cannot seek to the specified line number entries.

Note that the first section has an index of *one*.

The program must be loaded with the object file access routine library *libld.a*.

**SEE ALSO**

*ldclose(3X)*, *ldopen(3X)*, *ldhread(3X)*, *ldfcn(4)*.

**NAME**

ldohseek — seek to the optional file header of a common object file

**SYNOPSIS**

```
#include <stdio.h>
#include <filehdr.h>
#include <ldfcn.h>

int ldohseek (ldptr)
LDFILE *ldptr;
```

**DESCRIPTION**

*Ldohseek* seeks to the optional file header of the common object file currently associated with *ldptr*.

*Ldohseek* returns SUCCESS or FAILURE. *Ldohseek* fails if the object file has no optional header or if it cannot seek to the optional header.

The program must be loaded with the object file access routine library libld.a.

**SEE ALSO**

ldclose(3X), ldopen(3X), ldhread(3X), ldfcn(4).

## NAME

*ldopen*, *ldaopen* — open a common object file for reading

## SYNOPSIS

```
#include <stdio.h>
#include <filehdr.h>
#include <ldfcn.h>

LDFILE *ldopen (filename, ldptr)
char *filename;
LDFILE *ldptr;

LDFILE *ldaopen (filename, oldptr)
char *filename;
LDFILE *oldptr;
```

## DESCRIPTION

*Ldopen* and *ldclose(3X)* are designed to provide uniform access to both simple object files and object files that are members of archive files. Thus, an archive of common object files can be processed as if it were a series of simple common object files.

If *ldptr* has the value NULL, *ldopen* opens *filename*, allocates and initializes the LDFILE structure, and returns a pointer to the structure to the calling program.

If *ldptr* is valid and *TYPE(ldptr)* is the archive magic number, *ldopen* reinitializes the LDFILE structure for the next archive member of *filename*.

*Ldopen* and *ldclose* are designed to work in concert. *Ldclose* returns FAILURE only when *TYPE(ldptr)* is the archive magic number and there is another file in the archive to be processed. Only then should *ldopen* be called with the current value of *ldptr*. In all other cases, in particular whenever a new *filename* is opened, *ldopen* should be called with a NULL *ldptr* argument.

The following is a prototype for the use of *ldopen* and *ldclose*.

```
/* for each filename to be processed */
ldptr = NULL;
do
    if ( (ldptr = ldopen(filename, ldptr)) != NULL )
    {
        /* check magic number */
        /* process the file */
    }
} while (ldclose(ldptr) == FAILURE);
```

If the value of *oldptr* is not NULL, *ldaopen* opens *filename* anew and allocates and initializes a new LDFILE structure, copying the TYPE, OFFSET, and HEADER fields from *oldptr*. *Ldaopen* returns a pointer to the new LDFILE structure. This new pointer is independent of the old pointer, *oldptr*. The two pointers may be used concurrently to read separate parts of the object file. For example, one pointer may be used to step sequentially through the relocation information, while the other is used to read indexed symbol table entries.

Both *ldopen* and *ldaopen* open *filename* for reading. Both functions return NULL if *filename* cannot be opened or if memory for the LDFILE structure cannot be allocated. A successful open does not insure that the given file is a common object file or an archived object file.

The program must be loaded with the object file access routine library **libld.a**.

**SEE ALSO**

**fopen(3S)**, **ldclose(3X)**, **ldfcn(4)**.

**NAME**

`ldrseek`, `ldnrseek` -- seek to relocation entries of a section of a common object file

**SYNOPSIS**

```
#include <stdio.h>
#include <filehdr.h>
#include <ldfcn.h>

int ldrseek (ldptr, sectindx)
LDFILE *ldptr;
unsigned short sectindx;

int ldnrseek (ldptr, sectname)
LDFILE *ldptr;
char *sectname;
```

**DESCRIPTION**

*Ldrseek* seeks to the relocation entries of the section specified by *sectindx* of the common object file currently associated with *ldptr*.

*Ldnrseek* seeks to the relocation entries of the section specified by *sectname*.

*Ldrseek* and *ldnrseek* return **SUCCESS** or **FAILURE**. *Ldrseek* fails if *sectindx* is greater than the number of sections in the object file; *ldnrseek* fails if there is no section name corresponding with *sectname*. Either function fails if the specified section has no relocation entries or if it cannot seek to the specified relocation entries.

Note that the first section has an index of *one*.

The program must be loaded with the object file access routine library `libld.a`.

**SEE ALSO**

`ldclose(3X)`, `ldopen(3X)`, `ldshread(3X)`, `ldfcn(4)`.

**NAME**

`ldshread`, `ldnshread` — read an indexed/named section header of a common object file

**SYNOPSIS**

```
#include <stdio.h>
#include <filehdr.h>
#include <scnhdr.h>
#include <ldfcn.h>

int ldshread (ldptr, sectindx, secthead)
LDFILE *ldptr;
unsigned short sectindx;
SCNHDR *secthead;

int ldnshread (ldptr, sectname, secthead)
LDFILE *ldptr;
char *sectname;
SCNHDR *secthead;
```

**DESCRIPTION**

*Ldshread* reads the section header specified by *sectindx* of the common object file currently associated with *ldptr* into the area of memory beginning at *secthead*.

*Ldnshread* reads the section header specified by *sectname* into the area of memory beginning at *secthead*.

*Ldshread* and *ldnshread* return SUCCESS or FAILURE. *Ldshread* fails if *sectindx* is greater than the number of sections in the object file; *ldnshread* fails if there is no section name corresponding with *sectname*. Either function fails if it cannot read the specified section header.

Note that the first section header has an index of *one*.

The program must be loaded with the object file access routine library `libld.a`.

**SEE ALSO**

`ldclose(3X)`, `ldopen(3X)`, `ldfcn(4)`.



**NAME**

*ldsseek*, *ldnsseek* — seek to an indexed/named section of a common object file

**SYNOPSIS**

```
#include <stdio.h>
#include <filehdr.h>
#include <ldfcn.h>

int ldsseek (ldptr, sectindx)
LDFILE *ldptr;
unsigned short sectindx;

int ldnsseek (ldptr, sectname)
LDFILE *ldptr;
char *sectname;
```

**DESCRIPTION**

*Ldsseek* seeks to the section specified by *sectindx* of the common object file currently associated with *ldptr*.

*Ldnsseek* seeks to the section specified by *sectname*.

*Ldsseek* and *ldnsseek* return SUCCESS or FAILURE. *Ldsseek* fails if *sectindx* is greater than the number of sections in the object file; *ldnsseek* fails if there is no section name corresponding with *sectname*. Either function fails if there is no section data for the specified section or if it cannot seek to the specified section.

Note that the first section has an index of *one*.

The program must be loaded with the object file access routine library *libld.a*.

**SEE ALSO**

*ldclose(3X)*, *ldopen(3X)*, *ldshread(3X)*, *ldfcn(4)*.

**NAME**

*ldtbindex* - compute the index of a symbol table entry of a common object file

**SYNOPSIS**

```
#include <stdio.h>
#include <filehdr.h>
#include <syms.h>
#include <ldfcn.h>

long ldtbindex (ldptr)
LDFILE *ldptr;
```

**DESCRIPTION**

*Ldtbindex* returns the (long) index of the symbol table entry at the current position of the common object file associated with *ldptr*.

The index returned by *ldtbindex* may be used in subsequent calls to *ldtbread*(3X). However, since *ldtbindex* returns the index of the symbol table entry that begins at the current position of the object file, if *ldtbindex* is called immediately after a particular symbol table entry has been read, it returns the the index of the next entry.

*Ldtbindex* fails if there are no symbols in the object file or if the object file is not positioned at the beginning of a symbol table entry.

Note that the first symbol in the symbol table has an index of zero.

The program must be loaded with the object file access routine library *libld.a*.

**SEE ALSO**

*ldclose*(3X), *ldopen*(3X), *ldtbread*(3X), *ldtbseek*(3X), *ldfcn*(4).

**NAME**

ldtbread — read an indexed symbol table entry of a common object file

**SYNOPSIS**

```
#include <stdio.h>
#include <filehdr.h>
#include <syms.h>
#include <ldfcn.h>
```

```
int ldtbread (ldptr, symindex, symbol)
LDFILE *ldptr;
long symindex;
SYMENT *symbol;
```

**DESCRIPTION**

*Ldtbread* reads the symbol table entry specified by *symindex* of the common object file currently associated with *ldptr* into the area of memory beginning at *symbol*.

*Ldtbread* returns **SUCCESS** or **FAILURE**. *Ldtbread* fails if *symindex* is greater than the number of symbols in the object file or if it cannot read the specified symbol table entry.

Note that the first symbol in the symbol table has an index of zero.

The program must be loaded with the object file access routine library **libld.a**.

**SEE ALSO**

**ldclose(3X)**, **ldgetname(3X)**, **ldopen(3X)**, **ldtbseek(3X)**, **ldgetname(3X)**, **ldfcn(4)**.

**NAME**

ldtbseek — seek to the symbol table of a common object file

**SYNOPSIS**

```
#include <stdio.h>
#include <filehdr.h>
#include <ldfcn.h>

int ldtbseek (ldptr)
LDFILE *ldptr;
```

**DESCRIPTION**

*Ldtbseek* seeks to the symbol table of the object file currently associated with *ldptr*.

*Ldtbseek* returns SUCCESS or FAILURE. *Ldtbseek* fails if the symbol table has been stripped from the object file or if it cannot seek to the symbol table.

The program must be loaded with the object file access routine library *libld.a*.

**SEE ALSO**

ldclose(3X), ldopen(3X), ldtbread(3X), ldfcn(4).

**LEN (3F)**

**LEN (3F)**

**NAME**

len — return length of Fortran string

**SYNOPSIS**

character\*N ch

integer i

i = len(ch)

**DESCRIPTION**

*Len* returns the length of string *ch*.

## NAME

lockf – record locking on files

## SYNOPSIS

```
# include <unistd.h>
```

```
lockf (fildes, function, size) long size;
function;
```

## DESCRIPTION

The *lockf* call will allow sections of a file to be locked (advisory write locks; mandatory or enforcement mode record locks are not currently available). Locking calls from other processes which attempt to lock the locked file section will either return an error value or be put to sleep until the resource becomes unlocked. All the locks for a process are removed when the process terminates. [See *fcntl(2)* for more information about record locking.]

*Fildes* is an open file descriptor. The file descriptor must have O\_WRONLY or O\_RDWR permission in order to establish lock with this function call.

*Function* is a control value which specifies the action to be taken. The permissible values for *function* are defined in <unistd.h> as follows:

```
#define F_ULOCK    0    /* Unlock a previously locked section */
#define F_LOCK     1    /* Lock a section for exclusive use */
#define F_TLOCK    2    /* Test and lock a section for exclusive use */
#define F_TEST     3    /* Test section for other processes locks */
```

All other values of *function* are reserved for future extensions and will result in an error return if not implemented.

F\_TEST is used to detect if a lock by another process is present on the specified section. F\_LOCK and F\_TLOCK both lock a section of a file if the section is available. F\_ULOCK removes locks from a section of the file.

*Size* is the number of contiguous bytes to be locked or unlocked. The resource to be locked starts at the current offset in the file and extends forward for a positive size and backward for a negative size. If *size* is zero, the section from the current offset through the largest file offset is locked (i.e., from the current offset through the present or any future end-of-file). An area need not be allocated to the file in order to be locked, as such locks may exist past the end-of-file.

The sections locked with `F_LOCK` or `F_TLOCK` may, in whole or in part, contain or be contained by a previously locked section for the same process. When this occurs, or if adjacent sections occur, the sections are combined into a single section. If the request requires that a new element be added to the table of active locks and this table is already full, an error is returned, and the new section is not locked.

`F_LOCK` and `F_TLOCK` requests differ only by the action taken if the resource is not available. `F_LOCK` will cause the calling process to sleep until the resource is available. `F_TLOCK` will cause the function to return a `-1` and set *errno* to `[EACCES]` error if the section is already locked by another process.

`F_ULOCK` requests may, in whole or in part, release one or more locked sections controlled by the process. When sections are not fully released, the remaining sections are still locked by the process. Releasing the center section of a locked section requires an additional element in the table of active locks. If this table is full, an `[EDEADLK]` error is returned and the requested section is not released.

A potential for deadlock occurs if a process controlling a locked resource is put to sleep by accessing another process's locked resource. Thus calls to *lock* or *fcntl* scan for a deadlock prior to sleeping on a locked resource. An error return is made if sleeping on the locked resource would cause a deadlock.

Sleeping on a resource is interrupted with any signal. The *alarm(2)* command may be used to provide a timeout facility in applications which require this facility.

#### ERRORS

The *lockf* utility will fail if one or more of the following are true:

- |                        |   |
|------------------------|---|
| <code>[EBADF]</code>   | <i>Fildes</i> is not a valid open descriptor.   |
| <code>[EACCES]</code>  | <i>Cmd</i> is <code>F_TLOCK</code> or <code>F_TEST</code> and the section is already locked by another process.   |
| <code>[EDEADLK]</code> | <i>Cmd</i> is <code>F_LOCK</code> or <code>F_TLOCK</code> and a deadlock would occur. Also the <i>cmd</i> is either of the above or <code>F_ULOCK</code> and the number of entries in the lock table would exceed the number allocated on the system. |

#### 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.

**CAVEATS**

Unexpected results may occur in processes that do buffering in the user address space. The process may later read/write data which is/was locked. The standard I/O package is the most common source of unexpected buffering.

**SEE ALSO**

close(2), creat(2), fcntl(2), intro(2), open(2), read(2), write(2).



**NAME**

*log*, *alog*, *dlog*, *clog* — Fortran natural logarithm intrinsic function

**SYNOPSIS**

**real** *r1*, *r2*  
**double precision** *dp1*, *dp2*  
**complex** *cx1*, *cx2*  
*r2* = **alog**(*r1*)  
*r2* = **log**(*r1*)  
*dp2* = **dlog**(*dp1*)  
*dp2* = **log**(*dp1*)  
*cx2* = **clog**(*cx1*)  
*cx2* = **log**(*cx1*)

**DESCRIPTION**

*Alog* returns the real natural logarithm of its real argument. *Dlog* returns the double-precision natural logarithm of its double-precision argument. *Clog* returns the complex logarithm of its complex argument. The generic function *log* becomes a call to *alog*, *dlog*, or *clog* depending on the type of its argument.

**SEE ALSO**

*exp*(3M).

**NAME**

log10, alog10, dlog10 — Fortran common logarithm intrinsic function

**SYNOPSIS**

```
real r1, r2
double precision dp1, dp2
r2 = alog10(r1)
r2 = log10(r1)
dp2 = dlog10(dp1)
dp2 = log10(dp1)
```

**DESCRIPTION**

*Alog10* returns the real common logarithm of its real argument. *Dlog10* returns the double-precision common logarithm of its double-precision argument. The generic function *log10* becomes a call to *alog10* or *dlog10* depending on the type of its argument.

**SEE ALSO**

exp(3M).

**NAME**

logname — return login name of user

**SYNOPSIS**

char \*logname( )

**DESCRIPTION**

*Logname* returns a pointer to the null-terminated login name; it extracts the **\$LOGNAME** variable from the user's environment.

This routine is kept in **/lib/libPW.a**.

**FILES**

**/etc/profile**

**SEE ALSO**

**env(1), login(1), profile(4), environ(5).**

**BUGS**

The return values point to static data whose content is overwritten by each call.

This method of determining a login name is subject to forgery.

## NAME

`lsearch`, `lfind` — linear search and update

## SYNOPSIS

```
#include <stdio.h>
#include <search.h>

char *lsearch ((char *)key, (char *)base, nelp, width, compar)
unsigned *nelp, width;
int (*compar)();

char *lfind ((char *)key, (char *)base, nelp, width compar)
unsigned *nelp, width;
int (*compar)();
```

## DESCRIPTION

*lsearch* is a linear search routine generalized from Knuth (6.1) Algorithm S. It returns a pointer into a table indicating where a datum may be found. If the datum does not occur, it is added at the end of the table. **Key** points to the datum to be sought in the table. **Base** points to the first element in the table. **Nelp** points to an integer containing the current number of elements in the table. The integer is incremented if the datum is added to the table. **Width** is the width of an element in bytes; *sizeof (\*key)* should be used. **Compar** is the name of the comparison function which the user must supply (*strcmp*, for example). It is called with two arguments that point to the elements being compared. The function must return zero if the elements are equal and non-zero otherwise.

*lfind* is the same as *lsearch* except that if the datum is not found, it is not added to the table. Instead, a `-1` pointer is returned.

## NOTES

The pointers to the key and the element at the base of the table should be of type pointer-to-element, and cast to type pointer-to-character. The comparison function need not compare every byte, so arbitrary data may be contained in the elements in addition to the values being compared. Although declared as type pointer-to-character, the value returned should be cast into type pointer-to-element.

## EXAMPLE

This fragment will read in  $\leq$  TABSIZE strings of length  $\leq$  ELSIZE and store them in a table, eliminating duplicates.

```
#include <stdio.h>
#include <search.h>

#define TABSIZE 50
#define ELSIZE 120

char line[ELSIZE], tab[TABSIZE][ELSIZE], *lsearch ();
unsigned nel = 0;
int strcmp ();
...
while (fgets(line, ELSIZE, stdin) != NULL &&
      nel < TABSIZE)
    (void) lsearch(line, (char *)tab, &nel,
                  ELSIZE, strcmp);
...
```

**SEE ALSO**

bsearch(3C), hsearch(3C), tsearch(3C).

**DIAGNOSTICS**

If the searched for datum is found, both *lsearch* and *lfind* return a pointer to it. Otherwise, *lfind* returns NULL and *lsearch* returns a pointer to the newly added element.

**BUGS**

Undefined results can occur if there is not enough room in the table to add a new item.

**NAME**

`malloc`, `free`, `realloc`, `calloc` — main memory allocator

**SYNOPSIS**

```
char *malloc (size)
unsigned size;

void free (ptr)
char *ptr;

char *realloc (ptr, size)
char *ptr;
unsigned size;

char *calloc (nelem, elsize)
unsigned nelem, elsize;

cfree (ptr, nelem, elsize)
char *ptr,
unsigned nelem, elsize;
```

**DESCRIPTION**

*Malloc* and *free* provide a simple general-purpose memory allocation package. *Malloc* returns a pointer to a block of at least *size* bytes suitably aligned for any use.

The argument to *free* is a pointer to a block previously allocated by *malloc*; after *free* is performed this space is made available for further allocation, but its contents are left undisturbed.

Undefined results occur if the space assigned by *malloc* is overrun or if some random number is handed to *free*.

*Malloc* allocates the first contiguous reach of free space of sufficient size found in a circular search from the last block allocated or freed; it coalesces adjacent free blocks as it searches. It calls *sbrk* (see *brk(2)*) to get more memory from the system when there is no suitable space already free.

*Realloc* changes the size of the block pointed to by *ptr* to *size* bytes and returns a pointer to the (possibly moved) block. The contents are unchanged up to the lesser of the new and old sizes. If no free block of *size* bytes is available in the storage arena, *realloc* asks *malloc* to enlarge the arena by *size* bytes and then moves the data to the new space.

*Realloc* also works if *ptr* points to a block freed since the last call of *malloc*, *realloc*, or *calloc*; thus sequences of *free*, *malloc*, and *realloc* can exploit the search strategy of *malloc* to do storage compaction.

*Calloc* allocates space for an array of *nelem* elements of size *elsize*. The space is initialized to zeros.

The arguments to *cfree* are the pointer to a block previously allocated by *calloc* plus the parameters to *calloc*.

Each of the allocation routines returns a pointer to space suitably aligned (after possible pointer coercion) for storage of any type of object.

**DIAGNOSTICS**

*Malloc*, *realloc*, and *calloc* return a NULL pointer if there is no available memory or if the arena has been detectably corrupted by storing outside the bounds of a block. When this happens the block pointed to by *ptr* may be destroyed.

**NOTE**

Search time increases when many objects have been allocated; i.e., if a program allocates space but never frees it, each successive allocation takes longer.

**SEE ALSO**

*brk(2)*, *malloc(3X)*. For an alternate, more flexible implementation, see *malloc(3X)*.

## NAME

`malloc`, `free`, `realloc`, `calloc`, `malloc`, `mallinfo` — fast main memory allocator

## SYNOPSIS

```
#include <malloc.h>
char *malloc (size)
unsigned size;
void free (ptr)
char *ptr;
char *realloc (ptr, size)
char *ptr;
unsigned size;
char *calloc (nelem, elsize)
unsigned nelem, elsize;
int malloc (cmd, value)
int cmd, value;
struct mallinfo mallinfo (max)
int max;
```

## DESCRIPTION

*Malloc* and *free* provide a simple general-purpose memory allocation package, which runs considerably faster than the *malloc(3C)* package. It is found in the library “*malloc*”, and is loaded if the option “*-lmalloc*” is used with *cc(1)* or *ld(1)*.

*Malloc* returns a pointer to a block of at least *size* bytes suitably aligned for any use.

The argument to *free* is a pointer to a block previously allocated by *malloc*; after *free* is performed this space is made available for further allocation, and its contents have been destroyed (but see *malloc* below for a way to change this behavior).

Undefined results will occur if the space assigned by *malloc* is overrun or if some random number is handed to *free*.

*Realloc* changes the size of the block pointed to by *ptr* to *size* bytes and returns a pointer to the (possibly moved) block. The contents will be unchanged up to the lesser of the new and old sizes.

*Calloc* allocates space for an array of *nelem* elements of size *elsize*. The space is initialized to zeros.

*Malloc* provides for control over the allocation algorithm. The available values for *cmd* are:

- |          |  |
|----------|--|
| M_MXFAST | Set <i>maxfast</i> to <i>value</i> . The algorithm allocates all blocks below the size of <i>maxfast</i> in large groups and then does them out very quickly. The default value for <i>maxfast</i> is 0. |
| M_NLBLKS | Set <i>numlblks</i> to <i>value</i> . The above mentioned “large groups” each contain <i>numlblks</i> blocks. <i>Numlblks</i> must be greater than 0. The default value for <i>numlblks</i> is 100.      |
| M_GRAIN  | Set <i>grain</i> to <i>value</i> . The sizes of all blocks smaller than <i>maxfast</i> are considered to be rounded up to the nearest multiple   |



of *grain*. *Grain* must be greater than 0. The default value of *grain* is the smallest number of bytes which will allow alignment of any data type. Value will be rounded up to a multiple of the default when *grain* is set.

**M\_KEEP** Preserve data in a freed block until the next *malloc*, *realloc*, or *calloc*. This option is provided only for compatibility with the old version of *malloc* and is not recommended.

These values are defined in the `<malloc.h>` header file.

*Mallot* may be called repeatedly, but may not be called after the first small block is allocated.

*Mallinfo* provides instrumentation describing space usage. It returns the structure:

```
struct mallinfo {
    int arena;          /* total space in arena */
    int ordblks;       /* number of ordinary blocks */
    int smblks;        /* number of small blocks */
    int hblkhd;        /* space in holding block headers */
    int hblks;         /* number of holding blocks */
    int usmblks;       /* space in small blocks in use */
    int fsmblks;       /* space in free small blocks */
    int uordblks;      /* space in ordinary blocks in use */
    int fordblks;      /* space in free ordinary blocks */
    int keepcost;      /* space penalty if keep option */
                      /* is used */
}
```

This structure is defined in the `<malloc.h>` header file.

Each of the allocation routines returns a pointer to space suitably aligned (after possible pointer coercion) for storage of any type of object.

**SEE ALSO**

`brk(2)`, `malloc(3C)`.

**DIAGNOSTICS**

*Malloc*, *realloc* and *calloc* return a NULL pointer if there is not enough available memory. When *realloc* returns NULL, the block pointed to by *ptr* is left intact. If *mallot* is called after any allocation or if *cmd* or *value* are invalid, non-zero is returned. Otherwise, it returns zero.

**WARNINGS**

This package usually uses more data space than *malloc(3C)*.

The code size is also bigger than *malloc(3C)*.

Note that unlike *malloc(3C)*, this package does not preserve the contents of a block when it is freed, unless the M\_KEEP option of *mallot* is used.

Undocumented features of *malloc(3C)* have not been duplicated.

## NAME

`matherr` — error-handling function

## SYNOPSIS

```
#include <math.h>

int matherr (x)
struct exception *x;
```

## DESCRIPTION

*Matherr* is invoked by functions in the Math Library when errors are detected. Users may define their own procedures for handling errors, by including a function named *matherr* in their programs. *Matherr* must be of the form described above. When an error occurs, a pointer to the exception structure *x* will be passed to the user-supplied *matherr* function. This structure, which is defined in the *<math.h>* header file, is as follows:

```
struct exception {
    int type;
    char *name;
    double arg1, arg2, retval;
};
```

The element *type* is an integer describing the type of error that has occurred, from the following list of constants (defined in the header file):

DOMAIN	argument domain error
SING	argument singularity
OVERFLOW	overflow range error
UNDERFLOW	underflow range error
TLOSS	total loss of significance
PLOSS	partial loss of significance

The element *name* points to a string containing the name of the function that incurred the error. The variables *arg1* and *arg2* are the arguments with which the function was invoked. *Retval* is set to the default value that will be returned by the function unless the user's *matherr* sets it to a different value.

If the user's *matherr* function returns non-zero, no error message will be printed, and *errno* will not be set.

If *matherr* is not supplied by the user, the default error-handling procedures, described with the math functions involved, will be invoked upon error. These procedures are also summarized in the table below. In every case, *errno* is set to EDOM or ERANGE and the program continues.

## EXAMPLE

```
#include <math.h>

int
matherr(x)
register struct exception *x;
{
    switch (x->type) {
        case DOMAIN:
            /* change sqrt to return sqrt(-arg1), not 0 */
```

```

if (!strcmp(x->name, "sqrt")) {
    x->retval = sqrt(-x->arg1);
    return (0); /* print message and set errno */
}
case SING:
    /* all other domain or sing errors, print message and abort */
    fprintf(stderr, "domain error in %s\n", x->name);
    abort();
case PLOSS:
    /* print detailed error message */
    fprintf(stderr, "loss of significance in %s(%g) = %g\n",
        x->name, x->arg1, x->retval);
    return (1); /* take no other action */
}
return (0); /* all other errors, execute default procedure */
    
```

**DEFAULT ERROR HANDLING PROCEDURES**

type	Types of Errors					
	DOMAIN	SING	OVERFLOW	UNDERFLOW	TLOSS	PLOSS
<i>errno</i>	EDOM	EDOM	ERANGE	ERANGE	ERANGE	ERANGE
BESSEL:	-	-	-	-	M, 0	*
y0, y1, yn (arg ≤ 0)	M, -H	-	-	-	-	-
EXP:	-	-	II	0	-	-
LOG, LOG10:						
(arg < 0)	M, -H	-	-	-	-	-
(arg = 0)	-	M, -H	-	-	-	-
POW:						
neg ** non-int	M, 0	-	±H	0	-	-
0 ** non-pos	-	-	-	-	-	-
SQRT:	M, 0	-	-	-	-	-
GAMMA:	-	M, H	H	-	-	-
HYPOT:	-	-	H	-	-	-
SINH:	-	-	±H	-	-	-
COSH:	-	-	H	-	-	-
SIN, COS, TAN:	-	-	-	-	M, 0	*
ASIN, ACOS, ATAN2:	M, 0	-	-	-	-	-

**ABBREVIATIONS**

- \* As much as possible of the value is returned.
- M Message is printed (EDOM error).
- H HUGE is returned.
- H -HUGE is returned.
- ±H HUGE or -HUGE is returned.
- 0 0 is returned.

## NAME

max, max0, amax0, max1, amax1, dmax1 — Fortran maximum-value functions

## SYNOPSIS

```
Integer i, j, k, l  
real a, b, c, d  
double precision dp1, dp2, dp3  
  
l = max(i, j, k)  
c = max(a, b)  
dp = max(a, b, c)  
k = max0(i, j)  
a = amax0(i, j, k)  
i = max1(a, b)  
d = amax1(a, b, c)  
dp3 = dmax1(dp1, dp2)
```

## DESCRIPTION

The maximum-value functions return the largest of their arguments; there may be any number of arguments. *Max* is the generic form which can be used for all data types and takes its return type from that of its arguments. All arguments must be of the same type. *Max0* returns the integer form of the maximum value of its integer arguments; *amax0*, the real form of its integer arguments; *max1*, the integer form of its real arguments; *amax1*, the real form of its real arguments; and *dmax1*, the double-precision form of its double-precision arguments.

## SEE ALSO

min(3F).

**NAME**

*mclock* — return Fortran time accounting

**SYNOPSIS**

integer i

i = *mclock* ( )

**DESCRIPTION**

*Mclock* returns time accounting information about the current process and its child processes. The value returned is the sum of the current process's user time and the user and system times of all child processes.

**SEE ALSO**

*times*(2), *clock*(3C), *system*(3F).

## NAME

memccpy, memchr, memcmp, memcpy, memset — memory operations

## SYNOPSIS

```
#include <memory.h>

char *memccpy (s1, s2, c, n)
char *s1, *s2;
int c, n;

char *memchr (s, c, n)
char *s;
int c, n;

int memcmp (s1, s2, n)
char *s1, *s2;
int n;

char *memcpy (s1, s2, n)
char *s1, *s2;
int n;

char *memset (s, c, n)
char *s;
int c, n;
```

## DESCRIPTION

These functions operate efficiently on memory areas (arrays of characters bounded by a count, not terminated by a null character). They do not check for the overflow of any receiving memory area.

*Memccpy* copies characters from memory area *s2* into *s1*, stopping after the first occurrence of character *c* has been copied or after *n* characters have been copied, whichever comes first. It returns either a pointer to the character after the copy of *c* in *s1* or a NULL pointer if *c* was not found in the first *n* characters of *s2*.

*Memchr* returns either a pointer to the first occurrence of character *c* in the first *n* characters of memory area *s* or a NULL pointer if *c* does not occur.

*Memcmp* compares its arguments, looking at the first *n* characters only. It returns an integer less than, equal to, or greater than 0, depending on whether *s1* is lexicographically less than, equal to, or greater than *s2*.

*Memcpy* copies *n* characters from memory area *s2* to *s1*. It returns *s1*.

*Memset* sets the first *n* characters in memory area *s* to the value of character *c*. It returns *s*.

## NOTE

For user convenience, all these functions are declared in the optional `<memory.h>` header file.

## BUGS

*Memcmp* uses native character comparison.

Because character movement is performed differently in different implementations, overlapping moves may yield unexpected results.

**NAME**

*min*, *min0*, *amin0*, *min1*, *amin1*, *dmin1* - Fortran minimum-value functions

**SYNOPSIS**

```
integer i, j, k, l
real a, b, c, d
double precision dp1, dp2, dp3

l = min(i, j, k)
c = min(a, b)
dp = min(a, b, c)
k = min0(i, j)
a = amin0(i, j, k)
i = min1(a, b)
d = amin1(a, b, c)
dp3 = dmin1(dp1, dp2)
```

**DESCRIPTION**

The minimum-value functions return the minimum of their arguments. There may be any number of arguments. *Min* is the generic form which can be used for all data types. It takes its return type from that of its arguments, which must all be of the same type. *Min0* returns the integer form of the minimum value of its integer arguments; *amin0*, the real form of its integer arguments; *min1*, the integer form of its real arguments; *amin1*, the real form of its real arguments; and *dmin1*, the double-precision form of its double-precision arguments.

**SEE ALSO**

*max*(3F).

**NAME**

`mktemp` — make a unique filename

**SYNOPSIS**

```
char *mktemp (template)
char *template;
```

**DESCRIPTION**

*Mktemp* replaces the contents of the string pointed to by *template* with a unique filename; it returns the address of *template*. The string in *template* should look like a filename with six trailing Xs; *mktemp* replaces the Xs with a letter and the current process ID. The letter is chosen so that the resulting name does not duplicate an existing file.

**SEE ALSO**

`getpid(2)`, `tmpfile(3S)`, `tmpnam(3S)`.

**BUGS**

It is possible to run out of letters.



**NAME**

mod, amod, dmod — Fortran remaindering intrinsic functions

**SYNOPSIS**

integer i, j, k  
real r1, r2, r3  
double precision dp1, dp2, dp3  
k = mod(i, j)  
r3 = amod(r1, r2)  
r3 = mod(r1, r2)  
dp3 = dmod(dp1, dp2)  
dp3 = mod(dp1, dp2)

**DESCRIPTION**

*Mod* returns the integer remainder of its first argument divided by its second argument. *Amod* and *dmod* return, respectively, the real and double-precision whole number remainder of the integer division of their two arguments. The generic version *mod* returns the data type of its arguments.

**NAME**

monitor — prepare execution profile

**SYNOPSIS**

```
#include <mon.h>

void monitor (lowpc, highpc, buffer, bufsize, nfunc)
int (*lowpc)(), (*highpc)();
WORD *buffer;
int bufsize, nfunc;
```

**DESCRIPTION**

An executable program created by `cc -p` automatically includes calls for *monitor* with default parameters; *monitor* needn't be called explicitly except to gain fine control over profiling.

*Monitor* is an interface to *profil(2)*. *Lowpc* and *highpc* are the addresses of two functions; *buffer* is the address of a (user supplied) array of *bufsize* WORD (defined in the `<mon.h>` header file). *Monitor* arranges to record a histogram in the buffer. This histogram shows periodically sampled values of the program counter and counts of calls of certain functions. The lowest address sampled is that of *lowpc*; the highest address is just below *highpc*. *Lowpc* may not equal 0 for this use of *monitor*. *Nfunc* is the maximum number of call counts that can be kept; only calls of functions compiled with the profiling option `-p` of `cc(1)` are recorded. (The C Library and Math Library supplied when `cc -p` is used also have call counts recorded.) For the results to be significant, especially where there are small, heavily used routines, it is suggested that the buffer be no more than a few times smaller than the range of locations sampled.

To profile the entire program, it is sufficient to use

```
extern etext;
...
monitor ((int (*)())2, etext, buf, bufsize, nfunc);
```

*Etext* lies just above all the program text; see *end(3C)*.

To stop execution monitoring and write the results on the file `mon.out`, use

```
monitor ((int (*)())0, 0, 0, 0, 0);
```

*Prof(1)* can then be used to examine the results.

**FILES**

```
mon.out
/lib/libp/libc.a
/lib/libp/libm.a
```

**SEE ALSO**

`cc(1)`, `prof(1)`, `profil(2)`, `end(3C)`.

**NAME**

*nlist* - get entries from name list

**SYNOPSIS**

```
#include <a.out.h>

int nlist (filename, nl)
char *filename;
struct nlist nl
```

**DESCRIPTION**

*Nlist* examines the name list in the executable file whose name is pointed to by *filename*; it selectively extracts a list of values and puts them in the array of *nlist* structures pointed to by *nl*. The name list *nl* consists of an array of structures containing names of variables, types, and values. The list is terminated with a null name; i.e., a null string is in the name position of the structure. Each variable name is looked up in the name list of the file. If the name is found, the type and value of the name are inserted in the next two fields. The type field will be set to 0 unless the file was compiled with the *-g* option. If the name is not found, both entries are set to 0. See *a.out(4)* for a discussion of the symbol table structure.

This function is useful for examining the system name list kept in the file */unix*. In this way programs can obtain system addresses that are up to date.

**SEE ALSO**

*a.out(4)*.

**DIAGNOSTICS**

All value entries are set to 0 if the file cannot be read or if it does not contain a valid name list.

*Nlist* returns *-1* upon error; otherwise it returns 0.

**NAME**

`perror`, `errno`, `sys_errlist`, `sys_nerr` — system error messages

**SYNOPSIS**

```
void perror (s)
char *s;
extern int errno;
extern char *sys_errlist[];
extern int sys_nerr;
```

**DESCRIPTION**

*Perror* produces a message on the standard error output, describing the last error encountered during a call to a system or library function. The argument string *s* is printed first, then a colon and a blank, then the message and a new-line. To be of most use, the argument string should include the name of the program that incurred the error. The error number is taken from the external variable *errno*, which is set when errors occur but not cleared when non-erroneous calls are made.

To simplify variant formatting of messages, the array of message strings *sys\_errlist* is provided; *errno* can be used as an index in this table to get the message string without the new-line. *sys\_nerr* is the largest message number provided for in the table; it should be checked because new error codes may be added to the system before they are added to the table.

**SEE ALSO**

`intro(2)`.

## NAME

plot — graphics interface subroutines

## SYNOPSIS

```

openpl ()
erase ()
label (s)
char *s;
line (x1, y1, x2, y2)
int x1, y1, x2, y2;
circle (x, y, r)
int x, y, r;
arc (x, y, x0, y0, x1, y1)
int x, y, x0, y0, x1, y1;
move (x, y)
int x, y;
cont (x, y)
int x, y;
point (x, y)
int x, y;
linemod (s)
char *s;
space (x0, y0, x1, y1)
int x0, y0, x1, y1;
closepl ()

```

## DESCRIPTION

These subroutines generate graphic output in a relatively device-independent manner. *Space* must be used before any of these functions to declare the amount of space necessary; see *plot(4)*. *Openpl* must be used before any of the others to open the device for writing. *Closepl* flushes the output.

*Circle* draws a circle of radius *r* with center at the point  $(x, y)$ .

*Arc* draws an arc of a circle with center at the point  $(x, y)$  between the points  $(x0, y0)$  and  $(x1, y1)$ .

String arguments to *label* and *linemod* are terminated by nulls and do not contain new-lines.

See *plot(4)* for a description of the effect of the remaining functions.

The library files listed below provide several variations of these routines.

## FILES

/usr/lib/libplot.a	produces output for <i>tplot(1G)</i> filters
/usr/lib/lib300.a	for DASI 300
/usr/lib/lib300s.a	for DASI 300s
/usr/lib/lib450.a	for DASI 450
/usr/lib/lib4014.a	for Tektronix 4014

## WARNINGS

To compile a program containing these functions in *file.c*, use `cc file.c -lplot`

**PLOT(3X)**

**PLOT(3X)**

To execute it, use `a.out | tplot`.

The above routines use `<stdio.h>`. Therefore, the size of programs not otherwise using standard I/O is increased more than might be expected.

**SEE ALSO**

`tplot(1G)`, `plot(4)`.

**NAME**

*popen*, *pclose* — initiate pipe to/from a process

**SYNOPSIS**

```
#include <stdio.h>
FILE *popen (command, type)
char *command, *type;
int pclose (stream)
FILE *stream;
```

**DESCRIPTION**

The arguments to *popen* are pointers to null-terminated strings; one string contains a shell command line and the other contains an I/O mode. The mode may be either *r* for reading or *w* for writing. *Popen* creates a pipe between the calling program and the command to be executed. The value returned is a stream pointer. If the I/O mode is *w*, one can write to the standard input of the command by writing to the file *stream*; if the I/O mode is *r*, one can read from the standard output of the command, by reading from the file *stream*.

A stream opened by *popen* should be closed by *pclose*, which waits for the associated process to terminate and returns the exit status of the command.

Because open files are shared, a type *r* command may be used as an input filter and a type *w* as an output filter.

**SEE ALSO**

*pipe(2)*, *wait(2)*, *fclose(3S)*, *fopen(3S)*, *system(3S)*.

**DIAGNOSTICS**

*Popen* returns a NULL pointer if files or processes cannot be created or if the shell cannot be accessed.

*Pclose* returns *-1* if *stream* is not associated with a command opened by *popen*.

**BUGS**

If the original processes and processes opened by *popen* concurrently read or write a common file, neither should use buffered I/O, because the buffering gets all mixed up. Problems with an output filter may be forestalled by careful buffer flushing, e.g., by using *fflush*; see *fclose(3S)*.

If an illegal type is passed, *popen* will fork and exec the command line passed to it before it discovers that the type was illegal. This will result in a NULL pointer being returned and a broken pipe (with the command executing in the background).

## NAME

printf, fprintf, sprintf — print formatted output

## SYNOPSIS

```
#include <stdio.h>

int printf (format [ , arg ] ... )
char *format;

int fprintf (stream, format [ , arg ] ... )
FILE *stream;
char *format;

int sprintf (s, format [ , arg ] ... )
char *s, format;
```

## DESCRIPTION

*Printf* places output on the standard output stream *stdout*. *Fprintf* places output on the named output *stream*. *Sprintf* places “output,” followed by the null character (\0) in consecutive bytes starting at \*s; it is the user’s responsibility to ensure that enough storage is available. Each function returns the number of characters transmitted (not including the \0 in the case of *sprintf*), or a negative value if an output error was encountered.

Each of these functions converts, formats, and prints its *args* under control of the *format*. The *format* is a character string that contains two types of objects: plain characters, which are simply copied to the output stream, and conversion specifications, each of which results in fetching zero or more *args*. The results are undefined if there are insufficient *args* for the format. If the format is exhausted while *args* remain, the excess *args* are simply ignored.

Each conversion specification is introduced by the character %. After the %, the following appear in sequence:

Zero or more *flags*, which modify the meaning of the conversion specification.

An optional decimal digit string specifying a minimum *field width*. If the converted value has fewer characters than the field width, it will be padded to the field width on the left (default) or right (if the left-adjustment flag ‘-’ has been given); see below for flag specification. If the field width for an *s* conversion is preceded by a 0, the string is right adjusted with zero padding on the left.

A *precision* that gives the minimum number of digits to appear for the *d*, *o*, *u*, *x*, or *X* conversions, the number of digits to appear after the decimal point for the *e* and *f* conversions, the maximum number of significant digits for the *g* conversion, or the maximum number of characters to be printed from a string in *s* conversion. The format of the precision is a period (.) followed by a decimal digit string; a null digit string is treated as zero.

An optional *l* (*ell*) specifying that a following *d*, *o*, *u*, *x*, or *X* conversion character applies to a long integer *arg*. An *l* before any other conversion character is ignored.

A character that indicates the type of conversion to be applied.

A field width or precision may be indicated by an asterisk (\*) instead of a digit string. In this case, an integer *arg* supplies the field width or



precision. The *arg* that is actually converted is not fetched until the conversion letter is seen; therefore, the *args* specifying field width or precision must appear *before* the *arg* (if any) to be converted.

The flag characters and their meanings are:

- The result of the conversion will be left-justified within the field.
- + The result of a signed conversion will always begin with a sign (+ or -).
- blank If the first character of a signed conversion is not a sign, a blank will be prefixed to the result. This implies that if the blank and + flags both appear, the blank flag will be ignored.
- # This flag specifies that the value is to be converted to an "alternate form." For c, d, s, and u conversions, the flag has no effect. For o conversion, it increases the precision to force the first digit of the result to be a zero. For x (X) conversion, a non-zero result will have 0x (0X) prefixed to it. For e, E, f, g, and G conversions, the result will always contain a decimal point, even if no digits follow the point (normally, a decimal point appears in the result of these conversions only if a digit follows it). For g and G conversions, trailing zeroes will *not* be removed from the result (which they normally are).

The conversion characters and their meanings are:

- d,o,u,x,X The integer *arg* is converted to signed decimal, unsigned octal, decimal, or hexadecimal notation (x and X), respectively; the letters abcdef are used for x conversion and the letters ABCDEF for X conversion. The precision specifies the minimum number of digits to appear; if the value being converted can be represented in fewer digits, it will be expanded with leading zeroes. (For compatibility with older versions, padding with leading zeroes may alternatively be specified by prepending a zero to the field width.) This does not imply an octal value for the field width. The default precision is 1. The result of converting a zero value with a precision of zero is a null string.
- f The float or double *arg* is converted to decimal notation in the style "[-]ddd.ddd", where the number of digits after the decimal point is equal to the precision specification. If the precision is missing, 6 digits are output; if the precision is explicitly 0, no decimal point appears.
- e,E The float or double *arg* is converted in the style "[-]d.ddde±dd", where there is one digit before the decimal point and the number of digits after it is equal to the precision; when the precision is missing, 6 digits are produced; if the precision is zero, no decimal point appears. The E format code produces a number with E instead of e introducing the exponent. The exponent always contains at least two digits.

- g,G** The float or double *arg* is printed in style *f* or *e* (or in style *E* in the case of a *G* format code), with the precision specifying the number of significant digits. The style used depends on the value converted: style *e* is used only if the exponent resulting from the conversion is less than  $-4$  or greater than the precision. Trailing zeroes are removed from the result; a decimal point appears only if it is followed by a digit.
- c** The character *arg* is printed.
- s** The *arg* is taken to be a string (character pointer) and characters from the string are printed until a null character ( $\backslash 0$ ) is encountered or the number of characters indicated by the precision specification is reached. If the precision is missing, it is taken to be infinite, so all characters up to the first null character are printed. A NULL value for *arg* yields undefined results.
- %** Print a %; no argument is converted.

In no case does a non-existent or small field width cause truncation of a field; if the result of a conversion is wider than the field width, the field is simply expanded to contain the conversion result. Characters generated by *printf* and *fprintf* are printed as if *putc*(3S) had been called.

#### EXAMPLES

To print a date and time in the form "Sunday, July 3, 10:02", where *weekday* and *month* are pointers to null-terminated strings:

```
printf("%s, %s %d, %.2d:%.2d", weekday, month, day, hour, min);
```

To print *pi* to 5 decimal places:

```
printf("pi = %.5f", 4*atan(1.0));
```

#### SEE ALSO

*ecvt*(3C), *putc*(3S), *scanf*(3S), *stdio*(3S).

**NAME**

*putc*, *putchar*, *fputc*, *putw* - put character or word on a stream

**SYNOPSIS**

```
#include <stdio.h>

int putc (c, stream)
int c;
FILE *stream;

int putchar (c)
int c;

int fputc (c, stream)
int c;
FILE *stream;

int putw (w, stream)
int w;
FILE *stream;
```

**DESCRIPTION**

*Putc* writes the character *c* onto the output *stream* at the position where the file pointer, if defined, is pointing. *Putchar(c)* is defined as *putc(c, stdout)*. *Putc* and *putchar* are macros.

*Fputc* behaves like *putc*, but is a function rather than a macro. *Fputc* runs more slowly than *putc*, but it takes less space per invocation and its name can be passed as an argument to a function.

*Putw* writes the word (32-bit integer on the 68000) *w* to the output *stream* at the position at which the file pointer, if defined, is pointing. *Putw* neither assumes nor causes special alignment in the file.

Output streams, with the exception of the standard error stream *stderr*, are by default buffered if the output refers to a file and line-buffered if the output refers to a terminal. The standard error output stream *stderr* is by default unbuffered, but use of *freopen*(see *fopen*(3S)) causes it to become buffered or line-buffered. When an output stream is unbuffered information, it is queued for writing on the destination file or terminal as soon as written; when it is buffered, many characters are saved up and written as a block; when it is line-buffered, each line of output is queued for writing on the destination terminal as soon as the line is completed (i.e., as soon as a new-line character is written or terminal input is requested). *Setbuf*(3S) may be used to change the stream's buffering strategy.

**SEE ALSO**

*fclose*(3S), *ferror*(3S), *fopen*(3S), *fread*(3S), *printf*(3S), *puts*(3S), *setbuf*(3S).

**DIAGNOSTICS**

On success, these functions each return the value they have written. On failure, they return the constant EOF. This occurs if the file *stream* is not open for writing or if the output file cannot be grown. Because EOF is a valid integer, *ferror*(3S) should be used to detect *putw* errors.

**BUGS**

Because it is implemented as a macro, *putc* treats incorrectly a *stream* argument with side effects. In particular, *putc(c, \*f++)*; doesn't work sensibly. *Fputc* should be used instead.

Because of possible differences in word length and byte ordering, files written using *putw* are machine-dependent and may not be read using *getw* on a different processor.

**NAME**

putenv — change or add value to environment

**SYNOPSIS**

```
int putenv (string)
char *string;
```

**DESCRIPTION**

*String* points to a string of the form "*name=value*." *Putenv* makes the value of the environment variable *name* equal to *value* by altering an existing variable or creating a new one. In either case, the string pointed to by *string* becomes part of the environment, so altering the string will change the environment. The space used by *string* is no longer used once a new string-defining *name* is passed to *putenv*.

**DIAGNOSTICS**

*Putenv* returns non-zero if it was unable to obtain enough space via *malloc* for an expanded environment, otherwise zero.

**SEE ALSO**

exec(2), getenv(3C), malloc(3C), environ(5).

**WARNINGS**

*Putenv* manipulates the environment pointed to by *environ*, and can be used in conjunction with *getenv*. However, *envp* (the third argument to *main*) is not changed.

This routine uses *malloc(3C)* to enlarge the environment.

After *putenv* is called, environmental variables are not in alphabetical order. A potential error is to call *putenv* with an automatic variable as the argument, then exit the calling function while *string* is still part of the environment.

**NAME**

putpwent — write password file entry

**SYNOPSIS**

```
#include <pwd.h>
int putpwent (p, f)
struct passwd *p;
FILE *f;
```

**DESCRIPTION**

*Putpwent* is the inverse of *getpwent(3C)*. Given a pointer to a *passwd* structure created by *getpwent* (or *getpwuid* or *getpwnam*), *putpwuid* writes a line on the stream *f* which matches the format of */etc/passwd*.

The *<pwd.h>* header file is described in *getpwent(3C)*.

**SEE ALSO**

*getpwent(3C)*.

**DIAGNOSTICS**

*Putpwent* returns non-zero if an error was detected during its operation; otherwise it returns zero.

**SEE ALSO**

*getpwent(3C)*.

**WARNING**

The above routine uses *<stdio.h>*. Therefore, the size of programs not otherwise using standard I/O is increased more than might be expected.

**NAME**

puts, fputs — put a string on a stream

**SYNOPSIS**

```
#include <stdio.h>
```

```
int puts (s)
```

```
char *s;
```

```
int fputs (s, stream)
```

```
char *s;
```

```
FILE *stream;
```

**DESCRIPTION**

*Puts* writes the null-terminated string pointed to by *s*, followed by a new-line character, to the standard output stream *stdout*.

*Fputs* writes the null-terminated string pointed to by *s* to the named output *stream*.

Neither function writes the terminating null character.

**SEE ALSO**

ferror(3S), fopen(3S), fread(3S), printf(3S), putc(3S).

**DIAGNOSTICS**

Both routines return EOF on error. This occurs if the routines try to write on a file that has not been opened for writing.

**NOTES**

*Puts* appends a new-line character while *fputs* does not.

**NAME**

qsort — quicker sort

**SYNOPSIS**

```
void qsort ((char *) base, nel, width, compar)
unsigned nel, width;
int (*compar)();
```

**DESCRIPTION**

*Qsort* is an implementation of the quicker-sort algorithm. It sorts a table of data in place.

*Base* points to the element at the base of the table. *Nel* is the number of elements in the table. *Width* is the width of an element in bytes; *sizeof (base)* should be used. *Compar* is the name of the comparison function, which is called with two arguments that point to the elements being compared. The function must return an integer less than, equal to, or greater than zero according as the first argument is to be considered less than, equal to, or greater than the second.

**NOTES**

The pointer to the base of the table should be of type pointer-to-element, and cast to type pointer-to-character.

The comparison function need not compare every byte, so arbitrary data may be contained in the elements in addition to the values being compared. The order in the output of the two items which compare as equal is unpredictable.

**EXAMPLE**

```
struct entry {
    char *name;
    int flags;
};

main()
{
    struct entry hp[100];
    int entcmp();
    int i, count;
    for (i = 0; i < (count = 100); i++) {
        /* fill the structure with the name and flags */
        :
        :
    }
    qsort( (char *) hp, count, sizeof (hp[0]), entcmp);
}

entcmp(ep,ep2)
struct entry *ep, *ep2;
{
    return (strcmp(ep->name, ep2->name));
}
```

will sort a set of names with associated flags in ASCII order.

**SEE ALSO**

sort(1), bsearch(3C), lsearch(3C), string(3C).



**NAME**

*rand*, *srand* — simple random-number generator

**SYNOPSIS**

```
int rand ( )  
void srand (seed)  
unsigned seed;
```

**DESCRIPTION**

*Rand* uses a multiplicative congruential random-number generator with period  $2^{32}$  that returns successive pseudo-random numbers in the range from 0 to  $2^{15}-1$ .

*Srand* can be called at any time to reset the random-number generator to a random starting point. The generator is initially seeded with a value of 1.

**NOTE**

The spectral properties of *rand* leave a great deal to be desired. *Drand48(3C)* provides a much better, though more elaborate, random-number generator.

**SEE ALSO**

*drand48(3C)*.

**NAME**

irand, srand, rand — Fortran uniform random-number generator

**SYNOPSIS**

call srand(iseed)

i = irand()

x = rand()

**DESCRIPTION**

*Irand* generates successive pseudo-random numbers in the range from 0 to  $2^{**}15-1$ . *Rand* generates pseudo-random numbers distributed in (0, 1.0). *Srand* uses its integer argument to re-initialize the seed for successive invocations of *irand* and *rand*.

**SEE ALSO**

rand(3C).

## NAME

*rcmd*, *rresvport*, *ruserok* — routines for returning a stream to a remote command

## SYNOPSIS

```
rem = rcmd(ahost, inport, locuser, remuser, cmd, fd2p);
char **ahost;
u_short inport;
char *locuser, *remuser, *cmd;
int *fd2p;

s = rresvport(port);
int *port;

ruserok(rhost, superuser, ruser, luser);
char *rhost;
int superuser;
char *ruser, *luser;

cc ... -lnet
```

## DESCRIPTION

*Rcmd* is a routine used by the super-user to execute a command on a remote machine using an authentication scheme based on reserved port numbers. *Rresvport* is a routine which returns a descriptor to a socket with an address in the privileged port space. *Ruserok* is a routine used by servers to authenticate clients requesting service with *rcmd*. All three functions are present in the same file and are used by the *remshd*(8N) server (among others).

*Rcmd* looks up the host *\*ahost* using *gethostent*(3N), returning -1 if the host does not exist. Otherwise *\*ahost* is set to the standard name of the host and a connection is established to a server residing at the well-known Internet port *inport*.

If the call succeeds, a socket of type `SOCK_STREAM` is returned to the caller, and given to the remote command as `stdin` and `stdout`. If *fd2p* is non-zero, then an auxiliary channel to a control process will be set up, and a descriptor for it will be placed in *\*fd2p*. The control process will return diagnostic output from the command (unit 2) on this channel, and will also accept bytes on this channel as being UNIX signal numbers, to be forwarded to the process group of the command. If *fd2p* is 0, then the `stderr` (unit 2 of the remote command) will be made the same as the `stdout` and no provision is made for sending arbitrary signals to the remote process, although you may be able to get its attention by using out-of-band data.

The protocol is described in detail in *remshd*(8N).

The *rresvport* routine is used to obtain a socket with a privileged address bound to it. This socket is suitable for use by *rcmd* and several other routines. Privileged addresses consist of a port in the range 0 to 1023. Only the super-user is allowed to bind an address of this sort to a socket.

*Ruserok* takes a remote host's name, as returned by a *gethostent*(3N) routine, two user names and a flag indicating if the local user's name is the super-user. It then checks the files */etc/hosts.equiv* and, possibly, *.rhosts* in the current working directory (normally the local user's home directory) to

see if the request for service is allowed. A 1 is returned if the machine name is listed in the "hosts.equiv" file, or the host and remote user name are found in the ".rhosts" file; otherwise *ruserok* returns 0. If the *superuser* flag is 1, the checking of the "host.equiv" file is bypassed.

**LINKING**

This library is accessed by specifying `-lnet` as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**SEE ALSO**

`rlogin(1N)`, `remsh(1N)`, `rexec(3N)`, `rexecd(8N)`, `rlogind(8N)`, `remshd(8N)`

**BUGS**

There is no way to specify options to the *socket* call which *rcmd* makes.

**NAME**

`readv` — read from file

**SYNOPSIS**

```
#include <sys/types.h>
#include <sys/uio.h>
```

```
cc = readv(d,iov,iovcnt)
int cc, d;
struct iovec *iov;
int iovcnt;
```

**DESCRIPTION**

*fdes* is a file descriptor obtained from a *creat*, *open*, *dup*, *fcntl*, *pipe*, or *socket* system call.

*Readv* attempts to read *nbyte* bytes from the file associated with *fdes* and scatters the input data into the *iovcnt* buffers specified by the members of the *iovec* array: *iov*[0], *iov*[1], ..., *iov*[*iovcnt* - 1].

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 devices capable of seeking, the *readv* starts at a position in the file given by the file pointer associated with *fdes*. Upon return from *readv*, the file pointer is incremented by the number of bytes actually read.

Devices that are incapable of seeking always read from the current position. The value of a file pointer associated with such a file is undefined.

Upon successful completion, *readv* returns the number of bytes actually read and placed in the buffer; this number may be less than *nbyte* if the file is associated with a communication line (see *ioctl*(2), *socket*(2N), and *termio*(7)), or if the number of bytes left in the file is less than *nbyte* bytes. A value of 0 is returned when an end-of-file has been reached.

When attempting to read from an empty pipe (or FIFO):

If *O\_NDELAY* is set, the read will return a 0.

If *O\_NDELAY* is clear, the read will block until data is written to the file or the file is no longer open for writing.

When attempting to read a file associated with a *tty* that has no data currently available:

If *O\_NDELAY* is set, the read will return a 0.

If *O\_NDELAY* is clear, the read will block until data becomes available.

*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 signal was caught during the *read* system call.

In addition, *readv* may return one of the following errors:

[EINVAL] *iovent* 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.

#### RETURN VALUE

Upon successful completion a non-negative integer is returned indicating the number of bytes actually read. Otherwise, a  $-1$  is returned and *errno* is set to indicate the error.

#### SEE ALSO

*creat*(2), *fcntl*(2), *ioctl*(2), *open*(2), *pipe*(2), *socket*(2N), *termio*(7) in the *Administrator Reference Manual*.

## NAME

regcmp, regex — compile and execute a regular expression

## SYNOPSIS

```
char *regcmp(string1 |, string2, ...), (char *)0)
char *string1, *string2, ...;
char *regex(re, subject1, ret0, ...)
char *re, *subject, *ret0, ...;
extern char *loc1;
```

## DESCRIPTION

*Regcmp* compiles a regular expression and returns a pointer to the compiled form. *Malloc*(3C) is used to create space for the vector. It is the user's responsibility to free unneeded space that has been allocated by *malloc*. A NULL return from *regcmp* indicates an incorrect argument. *Regcmp*(1) has been written to generally preclude the need for this routine at execution time.

*Regex* executes a compiled pattern against the subject string. Additional arguments are passed to receive values back. *Regex* returns NULL on failure or a pointer to the next unmatched character on success. A global character pointer *loc1* points to where the match began. *Regcmp* and *regex* were mostly borrowed from the editor, *ed*(1); however, the syntax and semantics have been changed slightly. The following are the valid symbols and their associated meanings.

- | | \* . These symbols retain their current meaning.
- \$ This symbol matches the end of the string; \n matches the new-line.
- Within brackets the minus means "through". For example, [a-z] is equivalent to [abcd...xyz]. The - can appear as itself only if used as the last or first character. For example, the character class expression [| - | matches the characters | and -.
- + A regular expression followed by + means "one or more times". For example, [0-9]+ is equivalent to [0-9][0-9]\*.
- {m} {m,u} Integer values enclosed in {} indicate the number of times the preceding regular expression is to be applied. The minimum number is *m* and the maximum number is *u*, which must be less than 256. If only *m* is present (e.g., {m}), it indicates the exact number of times the regular expression is to be applied. {m} is analogous to {m,infinity}. The plus (+) and star (\*) operations are equivalent to {1,} and {0,}, respectively.
- (...)\$n The value of the enclosed regular expression is to be returned. The value will be stored in the (n+1)th argument following the subject argument. At present, at most 10 enclosed regular expressions are allowed. *Regex* makes its assignments unconditionally.
- (...) Parentheses are used for grouping. An operator (e.g., \*, +, {}) can work on a single character or a regular expression enclosed in parentheses. For example, (a\*(cb+)\*)\$0.

By necessity, all the above defined symbols are special. They must, therefore, be escaped to be used as themselves.

#### EXAMPLES

Example 1:

```
char *cursor, *newcursor, *ptr;
...
newcursor = regex((ptr = regcmp("\n", 0)), cursor);
free(ptr);
```

This example will match a leading new-line in the subject string pointed at by cursor.

Example 2:

```
char ret0[9];
char *newcursor, *name;
...
name = regcmp("[A-Za-z][A-Za-z0-9_]{0,7}$0", 0);
newcursor = regex(name, "123Testing321", ret0);
```

This example will match through the string "Testing3" and will return the address of the character after the last matched character (cursor+11). The string "Testing3" will be copied to the character array *ret0*.

Example 3:

```
#include "file.i"
char *string, *newcursor;
...
newcursor = regex(name, string);
```

This example applies a precompiled regular expression in *file.i* (see *regcmp(1)*) against *string*.

This routine is kept in */lib/libPW.a*.

#### SEE ALSO

*ed(1)*, *regcmp(1)*, *malloc(3C)*.

#### BUGS

The user program may run out of memory if *regcmp* is called iteratively without freeing the vectors no longer required. The following user-supplied replacement for *malloc(3C)* reuses the same vector, saving time and space:

```
/* user's program */
...
char *
malloc(n)
unsigned n;
{
    static char rebuf[512];
    return (n <= sizeof rebuf) ? rebuf : NULL;
}
```



**NAME**

`rexec` — return stream to a remote command

**SYNOPSIS**

```
rem = rexec(ahost, inport, user, passwd, cmd, fd2p);
char **ahost;
u_short inport;
char *user, *passwd, *cmd;
int *fd2p;
cc ... -lnet
```

**DESCRIPTION**

*Rexec* looks up the host *\*ahost* using *gethostent(3N)*, returning `-1` if the host does not exist. Otherwise *\*ahost* is set to the standard name of the host. If a username and password are both specified, then these are used to authenticate to the foreign host; otherwise the environment and then the user's *.netrc* file in his home directory are searched for appropriate information. If all this fails, the user is prompted for the information.

The port *inport* specifies which well-known DARPA Internet port to use for the connection; it will normally be the value returned from the call `"getservbyname("exec", "tcp")"` (see *getservent(3N)*). The protocol for connection is described in detail in *rexecd(8N)*.

If the call succeeds, a socket of type `SOCK_STREAM` is returned to the caller, and given to the remote command as `stdin` and `stdout`. If *fd2p* is non-zero, then a auxiliary channel to a control process will be setup, and a descriptor for it will be placed in *\*fd2p*. The control process will return diagnostic output from the command (unit 2) on this channel, and will also accept bytes on this channel as being UNIX signal numbers, to be forwarded to the process group of the command. If *fd2p* is 0, then the `stderr` (unit 2 of the remote command) will be made the same as the `stdout` and no provision is made for sending arbitrary signals to the remote process, although you may be able to get its attention by using out-of-band data.

**LINKING**

This library is accessed by specifying `-lnet` as the last argument to the compile line, e.g.:

```
cc -o prog prog.c -lnet
```

**SEE ALSO**

`rcmd(3N)`, `rexecd(8N)`

**BUGS**

There is no way to specify options to the *socket* call which *rexec* makes.

**NAME**

**anint, dnint, nint, idnint** — Fortran nearest integer functions

**SYNOPSIS**

```
integer i
real r1, r2
double precision dp1, dp2

r2 = anint(r1)
i = nint(r1)

dp2 = anint(dp1)
dp2 = dnint(dp1)

i = nint(dp1)
i = idnint(dp1)
```

**DESCRIPTION**

*Anint* returns the nearest whole real number to its real argument (i.e.,  $\text{int}(a+0.5)$  if  $a \geq 0$ ,  $\text{int}(a-0.5)$  otherwise). *Dnint* does the same for its double-precision argument. *Nint* returns the nearest integer to its real argument. *Idnint* is the double-precision version. *Anint* is the generic form of *anint* and *dnint*, performing the same operation and returning the data type of its argument. *Nint* is also the generic form of *idnint*.

**NAME**

`scanf`, `fscanf`, `sscanf` — convert formatted input

**SYNOPSIS**

```
#include <stdio.h>

int scanf (format [ , pointer [ ... ] )
char *format;

int fscanf (stream, format [ , pointer [ ... ] )
FILE *stream;
char *format;

int sscanf (s, format [ , pointer [ ... ] )
char *s, *format;
```

**DESCRIPTION**

*Scanf* reads from the standard input stream *stdin*. *Fscanf* reads from the named input *stream*. *Sscanf* reads from the character string *s*. Each function reads characters, interprets them according to *format*, and stores the results in its arguments. Each function expects two arguments: a control string *format* (described below) and a set of *pointer* arguments indicating where the converted input should be stored.

The control string usually contains conversion specifications, which are used to direct interpretation of input sequences. The control string may contain:

1. White-space characters (blanks and tabs) which, except in two cases described below, cause input to be read up to the next non-white-space character.
2. An ordinary character (not %), which must match the next character of the input stream.
3. Conversion specifications, consisting of the character %, an optional assignment suppression character \*, an optional numerical maximum field width, an optional l (ell) or h indicating the size of the receiving variable, and a conversion code.

A conversion specification directs the conversion of the next input field; the result is placed in the variable pointed to by the corresponding argument, unless assignment suppression has been indicated by \*. The suppression of assignment provides a way of describing an input field which is to be skipped. An input field is defined as a string of non-white-space characters; it extends to the next inappropriate character or until the field width, if specified, is exhausted. For all descriptors except “[” and “c”, white space leading an input field is ignored.

The conversion code indicates the interpretation of the input field; the corresponding pointer argument must usually be of a restricted type. For a suppressed field, no pointer argument should be given. The following conversion codes are legal:

- % A single % is expected in the input at this point; no assignment is done.
- d A decimal integer is expected; the corresponding argument should be an integer pointer.

- u** An unsigned decimal integer is expected; the corresponding argument should be an unsigned integer pointer.
- o** An octal integer is expected; the corresponding argument should be an integer pointer.
- x** A hexadecimal integer is expected; the corresponding argument should be an integer pointer.
- e,f,g** A floating point number is expected; the next field is converted accordingly and stored through the corresponding argument, which should be a pointer to a *float*. The input format for floating point numbers is an optionally signed string of digits, possibly containing a decimal point, followed by an optional exponent field consisting of an E or an e, followed by an optional +, -, or space followed by an integer.
- s** A character string is expected; the corresponding argument should be a character pointer to an array of characters large enough to accept the string and a terminating `\0`, which will be added automatically. The input field is terminated by a white-space character.
- c** A character is expected; the corresponding argument should be a character pointer. The normal skip over white space is suppressed in this case; to read the next non-space character, use `%1s`. If a field width is given, the corresponding argument should refer to a character array; the indicated number of characters is read.
- l** String data and the normal skip over leading white space is suppressed. The left bracket is followed by a set of characters (the *scanset*) and a right bracket; the input field is the maximal sequence of input characters consisting entirely of characters in the *scanset*. The circumflex, (`^`), when it appears as the first character in the *scanset*, serves as a complement operator and redefines the *scanset* as the set of all characters *not* contained in the remainder of the *scanset* string. There are some conventions used in the construction of the *scanset*. A range of characters may be represented by the construct *first-last*; thus, `[0123456789]` may be expressed `[0-9]`. Using this convention, *first* must be lexically less than or equal to *last*, or else the dash will stand for itself. The dash will also stand for itself whenever it is the first or the last character in the *scanset*. To include the right square bracket as an element of the *scanset*, it must appear as the first character (possibly preceded by a circumflex) of the *scanset*; otherwise it will be interpreted syntactically as the closing bracket. The corresponding argument must point to a character array large enough to hold the data field and the terminating `\0`, which will be added automatically. At least one character must match for this conversion to be considered successful.

The conversion characters **d**, **u**, **o**, and **x** may be preceded by **l** or **h** to indicate that a pointer to **long** or **short**, rather than **int**, is in the argument list. Similarly, the conversion characters **e**, **f**, and **g** may be preceded by **l** to indicate that a pointer to **double**, rather than **float**, is in the argument list.

The **l** or **h** modifier is ignored for other conversion characters. *Scanf* conversion terminates at EOF, at the end of the control string, or when an

input character conflicts with the control string. In the latter case, the offending character is left unread in the input stream.

*Scanf* returns the number of successfully matched and assigned input items; this number can be zero when an early conflict between an input character and the control string occurs. If the input ends before the first conflict or conversion, EOF is returned.

#### EXAMPLES

The call

```
int i; float x; char name[50];
n = scanf ("%d%f%s", &i, &x, name);
```

with the input line

```
25 54.32E-1 thompson
```

will assign the value 3 to *n*, the value 25 to *i*, and the value 5.432 to *x*; *name* will contain *thompson\0*.

The call

```
int i; float x; char name[50];
(void) scanf ("%2d%f%d %|0-9|", &i, &x, name);
```

with input

```
56789 0123 56a72
```

will assign 56 to *i*, 789.0 to *x*, skip 0123, and place the string 56\0 in *name*. The next call to *getchar* (see *getc*(3S)) will return a.

#### SEE ALSO

*getc*(3S), *printf*(3S), *strtod*(3C), *strtol*(3C).

#### NOTE

Trailing white space is left unread unless matched in the control string.

#### DIAGNOSTICS

These functions return EOF on end of input and a short count for missing or illegal data items.

#### BUGS

The success of literal matches and suppressed assignments is not directly determinable.

## NAME

setbuf, setvbuf – assign buffering to a stream

## SYNOPSIS

```
#include <stdio.h>

void setbuf (stream, buf)
FILE *stream;
char *buf;

int setvbuf (stream, buf, type, size)
FILE *stream;
char *buf;
int type, size;
```

## DESCRIPTION

*Setbuf* may be used after a stream has been opened but before it is read or written. It causes the array pointed to by *buf* to be used instead of an automatically allocated buffer. If *buf* is the NULL pointer input/output will be completely unbuffered.

A constant BUFSIZ, defined in the <stdio.h> header file, tells how big an array is needed:

```
char buf[BUFSIZ];
```

*Setvbuf* may be used after a stream has been opened but before it is read or written. *Type* determines how *stream* will be buffered. Legal values for *type* (defined in *stdio.h*) are:

**\_IOFBF** causes input/output to be fully buffered.  
**\_IOLBF** causes output to be line buffered; the buffer will be flushed when a newline is written, the buffer is full, or input is requested.  
**\_IONBF** causes input/output to be completely unbuffered.

If *buf* is not the NULL pointer, the array it points to will be used for buffering, instead of an automatically allocated buffer. *Size* specifies the size of the buffer to be used. The constant BUFSIZ in <stdio.h> is suggested as a good buffer size. If input/output is unbuffered, *buf* and *size* are ignored.

By default, output to a terminal is line buffered and all other input/output is fully buffered.

## SEE ALSO

fopen(3S), getc(3S), malloc(3C), putc(3S), stdio(3S).

**DIAGNOSTICS**

If an illegal value for *type* or *size* is provided, *setvbuf* returns a non-zero value. Otherwise, the value returned will be zero.

**NOTE**

A common source of error is allocating buffer space as an "automatic" variable in a code block, and then failing to close the stream in the same block.

*Setbuf* allows assignment of a new I/O buffer after the stream has been read (written), and if unflushed data remains in the original buffer. This could lead to a loss of data error.

**NAME**

setjmp, longjmp — non-local goto

**SYNOPSIS**

```
#include <setjmp.h>
int setjmp (env)
jmp_buf env;
void longjmp (env, val)
jmp_buf env;
int val;
```

**DESCRIPTION**

These functions are useful for dealing with errors and interrupts encountered in a low-level subroutine of a program.

*Setjmp* saves its stack environment in *env* for later use by *longjmp*. The environment type *jmp\_buf* is defined in the `<setjmp.h>` header file. *Setjmp* returns the value 0.

*Longjmp* restores the environment saved by the last call of *setjmp* with the corresponding *env* argument. After *longjmp* is completed, program execution continues as if the corresponding call of *setjmp* (which must not itself have returned in the interim) had just returned the value *val*. *Longjmp* cannot cause *setjmp* to return the value 0. If *longjmp* is invoked with a second argument of 0, *setjmp* will return 1. All accessible data have values as of the time *longjmp* was called.

**SEE ALSO**

signal(2).

**WARNING**

*Longjmp* fails if it is called when *env* was never primed by a call to *setjmp* or when the last such call is in a function which has since returned.



**NAME**

sign, isign, dsign — Fortran transfer-of-sign intrinsic function

**SYNOPSIS**

integer i, j, k  
real r1, r2, r3  
double precision dp1, dp2, dp3  
k = isign(i, j)  
k = sign(i, j)  
r3 = sign(r1, r2)  
dp3 = dsign(dp1, dp2)  
dp3 = sign(dp1, dp2)

**DESCRIPTION**

*Isign* returns the magnitude of its first argument with the sign of its second argument. *Sign* and *dsign* are its real and double-precision counterparts, respectively. The generic version is *sign*, which devolves to the appropriate type depending on its arguments.

**SIGNAL (3F)****SIGNAL (3F)****NAME**

signal — specify Fortran action on receipt of a system signal

**SYNOPSIS**

integer i  
external integer intfnc  
call signal(i, intfnc)

**DESCRIPTION**

*Signal* allows a process to specify a function to be invoked upon receipt of a specific signal. The first argument specifies a fault or exception; the second argument specifies the function to be invoked.

**SEE ALSO**

kill(2), signal(2).

**NAME**

*sin*, *dsin*, *csin* — Fortran sine intrinsic function

**SYNOPSIS**

**real** r1, r2  
**double precision** dp1, dp2  
**complex** cx1, cx2  
r2 = *sin*(r1)  
dp2 = *dsin*(dp1)  
dp2 = *sin*(dp1)  
cx2 = *csin*(cx1)  
cx2 = *sin*(cx1)

**DESCRIPTION**

*Sin* returns the real sine of its real argument. *Dsin* returns the double-precision sine of its double-precision argument. *Csin* returns the complex sine of its complex argument. The generic *sin* function becomes *dsin* or *csin* as required by argument type.

**SEE ALSO**

trig(3M).

**NAME**

sinh, dsinh — Fortran hyperbolic sine intrinsic function

**SYNOPSIS**

```
real r1, r2
double precision dp1, dp2
r2 = sinh(r1)
dp2 = dsinh(dp1)
dp2 = sinh(dp1)
```

**DESCRIPTION**

*Sinh* returns the real hyperbolic sine of its real argument. *Dsinh* returns the double-precision hyperbolic sine of its double-precision argument. The generic form *sinh* may be used to return a double-precision value given a double-precision argument.

**SEE ALSO**

sinh(3M).

**NAME**

*sinh*, *cosh*, *tanh* — hyperbolic functions

**SYNOPSIS**

```
#include <math.h>
double sinh (x)
double x;
double cosh (x)
double x;
double tanh (x)
double x;
```

**DESCRIPTION**

*Sinh*, *cosh*, and *tanh* return, respectively, the hyperbolic sine, cosine, and tangent of their argument.

**DIAGNOSTICS**

*Sinh* and *cosh* return HUGE (and *sinh* may return -HUGE for negative *x*) when the correct value would overflow and set *errno* to ERANGE.

These error-handling procedures may be changed with the function *matherr*(3M).

**SEE ALSO**

*matherr*(3M).

**NAME**

*sleep* — suspend execution for interval

**SYNOPSIS**

**unsigned sleep (seconds)**  
**unsigned seconds;**

**DESCRIPTION**

*Sleep* suspends the current process from execution for the number of *seconds* specified by the argument. The actual suspension time may be less than that requested for two reasons: (1) scheduled wakeups occur at fixed 1-second intervals, (on the second, according to an internal clock) and (2) any caught signal will terminate *sleep* following execution of the signal catching routine. The suspension time may be longer than requested by an arbitrary amount, due to the scheduling of other activity in the system. The value returned by *sleep* is the "unslept" amount (the requested time minus the time actually slept) in case the caller had an alarm set to go off earlier than the end of the requested *sleep* time or in case there is premature arousal due to another caught signal.

The routine is implemented by setting an alarm signal and pausing until it (or some other signal) occurs. The previous state of the alarm signal is saved and restored. The calling program may have set up an alarm signal before calling *sleep*. If the *sleep* time exceeds the time before the alarm signal, the process sleeps only until the alarm signal would have occurred and the caller's alarm catch routine is executed just before the *sleep* routine returns. If the *sleep* time is less than the time before the calling program's alarm, the prior alarm time is reset to go off at the same time it would have without the intervening *sleep*.

**SEE ALSO**

*alarm(2)*, *pause(2)*, *signal(2)*.

**NAME**

*sputl*, *sgetl* — access long integer data in a machine independent fashion.

**SYNOPSIS**

*sputl* (*value*, *buffer*)

long *value*;

char \**buffer*;

long *sgetl* (*buffer*)

char \**buffer*;

**DESCRIPTION**

*Sputl* takes the 4 bytes of the long integer *value* and places them in memory, starting at the address pointed to by *buffer*. The ordering of the bytes is the same across all machines.

*Sgetl* retrieves the 4 bytes in memory, starting at the address pointed to by *buffer*, and returns the long integer value in the byte ordering of the host machine.

Use of *sputl* and *sgetl* provide a machine independent way of storing long numeric data in a file in binary form without conversion to characters.

A program that uses these functions must be loaded with the object file access routine library *libld.a*.

**SEE ALSO**

*ar*(4).

**NAME**

*sqrt*, *dsqrt*, *csqrt* — Fortran square root intrinsic function

**SYNOPSIS**

**real** *r1*, *r2*  
**double precision** *dp1*, *dp2*  
**complex** *cx1*, *cx2*  
*r2* = *sqrt*(*r1*)  
*dp2* = *dsqrt*(*dp1*)  
*dp2* = *sqrt*(*dp1*)  
*cx2* = *csqrt*(*cx1*)  
*cx2* = *sqrt*(*cx1*)

**DESCRIPTION**

*Sqrt* returns the real square root of its real argument. *Dsqrt* returns the double-precision square root of its double-precision argument. *Csqrt* returns the complex square root of its complex argument. *Sqrt*, the generic form, will become *dsqrt* or *csqrt* as required by its argument type.

**SEE ALSO**

*exp*(3M).



**NAME**

ssignal, gsignal – software signals

**SYNOPSIS**

```
#include <signal.h>

int (*ssignal (sig, action))( )
int sig, (*action)( );

int gsignal (sig)
int sig;
```

**DESCRIPTION**

*Ssignal* and *gsignal* implement a software facility similar to *signal(2)*. This facility is used by the Standard C Library to enable users to indicate the disposition of error conditions; it is also made available to users for their own purposes.

Software signals made available to users are associated with integers in the inclusive range 1 through 15. A call to *ssignal* associates a procedure, *action*, with the software signal, *sig*; the software signal, *sig*, is raised by a call to *gsignal*. Raising a software signal causes the action established for that signal to be taken.

The first argument to *ssignal* is a number identifying the type of signal for which an action is to be established. The second argument defines the action; it is either the name of a user-defined *action* function or one of the manifest constants SIG\_DFL (default) or SIG\_IGN (ignore). *Ssignal* returns the action previously established for that signal type; if no *action* has been established or the signal number (*sig*) is illegal, *ssignal* returns SIG\_DFL.

*Gsignal* raises the signal identified by its argument, *sig*:

If an *action* function has been established for *sig*, then that *action* is reset to SIG\_DFL and the *action* function is entered with argument *sig*. *Gsignal* returns the value returned to it by the *action* function.

If the *action* for *sig* is SIG\_IGN, *gsignal* returns the value 1 and takes no other action.

If the *action* for *sig* is SIG\_DFL, *gsignal* returns the value 0 and takes no other action.

If *sig* has an illegal value or no *action* was ever specified for *sig*, *gsignal* returns the value 0 and takes no other action.

**SEE ALSO**

signal(2).

**NOTES**

There are some additional signals with numbers outside the range 1 through 15 which are used by the Standard C Library to indicate error conditions. Thus, some signal numbers outside the range 1 through 15 are legal, although their use may interfere with the operation of the Standard C Library.

**NAME**

**stdio** — standard buffered input/output package

**SYNOPSIS**

```
#include <stdio.h>
```

```
FILE *stdin, *stdout, *stderr;
```

**DESCRIPTION**

The functions described in the entries of sub-class 3S of this manual constitute an efficient, user-level I/O buffering scheme. The input/output function may be grouped into the following categories: file access, file status, input, output, miscellaneous. For lists of the functions in each category, refer to the "Libraries" section of the *Programming Guide*. The in-line macros *getc*(3S) and *putc*(3S) handle characters quickly. The macros *getchar* and *putchar*, and the higher-level routines *fgetc*, *fgets*, *sprintf*, *putc*, *fputs*, *fread*, *fscanf*, *fwrite*, *gets*, *getw*, *printf*, *puts*, *putw*, and *scanf* all use *getc* and *putc*; they can be freely intermixed.

A file with associated buffering is called a *stream* and is declared to be a pointer to a defined type FILE. *Fopen*(3S) creates certain descriptive data for a stream and returns a pointer to designate the stream in all further transactions. Normally, there are three open streams with constant pointers declared in the <stdio.h> header file and associated with the standard open files:

<b>stdin</b>	standard input file
<b>stdout</b>	standard output file
<b>stderr</b>	standard error file.

A constant NULL (0) designates a nonexistent pointer.

An integer constant EOF (-1) is returned upon end-of-file or error by most integer functions that deal with streams (see the individual descriptions for details).

An integer constant BUFSIZ specifies the size of the buffers used by the particular implementation.

Any program that uses this package must include the header file of pertinent macro definitions, as follows:

```
#include <stdio.h>
```

The functions and constants mentioned in the entries of sub-class 3S of this manual are declared in that header file and need no further declaration. The constants and the following functions are implemented as macros: *getc*, *getchar*, *putc*, *putchar*, *feof*, *ferror*, *clearerr*, and *fileno*. Redclaration of these names is perilous.

The <stdio.h> file is illustrated in the "Libraries" section of the *Programming Guide*.

**SEE ALSO**

*open*(2), *close*(2), *lseek*(2), *pipe*(2), *read*(2), *ctermid*(3S), *cuserid*(3S), *fclose*(3S), *ferror*(3S), *fopen*(3S), *fread*(3S), *fseek*(3S), *getc*(3S), *gets*(3S), *popen*(3S), *printf*(3S), *putc*(3S), *puts*(3S), *scanf*(3S), *setbuf*(3S), *system*(3S), *tmpfile*(3S), *tmpnam*(3S), *ungetc*(3S), *write*(3).

**DIAGNOSTICS**

Invalid *stream* pointers cause serious errors, possibly including program termination. Individual function descriptions describe the possible error conditions.

**NAME**

*ftok* — standard interprocess communication package

**SYNOPSIS**

```
#include <sys/types.h>
#include <sys/ipc.h>

key_t ftok(path, id)
char *path;
char id;
```

**DESCRIPTION**

All interprocess communication facilities require the user to supply a key to be used by the *msgget(2)*, *semget(2)*, and *shmget(2)* system calls to obtain interprocess communication identifiers. One method for forming a key is to use the *ftok* subroutine described below. Another way to compose keys is to include the project ID in the most significant byte and to use the remaining portion as a sequence number. There are many other ways to form keys, but it is necessary for each system to define standards for forming them. If a standard is not adhered to, unrelated processes may interfere with each other's operation. Therefore, it is strongly suggested that the most significant byte of a key in some sense refer to a project so that keys do not conflict across a given system.

*ftok* returns a key based on *path* and *id* that is usable in subsequent *msgget*, *semget*, and *shmget* system calls. *Path* must be the pathname of an existing file that is accessible to the process. *Id* is a character that uniquely identifies a project. *ftok* returns the same key for linked files when called with the same *id*; it returns different keys when called with the same filename but different *ids*.

**SEE ALSO**

*intro(2)*, *msgget(2)*, *semget(2)*, *shmget(2)*.

**DIAGNOSTICS**

*ftok* returns (key\_t) -1 if *path* does not exist or if it is not accessible to the process.

**WARNING**

If the file whose *path* is passed to *ftok* is removed when keys still refer to the file, future calls to *ftok* with the same *path* and *id* will return an error. If the same file is recreated, *ftok* is likely to return a different key than it did the original time it was called.

**NAME**

lge, lgt, lle, llt — string comparison intrinsic functions

**SYNOPSIS**

character\*N a1, a2  
logical l

l = lge (a1,a2)

l = lgt (a1,a2)

l = lle (a1,a2)

l = llt (a1,a2)

**DESCRIPTION**

These functions return .TRUE. if the inequality holds and .FALSE. otherwise.

## NAME

strcat, strncat, strcmp, strncmp, strcpy, strncpy, strlen, strchr, strrchr, strpbrk, strspn, strespn, strtok - string operations

## SYNOPSIS

```
#include <string.h>
char *strcat (s1, s2)
char *s1, *s2;
char *strncat (s1, s2, n)
char *s1, *s2;
int n;
int strcmp (s1, s2)
char *s1, *s2;
int strncmp (s1, s2, n)
char *s1, *s2;
int n;
char *strcpy (s1, s2)
char *s1, *s2;
char *strncpy (s1, s2, n)
char *s1, *s2;
int n;
int strlen (s)
char *s;
char *strchr (s, c)
char *s;
int c;
char *strrchr (s, c)
char *s;
int c;
char *strpbrk (s1, s2)
char *s1, *s2;
int strspn (s1, s2)
char *s1, *s2;
int strespn (s1, s2)
char *s1, *s2;
char *strtok (s1, s2)
char *s1, *s2;
```

## DESCRIPTION

The arguments *s1*, *s2*, and *s* point to strings (arrays of characters terminated by a null character). The functions *strcat*, *strncat*, *strcpy*, and *strncpy* all alter *s1*. These functions do not check for overflow of the array pointed to by *s1*.

*Strcat* appends a copy of string *s2* to the end of string *s1*. *Strncat* appends at most *n* characters. Each function returns a pointer to the null-terminated result.

*Strcmp* performs a lexicographical comparison of its arguments and returns an integer less than, equal to, or greater than 0, when *s1* is less than, equal

to, or greater than *s2*, respectively. *Strncmp* makes the same comparison but looks at a maximum of *n* characters.

*Strncpy* copies string *s2* to string *s1*, stopping after the null character has been copied. *Strncpy* copies exactly *n* characters, truncating *s2* or adding null characters to *s1* if necessary. The result is not null-terminated if the length of *s2* is *n* or more. Each function returns *s1*.

*Strlen* returns the number of characters in *s*, not including the terminating null character.

*Strchr* (*strchr*) returns a pointer to the first (last) occurrence of character *c* in string *s*, or a NULL pointer if *c* does not occur in the string. The null character terminating a string is considered to be part of the string.

*Strpbrk* returns a pointer to the first occurrence in string *s1* of any character from string *s2*, or a NULL pointer if no character from *s2* exists in *s1*.

*Strspn* (*strcspn*) returns the length of the initial segment of string *s1* which consists entirely of characters from (not from) string *s2*.

*Strtok* considers the string *s1* to consist of a sequence of zero or more text tokens separated by spans of one or more characters from the separator string *s2*. The first call (with pointer *s1* specified) returns a pointer to the first character of the first token, and writes a null character into *s1* immediately following the returned token. The function keeps track of its position in the string between separate calls, so that on subsequent calls (which must be made with a NULL pointer as the first argument) it works through the string *s1* immediately following that token. This can be continued until no tokens remain. The separator string *s2* may be different from call to call. When no token remains in *s1*, a NULL pointer is returned.

**NOTE**

For user convenience, all these functions are declared in the optional `<string.h>` header file.

**BUGS**

*Strcmp* use native character comparison. Thus the sign of the value returned when one of the characters has its high-order bit set is implementation-dependent.

All string movement is performed character by character starting at the left. Thus overlapping moves toward the left will work as expected, but overlapping moves to the right may yield surprises.

**NAME**

`strtod`, `atof` — convert string to double-precision number

**SYNOPSIS**

`double strtod (str, ptr)`

`char *str, **ptr;`

`double atof (str)`

`char *str;`

**DESCRIPTION**

*Strtod* returns as a double-precision floating-point number the value represented by the character string pointed to by *str*. The string is scanned up to the first unrecognized character.

*Strtod* recognizes an optional string of "white-space" characters (as defined by *isspace* in *cctype*(3C)), then an optional sign, then a string of digits optionally containing a decimal point, then an optional e or E followed by an optional sign or space, followed by an integer.

If the value of *ptr* is not `(char **)NULL`, a pointer to the character terminating the scan is returned in the location pointed to by *ptr*. If no number can be formed, *\*ptr* is set to *str*, and zero is returned.

*Atof(str)* is equivalent to *strtod(str, (char \*\*)NULL)*.

**SEE ALSO**

*cctype*(3C), *scanf*(3S), *strtol*(3C).

**DIAGNOSTICS**

If the correct value would cause overflow, `±HUGE` is returned (according to the sign of the value), and *errno* is set to `ERANGE`.

If the correct value would cause underflow, zero is returned and *errno* is set to `ERANGE`.



**NAME**

`strtol`, `atol`, `atoi` — convert string to integer

**SYNOPSIS**

```
long strtol (str, ptr, base)
```

```
char *str,** ptr;
```

```
int base;
```

```
long atol (str)
```

```
char *str;
```

```
int atoi (str)
```

```
char *str;
```

**DESCRIPTION**

*Strtol* returns as a long integer the value represented by the character string pointed to by *str*. The string is scanned up to the first character inconsistent with the base. Leading white-space characters (blanks and tabs) are ignored.

If the value of *ptr* is not `(char **)NULL`, a pointer to the character terminating the scan is returned in the location pointed to by *ptr*. If no integer can be formed, zero is returned.

If *base* is positive (and not greater than 36), it is used as the base for conversion. After an optional leading sign, leading zeros are ignored; a leading `0x` or `0X` is ignored if *base* is 16.

If *base* is zero, the string itself determines the base. After an optional leading sign, a leading zero indicates octal conversion and a leading `0x` or `0X` indicates hexadecimal conversion; otherwise, decimal conversion is used.

Truncation from long to int can take place upon assignment or by an explicit cast.

*Atol(str)* is equivalent to `strtol(str, (char **)NULL, 10)`.

*Atoi(str)* is equivalent to `(int) strtol(str, (char **)NULL, 10)`.

**SEE ALSO**

`ctype(3C)`, `scanf(3S)`, `strtod(3C)`.

**BUGS**

Overflow conditions are ignored.

**NAME**

swab - swap bytes

**SYNOPSIS**

```
void swab (from, to, nbytes)
char *from, *to;
int nbytes;
```

**DESCRIPTION**

*Swab* copies *nbytes* bytes pointed to by *from* to the array pointed to by *to*, exchanging adjacent even and odd bytes. It is useful for carrying binary data between PDP-11s and other machines. *Nbytes* should be even and non-negative. If *nbytes* is odd and positive, *swab* uses *nbytes*-1 instead. If *nbytes* is negative, *swab* does nothing.

**NAME**

system — issue a shell command from Fortran

**SYNOPSIS**

character \*N c

call system(c)

**DESCRIPTION**

*System* causes its character argument to be given to *sh(1)* as input, as if the string had been typed at a terminal. The current process waits until the shell has completed.

**SEE ALSO**

*sh(1)*, *exec(2)*, *system(3S)*.

**NAME**

`system` — issue a shell command

**SYNOPSIS**

```
#include <stdio.h>

int system (string)
char *string;
```

**DESCRIPTION**

*System* causes *string* to be given to *sh*(1) as input, as if the string had been typed as a command at a terminal. The current process waits until the shell has completed, then returns the exit status of the shell.

**FILES**

/bin/sh

**SEE ALSO**

*sh*(1), *exec*(2).

**DIAGNOSTICS**

*System* forks to create a child process that in turn performs *exec*(2) on */bin/sh* in order to execute *string*. If the fork or exec fails, *system* returns a negative value and sets *errno*.

**NAME**

tan, dtan - Fortran tangent intrinsic function

**SYNOPSIS**

real r1, r2  
double precision dp1, dp2

r2 = tan(r1)

dp2 = dtan(dp1)

dp2 = tan(dp1)

**DESCRIPTION**

*Tan* returns the real tangent of its real argument. *Dtan* returns the double-precision tangent of its double-precision argument. The generic *tan* function becomes *dtan* as required with a double-precision argument.

**SEE ALSO**

trig(3M).

**NAME**

*tanh*, *dtanh* — Fortran hyperbolic tangent intrinsic function

**SYNOPSIS**

real *r1*, *r2*

double precision *dp1*, *dp2*

*r2* = *tanh*(*r1*)

*dp2* = *dtanh*(*dp1*)

*dp2* = *tanh*(*dp1*)

**DESCRIPTION**

*Tanh* returns the real hyperbolic tangent of its real argument. *Dtanh* returns the double-precision hyperbolic tangent of its double precision argument. The generic form *tanh* may be used to return a double-precision value given a double-precision argument.

**SEE ALSO**

*sinh*(3M).

## NAME

`tgetent`, `tgetnum`, `tgetflag`, `tgetstr`, `tgoto`, `tputs` – terminal independent operation routines

## SYNOPSIS

```
char PC;
char *BC;
char *UP;
short ospeed;

tgetent(bp, name)
char *bp, *name;

tgetnum(id)
char *id;

tgetflag(id)
char *id;

char *
tgetstr(id, area)
char *id, **area;

char *
tgoto(cm, destcol, destline)
char *cm;

tputs(cp, affcnt, outc)
register char *cp;
int affcnt;
int *outc);
```

## DESCRIPTION

These functions extract and use capabilities from the terminal capability data base *termcap*(5). Note that these are low-level routines.

*Tgetent* extracts the entry for terminal *name* into the buffer at *bp*. *Bp* should be a character buffer of size 1024 and must be retained through all subsequent calls to *tgetnum*, *tgetflag*, and *tgetstr*. *Tgetent* returns -1 if it cannot open the *termcap* file, 0 if the terminal name given does not have an entry, and 1 if successful. It looks in the environment for a *TERMCAP* variable. If a variable is found whose value does not begin with a slash and the terminal type *name* is the same as the environment string *TERM*, the *TERMCAP* string is used instead of reading the *termcap* file. If the value does begin with a slash, the

string is used as a pathname rather than */etc/termcap*. This can speed up entry into programs that call *tgetent*. It can also help debug new terminal descriptions or be used to make one for your terminal if you can't write the file */etc/termcap*.

*Tgetnum* gets the numeric value of capability *id*, returning *-1* if is not given for the terminal. *Tgetflag* returns 1 if the specified capability is present in the terminal's entry, 0 if it is not. *Tgetstr* gets the string value of capability *id*, placing it in the buffer at *area*, advancing the *area* pointer. It decodes the abbreviations for this field described in *termcap(5)*, except for cursor addressing and padding information.

*Tgoto* returns a cursor addressing string decoded from *cm* to go to column *destcol* in line *destline*. It uses the external variables UP (from the up capability) and BC (if bc is given rather than bs) if necessary to avoid placing *\n*, *^D* or *^@* in the returned string. (Programs that call *tgoto* should be sure to turn off the XTABS bit(s), since *tgoto* may now output a tab. Note that programs using *termcap* should in general turn off XTABS anyway since some terminals use control-I for other functions, such as nondestructive space.) If a % sequence is given which is not understood, then *tgoto* returns OOPS.

*Tputs* decodes the leading padding information of the string *cp*; *affcnt* gives the number of lines affected by the operation, or 1 if this is not applicable; *putc* is a routine that is called with each character in turn. The external variable *ospeed* should contain the output speed of the terminal as encoded by *stty(1)*. The external variable PC should contain a pad character to be used (from the pc capability) if a null (*^@*) is inappropriate.

#### FILES

*/usr/lib/libtermcap.a* *-termcap* library  
*/etc/termcap* data base

#### SEE ALSO

*ex(1)*, *termcap(5)*



**NAME**

`tmpfile` — create a temporary file

**SYNOPSIS**

```
#include <stdio.h>
```

```
FILE *tmpfile ()
```

**DESCRIPTION**

*Tmpfile* creates a temporary file using a name generated by *tmpnam(3S)*, and returns a corresponding FILE pointer. If the file cannot be opened, an error message is printed using *perror(3C)*, and a NULL pointer is returned. The file is automatically deleted when the process using it terminates. The file is opened for update ("w+"). *Tmpfile* calls *fopen* and so returns any error code passed to it from *fopen*.

**SEE ALSO**

`creat(2)`, `unlink(2)`, `fopen(3S)`, `mktemp(3C)`, `perror(3C)`, `tmpnam(3S)`.

**NAME**

`tmpnam`, `tempnam` — create a name for a temporary file

**SYNOPSIS**

```
#include <stdio.h>

char *tmpnam (s)
char *s;

char *tempnam (dir, pfx)
char *dir, *pfx;
```

**DESCRIPTION**

These functions generate filenames that can safely be used for a temporary file.

*tmpnam* always generates a filename using the pathname defined as *P\_tmpdir* in the `<stdio.h>` header file. If *s* is NULL, *tmpnam* leaves its result in an internal static area and returns a pointer to that area. The next call to *tmpnam* will destroy the contents of the area. If *s* is not NULL, it is assumed to be the address of an array of at least *L\_tmpnam* bytes, where *L\_tmpnam* is a constant defined in `<stdio.h>`; *tmpnam* places its result in that array and returns *s*.

*Tempnam* allows the user to control the choice of a directory. The argument *dir* points to the pathname of the directory in which the file is to be created. If *dir* is NULL or points to a string which is not a pathname for an appropriate directory, the pathname defined as *P\_tmpdir* in the `<stdio.h>` header file is used. If that pathname is not accessible, `/tmp` will be used as a last resort. This entire sequence can be upstaged by providing an environment variable `TMPDIR` in the user's environment, whose value is a pathname for the desired temporary-file directory.

Many applications prefer that names of temporary files contain favorite initial letter sequences. Use the *pfx* argument for this. This argument may be NULL or point to a string of up to 5 characters to be used as the first few characters of the name of the temporary file.

*Tempnam* uses `malloc(3C)` to get space for the constructed filename and returns a pointer to this area. Thus, any pointer value returned from *tempnam* may serve as an argument to `free` (see `malloc(3C)`). If *tempnam* cannot return the expected result for any reason (i.e., `malloc` failed or attempts to find an appropriate directory were unsuccessful), a NULL pointer will be returned.

**NOTES**

These functions generate a different filename each time they are called.

Files created using these functions and either `fopen(3S)` or `creat(2)` are temporary only in the sense that they reside in a directory intended for temporary use and their names are unique. It is the user's responsibility to use `unlink(2)` to remove the file when its use is ended.

**SEE ALSO**

`creat(2)`, `unlink(2)`, `fopen(3S)`, `malloc(3C)`, `mktemp(3C)`, `tmpfile(3S)`.

**BUGS**

If called more than 17,576 times in a single process, *tmpnam* and *tempnam*

will start recycling previously used names.

Between the time a filename is created and the file is opened, it is possible for some other process to create a file with the same name. This can never happen if that other process is using *tmpnam*, *tempnam*, or *mktemp(3C)* and the filenames are chosen carefully to avoid duplication by other means.

## NAME

sin, cos, tan, asin, acos, atan, atan2 — trigonometric functions

## SYNOPSIS

```
#include <math.h>

double sin (x)
double x;

double cos (x)
double x;

double tan (x)
double x;

double asin (x)
double x;

double acos (x)
double x;

double atan (x)
double x;

double atan2 (y, x)
double x, y;
```

## DESCRIPTION

*Sin*, *cos*, and *tan* return, respectively, the sine, cosine, and tangent of their argument, which is in radians.

*Asin* returns the arcsine of  $x$ , in the range  $-\pi/2$  to  $\pi/2$ .

*Acos* returns the arccosine of  $x$ , in the range 0 to  $\pi$ .

*Atan* returns the arctangent of  $x$ , in the range  $-\pi/2$  to  $\pi/2$ .

*Atan2* returns the arctangent of  $y/x$ , in the range  $-\pi$  to  $\pi$ , using the signs of both arguments to determine the quadrant of the return value.

## DIAGNOSTICS

*Sin*, *cos*, and *tan* lose accuracy when their argument is far from zero. For arguments sufficiently large, these functions return 0 when there would otherwise be a complete loss of significance. In this case a message indicating TLOSS error is printed on the standard error output. For less extreme arguments, a PLOSS error is generated but no message is printed. In both cases, *errno* is set to ERANGE.

If the magnitude of the argument of *asin* or *acos* is greater than one, or if both arguments of *atan2* are zero, zero is returned and *errno* is set to EDOM. In addition, a message indicating DOMAIN error is printed on the standard error output.

These error-handling procedures may be changed with the function *matherr*(3M).

## SEE ALSO

*matherr*(3M).

## NAME

*tsearch*, *tfind*, *tdelete*, *twalk* — manage binary search trees

## SYNOPSIS

```
#include <search.h>

char *tsearch ((char *) key, (char **) rootp, compar)
int (*compar) ( );

char *tfind ((char *) key, (char **) rootp, compar)
int (*compar) ( );

char *tdelete ((char *) key, (char **) rootp, compar)
int (*compar) ( );

void twalk ((char *) root, action)
void (*action) ( );
```

## DESCRIPTION

*Tsearch*, *tfind*, *tdelete*, and *twalk* are routines for manipulating binary search trees. They are generalized from Knuth (6.2.2) Algorithms T and D. All comparisons are done with a user-supplied routine. This routine is called with two arguments, the pointers to the elements being compared. It returns an integer less than, equal to, or greater than 0, according to whether the first argument is to be considered less than, equal to or greater than the second argument. The comparison function need not compare every byte, so arbitrary data may be contained in the elements in addition to the values being compared.

*Tsearch* is used to build and access the tree. *Key* is a pointer to a datum to be accessed or stored. If there is a datum in the tree equal to *\*key* (the value pointed to by *key*), a pointer to this found datum is returned. Otherwise, *\*key* is inserted, and a pointer to it returned. Only pointers are copied, so the calling routine must store the data. *Rootp* points to a variable that points to the root of the tree. A NULL value for the variable pointed to by *rootp* denotes an empty tree; in this case, the variable will be set to point to the datum which will be at the root of the new tree.

Like *tsearch*, *tfind* will search for a datum in the tree, returning a pointer to it if found. However, if it is not found, *tfind* will return a NULL pointer. The arguments for *tfind* are the same as for *tsearch*.

*Tdelete* deletes a node from a binary search tree. The arguments are the same as for *tsearch*. The variable pointed to by *rootp* will be changed if the deleted node was the root of the tree. *Tdelete* returns a pointer to the parent of the deleted node, or a NULL pointer if the node is not found.

*Twalk* traverses a binary search tree. *Root* is the root of the tree to be traversed. (Any node in a tree may be used as the root for a walk below that node.) *Action* is the name of a routine to be invoked at each node. This routine is, in turn, called with three arguments. The first argument is the address of the node being visited. The second argument is a value from an enumeration data type `typedef enum { preorder, postorder, endorder, leaf } VISIT;` (defined in the `<search.h>` header file), depending on whether this is the first, second or third time that the node has been visited (during a depth-first, left-to-right traversal of the tree), or whether the node is a leaf. The third argument is the level of the node in the tree, with the root being level zero.

The pointers to the key and the root of the tree should be of type pointer-to-element, and cast to type pointer-to-character. Similarly, although declared as type pointer-to-character, the value returned should be cast into type pointer-to-element.

**EXAMPLE**

The following code reads in strings and stores structures containing a pointer to each string and a count of its length. It then walks the tree, printing out the stored strings and their lengths in alphabetical order.

```
#include <search.h>
#include <stdio.h>

struct node {          /* pointers to these are stored in the tree */
    char *string;
    int length;
};
char string_space[10000]; /* space to store strings */
struct node nodes[500];   /* nodes to store */
struct node *root = NULL; /* this points to the root */

main( )
{
    char *strptr = string_space;
    struct node *nodeptr = nodes;
    void print_node( ), twalk( );
    int i = 0, node_compare( );

    while (gets(strptr) != NULL && i++ < 500) {
        /* set node */
        nodeptr->string = strptr;
        nodeptr->length = strlen(strptr);
        /* put node into the tree */
        (void) tsearch((char *)nodeptr, &root,
                       node_compare);
        /* adjust pointers, so we don't overwrite tree */
        strptr += nodeptr->length + 1;
        nodeptr++;
    }
    twalk(root, print_node);
}
/*
   This routine compares two nodes, based on an
   alphabetical ordering of the string field.
*/
int
node_compare(node1, node2)
struct node *node1, *node2;
{
    return strcmp(node1->string, node2->string);
}
/*
   This routine prints out a node, the first time
   twalk encounters it.
*/
```

```

*/
void
print_node(node, order, level)
struct node **node;
VISIT order;
int level;
{
    if (order == preorder || order == leaf) {
        (void)printf("string = %20s, length = %d\n",
            (*node)--> string, (*node)--> length);
    }
}

```

**SEE ALSO**

bsearch(3C), hsearch(3C), lsearch(3C).

**DIAGNOSTICS**

A NULL pointer is returned by *tsearch* if there is not enough space available to create a new node.

A NULL pointer is returned by *tsearch*, *tfind* and *tdelete* if **rootp** is NULL on entry.

If the datum is found, both *tsearch* and *tfind* return a pointer to it. If not, *tfind* returns NULL, and *tsearch* returns a pointer to the inserted item.

**WARNINGS**

The **root** argument to *twalk* is one level of indirection less than the **rootp** arguments to *tsearch* and *tdelete*.

There are two nomenclatures used to refer to the order in which tree nodes are visited. *Tsearch* uses *preorder*, *postorder* and *endorder* to respectively refer to visiting a node before any of its children, after its left child and before its right, and after both its children. The alternate nomenclature uses *preorder*, *inorder* and *postorder* to refer to the same visits, which could result in some confusion over the meaning of *postorder*.

**BUGS**

If the calling function alters the pointer to the root, results are unpredictable.

**NAME**

*ttyname*, *isatty* — find name of a terminal

**SYNOPSIS**

```
char *ttyname (fildes)
int isatty (fildes)
int fildes;
```

**DESCRIPTION**

*Ttyname* returns a pointer to a string containing the null-terminated path-name of the terminal device associated with file descriptor *fildes*.

*Isatty* returns 1 if *fildes* is associated with a terminal device; otherwise, it returns 0.

**FILES**

/dev/\*

**DIAGNOSTICS**

*Ttyname* returns a NULL pointer if *fildes* does not describe a terminal device in directory /dev.

**BUGS**

The return value points to static data whose content is overwritten by each call.



**NAME**

ttyslot — find the slot in the utmp file of the current user

**SYNOPSIS**

int ttyslot ( )

**DESCRIPTION**

*Ttyslot* returns the index of the current user's entry in the */etc/utmp* file. This is accomplished by scanning the file */etc/inittab* for the name of the terminal device associated with the standard input, the standard output, or the error output (0, 1, or 2).

**FILES**

*/etc/inittab*  
*/etc/utmp*

**SEE ALSO**

getut(3C), ttyname(3C).

**DIAGNOSTICS**

A value of 0 is returned if an error is encountered while searching for the terminal name or if none of the above file descriptors is associated with a terminal device.

**NAME**

`ungetc` — push character back into input stream

**SYNOPSIS**

```
#include <stdio.h>
int ungetc (c, stream)
char c;
FILE *stream;
```

**DESCRIPTION**

*Ungetc* inserts the character *c* into the buffer associated with an input *stream*. That character, *c*, will be returned by the next *getc* call on that *stream*. *Ungetc* returns *c* and leaves the file *stream* unchanged.

One character of pushback is guaranteed provided something has been read from the stream and the stream is actually buffered. In the case that *stream* is *stdin*, one character may be pushed back onto the buffer without a previous read statement.

If *c* equals EOF, *ungetc* does nothing to the buffer and returns EOF.

*Fseek*(3S) erases all memory of inserted characters.

**SEE ALSO**

*fseek*(3S), *getc*(3S), *setbuf*(3S).

**DIAGNOSTICS**

*Ungetc* returns EOF if it can't insert the character.

**NAME**

*vprintf*, *vfprintf*, *vsprintf* — print formatted output of a *varargs* argument list

**SYNOPSIS**

```
#include <stdio.h>
#include <varargs.h>

int vprintf (format, ap)
char *format;
va_list ap;

int fprintf (stream, format, ap)
FILE *stream;
char *format;
va_list ap;

int sprintf (s, format, ap);
char *s, *format;
va_list ap;
```

**DESCRIPTION**

*vprintf*, *vfprintf*, and *vsprintf* are the same as *printf*, *fprintf*, and *sprintf* respectively, except that instead of being called with a variable number of arguments, they are called with an argument list as defined by *varargs*(5).

**EXAMPLE**

The following demonstrates how *vfprintf* could be used to write an error routine.

```
#include <stdio.h>
#include <varargs.h>
.
.
.
/*
 * error should be called like
 * error(function_name, format, arg1, arg2...);
 */
/*VARARGS0*/
void
error(va_list)
/* Note that the function_name and format arguments cannot be
 * separately declared because of the definition of varargs.
 */
va_dcl
{
    va_list args;
    char *fmt;

    va_start(args);
    /* print out name of function causing error */
    (void)fprintf(stderr, "ERROR in %s: ", va_arg(args, char *));
    fmt = va_arg(args, char *);
    /* print out remainder of message */
    (void)vfprintf(fmt, args);
    va_end(args);
}
```

**VPRINTF(3S)**

**VPRINTF(3S)**

```
        (void)abort( );  
    }
```

**SEE ALSO**

`vprintf(3X)`, `varargs(5)`.

## NAME

*vprintf*, *vfprintf*, *vsprintf* - print formatted output of a *varargs* argument list

## SYNOPSIS

```
#include <stdio.h>
#include <varargs.h>

int vprintf (format, ap)
char *format;
va_list ap;

int vfprintf (stream, format, ap)
FILE *stream;
char *format;
va_list ap;

int vsprintf (s, format, ap)
char *s, *format;
va_list ap;
```

## DESCRIPTION

*vprintf*, *vfprintf*, and *vsprintf* are the same as *printf*, *fprintf*, and *sprintf* respectively, except that instead of being called with a variable number of arguments, they are called with an argument list as defined by *varargs(5)*.

## EXAMPLE

The following demonstrates how *vfprintf* could be used to write an error routine.

```
#include <stdio.h>
#include <varargs.h>
.
.
.
/*
 *   error should be called like
 *       error(function_name, format, arg1, arg2...);
 */
/*VARARGS0*/
void
error(va_list)
/* Note that the function_name and format arguments cannot be
 *   separately declared because of the definition of varargs.
 */
va_dcl
{
    va_list args;
    char *fmt;

    va_start(args);
    /* print out name of function causing error */
    (void)fprintf(stderr, "ERROR in %s: ", va_arg(args, char *));
    fmt = va_arg(args, char *);
    /* print out remainder of message */
    (void)vfprintf(fmt, args);
    va_end(args);
}
```

**VPRINTF (3X)**

```
        (void)abort( );  
    }
```

**SEE ALSO**

printf(3S), varargs(5).

**VPRINTF (3X)**

**NAME**

write — write on a file

**SYNOPSIS**

```
int write (fildes, buf, nbyte)
int fildes;
char *buf;
unsigned nbyte;
```

**DESCRIPTION**

*Fildes* is a file descriptor obtained from a *creat*, *open*, *dup*, *fcntl*, *pipe*, or *socket* system call.

*Write* attempts to write *nbyte* bytes from the buffer pointed to by *buf* to the file associated with the *fildes*.

On devices capable of seeking, the actual writing of data proceeds from the position in the file indicated by the file pointer. Upon return from *write*, the file pointer is incremented by the number of bytes actually written.

On devices incapable of seeking, writing always takes place starting at the current position. The value of a file pointer associated with such a device is undefined.

If the *O\_APPEND* flag of the file status flags is set, the file pointer will be set to the end of the file prior to each write.

*Write* will fail and the file pointer will remain unchanged if one or more of the following are true:

- [EBADF] *Fildes* is not a valid file descriptor open for writing.
- [EPIPE and SIGPIPE signal] 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 pipe that is not open for reading by any process.
- [EFBIG] An attempt was made to write a file that exceeds the process's file size limit or the maximum file size. See *ulimit(2)*.
- [EFAULT] Part of *iov* or data to be written to the file points outside the process's allocated address space.
- [EFAULT] *Buf* points outside the process's allocated address space.
- [EINTR] A signal was caught during the *write* system call.

If a *write* requests that more bytes be written than there is room for (e.g., the *ulimit* (see *ulimit(2)*) or the physical end of a medium), only as many bytes as there is room for will be written. For example, suppose there is space for 20 bytes more in a file before reaching a limit. A write of 512 bytes will return 20. The next write of a non-zero number of bytes will give a failure return (except as noted below).

If the file being written is a pipe (or FIFO) and the *O\_NDELAY* flag of the file flag word is set, then write to a full pipe (or FIFO) will return a count of 0. Otherwise (*O\_NDELAY* clear), writes to a full pipe (or FIFO) will block until space becomes available.

**WRITE(3)**

**WRITE(3)**

**RETURN VALUE**

Upon successful completion the number of bytes actually written is returned. Otherwise, -1 is returned and *errno* is set to indicate the error.

**SEE ALSO**

`creat(2)`, `lseek(2)`, `open(2)`, `pipe(2)`, `socket(2N)`, `ulimit(2)`.



## NAME

writev — write on a file

## SYNOPSIS

```
#include <sys/types.h>
#include <sys/uio.h>
```

```
writev(d, iov, iovectlen)
int d;
struct iovec *iov;
int iovectlen;
```

## DESCRIPTION

*Fildes* is a file descriptor obtained from a *creat*, *open*, *dup*, *fcntl*, *pipe*, or *socket* system call.

*Writev* attempts to write *nbyte* bytes to the file associated with the *fildes* and gathers the output data from the *iovlen* buffers specified by the members of the *iovec* array: *iov[0]*, *iov[1]*, etc.

On devices capable of seeking, the actual writing of data proceeds from the position in the file indicated by the file pointer. Upon return from *writev*, the file pointer is incremented by the number of bytes actually written.

On devices incapable of seeking, writing always takes place starting at the current position. The value of a file pointer associated with such a device is undefined.

If the *O\_APPEND* flag of the file status flags is set, the file pointer will be set to the end of the file prior to each write.

*Writev* will fail and the file pointer will remain unchanged if one or more of the following are true:

- [EBADF] *Fildes* is not a valid file descriptor open for writing.
- [EPIPE and SIGPIPE signal] 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 pipe that is not open for reading by any process.
- [EFBIG] An attempt was made to write a file that exceeds the process's file size limit or the maximum file size. See *ulimit(2)*.
- [EFAULT] Part of *iov* or data to be written to the file points outside the process's allocated address space.
- [EFAULT] *Buf* points outside the process's allocated address space.
- [EINTR] A signal was caught during the *writev* system call.

If a *writev* requests that more bytes be written than there is room for (e.g., the *ulimit* (see *ulimit(2)*) or the physical end of a medium), only as many bytes as there is room for will be written. For example, suppose there is space for 20 bytes more in a file before reaching a limit. A write of 512 bytes will return 20. The next write of a non-zero number of bytes will give a failure return (except as noted below).

If the file being written is a pipe (or FIFO) and the `O_NDELAY` flag of the file flag word is set, then write to a full pipe (or FIFO) will return a count of 0. Otherwise (`O_NDELAY` clear), writes to a full pipe (or FIFO) will block until space becomes available.

**RETURN VALUE**

Upon successful completion the number of bytes actually written is returned. Otherwise, `-1` is returned and *errno* is set to indicate the error.

**SEE ALSO**

`creat(2)`, `lseek(2)`, `open(2)`, `pipe(2)`, `socket(2N)`, `ulimit(2)`.

**NAME**

intro - introduction to file formats

**DESCRIPTION**

This section outlines the formats of various files. The C struct declarations for the file formats are given where applicable. Usually, these structures can be found in the directories `/usr/include` or `/usr/include/sys`.

References of the type *name(1M)* refer to entries found in Section 1 of the *UniPlus+ Administrator's Manual*.

## NAME

acct — per-process accounting file format

## SYNOPSIS

```
#include <sys/acct.h>
```

## DESCRIPTION

Files produced as a result of calling *acct(2)* have records in the form defined by *<sys/acct.h>*, whose contents are:

```
typedef  ushort comp_t; /* "floating point" */
                /* 13-bit fraction, 3-bit exponent */

struct   acct {
    char   ac_flag;      /* Accounting flag */
    char   ac_stat;     /* Exit status */
    ushort ac_uid;      /* Accounting user ID */
    ushort ac_gid;      /* Accounting group ID */
    dev_t  ac_tty;      /* control typewriter */
    time_t ac_btime;    /* Beginning time */
    comp_t ac_utime;     /* acctng user time in clock ticks */
    comp_t ac_stime;     /* acctng system time in clock ticks */
    comp_t ac_etime;     /* acctng elapsed time in clock ticks */
    comp_t ac_mem;      /* memory usage in clicks */
    comp_t ac_io;        /* chars transfd by read/write */
    comp_t ac_rw;        /* number of block reads/writes */
    char   ac_comm[8];  /* command name */
};

extern   struct acct acctbuf;
extern   struct inode *acctp; /* inode of accounting file */

#define  AFORK  01      /* has executed fork, but no exec */
#define  ASU    02      /* used super-user privileges */
#define  ACCTF  0300    /* record type: 00 = acct */
```

In *ac\_flag*, the AFORK flag is turned on by each *fork(2)* and turned off by an *exec(2)*. The *ac\_comm* field is inherited from the parent process and is reset by any *exec*. Each time the system charges the process with a clock tick, it also adds to *ac\_mem* the current process size, computed as follows:

$$(\text{data size}) + (\text{text size}) / (\text{number of in-core processes using text})$$

The value of  $ac\_mem / (ac\_stime + ac\_utime)$  can be viewed as an approximation to the mean process size, as modified by text-sharing.

The structure `tacct`, which resides with the source files of the accounting commands, represents the total accounting format used by the various accounting commands:

```

/*
 * total accounting (for acct period), also for day
 */
struct  tacct {
    uid_t    ta_uid;        /* userid */
    char     ta_name[8];   /* login name */
    float    ta_cpu[2];    /* cum. cpu time, p/np (mins) */
    float    ta_kcore[2]; /* cum kcore-minutes, p/np */
    float    ta_con[2];   /* cum. connect time, p/np, mins */
    float    ta_du;       /* cum. disk usage */
    long     ta_pc;       /* count of processes */
    unsigned short ta_sc; /* count of login sessions */
    unsigned short ta_dc; /* count of disk samples */
    unsigned short ta_fee; /* fee for special services */
};

```

#### SEE ALSO

`acct(1M)`, `acctcom(1)`, `acct(2)`, `exec(2)`, `fork(2)`.

#### BUGS

The `ac_mem` value for a short-lived command gives little information about the actual size of the command, because `ac_mem` may be incremented while a different command (e.g., the shell) is being executed by the process.

**NAME**

aliases – aliases file for delivermail

**SYNOPSIS**

`/usr/lib/aliases`

**DESCRIPTION**

This file describes user ID aliases that will be used by `/etc/delivermail`. It is formatted as a series of lines of the form

`name:addr1,addr2,...addrn`

The *name* is the name to alias, and the *addr*i** are the addresses to send the message to. Lines beginning with white space are continuation lines. Lines beginning with '#' are comments.

Aliasing occurs only on local names. Loops cannot occur since no message will be sent to any person more than once.

**SEE ALSO**

`delivermail(8N)`.

**NAME**

`altblk` — alternate block information for bad block handling

**SYNOPSIS**

```
#include <altblk.h>
```

**DESCRIPTION**

*Altblk* is the data structure used by *badblk(1M)* to handle bad blocks for disk drives that support soft sector bad block remapping.

The layout of this structure is as follows:

```
#define MAXALT 50 /* max alternate disk blocks */
#define ALTMAGIC 0xDBDF /* bad block information is valid flag */

/*
 * structure for alternate block mapping
 */
struct a_map {
    long a_altblk; /* bad block */
    long a_index; /* relative bad block index */
};

/*
 * disk header block format for alternate block mapping
 */
struct altblk {
    char a_fill[BSIZE-sizeof(struct a_map)-4*sizeof(long)];
    /* fill to make structure BSIZE bytes long */
    struct a_map a_map[1]; /* mapping */
    long a_magic; /* verification code (ALTMAGIC) */
    long a_count; /* bad block count */
    long a_nicbad; /* max number of bad blocks */
    long a_maxalt; /* max alt block used so far */
};
```

This structure describes the upper portion of block 0 of each physical disk. The array `a_map` is inverted (i.e., it is indexed backwards). The specific fields in *altblk* are:

`a_maxalt` — the next usable block in bad block area relative to the start of the bad block area  
`a_nicbad` — the maximum number of elements in the `a_map` structure  
`a_count` — the number of bad blocks currently remapped on the disk  
`a_magic` — a magic number for verification  
`a_map` — bad block remap information

**SEE ALSO**

*badblk(1M)*

**NAME**

a.out — common assembler and link editor output

**DESCRIPTION**

**A.out** is the output file from the assembler *as(1)* and the link editor *ld(1)*. *A.out* can be executed on the target machine if there were no errors in assembling or linking and no unresolved external references.

The object file format supports user-defined sections and contains extensive information for symbolic software testing. A common object file consists of a file header, an optional aout header, a table of section headers, relocation information, (optional) line numbers, and a symbol table. The order is given below.

```

File header.
Optional aout header.
Section 1 header.
...
Section n header.
Section 1 data.
...
Section n data.
Section 1 relocation.
...
Section n relocation.
Section 1 line numbers.
...
Section n line numbers.
Symbol table.
String table.

```

The last four sections (relocation, line numbers, symbol table, and string table) may be missing if the program was linked with the *-s* option of *ld(1)* or if the symbol table and relocation bits were removed by *strip(1)*. Also note that if the program was linked without the *-r* option, the relocation information will be absent. The string table exists only if necessary.

When an **a.out** file is loaded into memory for execution, three logical segments are set up: the text segment, the data segment (initialized data followed by uninitialized data, the latter actually being initialized to all 0's), and a stack. The text segment begins at location 0 in the core image; the header is not loaded. If the magic number (the first field in the optional aout header) is 407 (octal), it indicates that the text segment is not to be write-protected or shared, so the data segment will be contiguous with the text segment. If the magic number is 410 (octal), the data segment begins at the next segment boundary following the text segment, and the text segment is not writable by the program. If other processes are executing the same **a.out** file, they will share a single text segment.

On the M68000 family of processors the stack begins at the end of memory and grows toward lower addresses. The stack is automatically extended as required. The data segment is extended only as requested by the *brk(2)* and *sbrk(2)* system calls.

The value of a word in the text or data portions that is not a reference to an undefined external symbol is exactly the value that will appear in



memory when the file is executed. If a word in the text involves a reference to an undefined external symbol, the storage class of the symbol-table entry for that word will be marked as an "external symbol", and the section number will be set to 0. When the file is processed by the link editor and the external symbol becomes defined, the value of the symbol will be added to the word in the file.

See *aouthdr(4)*, *filehdr(4)*, *linenum(4)*, *scnhdr(4)*, *reloc(4)*, and *syms(4)* for descriptions of the individual parts. Every section created by *as(1)* contains a multiple-of-four number of bytes; directives to *ld(1)* can create a section with an odd number of bytes.

**SEE ALSO**

*as(1)*, *cc(1)*, *ld(1)*, *aouthdr(4)*, *filehdr(4)*, *ldfcn(4)*, *linenum(4)*, *reloc(4)*, *scnhdr(4)*, *syms(4)*.

**NAME**

a.out5.0 — assembler and link editor output

**SYNOPSIS**

#include &lt;a.out.h&gt;

**DESCRIPTION**

*A.out5.0* is the output file of the assembler *as5.0(1)* and the link loader *ld5.0(1)*. *Ld5.0(1)* makes *a.out5.0* executable if there were no errors and no unresolved external references. Layout information as given in the include file for the 68000 is:

```

/*
 * Layout of a.out file:
 *
 * header of 8 longs   magic number 405, 407, 410, 411
 *                    text size
 *                    data size
 *                    bss size
 *                    symbol table size
 *                    text relocation size
 *                    data relocation size
 *                    entry point
 *
 * header:           0
 * text:             32
 * data:             32+textsize
 * symbol table:     32+textsize+datasize
 * text relocation:  32+textsize+datasize+symsize
 * data relocation:  32+textsize+datasize+symsize+rtxtsize
 *
 */

/* various parameters */
#define SYMLENGTH 50 /* maximum length of a symbol */

/* types of files */
#define ARCMAGIC 0177545 /* ar files */
#define FMAGIC 0407 /* standard executable */
#define NMAGIC 0410 /* shared text executable */

/* symbol types */
#define EXTERN 040 /* external */
#define UNDEF 00 /* undefined */
#define ABS 01 /* absolute */
#define TEXT 02 /* text */
#define DATA 03 /* data */
#define BSS 04 /* bss */
#define COMM 05 /* internal use only */
#define REG 06 /* register name */

/* relocation regions */
#define RTEXT 00
#define RDATA 01
#define RBSS 02
#define REXT 03

```

```

/* relocation sizes */
#define RBYTE      00
#define RWORD     01
#define RLONG     02

/* macros which define various positions in file based on a bhdr, filhdr */
#define TEXTPOS    ((long) sizeof(filhdr))
#define DATAPOS  (TEXTPOS + filhdr.tsize)
#define SYMPOS    (DATAPOS + filhdr.dsize)
#define RTEXTPOS  (SYMPOS + filhdr.ssize)
#define RDATAPOS  (RTEXTPOS + filhdr.rtsize)
#define ENDPOS    (RDATAPOS + filhdr.rdsiZe)

/* header of a.out files */
struct bhdr {
    long    fmagic;
    long    tsize;
    long    dsize;
    long    bsize;
    long    ssize;
    long    rtsize;
    long    rdsiZe;
    long    entry;
};

/* symbol management */
struct sym {
    char    stype;        /* symbol type */
    char    sympad;      /* pad to short align */
    long    sval;        /* value */
};

/* relocation commands */
struct reloc {
    unsigned rsegment:2; /* RTEXT, RDATA, RBSS, or REXTERN */
    unsigned rsize:2;    /* RBYTE, RWORD, or RLONG */
    unsigned rdisp:1;    /* 1 => a displacement */
    unsigned relpad1:3;  /* pad 1 */
    char     relpad2;    /* pad 2 */
    short    rsymbol;    /* id of the symbol of external relocations */
    long     rpos;       /* position of relocation in segment */
};

/* symbol table entry */
struct nlist {
    char     n_name[8];  /* symbol name */
    int      n_type;     /* type flag */
    unsigned n_value;    /* value */
};

```

```

/* values for type flag */
#define N_UNDF 0 /* undefined */
#define N_ABS 01 /* absolute */
#define N_TEXT 02 /* text symbol */
#define N_DATA 03 /* data symbol */
#define N_BSS 04 /* bss symbol */
#define N_TYPE 037
#define N_REG 024 /* register name */
#define N_FN 037 /* file name symbol */
#define N_EXT 040 /* external bit, or'ed in */
#define FORMAT "%06o" /* to print a value */

```

The file has four sections: a header, the program and data text, a symbol table, and relocation information. The last two may be empty if the program was loaded with the `-s` option of `ld5.0` or if the symbols and relocation have been removed by `strip(1)`.

In the header the sizes of each section are given in bytes, but are even. The size of the header is not included in any of the other sizes.

When an `a.out5.0` file is loaded into core for execution, three logical segments are set up: the text segment, the data segment (with uninitialized data, which starts off as all 0, following initialized data), and a stack. The text segment begins at the user program start address in the core image; the header is not loaded. If the magic number in the header is `FMAGIC`, it indicates that the text segment is not to be write-protected and shared, so the data segment is immediately contiguous with the text segment. If the magic number is `NMAGIC`, the data segment begins at the next segment boundary following the text segment, and the text segment is not writable by the program; if other processes are executing the same file, they will share the text segment.

The stack will occupy the highest possible user program locations in the core image and will grow downwards. The stack is automatically extended as required. The data segment is only extended as requested by `brk(2)`.

The start of the text segment in the file is `32(10)`; the start of the data segment is `32+St` (the size of the text) the start of the relocation information is `32+St+Sd`; the start of the symbol table is `32+2(St+Sd)` if the relocation information is present, `32+St+Sd` if not.

The layout of a symbol table entry and the principal flag values that distinguish symbol types are given in the include file.

If a symbol's type is undefined external, and the value field is non-zero, the symbol is interpreted by the loader `ld` as the name of a common region whose size is indicated by the value of the symbol.

The value of a word in the text or data portions which is not a reference to an undefined external symbol is exactly that value which will appear in core when the file is executed. If a word in the text or data portion involves a reference to an undefined external symbol, as indicated by the relocation information for that word, then the value of the word as stored in the file is an offset from the associated external symbol. When the file is processed by the link editor and the external symbol becomes defined, the value of the symbol will be added into the word in the file.

If relocation information is present, it will appear in the form of the structure shown above.

**SEE ALSO**

as5.0(1), ld5.0(1), nm5.0(1)

**NAME**

aouthdr.h - a.out header for common object files

**DESCRIPTION**

Optional a.out header for common object files. The C structure appears below.

```

/*
 * static char ID_aouth[] = "@(#)aouthdr.h 2.1 ";
 */

typedef struct aouthdr {
    short magic; /* see magic.h */
    short vstamp; /* version stamp */
    long tsize; /* text size in bytes, padded to FW
                bdry */
    long dsize; /* initialized data " */
    long bsize; /* uninitialized data " */
#ifdef u3b
    long dum1;
    long dum2; /*Pad to entry point*/
#endif
    long entry; /* entry pt. */
    long text_start; /* base of text used for this file*/
    long data_start; /* base of data used for this file*/
} AOUTHDR;

```

**SEE ALSO**

a.out(4).

## NAME

ar — common archive file format

## DESCRIPTION

The archive command *ar* is used to combine several files into one. Archives are used mainly as libraries to be searched by the link editor *ld(1)*.

Each archive begins with the archive magic string.

```
#define ARMAG "!<arch>\n" /* magic string */
#define SARMAG 8 /* length of magic string */
```

Each archive which contains common object files (see *a.out(4)*) includes an archive symbol table. This symbol table is used by the link editor *ld(1)* to determine which archive members must be loaded during the link edit process. The archive symbol table (if it exists) is always the first file in the archive (but is never listed) and is automatically created and/or updated by *ar*.

Following the archive magic string are the archive file members. Each file member is preceded by a file member header which is of the following format:

```
#define ARFMAG "\n" /* header trailer string */

struct ar_hdr /* file member header */
{
    char ar_name[16]; /* '/' terminated file member name */
    char ar_date[12]; /* file member date */
    char ar_uid[6]; /* file member user identification */
    char ar_gid[6]; /* file member group identification */
    char ar_mode[8]; /* file member mode */
    char ar_size[10]; /* file member size */
    char ar_fm[2]; /* header trailer string */
};
```

All information in the file member headers is in printable ASCII. The numeric information contained in the headers is stored as decimal numbers (except for *ar\_mode* which is in octal). Thus, if the archive contains printable files, the archive itself is printable.

The *ar\_name* field is blank-padded and slash (/) terminated. The *ar\_date* field is the modification date of the file at the time of its insertion into the archive. Common format archives can be moved from system to system as long as the portable archive command *ar(1)* is used.

Each archive file member begins on an even byte boundary; a newline is inserted between files if necessary. Nevertheless, the size given reflects the actual size of the file exclusive of padding.

Notice there is no provision for empty areas in an archive file.

If the archive symbol table exists, the first file in the archive has a zero length name (i.e., *ar\_name[0] = '/'*). The contents of this file are as follows:

- The number of symbols. Length: 4 bytes.

- The array of offsets into the archive file. Length: 4 bytes \* "the number of symbols".
- The name string table. Length: *ar\_size* - (4 bytes \* ("the number of symbols" + 1)). The number of symbols and the array of offsets are managed with *sgetl* and *sputl*. The string table contains exactly as many null terminated strings as there are elements in the offsets array. Each offset from the array is associated with the corresponding name from the string table (in order). The names in the string table are all the defined global symbols found in the common object files in the archive. Each offset is the location of the archive header for the associated symbol.

**SEE ALSO**

*ar*(1), *ld*(1), *strip*(1), *sputl*(3X), *a.out*(4).

**WARNINGS**

*Strip*(1) will remove all archive symbol entries from the header. The archive symbol entries must be restored via the *s* option of the *ar*(1) command before the archive can be used with the link editor *ld*(1).



**NAME**

ar5.0 — archive (library) file format

**SYNOPSIS**

#include &lt;ar.h&gt;

**DESCRIPTION**

The archive command *ar5.0* is used to combine several files into one. Archives are used mainly as libraries to be searched by the link-editor *ld5.0*.

A file produced by *ar5.0* has a magic number at the start, followed by the constituent files, each preceded by a file header. The magic number and header layout as described in the include file are:

```
#define ARFMAG 0177545

struct ar_hdr {
    char    ar_name[14];
    long    ar_date;
    short   ar_uid;
    short   ar_gid;
    short   ar_mode;
    long    ar_size;
};
```

The "ar\_fm<sub>ag</sub>" field contains the 32-bit number ARFMAG to help verify the presence of a header. The name is a blank padded string. The other fields are left-adjusted, blank-padded numbers. They are decimal except for "ar\_mode", which is octal. The date is the modification date of the file at the time of its insertion into the archive.

Each file begins on an even (0 mod 2) boundary; a new-line is inserted between files if necessary. Nevertheless the size given reflects the actual size of the file exclusive of padding.

There is no provision for empty areas in an archive file.

**SEE ALSO**

ar5.0(1), ld5.0(1), nm5.0(1)

**BUGS**

File names lose trailing blanks. Most software dealing with archives takes even an included blank as a name terminator.

**NAME**

checklist — list of file systems processed by fsck

**DESCRIPTION**

*Checklist* resides in directory */etc* and contains a list of at most 15 *special* filenames. Each *special* filename is contained on a separate line and corresponds to a file system. If no *file-system* argument is provided to *fsck(1M)*, each file listed in */etc/checklist* is automatically read and checked for inconsistencies.

**SEE ALSO**

*fsck(1M)*.

**NAME**

core — format of core image file

**DESCRIPTION**

The UNIX System writes out a core image of a terminated process when any of various errors occur. See *signal(2)* for the list of reasons; the most common are memory violations, illegal instructions, bus errors, and user-generated quit signals. The core image is called *core* and is written in the process's working directory (provided it can be; normal access controls apply). A process with an effective user ID different from the real user ID will not produce a core image.

The first section of the core image is a copy of the system's per-user data for the process, including the registers as they were at the time of the fault. The size of this section depends on the parameter *USIZE*, which is defined in */usr/include/sys/param.h*. The remainder represents the actual contents of the user's core area when the core image was written. If the text segment is read-only and shared, or separated from data space, it is not dumped.

The format of the information in the first section is described by the *user* structure of the system, defined in */usr/include/sys/user.h*. The important stuff not detailed therein is the locations of the registers, which are outlined in */usr/include/sys/reg.h*.

**SEE ALSO**

*setuid(2)*, *signal(2)*.

## NAME

cpio — format of cpio archive

## DESCRIPTION

The *header* structure, when the *-c* option of *cpio(1)* is not used, is:

```
struct {
    short   h_magic,
           h_dev;
    ushort  h_ino,
           h_mode,
           h_uid,
           h_gid;
    short   h_nlink,
           h_rdev,
           h_mtime[2],
           h_namesize,
           h_filesiz[2];
    char    h_name[h_namesize rounded to word];
} Hdr;
```

When the *-c* option is used, the *header* information is described by:

```
sscanf(Chdr,"%6o%6o%6o%6o%6o%6o%6o%6o%11lo%6o%11lo%s",
        &Hdr.h_magic, &Hdr.h_dev, &Hdr.h_ino, &Hdr.h_mode,
        &Hdr.h_uid, &Hdr.h_gid, &Hdr.h_nlink, &Hdr.h_rdev,
        &Longtime, &Hdr.h_namesize,&Longfile,Hdr.h_name);
```

*Longtime* and *Longfile* are equivalent to *Hdr.h\_mtime* and *Hdr.h\_filesiz*, respectively. The contents of each file are recorded in an element of the array of varying length structures, *archive*, together with other items describing the file. Every instance of *h\_magic* contains the constant 070707 (octal). The items *h\_dev* through *h\_mtime* have meanings explained in *stat(2)*. The length of the null-terminated path name *h\_name*, including the null byte, is given by *h\_namesize*.

The last record of the *archive* always contains the name TRAILER!!!. Special files, directories, and the trailer are recorded with *h\_filesiz* equal to zero.

## SEE ALSO

*cpio(1)*, *find(1)*, *stat(2)*.

**NAME**

dir — format of directories

**SYNOPSIS**

```
#include <sys/dir.h>
```

**DESCRIPTION**

A directory behaves exactly like an ordinary file, save that no user may write into a directory. The fact that a file is a directory is indicated by a bit in the flag word of its i-node entry (see *fs(4)*). The structure of a directory entry as given in the include file is:

```
#ifndef DIRSIZ
#define DIRSIZ 14
#endif
struct direct {
    ino_t    d_ino;
    char    d_name[DIRSIZ];
};
```

By convention, the first two entries in each directory are for . and .. The first is an entry for the directory itself. The second is for the parent directory. The meaning of .. is modified for the root directory of the master file system; there is no parent, so .. has the same meaning as ..

**SEE ALSO**

*fs(4)*.

**NAME**

errfile — error-log file format

**DESCRIPTION**

When hardware errors are detected by the system, an error record is generated and passed to the error-logging daemon for recording in the error log for later analysis. The default error log is `/usr/adm/errfile`.

The format of an error record depends on the type of error that was encountered. Every record, however, has a header with the following format:

```
struct errhdr {
    short      e_type;          /* record type */
    short      e_len;          /* bytes in record (inc hdr) */
    time_t     e_time;        /* time of day */
};
```

The permissible record types are as follows:

```
#define E_GOTS          010    /* start for the UNIX/TS*/
#define E_GORT          011    /* start for the UNIX/RT*/
#define E_STOP          012    /* stop */
#define E_TCRG          013    /* time change */
#define E_CCHG          014    /* configuration change */
#define E_BLK           020    /* block device error */
#define E_STRAY         030    /* stray interrupt */
#define E_PRTY          031    /* memory parity */
```

Some records in the error file are of an administrative nature. These include the startup record that is entered into the file when logging is activated, the stop record that is written if the daemon is terminated “gracefully”, and the time-change record that is used to account for changes in the system’s time-of-day. These records have the following formats:

```
struct estart {
    short      e_cpu;          /* CPU type */
    struct utsname e_name;     /* system names */
};
```

```
#define eend errhdr          /* record header */
```

```
struct etimchg {
    time_t     e_ntime;       /* new time */
};
```

Stray interrupts cause a record with the following format to be logged:

```
struct estray {
    uint       e_saddr;       /* stray loc or device addr */
};
```

Generation of memory subsystem errors is not supported in this release.

Error records for block devices have the following format:

**NAME**

dir — format of directories

**SYNOPSIS**

```
#include <sys/dir.h>
```

**DESCRIPTION**

A directory behaves exactly like an ordinary file, save that no user may write into a directory. The fact that a file is a directory is indicated by a bit in the flag word of its i-node entry (see *fs(4)*). The structure of a directory entry as given in the include file is:

```
#ifndef DIRSIZ
#define DIRSIZ 14
#endif
struct direct {
    ino_t    d_ino;
    char    d_name[DIRSIZ];
};
```

By convention, the first two entries in each directory are for . and .. The first is an entry for the directory itself. The second is for the parent directory. The meaning of .. is modified for the root directory of the master file system; there is no parent, so .. has the same meaning as ..

**SEE ALSO**

*fs(4)*.

**NAME**

errfile — error-log file format

**DESCRIPTION**

When hardware errors are detected by the system, an error record is generated and passed to the error-logging daemon for recording in the error log for later analysis. The default error log is `/usr/adm/errfile`.

The format of an error record depends on the type of error that was encountered. Every record, however, has a header with the following format:

```
struct errhdr {
    short      e_type;          /* record type */
    short      e_len;          /* bytes in record (inc hdr) */
    time_t     e_time;         /* time of day */
};
```

The permissible record types are as follows:

```
#define E_GOTS          010    /* start for the UNIX/TS */
#define E_GORT          011    /* start for the UNIX/RT */
#define E_STOP          012    /* stop */
#define E_TCHG          013    /* time change */
#define E_CCHG          014    /* configuration change */
#define E_BLK           020    /* block device error */
#define E_STRAY         030    /* stray interrupt */
#define E_PRTY         031    /* memory parity */
```

Some records in the error file are of an administrative nature. These include the startup record that is entered into the file when logging is activated, the stop record that is written if the daemon is terminated “gracefully”, and the time-change record that is used to account for changes in the system’s time-of-day. These records have the following formats:

```
struct estart {
    short      e_cpu;          /* CPU type */
    struct utmsname e_name;    /* system names */
};
```

```
#define eend errhdr          /* record header */
```

```
struct etimchg {
    time_t     e_nptime;      /* new time */
};
```

Stray interrupts cause a record with the following format to be logged:

```
struct estray {
    uint       e_saddr;        /* stray loc or device addr */
};
```

Generation of memory subsystem errors is not supported in this release.

Error records for block devices have the following format:



```

struct eblock {
    dev_t e_dev;          /* "true" major + minor dev no */
    physadr e_regloc;     /* controller address */
    short e_bacty;       /* other block I/O activity */
    struct iostat {
        long io_ops;     /* number read/writes */
        long io_misc;    /* number "other" operations */
        ushort io_unlog; /* number unlogged errors */
    } e_stats;
    short e_bflags;      /* read/write, error, etc */
    short e_cylloff;     /* logical dev start cyl */
    daddr_t e_bnum;      /* logical block number */
    ushort e_bytes;      /* number bytes to transfer */
    paddr_t e_memadd;    /* buffer memory address */
    ushort e_rtry;       /* number retries */
    short e_nreg;        /* number device registers */
};

```

The following values are used in the *e\_bflags* word:

```

#define E_WRITE      0      /* write operation */
#define E_READ       1      /* read operation */
#define E_NOD        02     /* no I/O pending */
#define E_PHYS       04     /* physical I/O */
#define E_FORMAT     010    /* Formatting Disk */
#define E_ERROR      020    /* I/O failed */

```

**SEE ALSO**  
errdemon(1M).

**NAME**

filehdr — file header for common object files

**SYNOPSIS**

```
#include <filehdr.h>
```

**DESCRIPTION**

Every common object file begins with a 20-byte header. The following C struct declaration is used:

```
struct filehdr
{
    unsigned short f_magic ; /* magic number */
    unsigned short f_nscns ; /* number of sections */
    long          f_tmdat ; /* time & date stamp */
    long          f_sympr ; /* file ptr to symtab */
    long          f_nsyms ; /* # symtab entries */
    unsigned short f_opthdr ; /* sizeof(opt hdr) */
    unsigned short f_flags ; /* flags */
};
```

*f\_sympr* is the byte offset into the file at which the symbol table can be found. Its value can be used as the offset in *fseek*(3S) to position an I/O stream to the symbol table. See *aouthdr*(4) for the structure of the optional *aout* header. The valid magic number is:

```
#define MC68MAGIC 0520 /* magic number */
```

The value in *f\_tmdat* is obtained from the *time*(2) system call. Flag bits currently defined are:

```
#define F_RELFLG 00001 /* relocation entries stripped */
#define F_EXEC 00002 /* file is executable */
#define F_LNNO 00004 /* line numbers stripped */
#define F_LSYMS 00010 /* local symbols stripped */
#define F_MINMAL 00020 /* minimal object file */
#define F_UPDATE 00040 /* update file, ogen produced */
#define F_SWABD 00100 /* file is "pre-swabbed" */
#define F_AR16WR 00200 /* 16-bit DEC host */
#define F_AR32WR 00400 /* 32-bit DEC host */
#define F_AR32W 01000 /* non-DEC host */
#define F_PATCH 02000 /* "patch" list in opt hdr */
```

**SEE ALSO**

*time*(2), *fseek*(3S), *a.out*(4), *aouthdr*(4).

## NAME

file system — format of system volume

## SYNOPSIS

```
#include <sys/filsys.h>
#include <sys/types.h>
#include <sys/param.h>
```

## DESCRIPTION

Every file system storage volume has a common format for certain vital information. Every such volume is divided into a certain number of 512-byte long sectors. Sector 0 is unused and is available to contain a bootstrap program or other information.

Sector 1 is the *superblock*. The format of a superblock is:

```
/*
 * Structure of the superblock
 */
struct filsys
{
    ushort    s_size;           /* size in blocks of i-list */
    daddr_t   s_fsize;         /* size in blocks of entire volume */
    short     s_nfree;         /* number of addresses in s_free */
    daddr_t   s_free[NICFREE]; /* free block list */
    short     s_ninode;        /* number of inodes in s_inode */
    ino_t     s_inode[NICINOD]; /* free inode list */
    char      s_lock;          /* lock during free list manipulation */
    char      s_i_lock;        /* lock during i-list manipulation */
    char      s_fmod;          /* superblock modified flag */
    char      s_ronly;         /* mounted read-only flag */
    time_t    s_time;          /* last superblock update */
    short     s_dinfo[4];      /* device information */
    daddr_t   s_ifree;         /* total free blocks */
    ino_t     s_tinode;        /* total free inodes */
    char      s_fname[6];      /* file system name */
    char      s_fpack[6];      /* file system pack name */
    long      s_fih[14];       /* ADJUST size of filsys to 512 */
    ino_t     s_lasti;         /* start place for circular search */
    ino_t     s_nbehind;       /* est # free inodes before s_lasti */
    long      s_magic;         /* magic number to indicate new file system */
    long      s_type;          /* type of new file system */
};

#define Fsmagic 0xfd187e20 /* s_magic number */
#define Fsb 1 /* 512-byte block */
#define Fs2b 2 /* 1024-byte block */
#define Fs4b 4 /* 2048-byte block */
```

*S\_type* indicates the file system type. Currently, two types of file systems are supported: the original 512-byte oriented and the new improved 1024-byte oriented. *S\_magic* is used to distinguish the original 512-byte oriented file systems from the newer file systems. If this field is not equal to the magic number, *Fsmagic*, the type is assumed to be *Fsb*, otherwise the *s\_type* field is used. In the following description, a block is then determined by the type. For the original 512-byte oriented file system, a block is 512

bytes. For the 1024-byte oriented file system, a block is 1024 bytes or two sectors. The operating system takes care of all conversions from logical block numbers to physical sector numbers.

$S\_isize$  is the address of the first data block after the i-list; the i-list starts just after the super-block, namely in block 2; thus the i-list is  $s\_isize-2$  blocks long.  $S\_fsize$  is the first block not potentially available for allocation to a file. These numbers are used by the system to check for bad block numbers; if an "impossible" block number is allocated from the free list or is freed, a diagnostic is written on the on-line console. Moreover, the free array is cleared, so as to prevent further allocation from a presumably corrupted free list.

The free list for each volume is maintained as follows. The  $s\_free$  array contains, in  $s\_free[1], \dots, s\_free[s\_nfree-1]$ , up to 49 numbers of free blocks.  $S\_free[0]$  is the block number of the head of a chain of blocks constituting the free list. The first long in each free-chain block is the number (up to 50) of free-block numbers listed in the next 50 longs of this chain member. The first of these 50 blocks is the link to the next member of the chain. To allocate a block: decrement  $s\_nfree$ , and the new block is  $s\_free[s\_nfree]$ . If the new block number is 0, there are no blocks left, so give an error. If  $s\_nfree$  became 0, read in the block named by the new block number, replace  $s\_nfree$  by its first word, and copy the block numbers in the next 50 longs into the  $s\_free$  array. To free a block, check if  $s\_nfree$  is 50; if so, copy  $s\_nfree$  and the  $s\_free$  array into it, write it out, and set  $s\_nfree$  to 0. In any event set  $s\_free[s\_nfree]$  to the freed block's number and increment  $s\_nfree$ .

$S\_tfree$  is the total free blocks available in the file system.

$S\_ninode$  is the number of free i-numbers in the  $s\_inode$  array. To allocate an inode: if  $s\_ninode$  is greater than 0, decrement it and return  $s\_inode[s\_ninode]$ . If it was 0, read the i-list and place the numbers of all free inodes (up to 100) into the  $s\_inode$  array, then try again. To free an inode, provided  $s\_ninode$  is less than 100, place its number into  $s\_inode[s\_ninode]$  and increment  $s\_ninode$ . If  $s\_ninode$  is already 100, do not bother to enter the freed inode into any table. This list of inodes is only to speed up the allocation process; the information as to whether the inode is really free or not is maintained in the inode itself.

$S\_tinode$  is the total free inodes available in the file system.

$S\_flock$  and  $s\_ilock$  are flags maintained in the core copy of the file system while it is mounted and their values on disk are immaterial. The value of  $s\_fmod$  on disk is likewise immaterial; it is used as a flag to indicate that the super-block has changed and should be copied to the disk during the next periodic update of file system information.

$S\_ronly$  is a read-only flag to indicate write-protection.

$S\_time$  is the last time the super-block of the file system was changed, and is the number of seconds that have elapsed since 00:00 Jan. 1, 1970 (GMT). During a reboot, the  $s\_time$  of the super-block for the root file system is used to set the system's idea of the time.

$S\_fname$  is the name of the file system and  $s\_fpack$  is the name of the pack.

I-numbers begin at 1, and the storage for inodes begins in block 2. Also, inodes are 64 bytes long. Inode 1 is reserved for future use. Inode 2 is

reserved for the root directory of the file system, but no other i-number has a built-in meaning. Each inode represents one file. For the format of an inode and its flags, see *inode(4)*.

**FILES**

`/usr/include/sys/files.h`  
`/usr/include/sys/stat.h`

**SEE ALSO**

*fsck(1M)*, *fsdb(1M)*, *mkfs(1M)*, *inode(4)*.

## NAME

fspec — format specification in text files

## DESCRIPTION

It is sometimes convenient to maintain text files on the UNIX System with non-standard tabs, (i.e., tabs which are not set at every eighth column). Such files must generally be converted to a standard format, frequently by replacing all tabs with the appropriate number of spaces, before they can be processed by UNIX System commands. A format specification occurring in the first line of a text file specifies how tabs are to be expanded in the remainder of the file.

A format specification consists of a sequence of parameters separated by blanks and surrounded by the brackets <: and :>. Each parameter consists of a keyletter, possibly followed immediately by a value. The following parameters are recognized:

**t***tabs* The **t** parameter specifies the tab settings for the file. The value of *tabs* must be one of the following:

1. a list of column numbers separated by commas, indicating tabs set at the specified columns;
2. a — followed immediately by an integer *n*, indicating tabs at intervals of *n* columns;
3. a — followed by the name of a “canned” tab specification.

Standard tabs are specified by **t-8**, or equivalently, **t1,9,17,25**, etc. The canned tabs which are recognized are defined by the *tabs(1)* command.

**s***size* The **s** parameter specifies a maximum line size. The value of *size* must be an integer. Size checking is performed after tabs have been expanded, but before the margin is prepended.

**m***margin* The **m** parameter specifies a number of spaces to be prepended to each line. The value of *margin* must be an integer.

**d** The **d** parameter takes no value. Its presence indicates that the line containing the format specification is to be deleted from the converted file.

**e** The **e** parameter takes no value. Its presence indicates that the current format is to prevail only until another format specification is encountered in the file.

Default values, which are assumed for parameters not supplied, are **t-8** and **m0**. If the **s** parameter is not specified, no size checking is performed. If the first line of a file does not contain a format specification, the above defaults are assumed for the entire file. The following is an example of a line containing a format specification:

```
* <:t5,10,15 s72:> *
```

If a format specification can be disguised as a comment, it is not necessary to code the **d** parameter.

Several UNIX System commands correctly interpret the format specification for a file. Among them is *gath* which may be used to convert files to a standard format acceptable to other UNIX System commands.

## SEE ALSO

*ed(1)*, *newform(1)*, *tabs(1)*.

## NAME

gettydefs — speed and terminal settings used by getty

## DESCRIPTION

The `/etc/gettydefs` file contains information used by `getty(1M)` to set up the speed and terminal settings for a line. It supplies information on what the `login` prompt should look like. It also supplies the speed to try next if the user indicates the current speed is not correct by typing a `<break>` character.

Each entry in `/etc/gettydefs` has the following format:

```
label# initial-flags # final-flags # login-prompt #next-label
```

Each entry is followed by a blank line. The various fields can contain quoted characters of the form `\b`, `\n`, `\c`, etc., as well as `\nnn`, where `nnn` is the octal value of the desired character. The various fields are:

- label* This is the string against which `getty` tries to match its second argument. It is often the speed, such as `1200`, at which the terminal is supposed to run, but it need not be (see below).
- initial-flags* These flags are the initial `ioctl(2)` settings to which the terminal is to be set if a terminal type is not specified to `getty`. The flags that `getty` understands are the same as the ones listed in `/usr/include/sys/termio.h` (see `termio(7)`). Normally only the speed flag is required in the *initial-flags*. `Getty` automatically sets the terminal to raw input mode and takes care of most of the other flags. The *initial-flag* settings remain in effect until `getty` executes `login(1)`.
- final-flags* These flags take the same values as the *initial-flags* and are set just prior to `getty` executing `login`. The speed flag is again required. The composite flags `SANE` or `SANE2` take care of most of the other flags that need to be set so that the processor and terminal are communicating in a rational fashion. The other two commonly specified *final-flags* are `TAB3`, so that tabs are sent to the terminal as spaces, and `HUPCL`, so that the line is hung up on the final close. Flag attributes are added from left to right, flags that start with a `^` are subtracted, e.g., `SANE ^PARENB`. `SANE` is defined to be `BRKINT IGNPAR ISTRIP ICRNL IXON OPOST ONLCR CS7 PARENB CREAD ISIG ICANON ECHO ECHOK`. `SANE2` is the same as `SANE` but with eight bits and no parity, e.g., `SANE2 = SANE ^CS7 CS8 ^PARENB`.
- login-prompt* This entire field is printed as the *login-prompt*. Unlike the above fields where white space is ignored (a space, tab or new-line), they are included in the *login-prompt* field.
- next-label* If this entry does not specify the desired speed, indicated by the user typing a `<break>` character, then `getty` will search for the entry with *next-label* as its *label* field and set up the terminal for those settings. Usually, a series of speeds are linked together in this fashion, into a closed set; For instance, `2400` linked to `1200`, which in turn is linked to `300`, which finally is linked to `2400`.

If *getty* is called without a second argument, then the first entry of */etc/gettydefs* is used, thus making the first entry of */etc/gettydefs* the default entry. It is also used if *getty* can not find the specified *label*. If */etc/gettydefs* itself is missing, there is one entry built into the command which will bring up a terminal at 300 baud.

It is strongly recommended that after making or modifying */etc/gettydefs*, it be run through *getty* with the *check* option to be sure there are no errors.

The following four symbols define the state.

```
# define ISANE (BRKINT|IGNPAR|ISTRIP|ICRNL|IXON)
# define OSANE (OPOST|ONLCR)
# define CSANE (CS7|PARENB|CREAD)
# define LSANE (ISIG|ICANON|ECHO|ECHOK)
```

**FILES**

*/etc/gettydefs*

**SEE ALSO**

*login(1)*, *ioctl(2)*.

*getty(1M)*, *termio(7)* in the *Administrator Reference Manual*.



**NAME**

group - group file

**DESCRIPTION**

*Group* contains for each group the following information:

- group name
- encrypted password
- numerical group ID
- comma-separated list of all user allowed in the group

This is an ASCII file. The fields are separated by colons; each group is separated from the next by a new-line. If the password field is null, no password is demanded.

This file resides in directory */etc*. Because of the encrypted passwords, it can and does have general read permission and can be used, for example, to map numerical group ID's to names.

**FILES**

*/etc/group*

**SEE ALSO**

*newgrp(1)*, *passwd(1)*, *crypt(3C)*, *passwd(4)*.

**NAME**

hosts - host name data base

**DESCRIPTION**

The *hosts* file contains information regarding the known hosts on the DARPA Internet. For each host a single line should be present with the following information:

official host name  
Internet address  
aliases

Items are separated by any number of blanks and/or tab characters. A “#” indicates the beginning of a comment; characters up to the end of the line are not interpreted by routines which search the file. This file is normally created from the official host data base maintained at the Network Information Control Center (NIC), though local changes may be required to bring it up to date regarding unofficial aliases and/or unknown hosts.

Network addresses are specified in the conventional “.” notation using the *inet\_addr()* routine from the Internet address manipulation library, *inet(3N)*. Host names may contain any printable character other than a field delimiter, newline, or comment character.

**FILES**

/etc/hosts

**SEE ALSO**

gethostent(3N)

**BUGS**

A name server should be used instead of a static file. A binary indexed file format should be available for fast access.

## NAME

inittab — script for the init process

## DESCRIPTION

The *inittab* file supplies the script to *init*'s role as a general process dispatcher. The process that constitutes the majority of *init*'s process dispatching activities is the line process */etc/getty* that initiates individual terminal lines. Other processes typically dispatched by *init* are daemons and the shell.

The *inittab* file is composed of entries that are position dependent and have the following format:

```
id:rstate:action:process
```

Each entry is delimited by a newline, however, a backslash (\) preceding a newline indicates a continuation of the entry. Up to 512 characters per entry are permitted. Comments may be inserted in the *process* field using the *sh*(1) convention for comments. Comments for lines that spawn *gettys* are displayed by the *who*(1) command. It is expected that they will contain some information about the line such as the location. There are no limits (other than maximum entry size) imposed on the number of entries within the *inittab* file. The entry fields are:

- id* This is one or two characters used to uniquely identify an entry.
- rstate* This defines the *run-level* in which this entry is to be processed. *Run-levels* effectively correspond to a configuration of processes in the system. That is, each process spawned by *init* is assigned a *run-level* or *run-levels* in which it is allowed to exist. The *run-levels* are represented by a number ranging from 0 through 6. As an example, if the system is in *run-level* 1, only those entries having a 1 in the *rstate* field will be processed. When *init* is requested to change *run-levels*, all processes which do not have an entry in the *rstate* field for the target *run-level* will be sent the warning signal (SIGTERM) and allowed a 20-second grace period before being forcibly terminated by a kill signal (SIGKILL). The *rstate* field can define multiple *run-levels* for a process by selecting more than one *run-level* in any combination from 0-6. If no *run-level* is specified, then the process is assumed to be valid at all *run-levels* 0-6. There are three other values, a, b and c, which can appear in the *rstate* field, even though they are not true *run-levels*. Entries which have these characters in the *rstate* field are processed only when the *telinit* (see *init*(IM)) process requests them to be run (regardless of the current *run-level* of the system). They differ from *run-levels* in that *init* can never enter *run-level* a, b or c. Also, a request for the execution of any of these processes does not change the current *run-level*. Furthermore, a process started by an a, b or c command is not killed when *init* changes levels. They are only killed if their line in */etc/inittab* is marked off in the *action* field, their line is deleted entirely from */etc/inittab*, or *init* goes into the *SINGLE USER* state.
- action* Key words in this field tell *init* how to treat the process specified in the *process* field. The actions recognized by *init* are as follows:
- respawn** If the process does not exist then start the process, do not wait for its termination (continue scanning

- the *inittab* file), and when it dies restart the process. If the process currently exists then do nothing and continue scanning the *inittab* file.
- wait** Upon *init*'s entering the *run-level* that matches the entry's *rstate*, start the process and wait for its termination. All subsequent reads of the *inittab* file while *init* is in the same *run-level* will cause *init* to ignore this entry.
- once** Upon *init*'s entering a *run-level* that matches the entry's *rstate*, start the process, do not wait for its termination. When it dies, do not restart the process. If upon entering a new *run-level*, where the process is still running from a previous *run-level* change, the program will not be restarted.
- boot** The entry is to be processed only at *init*'s boot-time read of the *inittab* file. *Init* is to start the process, not wait for its termination; and when it dies, not restart the process. In order for this instruction to be meaningful, the *rstate* should be the default or it must match *init*'s *run-level* at boot time. This action is useful for an initialization function following a hardware reboot of the system.
- bootwait** The entry is to be processed only at *init*'s boot-time read of the *inittab* file. *Init* is to start the process, wait for its termination and, when it dies, not restart the process.
- powerfail** Execute the process associated with this entry only when *init* receives a power fail signal (SIGPWR see *signal(2)*).
- powerwait** Execute the process associated with this entry only when *init* receives a power fail signal (SIGPWR) and wait until it terminates before continuing any processing of *inittab*.
- off** If the process associated with this entry is currently running, send the warning signal (SIGTERM) and wait 20 seconds before forcibly terminating the process via the kill signal (SIGKILL). If the process is nonexistent, ignore the entry.
- ondemand** This instruction is really a synonym for the **respawn** action. It is functionally identical to **respawn** but is given a different keyword in order to divorce its association with *run-levels*. This is used only with the a, b or c values described in the *rstate* field.
- initdefault** An entry with this *action* is only scanned when *init* initially invoked. *Init* uses this entry, if it exists, to determine which *run-level* to enter initially. It does this by taking the highest *run-level* specified in the *rstate* field and using that as its initial state. If the *rstate* field is empty, this is interpreted as 0123456 and so *init* will enter *run-level* 6. Also, the

**initdefault** entry cannot specify that *init* start in the *SINGLE USER* state. Additionally, if *init* does not find an **initdefault** entry in */etc/inittab*, then it will request an initial *run-level* from the user at reboot time.

**sysinit** Entries of this type are executed before *init* tries to access the console. It is expected that this entry will be only used to initialize devices on which *init* might try to ask the *run-level* question. These entries are executed and waited for before continuing.

**process** This is a *sh* command to be executed. The entire **process** field is prefixed with *exec* and passed to a forked *sh* as *sh -c 'exec command'*. For this reason, any legal *sh* syntax can appear in the **process** field. Comments can be inserted with the ; *#comment* syntax.

**FILES**

*/etc/inittab*

**SEE ALSO**

*sh*(1), *who*(1), *exec*(2), *open*(2), *signal*(2),  
*getty*(1M), *init*(1M) in the *Administrator Manual*.

**NAME**

inode — format of an inode

**SYNOPSIS**

```
#include <sys/types.h>
#include <sys/ino.h>
```

**DESCRIPTION**

An i-node for a plain file or directory in a file system has the following structure defined by `<sys/ino.h>`.

```
/* Inode structure as it appears on a disk block. */
struct dinode {
    ushort    di_mode;      /* mode and type of file */
    short     di_nlink;    /* number of links to file */
    ushort    di_uid;      /* owner's user id */
    ushort    di_gid;      /* owner's group id */
    off_t     di_size;     /* number of bytes in file */
    char      di_addr[40]; /* disk block addresses */
    time_t    di_atime;    /* time last accessed */
    time_t    di_mtime;    /* time last modified */
    time_t    di_ctime;    /* time created */
};
/*
 * the 40 address bytes:
 *   39 used; 13 addresses
 *   of 3 bytes each.
 */
```

For the meaning of the defined types `off_t` and `time_t` see `types(5)`.

**FILES**

`/usr/include/sys/ino.h`

**SEE ALSO**

`stat(2)`, `fs(4)`, `types(5)`.

**NAME**

issue — issue identification file

**DESCRIPTION**

The file `/etc/issue` contains the *issue* or project identification to be printed as a login prompt. This is an ASCII file which is read by program *getty* and then written to any terminal spawned or respawnd from the *lines* file.

**FILES**

`/etc/issue`

**SEE ALSO**

`login(1)`.

**NAME**

ldfcn — common object file access routines

**SYNOPSIS**

```
#include <stdio.h>
#include <filehdr.h>
#include <ldfcn.h>
```

**DESCRIPTION**

The common object file access routines are a collection of functions for reading an object file that is in common object file form. Although the calling program must know the detailed structure of the parts of the object file that it processes, the routines effectively insulate the calling program from knowledge of the overall structure of the object file.

The interface between the calling program and the object file access routines is based on the defined type **LDFILE** (defined as **struct ldfile**), which is declared in the header file **<ldfcn.h>**. The primary purpose of this structure is to provide uniform access to both simple object files and object files that are members of an archive file.

The function *ldopen(3X)* allocates and initializes the **LDFILE** structure and returns a pointer to the structure to the calling program. The fields of the **LDFILE** structure may be accessed individually through macros defined in **<ldfcn.h>** and contain the following information:

**LDFILE \*ldptr;**

**TYPE(ldptr)** The file magic number, used to distinguish between archive members and simple object files.

**IOPTR(ldptr)** The file pointer returned by *fopen(3S)* and used by the standard input/output functions.

**OFFSET(ldptr)** The file address of the beginning of the object file; the offset is non-zero if the object file is a member of an archive file.

**HEADER(ldptr)** The file header structure of the object file.

The object file access functions may be divided into four categories:

- (1) functions that open or close an object file

```
ldopen(3X) and ldaopen
    open a common object file
ldclose(3X) and ldaclose
    close a common object file
```

- (2) functions that read header or symbol table information

```
ldahread(3X)
    read the archive header of a member of an archive
    file
ldfhread(3X)
    read the file header of a common object file
ldshread(3X) and ldnshread
    read a section header of a common object file
ldtbread(3X)
    read a symbol table entry of a common object file
```



*ldgetname(3X)*

retrieve a symbol name from a symbol table entry or from the string table

(3) functions that position an object file at (seek to) the start of the section, relocation, or line number information for a particular section.

*ldohseek(3X)*

seek to the optional file header of a common object file

*ldsseek(3X)* and *ldnsseek*

seek to a section of a common object file

*ldrseek(3X)* and *ldnrseek*

seek to the relocation information for a section of a common object file

*ldlseek(3X)* and *ldnlseek*

seek to the line number information for a section of a common object file

*ldtbseek(3X)*

seek to the symbol table of a common object file

(4) the function *ldtbindex(3X)* which returns the index of a particular common object file symbol table entry

These functions are described in detail in the manual pages identified for each function.

All the functions except *ldopen*, *ldgetname(3X)*, *ldaopen*, and *ldtbindex* return either SUCCESS or FAILURE, which are constants defined in <ldfcn.h>. *ldopen* and *ldaopen* both return pointers to a LDFILE structure.

**MACROS**

Additional access to an object file is provided through a set of macros defined in <ldfcn.h>. These macros parallel the standard input/output file reading and manipulating functions, translating a reference of the LDFILE structure into a reference to its file descriptor field.

The following macros are provided:

```
GETC(ldptr)
FGETC(ldptr)
GETW(ldptr)
UNGETC(c, ldptr)
FGETS(s, n, ldptr)
FREAD((char *) ptr, sizeof (*ptr), nitems, ldptr)
FSEEK(ldptr, offset, ptrname)
FTELL(ldptr)
REWIND(ldptr)
FEOF(ldptr)
FERROR(ldptr)
FILENO(ldptr)
SETBUF(ldptr, buf)
STROFFSET(ldptr)
```

The STROFFSET macro calculates the address of the string table in a object file. See the manual entries for the corresponding standard input/output library functions for details on the use of these macros. (The

functions are identified as 3S in Section 3 of this manual.)

The program must be loaded with the object file access routine library `libld.a`.

#### WARNINGS

The macro `FSEEK` defined in the header file `<ldfcn.h>` translates into a call to the standard input/output function `fseek(3S)`. `FSEEK` should not be used to seek from the end of an archive file since the end of an archive file may not be the same as the end of one of its object file members.

#### SEE ALSO

`fopen(3S)`, `fseek(3S)`, `ldahread(3X)`, `ldclose(3X)`, `ldhread(3X)`, `ldgetname(3X)`, `ldhread(3X)`, `ldlseek(3X)`, `ldohseek(3X)`, `ldopen(3X)`, `ldrseek(3X)`, `ldlseek(3X)`, `ldhread(3X)`, `ldtbindex(3X)`, `ldtbread(3X)`, `ldtbseek(3X)`.

`COFF` in the *Programming Guide*.

**NAME**

linenum — line number entries in a common object file

**SYNOPSIS**

```
#include <linenum.h>
```

**DESCRIPTION**

The C compiler generates an entry in the object file for each C source line on which a breakpoint is possible (when invoked with the `-g` option; see `cc(1)`). Users can then reference line numbers when using the appropriate software test system (see `sdb(1)`). The structure of these line number entries appears below.

```
struct lineno
{
    union
    {
        long    l_symndx ;
        long    l_paddr ;
    }
    unsigned short l_inno ;
};
```

Numbering starts with one for each function. The initial line number entry for a function has `l_inno` equal to zero, and the symbol table index of the function's entry is in `l_symndx`. Otherwise, `l_inno` is non-zero, and `l_paddr` is the physical address of the code for the referenced line. Thus the overall structure is the following:

<code>l_addr</code>	<code>l_inno</code>
function symtab index	0
physical address	line
physical address	line
...	
function symtab index	0
physical address	line
physical address	line
...	

**SEE ALSO**

`cc(1)`, `sdb(1)`, `a.out(4)`.

**NAME**

master — master device information table

**DESCRIPTION**

This file is used by the *config(1M)* program to obtain device information that enables it to generate the configuration files. The file consists of 3 parts, each separated by a line with a dollar sign (\$) in column 1. Part 1 contains device information; part 2 contains names of devices that have aliases; part 3 contains tunable parameter information. Any line with an asterisk (\*) in column 1 is treated as a comment.

Part 1 contains lines consisting of at least 7 fields and at most 10 fields, with the fields delimited by tabs and/or blanks:

Field 1:	device name (8 chars. maximum).
Field 2:	handlers (9 character string). 's', INIT, xxxinit() 'o', OPEN, xxxopen() 'c', CLOSE, xxxclose() 'r', READ, xxxread() 'w', WRITE, xxxwrite() 'i', IOCTL, xxxioctl() 't', SELECT, xxxselect() '-', No handlers
Field 3:	device type indicator (9 character string): 't', TTYS, a tty device 'o', ONCE, can only be specified once 's', NOCNT, suppress count & other stuff 'r', REQ, required device 'b', BLOCK, a block device 'c', CHAR, a character device 'k', CLOCK, the clock device 'p', PECULIAR, peculiar (use devname, not prefix) 'f', FORCE, '-', No type indicators
Field 4:	handler prefix (4 chars. maximum).
Field 5:	major device number for block-type device (short decimal).
Field 6:	major device number for character-type device (short decimal).
Field 7:	maximum number of devices/lines (short decimal).
Fields 8-10:	optional structure declarations (40 chars. maximum).

Part 2 contains lines with 2 fields each:

Field 1:	alias name of device (20 chars. maximum).
Field 2:	reference name of device (20 chars. maximum; specified in part 1).

Part 3 contains lines with 2 or 3 fields each:

Field 1:	parameter name (as it appears in description file; 20 chars. maximum)
Field 2:	text form (as it appears in the <i>conf.c</i> file; 20 chars. maximum)
Field 3:	default parameter value (20 chars. maximum; specification in description file is required if this field

**MASTER (4)**

**MASTER (4)**

is omitted)

**SEE ALSO**  
config(1M).

**FILES**  
/etc/master

## NAME

master – master device information table

## DESCRIPTION

This file is used by the *config(1M)* program to obtain device information that enables it to generate the configuration files. The file consists of 3 parts, each separated by a line with a dollar sign (\$) in column 1. Part 1 contains device information; part 2 contains names of devices that have aliases; part 3 contains tunable parameter information. Any line with an asterisk (\*) in column 1 is treated as a comment.

Part 1 contains lines consisting of at least 7 fields and at most 10 fields, with the fields delimited by tabs and/or blanks:

Field 1:	device name (8 chars. maximum).
Field 2:	handlers (9 character string)
	's', INIT, xxxinit()
	'o', OPEN, xxxopen()
	'c', CLOSE, xxxclose()
	'r', READ, xxxread()
	'w', WRITE, xxxwrite()
	'i', IOCTL, xxxioctl()
	't', SELECT, xxxselect()
	'-', No handlers
Field 3:	device type indicator (9 character string):
	'm', SEMAS, define semaphores
	't', TTYs, a tty device
	'o', ONCE, can only be specified once
	's', NOCNT, suppress count & other stuff
	'r', REQ, required device
	'b', BLOCK, a block device
	'c', CHAR, a character device
	'k', CLOCK, the clock device
	'p', PECULLAR, peculiar (use devname, not prefix)
	'f', FORCE, Define count if not a tty
	'-', No type indicators
Field 4:	handler prefix (4 chars. maximum).
Field 5:	major device number for block-type device (short decimal).
Field 6:	major device number for character-type device (short decimal).
Field 7:	maximum number of devices/lines (short decimal).
Fields 8-10:	optional structure declarations (40 chars. maximum).

Part 2 contains lines with 2 fields each:

Field 1:	alias name of device (20 chars. maximum).
Field 2:	reference name of device (20 chars. maximum; specified in part 1).

Part 3 contains lines with 2 or 3 fields each:

Field 1:	parameter name (as it appears in description file; 20 chars. maximum)
Field 2:	text form (as it appears in the <i>conf.c</i> file; 20 chars. maximum)
Field 3:	default parameter value (20 chars. maximum; specification in description file is required if this field is omitted)

## SEE ALSO

*config(1M)*.

**MASTER (4)**

*(Virtual)*

**MASTER (4)**

**FILES**

*/etc/master*

**NAME**

mnttab — mounted file system table

**SYNOPSIS**

```
#include <mnttab.h>
```

**DESCRIPTION**

*Mnttab* resides in directory */etc* and contains a table of devices, mounted by the *mount(1M)* command, in the following structure as defined by *<mnttab.h>*:

```
struct mnttab {
    char    mt_dev[32];
    char    mt_filsys[32];
    short   mt_ro_flg;
    time_t  mt_time;
};
```

Each entry is 70 bytes in length; the first 32 bytes are the null-padded name of the place where the *special file* is mounted; the next 32 bytes represent the null-padded root name of the mounted special file; the remaining 6 bytes contain the mounted *special file*'s read/write permissions and the date on which it was mounted.

**SEE ALSO**

*df(1M)*, *mount(1M)*, *setmnt(1M)* in the *Administrator Reference Manual*.



**NAME**

networks — network name data base

**DESCRIPTION**

The *networks* file contains information regarding the known networks which comprise the DARPA Internet. For each network a single line should be present with the following information:

official network name  
network number  
aliases

Items are separated by any number of blanks and/or tab characters. A “#” indicates the beginning of a comment; characters up to the end of the line are not interpreted by routines which search the file. This file is normally created from the official network data base maintained at the Network Information Control Center (NIC), though local changes may be required to bring it up to date regarding unofficial aliases and/or unknown networks.

Network number may be specified in the conventional “.” notation using the *inet\_network()* routine from the Internet address manipulation library, *inet(3N)*. Network names may contain any printable character other than a field delimiter, newline, or comment character.

**FILES**

/etc/networks

**SEE ALSO**

getnetent(3N)

**BUGS**

A name server should be used instead of a static file. A binary indexed file format should be available for fast access.

**NAME**

passwd — password file

**DESCRIPTION***Passwd* contains for each user the following information:

login name  
 encrypted password  
 numerical user ID  
 numerical group ID  
 GCOS job number, box number, optional GCOS user ID  
 initial working directory  
 program to use as Shell

This is an ASCII file. Each field within each user's entry is separated from the next by a colon. The GCOS field is used only when communicating with that system, and in other installations can contain any desired information. Each user is separated from the next by a new-line. If the password field is null, no password is demanded; if the Shell field is null, the Shell itself is used.

This file resides in directory */etc*. Because of the encrypted passwords, it can and does have general read permission and can be used, for example, to map numerical user IDs to names.

The encrypted password consists of 13 characters chosen from a 64-character alphabet (*.*, */*, 0-9, A-Z, a-z), except when the password is null, in which case the encrypted password is also null. Password aging is effected for a particular user if his encrypted password in the password file is followed by a comma and a non-null string of characters from the above alphabet. (Such a string must be introduced in the first instance by the super-user.)

The first character of the age, *M* say, denotes the maximum number of weeks for which a password is valid. A user who attempts to login after his password has expired will be forced to supply a new one. The next character, *m* say, denotes the minimum period in weeks which must expire before the password may be changed. The remaining characters define the week (counted from the beginning of 1970) when the password was last changed. (A null string is equivalent to zero.) *M* and *m* have numerical values in the range 0-63 that correspond to the 64-character alphabet shown above (i.e., */* = 1 week; *z* = 63 weeks). If *m* = *M* = 0 (derived from the string *.* or *..*) the user will be forced to change his password the next time he logs in (and the "age" will disappear from his entry in the password file). If *m* > *M* (signified, e.g., by the string *./*) only the super-user will be able to change the password.

**FILES***/etc/passwd***SEE ALSO**

login(1), passwd(1), a64I(3C), crypt(3C), getpwent(3C), group(4).

**NAME**

plot - graphics interface

**DESCRIPTION**

Files of this format are produced by routines described in *plot(3X)* and are interpreted for various devices by commands described in *tplot(1G)*. A graphics file is a stream of plotting instructions. Each instruction consists of an ASCII letter usually followed by bytes of binary information. The instructions are executed in order. A point is designated by four bytes representing the *x* and *y* values; each value is a signed integer. The last designated point in an *l*, *m*, *n*, or *p* instruction becomes the "current point" for the next instruction.

Each of the following descriptions begins with the name of the corresponding routine in *plot(3X)*.

- m** move: The next four bytes give a new current point.
- n** cont: Draw a line from the current point to the point given by the next four bytes. See *tplot(1G)*.
- p** point: Plot the point given by the next four bytes.
- l** line: Draw a line from the point given by the next four bytes to the point given by the following four bytes.
- t** label: Place the following ASCII string so that its first character falls on the current point. The string is terminated by a new-line.
- e** erase: Start another frame of output.
- f** linemod: Take the following string, up to a new-line, as the style for drawing further lines. The styles are "dotted", "solid", "longdashed", "shortdashed", and "dotdashed". Effective only for the -T4014 and -Tver options of *tplot(1G)* (TEKTRONIX 4014 terminal and Versatec plotter).
- s** space: The next four bytes give the lower left corner of the plotting area; the following four give the upper right corner. The plot will be magnified or reduced to fit the device as closely as possible.

Space settings that exactly fill the plotting area with unity scaling appear below for devices supported by the filters of *tplot(1G)*. The upper limit is just outside the plotting area. In every case the plotting area is taken to be square; points outside may be displayable on devices whose face is not square.

DASI 300	space(0, 0, 4096, 4096);
DASI 300s	space(0, 0, 4096, 4096);
DASI 450	space(0, 0, 4096, 4096);
TEKTRONIX 4014	space(0, 0, 3120, 3120);
Versatec plotter	space(0, 0, 2048, 2048);

**SEE ALSO**

*tplot(1G)*, *plot(3X)*, *term(5)*.

**WARNING**

The plotting library *plot(3X)* and the curses library *curses(3X)* both use the names *erase()* and *move()*. The curses versions are macros. If you need

**PLOT(4)**

**PLOT(4)**

both libraries, put the *plot(3X)* code in a different source file than the *curses(3X)* code, and/or `#undef move()` and `erase()` in the *plot(3X)* code.

**NAME**

pnch — file format for card images

**DESCRIPTION**

The PNCH format is a convenient representation for files consisting of card images in an arbitrary code.

A PNCH file is a simple concatenation of card records. A card record consists of a single control byte followed by a variable number of data bytes. The control byte specifies the number (which must lie in the range 0-80) of data bytes that follow. The data bytes are 8-bit codes that constitute the card image. If there are fewer than 80 data bytes, it is understood that the remainder of the card image consists of trailing blanks.

**SEE ALSO**

send(2N).

**NAME**

profile — setting up an environment at login time

**DESCRIPTION**

If your login directory contains a file named `.profile`, that file will be executed (via the shell's `exec .profile`) before your session begins; `.profiles` are handy for setting exported environment variables and terminal modes. If the file `/etc/profile` exists, it will be executed for every user before the `.profile`. The following example is typical (except for the comments):

```
# Make some environment variables global
export MAIL PATH TERM
# Set file creation mask
umask 22
# Tell me when new mail comes in
MAIL=/usr/mail/myname
# Add my /bin directory to the shell search sequence
PATH=$PATH:$HOME/bin
# Set terminal type
echo "terminal: \c"
read TERM
case $TERM in
    300)      stty cr2 n10 tabs; tabs;;
    300s)     stty cr2 n10 tabs; tabs;;
    450)     stty cr2 n10 tabs; tabs;;
    hp)      stty cr0 n10 tabs; tabs;;
    745|735) stty cr1 n11 -tabs; TERM=745;;
    43)      stty cr1 n10 -tabs;;
    4014|tek) stty cr0 n10 -tabs ff1; TERM=4014; echo "\33";;
    *)      echo "$TERM unknown";;
esac
```

**FILES**

`$HOME/.profile`  
`/etc/profile`

**SEE ALSO**

`env(1)`, `login(1)`, `mail(1)`, `sh(1)`, `stty(1)`, `su(1)`, `environ(5)`, `term(5)`.

**NAME**

protocols — protocol name data base

**DESCRIPTION**

The *protocols* file contains information regarding the known protocols used in the DARPA Internet. For each protocol a single line should be present with the following information:

official protocol name  
protocol number  
aliases

Items are separated by any number of blanks and/or tab characters. A “#” indicates the beginning of a comment; characters up to the end of the line are not interpreted by routines which search the file.

Protocol names may contain any printable character other than a field delimiter, newline, or comment character.

**FILES**

/etc/protocols

**SEE ALSO**

getprotoent(3N)

**BUGS**

A name server should be used instead of a static file. A binary indexed file format should be available for fast access.

**NAME**

reloc — relocation information for a common object file

**SYNOPSIS**

```
#include <reloc.h>
```

**DESCRIPTION**

Object files have one relocation entry for each relocatable reference in the text or data. If relocation information is present, it will be in the following format.

```
struct   reloc
{
    long   r_vaddr ;    /* (virtual) address of reference */
    long   r_symndx ;   /* index into symbol table */
    short  r_type ;     /* relocation type */
};

/*
 * All generics
 *   reloc. already performed to symbol in the same section
 */
#define R_ABS          0

/*
 * DEC Processors VAX 11/780 and VAX 11/750
 *
 */
#define R_RELBYTE  017
#define R_RELWORD  020
#define R_RELLONG  021
#define R_PCRBYTE  022
#define R_PCRWORD  023
#define R_PCRLONG  024

/*
 * Motorola 68000 uses R_RELBYTE, R_RELWORD, R_RELLONG,
 * R_PCRBYTE, and R_PCRWORD as for DEC machines above.
 */
```

As the link editor reads each input section and performs relocation, the relocation entries are read. They direct how references found within the input section are treated.

**R\_ABS**           The reference is absolute, and no relocation is necessary. The entry will be ignored.

**R\_RELBYTE**       A direct 8-bit reference to a symbol's virtual address.

**R\_RELWORD**       A direct 16-bit reference to a symbol's virtual address.

**R\_RELLONG**       A direct 32-bit reference to a symbol's virtual address.

**R\_PCRBYTE**       A "PC-relative" 8-bit reference to a symbol's virtual address.



**R\_PCRWORD** A "PC-relative" 16-bit reference to a symbol's virtual address.

**R\_PCRLONG** A "PC-relative" 32-bit reference to a symbol's virtual address.

On the VAX processors, relocation of a symbol index of -1 indicates that the relative difference between the current segment's start address and the program's load address is added to the relocatable address.

Other relocation types will be defined as they are needed.

Relocation entries are generated automatically by the assembler and automatically utilized by the link editor. A link editor option exists for removing the relocation entries from an object file.

**SEE ALSO**

ld(1), strip(1), a.out(4), syms(4).

## NAME

sccsfile — format of SCCS file

## DESCRIPTION

An SCCS file is an ASCII file. It consists of six logical parts: the *checksum*, the *delta table* (contains information about each delta), *user names* (contains login names and/or numerical group IDs of users who may add deltas), *flags* (contains definitions of internal keywords), *comments* (contains arbitrary descriptive information about the file), and the *body* (contains the actual text lines intermixed with control lines).

Throughout an SCCS file there are lines which begin with the ASCII SOH (start of heading) character (octal 001). This character is hereafter referred to as *the control character* and will be represented graphically as @. Any line described below which is not depicted as beginning with the control character is prevented from beginning with the control character.

Entries of the form DDDDD represent a five-digit string (a number between 00000 and 99999).

Each logical part of an SCCS file is described in detail below.

*Checksum*

The checksum is the first line of an SCCS file. The form of the line is:

```
@hDDDDDD
```

The value of the checksum is the sum of all characters, except those of the first line. The @h provides a *magic number* of (octal) 064001.

*Delta table*

The delta table consists of a variable number of entries of the form:

```
@s DDDDD/DBDDDD/DDDDDD
@d <type> <SCCS ID> yr/mo/da hr:mi:se <pgmr> DDDDD DDDDD
@f DDDDD ...
@x DDDDD ...
@g DDDDD ...
@m <MR number>
.
.
@c <comments> ...
.
.
@e
```

The first line (@s) contains the number of lines inserted/deleted/unchanged, respectively. The second line (@d) contains the type of the delta (currently, normal: D, and removed: R), the SCCS ID of the delta, the date and time of creation of the delta, the login name corresponding to the real user ID at the time

the delta was created, and the serial numbers of the delta and its predecessor, respectively.

The @i, @x, and @g lines contain the serial numbers of deltas included, excluded, and ignored, respectively. These lines are optional.

The @m lines (optional) each contain one MR number associated with the delta; the @c lines contain comments associated with the delta.

The @e line ends the delta table entry.

#### *User names*

The list of login names and/or numerical group IDs of users who may add deltas to the file, separated by new-lines. The lines containing these login names and/or numerical group IDs are surrounded by the bracketing lines @u and @U. An empty list allows anyone to make a delta. Any line starting with a ! prohibits the succeeding group or user from making deltas.

#### *Flags*

Keywords used internally (see *admin(1)* for more information on their use). Each flag line takes the form:

```
@f <flag> <optional text>
```

The following flags are defined:

```
@f t <type of program>
@f v <program name>
@f i <keyword string>
@f b
@f m <module name>
@f f <floor>
@f c <ceiling>
@f d <default-sid>
@f n
@f j
@f l <lock-releases>
@f q <user defined>
@f z <reserved for use in interfaces>
```

The t flag defines the replacement for the %Y% identification keyword. The v flag controls prompting for MR numbers in addition to comments; if the optional text is present it defines an MR number validity checking program. The i flag controls the warning/error aspect of the "No id keywords" message. When the l flag is not present, this message is only a warning; when the i flag is present, this message will cause a "fatal" error (the file will not be gotten, or the delta will not be made). When the b flag is present the -b keyletter may be used on the *get* command to cause a branch in the delta tree. The m flag defines the first choice for the replacement

text of the %M% identification keyword. The *f* flag defines the "floor" release; the release below which no deltas may be added. The *c* flag defines the "ceiling" release; the release above which no deltas may be added. The *d* flag defines the default SID to be used when none is specified on a *get* command. The *n* flag causes *delta* to insert a "null" delta (a delta that applies *no* changes) in those releases that are skipped when a delta is made in a *new* release (e.g., when delta 5.1 is made after delta 2.7; releases 3 and 4 are skipped). The absence of the *n* flag causes skipped releases to be completely empty. The *j* flag causes *get* to allow concurrent edits of the same base SID. The *l* flag defines a *list* of releases that are *locked* against editing (*get*(1) with the *-e* keyletter). The *q* flag defines the replacement for the %Q% identification keyword. The *z* flag is used in certain specialized interface programs.

#### Comments

Arbitrary text is surrounded by the bracketing lines @t and @T. The comments section typically will contain a description of the file's purpose.

#### Body

The body consists of text lines and control lines. Text lines do not begin with the control character, control lines do. There are three kinds of control lines: *insert*, *delete*, and *end*, represented by:

```
@I DDDDD
@D DDDDD
@E DDDDD
```

respectively. The digit string is the serial number corresponding to the delta for the control line.

#### SEE ALSO

admin(1), delta(1), get(1), prs(1).  
SCCS in the *Programming Tools Guide*.

**NAME**

scnhdr — section header for a common object file

**SYNOPSIS**

```
#include <scnhdr.h>
```

**DESCRIPTION**

Every common object file has a table of section headers to specify the layout of the data within the file. Each section within an object file has its own header. The C structure appears below.

```
struct scnhdr
{
    char          s_name[SYMNMLEN]; /* section name */
    long          s_paddr; /* physical address */
    long          s_vaddr; /* virtual address */
    long          s_size; /* section size */
    long          s_scnptr; /* file ptr to raw data */
    long          s_relptr; /* file ptr to relocation */
    long          s_innoptr; /* file ptr to line numbers */
    unsigned short s_nreloc; /* # reloc entries */
    unsigned short s_nlnno; /* # line number entries */
    long          s_flags; /* flags */
};
```

File pointers are byte offsets into the file; they can be used as the offset in a call to *fseek(3S)*. If a section is initialized, the file contains the actual bytes. An uninitialized section is somewhat different. It has a size, symbols defined in it, and symbols that refer to it, but it can have no relocation entries, line numbers, or data. Consequently, an uninitialized section has no raw data in the object file, and the values for *s\_scnptr*, *s\_relptr*, *s\_innoptr*, *s\_nreloc*, and *s\_nlnno* are zero.

**SEE ALSO**

ld(1), *fseek(3S)*, *a.out(4)*.

**NAME**

services — service name data base

**DESCRIPTION**

The *services* file contains information regarding the known services available in the DARPA Internet. For each service a single line should be present with the following information:

official service name  
port number  
protocol name  
aliases

Items are separated by any number of blanks and/or tab characters. The port number and protocol name are considered a single *item*; a “/” is used to separate the port and protocol (e.g. “512/tcp”). A “#” indicates the beginning of a comment; characters up to the end of the line are not interpreted by routines which search the file.

Service names may contain any printable character other than a field delimiter, newline, or comment character.

**FILES**

/etc/services

**SEE ALSO**

getservent(3N)

**BUGS**

A name server should be used instead of a static file. A binary indexed file format should be available for fast access.

**NAME**

syms — common object file symbol table format

**SYNOPSIS**

#include &lt;syms.h&gt;

**DESCRIPTION**

Common object files contain information to support *symbolic* software testing (see *sdb(1)*). Line number entries, *linenum(4)*, and extensive symbolic information permit testing at the C *source* level. Every object file's symbol table is organized as shown below.

```

Filename 1.
    Function 1.
        Local symbols for function 1.
    Function 2
        Local symbols for function 2.
    ...
    Static externs for file 1.

Filename 2.
    Function 1.
        Local symbols for function 1.
    Function 2.
        Local symbols for function 2.
    ...
    Static externs for file 2.

...

Defined global symbols.
Undefined global symbols.

```

The entry for a symbol is a fixed-length structure. The members of the structure hold the name (null padded), its value, and other information. The C structure is given below.

```

#define SYMNMLEN      8
#define FILNMLEN     14

struct    syment
{
    union
    {
        char    _n_name[SYMNMLEN] /* names less than 8 chars. */
        struct
        {
            long    _n_zeroes; /* == 0L when in string table */
            long    _n_offset; /* location of name in table */
        } n_n;
        char    *n_nptr[2]; /* allows overlaying */
    } _n;
    long    n_value; /* value of symbol */
    short   n_snum; /* section number */
    unsigned short n_type; /* type and derived type */
    char    n_class; /* storage class */
    char    n_numaux; /* number of aux entries */

```

```

};
#define n_name      _n_n_name
#define n_zeroes   _n_n_n_n_zeroes
#define n_offset   _n_n_n_n_offset
#define n_nptr     _n_n_nptr[]

```

Meaningful values and explanations for them are given in both `syms.h` and *Common Object File Format*. Anyone who needs to interpret the entries should seek more information in these sources. Some symbols require more information than a single entry; they are followed by *auxiliary entries* that are the same size as a symbol entry. The format follows.

```

union auxent
{
    struct
    {
        long      x_tagndx;
        union
        {
            struct
            {
                unsigned short x_inno;
                unsigned short x_size;
            } x_insz;
            long      x_fsize;
        } x_misc;
        union
        {
            struct
            {
                long      x_innoptr;
                long      x_endndx;
            } x_fc;
            struct
            {
                unsigned short x_dimen[DIMNUM];
            } x_ary;
        } x_fcary;
        unsigned short x_tvndx;
    } x_sym;
    struct
    {
        char      x_fname[FILNMLEN];
    } x_file;
    struct
    {
        long      x_scnlen;
        unsigned short x_nreloc;
        unsigned short x_nliino;
    } x_scn;
    struct
    {
        unsigned short x_tvlen;

```



```
        unsigned short x_tvran[2];  
    }  
    x_tv;
```

Indexes of symbol table entries begin at zero.

**SEE ALSO**

sdb(1), a.out(4), linenum(4).

COFF in the *Programming Guide*.

**WARNING**

In machines in which longs are equivalent to ints (M68000 and VAX), the longs are converted to ints in the compiler to minimize the complexity of the compiler code generator. Thus, the information about which symbols are declared as longs and which as ints cannot be determined from the symbol table.

**NAME**

term — format of compiled term file.

**SYNOPSIS**

term

**DESCRIPTION**

Compiled terminfo descriptions are placed under the directory `/usr/lib/terminfo`. In order to avoid a linear search of a huge UNIX system directory, a two-level scheme is used: `/usr/lib/terminfo/c/name` where *name* is the name of the terminal, and *c* is the first character of *name*. Thus, `act4` can be found in the file `/usr/lib/terminfo/a/act4`. Synonyms for the same terminal are implemented by multiple links to the same compiled file.

The format has been chosen so that it will be the same on all hardware. An 8 or more bit byte is assumed, but no assumptions about byte ordering or sign extension are made.

The compiled file is created with the `compile` program, and read by the routine `setupterm`. Both of these pieces of software are part of `curses(3X)`. The file is divided into six parts: the header, terminal names, boolean flags, numbers, strings, and string table.

The header section begins the file. This section contains six short integers in the format described below. These integers are (1) the magic number (octal 0432); (2) the size, in bytes, of the names section; (3) the number of bytes in the boolean section; (4) the number of short integers in the numbers section; (5) the number of offsets (short integers) in the strings section; (6) the size, in bytes, of the string table.

Short integers are stored in two 8-bit bytes. The first byte contains the least significant 8 bits of the value, and the second byte contains the most significant 8 bits. (Thus, the value represented is  $256 \cdot \text{second} + \text{first}$ .) The value `-1` is represented by `0377, 0377`, other negative values are illegal. The `-1` generally means that a capability is missing from this terminal. Note that this format corresponds to the hardware of the VAX and PDP-11. Machines where this does not correspond to the hardware read the integers as two bytes and compute the result.

The terminal names section comes next. It contains the first line of the terminfo description, listing the various names for the terminal, separated by the `†` character. The section is terminated with an ASCII NUL character.

The boolean flags have one byte for each flag. This byte is either 0 or 1 as the flag is present or absent. The capabilities are in the same order as the file `<term.h>`.

Between the boolean section and the number section, a null byte will be inserted, if necessary, to ensure that the number section begins on an even byte. All short integers are aligned on a short word boundary.

The numbers section is similar to the flags section. Each capability takes up two bytes, and is stored as a short integer. If the value represented is `-1`, the capability is taken to be missing.

The strings section is also similar. Each capability is stored as a short integer, in the format above. A value of `-1` means the capability is missing. Otherwise, the value is taken as an offset from the beginning of the string table. Special characters in `\X` or `\c` notation are stored in their

interpreted form, not the printing representation. Padding information \$<nn> and parameter information %x are stored intact in uninterpreted form.

The final section is the string table. It contains all the values of string capabilities referenced in the string section. Each string is null terminated.

Note that it is possible for *setupterm* to expect a different set of capabilities than are actually present in the file. Either the database may have been updated since *setupterm* has been recompiled (resulting in extra unrecognized entries in the file) or the program may have been recompiled more recently than the database was updated (resulting in missing entries). The routine *setupterm* must be prepared for both possibilities — this is why the numbers and sizes are included. Also, new capabilities must always be added at the end of the lists of boolean, number, and string capabilities.

As an example, an octal dump of the description for the Microterm ACT 4 is included:

```
microtermact4|microterm act iv,
cr=^M, cudl=^J, ind=^J, bel=^G, am, cubl=^H,
ed=^_, ej=^^, clear=^L, cup=^T%p1%c%p2%c,
cols#80, lines#24, cuf1=^X, cuul=^Z, home=^_].
```

```
000 032 001  \0 025 \0 \b \0 212 \0  " \0 m i c r
020 o t e r m | a c t 4 | m i c r o
040 t e r m   a c t   i v \0 \001 \0 \0
060 \0 \0 \0 \0 \0 \0 \0 \0 \0 \0 \0 \0 \0 \0 \0 \0
100 \0 \0 P \0 377 377 030 \0 377 377 377 377 377 377 377 377
120 377 377 377 377 \0 \0 002 \0 377 377 377 377 004 \0 006 \0
140 \b \0 377 377 377 377 \n \0 026 \0 030 \0 377 377 032 \0
160 377 377 377 377 034 \0 377 377 036 \0 377 377 377 377 377 377
200 377 377 377 377 377 377 377 377 377 377 377 377 377 377 377
*
520 377 377 377 377  \0 377 377 377 377 377 377 377 377 377 377
540 377 377 377 377 377 377 007 \0 \r \0 \f \0 036 \0 037 \0
560 024 % p 1 % c % p 2 % c \0 \n \0 035 \0
600 \b \0 030 \0 032 \0 \n \0
```

Some limitations: total compiled entries cannot exceed 4096 bytes. The name field cannot exceed 128 bytes.

#### FILES

/usr/lib/terminfo/\*/\* compiled terminal capability data base

#### SEE ALSO

curses(3X), terminfo(4).

**NAME**

**terminfo** – terminal capability data base

**SYNOPSIS**

*/usr/lib/terminfo/\*/\**

**DESCRIPTION**

*Terminfo* is a data base describing terminals, used, *e.g.*, by *vi*(1) and *curses*(3X). Terminals are described in *terminfo* by giving a set of capabilities which they have, and by describing how operations are performed. Padding requirements and initialization sequences are included in *terminfo*.

Entries in *terminfo* consist of a number of ',' separated fields. White space after each ',' is ignored. The first entry for each terminal gives the names which are known for the terminal, separated by '|' characters. The first name given is the most common abbreviation for the terminal, the last name given should be a long name fully identifying the terminal, and all others are understood as synonyms for the terminal name. All names but the last should be in lower case and contain no blanks; the last name may well contain upper case and blanks for readability.

Terminal names (except for the last, verbose entry) should be chosen using the following conventions. The particular piece of hardware making up the terminal should have a root name chosen, thus "hp2621". This name should not contain hyphens, except that synonyms may be chosen that do not conflict with other names. Modes that the hardware can be in, or user preferences, should be indicated by appending a hyphen and an indicator of the mode. Thus, a vt100 in 132 column mode would be vt100-w. The following suffixes should be used where possible:

Suffix	Meaning	Example
-w	Wide mode (more than 80 columns)	vt100-w
-am	With auto. margins (usually default)	vt100-am
-nam	Without automatic margins	vt100-nam
-n	Number of lines on the screen	aaa-60
-na	No arrow keys (leave them in local)	c100-na
-np	Number of pages of memory	c100-4p
-rv	Reverse video	c100-rv

**CAPABILITIES**

The variable is the name by which the programmer (at the *terminfo* level) accesses the capability. The capname is the short name used in the text of the database, and is used by a person updating the database. The i.code is the two

letter internal code used in the compiled database, and always corresponds to the old termcap capability name.

Capability names have no hard length limit, but an informal limit of 5 characters has been adopted to keep them short and to allow the tabs in the source file caps to line up nicely. Whenever possible, names are chosen to be the same as or similar to the ANSI X3.64-1979 standard. Semantics are also intended to match those of the specification.

(P) indicates that padding may be specified

(G) indicates that the string is passed through `tparm` with `parms` as given (`#i`).

(\*) indicates that padding may be based on the number of lines affected

(#*i*) indicates the *i*<sup>th</sup> parameter.

Variable	Cap-name	L Code	Description
auto_left_margin,	bw	bw	cu1 wraps from column 0 to last column
auto_right_margin,	am	am	Terminal has automatic margins
beehive_glitch,	xsb	xb	Beehive ( <code>f1=escape</code> , <code>f2=ctrl C</code> )
ceol_standout_glitch,	xhp	xs	Standout not erased by overwriting ( <code>hp</code> )
eat_newline_glitch,	xenl	xn	newline ignored after 80 cols (Concept)
erase_overstrike,	eo	eo	Can erase overstrikes with a blank
generic_type,	gn	gn	Generic line type (e.g., dialup, switch).
hard_copy,	hc	hc	Hardcopy terminal
has_meta_key,	km	km	Has a meta key (shift, sets parity bit)
has_status_line,	hs	hs	Has extra "status line"
insert_null_glitch,	in	in	Insert mode distinguishes nulls
memory_above,	da	da	Display may be retained above the screen
memory_below,	db	db	Display may be retained below the screen
move_insert_mode,	mir	mi	Safe to move while in insert mode
move_standout_mode,	msgr	ms	Safe to move in standout modes
over_strike,	os	os	Terminal overstrikes

status_line_esc_ok,	eslok	es	Escape can be used on the status line
telera_y_glitch,	xt	xt	Tabs ruin, magic so char (Telera_y 1061)
tilde_glitch,	hz	hz	Hazeltine; can not print '~'s
transparent_underline,	ul	ul	underline character overstrikes
xon_xoff,	xon	xo	Terminal uses xon/xoff handshaking
<b>Numbers:</b>			
columns,	cols	co	Number of columns in a line
init_tabs,	it	it	Tabs initially every # spaces
lines,	lines	li	Number of lines on screen or page
lines_of_memory,	lm	lm	Lines of memory if > lines. 0 means varies
magic_cookie_glitch,	xmc	sg	Number of blank chars left by smso or rns0
padding_baud_rate,	pb	pb	Lowest baud where cr/nl padding is needed
virtual_terminal,	vt	vt	Virtual terminal number (UNIX system)
width_status_line,	wsl	ws	No. columns in status line
<b>Strings:</b>			
back_tab,	cbt	bt	Back tab (P)
bell,	bel	bl	Audible signal (bell) (P)
carriage_return,	cr	cr	Carriage return (P*)
change_scroll_region,	csr	cs	change to lines #1 through #2 (vt100) (PG)
clear_all_tabs,	tbc	ct	Clear all tab stops (P)
clear_screen,	clear	cl	Clear screen and home cursor (P*)
clr_eol,	el	ce	Clear to end of line (P)
clr_eos,	ed	cd	Clear to end of display (P*)
column_address,	hpa	ch	Set cursor column (PG)
command_character,	cmdch	CC	Term. settable cmd char in prototype
cursor_address,	cup	cm	Screen rel. cursor motion row #1 col #2 (PG)
cursor_down,	cuD1	do	Down one line
cursor_home,	home	ho	Home cursor (if no cup)
cursor_invisible,	civis	vi	Make cursor invisible
cursor_left,	cub1	le	Move cursor left one space
cursor_mem_address,	mrcup	CM	Memory relative cursor addressing

cursor_normal,	cnorm	ve	Make cursor appear normal (undo vs/vi)
cursor_right,	cuf1	nd	Non-destructive space (cursor right)
cursor_to_ll,	ll	ll	Last line, first column (if no cup)
cursor_up,	cuu1	up	Upline (cursor up)
cursor_visible,	cvvis	vs	Make cursor very visible
delete_character,	dch1	dc	Delete character (P*)
delete_line,	dll	dl	Delete line (P*)
dis_status_line,	dsl	ds	Disable status line
down_half_line,	hd	hd	Half-line down (forward 1/2 linefeed)
enter_alt_charset_mode,	smacs	as	Start alternate character set (P)
enter_blink_mode,	blink	mb	Turn on blinking
enter_bold_mode,	bold	md	Turn on bold (extra bright) mode
enter_ca_mode,	smcup	ti	String to begin programs that use cup
enter_delete_mode,	smdc	dm	Delete mode (enter)
enter_dim_mode,	dim	mh	Turn on half-bright mode
enter_insert_mode,	amir	im	Insert mode (enter);
enter_protected_mode,	prot	mp	Turn on protected mode
enter_reverse_mode,	rev	mr	Turn on reverse video mode
enter_secure_mode,	invis	mk	Turn on blank mode (chars invisible)
enter_standout_mode,	smso	so	Begin stand out mode
enter_underline_mode,	smul	us	Start underscore mode
erase_chars	ech	ec	Erase #1 characters (PG)
exit_alt_charset_mode,	rmacs	ae	End alternate character set (P)
exit_attribute_mode,	sgr0	me	Turn off all attributes
exit_ca_mode,	rmcup	te	String to end programs that use cup
exit_delete_mode,	rmdc	ed	End delete mode
exit_insert_mode,	rmir	ei	End insert mode
exit_standout_mode,	rmso	se	End stand out mode
exit_underline_mode,	rmul	ue	End underscore mode
flash_screen,	flash	vb	Visible bell (may not move cursor)
form_feed,	ff	ff	Hardcopy terminal page eject (P*)
from_status_line,	fsl	fs	Return from status line
init_lstring,	is1	i1	Terminal initialization string
init_2string,	is2	i2	Terminal initialization string
init_3string,	is3	i3	Terminal initialization string
init_file,	if	if	Name of file containing is
insert_character,	ich1	ic	Insert character (P)
insert_line,	il1	al	Add new blank line (P*)
insert_padding,	ip	ip	Insert pad after character inserted

key_backspace,	kbs	kb	(p*) Sent by backspace key
key_catab,	ktbc	ka	Sent by clear-all-tabs key
key_clear,	kclr	kC	Sent by clear screen or erase key
key_ctab,	kctab	kt	Sent by clear-tab key
key_dc,	kdch1	kD	Sent by delete character key
key_dl,	kdll	kL	Sent by delete line key
key_down,	kcud1	kd	Sent by terminal down arrow key
key_eic,	krmir	kM	Sent by rmir or smir in insert mode
key_eol,	kel	kE	Sent by clear-to-end-of-line key
key_eos,	ked	kS	Sent by clear-to-end-of-screen key
key_f0,	kf0	k0	Sent by function key f0
key_f1,	kf1	k1	Sent by function key f1
key_f10,	kf10	ka	Sent by function key f10
key_f2,	kf2	k2	Sent by function key f2
key_f3,	kf3	k3	Sent by function key f3
key_f4,	kf4	k4	Sent by function key f4
key_f5,	kf5	k5	Sent by function key f5
key_f6,	kf6	k6	Sent by function key f6
key_f7,	kf7	k7	Sent by function key f7
key_f8,	kf8	k8	Sent by function key f8
key_f9,	kf9	k9	Sent by function key f9
key_home,	khome	kh	Sent by home key
key_ic,	kich1	ki	Sent by ins char/enter ins mode key
key_il,	kill	kA	Sent by insert line
key_left,	kcub1	kl	Sent by terminal left arrow key
key_ll,	kl	kH	Sent by home-down key
key_npage,	knp	kN	Sent by next-page key
key_ppage,	kpp	kP	Sent by previous-page key
key_right,	kcuf1	kr	Sent by terminal right arrow key
key_sf,	kind	kF	Sent by scroll-forward/down key
key_sr,	kri	kR	Sent by scroll-backward/up key
key_stab,	khts	kT	Sent by set-tab key
key_up,	kcuu1	ku	Sent by terminal up arrow key
keypad_local,	rmkx	ke	Out of "keypad transmit" mode
keypad_xmit,	smkx	ks	Put terminal in "keypad transmit" mode
lab_f0,	lf0	l0	Labels on function key f0 if not f0
lab_f1,	lf1	l1	Labels on function key f1 if not f1
lab_f10,	lf10	la	Labels on function key f10 if not f10



lab_f2,	lf2	l2	Labels on function key f2 if not f2
lab_f3,	lf3	l3	Labels on function key f3 if not f3
lab_f4,	lf4	l4	Labels on function key f4 if not f4
lab_f5,	lf5	l5	Labels on function key f5 if not f5
lab_f6,	lf6	l6	Labels on function key f6 if not f6
lab_f7,	lf7	l7	Labels on function key f7 if not f7
lab_f8,	lf8	l8	Labels on function key f8 if not f8
lab_f9,	lf9	l9	Labels on function key f9 if not f9
meta_on,	amm	mm	Turn on "meta mode" (8th bit)
meta_off,	rmm	mo	Turn off "meta mode"
newline,	nel	nw	Newline (behaves like cr followed by lf)
pad_char,	pad	pc	Pad character (rather than null)
parm_dch,	dch	DC	Delete #1 chars (PG*)
parm_delete_line,	dl	DL	Delete #1 lines (PG*)
parm_down_cursor,	cud	DO	Move cursor down #1 lines (PG*)
parm_ich,	ich	IC	Insert #1 blank chars (PG*)
parm_index,	indn	SF	Scroll forward #1 lines (PG)
parm_insert_line,	il	AL	Add #1 new blank lines (PG*)
parm_left_cursor,	cub	LE	Move cursor left #1 spaces (PG)
parm_right_cursor,	cuf	RI	Move cursor right #1 spaces (PG*)
parm_rindex,	rin	SR	Scroll backward #1 lines (PG)
parm_up_cursor,	cuu	UP	Move cursor up #1 lines (PG*)
pkey_key,	pfkey	pk	Prog funct key #1 to type string #2
pkey_local,	pfloc	pl	Prog funct key #1 to execute string #2
pkey_xmit,	pfx	px	Prog funct key #1 to xmit string #2
print_screen,	mc0	ps	Print contents of the screen
ptr_off,	mc4	pf	Turn off the printer
ptr_on,	mc5	po	Turn on the printer
repeat_char,	rep	rp	Repeat char #1 #2 times. (PG*)
reset_1string,	rs1	r1	Reset terminal completely to sane modes.
reset_2string,	rs2	r2	Reset terminal completely to sane modes.
reset_3string,	rs3	r3	Reset terminal completely to sane modes.
reset_file,	rf	rf	Name of file containing reset string
restore_cursor,	rc	rc	Restore cursor to position of last sc
row_address,	vpa	cv	Vertical position absolute (set row) (PG)
save_cursor,	sc	sc	Save cursor position (P)
scroll_forward,	ind	sf	Scroll text up (P)

scroll_reverse,	ri	sr	Scroll text down (P)
set_attributes,	sgr	sa	Define the video attributes (PG9)
set_tab,	hts	st	Set a tab in all rows, current column
set_window,	wind	wi	Current window is lines #1-#2 cols #3-#4
tab,	ht	ta	Tab to next 8 space hardware tab stop
to_status_line,	tsl	ts	Go to status line, column #1
underline_char,	uc	uc	Underscore one char and move past it
up_half_line,	hu	hu	Half-line up (reverse 1/2 linefeed)
init_prog,	iprogram	iP	Path name of program for init
key_a1,	ka1	K1	Upper left of keypad
key_a3,	ka3	K3	Upper right of keypad
key_b2,	kb2	K2	Center of keypad
key_c1,	kc1	K4	Lower left of keypad
key_c3,	kc3	K5	Lower right of keypad
prtr_non,	mc5p	pO	Turn on the printer for #1 bytes

### A Sample Entry

The following entry, which describes the Concept-100, is among the more complex entries in the *terminfo* file as of this writing.

```
concept100|c100|concept|c104|c100-4p|concept 100,
am, bel='G, blank='\EH, blink='\EC, clear='L$<2*>, cnorm='\Ew,
cols#80, cr='M$<9>, cub1='H, cud1='J, cufl='\E=,
cup='\Ea%p1%' '%+%c%p2%' '%+%c,
cuu1='\E;, cvvis='\EW, db, dch1='\E"A$<16*>, dim='\EE, dl1='\E"B$<3*>,
ed='\E"C$<16*>, el='\E"U$<16>, eo, flash='\Ek$<20>\EK, ht='\t$<8>,
il1='\E"R$<3*>, in, ind='J, .ind='J$<9>, ip=$<16*>,
is2='\EU\EAE7\E5\E3\E1\ENH\EK\B\200\Eo&\200\Eo\47\E,
kbs='h, kcub1='\E>, kcud1='\E<, kcufl='\E=, kcuu1='\E;,
kf1='\E5, kf2='\E6, kf3='\E7, khoms='\E7,
lines#24, mir, pb#9600, prot='\EI, rep='\Er%p1%c%p2%' '%+%c$<2*>,
rev='\ED, rmcup='\Ev $<6>\Epr'n, rmir='\E\200, rmx='\Ex,
rmso='\Ed\Ee, rmul='\Eg, rmul='\Eg, sgr0='\EN\200,
smcup='\EU\Ev 8p\Epr, smir='\E"P, smkx='\EX, smso='\EE\ED,
smul='\EG, tabs, ul, vt#8, xenl,
```

Entries may continue onto multiple lines by placing white space at the beginning of each line except the first. Comments may be included on lines beginning with "#". Capabilities in *terminfo* are of three types: Boolean

capabilities which indicate that the terminal has some particular feature, numeric capabilities giving the size of the terminal or the size of particular delays, and string capabilities, which give a sequence which can be used to perform particular terminal operations.

### Types of Capabilities

All capabilities have names. For instance, the fact that the Concept has *automatic margins* (i.e., an automatic return and linefeed when the end of a line is reached) is indicated by the capability `am`. Hence the description of the Concept includes `am`. Numeric capabilities are followed by the character '#' and then the value. Thus `cols`, which indicates the number of columns the terminal has, gives the value '80' for the Concept.

Finally, string valued capabilities, such as `el` (clear to end of line sequence) are given by the two-character code, an '=', and then a string ending at the next following ','. A delay in milliseconds may appear anywhere in such a capability, enclosed in `$<...>` brackets, as in `el=EK$<3>`, and padding characters are supplied by `tputs` to provide this delay. The delay can be either a number, e.g., '20', or a number followed by an '\*', i.e., '3\*'. A '\*' indicates that the padding required is proportional to the number of lines affected by the operation, and the amount given is the per-affected-unit padding required. (In the case of insert character, the factor is still the number of *lines* affected. This is always one unless the terminal has `xenl` and the software uses it.) When a '\*' is specified, it is sometimes useful to give a delay of the form '3.5' to specify a delay per unit to tenths of milliseconds. (Only one decimal place is allowed.)

A number of escape sequences are provided in the string valued capabilities for easy encoding of characters there. Both `\E` and `\e` map to an ESCAPE character, `^x` maps to a control-`x` for any appropriate `x`, and the sequences `\n` `\r` `\t` `\b` `\f` `\s` give a newline, linefeed, return, tab, backspace, formfeed, and space. Other escapes include `\^` for `^`, `\` for `\`, `\,` for comma, `\;` for `;`, and `\0` for null. (`\0` will produce `\200`, which does not terminate a string but behaves as a null character on most terminals.) Finally, characters may be given as three octal digits after a `\`.

Sometimes individual capabilities must be commented out. To do this, put a period before the capability name. For example, see the second `ind` in the example above.

### Preparing Descriptions

We now outline how to prepare descriptions of terminals. The most effective way to prepare a terminal description is by imitating the description of a similar terminal in *terminfo* and to build up a description gradually, using partial descriptions with *vi* to check that they are correct. Be aware that a very unusual terminal may expose deficiencies in the ability of the *terminfo* file to describe it or bugs in *vi*. To easily test a new terminal description you can set the environment variable `TERMINFO` to a pathname of a directory containing the compiled description you are working on and programs will look there rather than in */usr/lib/terminfo*. To get the padding for insert line right (if the terminal manufacturer did not document it) a severe test is to edit */etc/passwd* at 9600 baud, delete 16 or so lines from the middle of the screen, then hit the 'u' key several times quickly. If the terminal messes up, more padding is usually needed. A similar test can be used for insert character.

### Basic Capabilities

The number of columns on each line for the terminal is given by the `cols` numeric capability. If the terminal is a CRT, then the number of lines on the screen is given by the `lines` capability. If the terminal wraps around to the beginning of the next line when it reaches the right margin, then it should have the `am` capability. If the terminal can clear its screen, leaving the cursor in the home position, then this is given by the `clear` string capability. If the terminal overstrikes (rather than clearing a position when a character is struck over) then it should have the `os` capability. If the terminal is a printing terminal, with no soft copy unit, give it both `hc` and `os`. (`os` applies to storage scope terminals, such as TEKTRONIX 4010 series, as well as hard copy and APL terminals.) If there is a code to move the cursor to the left edge of the current row, give this as `er`. (Normally this will be carriage return, control M.) If there is a code to produce an audible signal (bell, beep, etc) give this as `bel`.

If there is a code to move the cursor one position to the left (such as backspace) that capability should be given as `cub1`. Similarly, codes to move to the right, up, and down should be given as `cuf1`, `cuu1`, and `cmd1`. These local cursor motions should not alter the text they pass over, for example, you would not normally use `'cuf1= '` because the space would erase the character moved over.

A very important point here is that the local cursor motions encoded in *terminfo* are undefined at the left and top edges of a CRT terminal. Programs should never attempt to backspace around the left edge, unless `bw` is given, and never attempt to go up locally off the top. In order to scroll text up, a program will go

to the bottom left corner of the screen and send the `ind` (index) string.

To scroll text down, a program goes to the top left corner of the screen and sends the `ri` (reverse index) string. The strings `ind` and `ri` are undefined when not on their respective corners of the screen.

Parameterized versions of the scrolling sequences are `indn` and `rin` which have the same semantics as `ind` and `ri` except that they take one parameter, and scroll that many lines. They are also undefined except at the appropriate edge of the screen.

The `am` capability tells whether the cursor sticks at the right edge of the screen when text is output, but this does not necessarily apply to a `cuf1` from the last column. The only local motion which is defined from the left edge is if `bw` is given, then a `cub1` from the left edge will move to the right edge of the previous row. If `bw` is not given, the effect is undefined. This is useful for drawing a box around the edge of the screen, for example. If the terminal has switch selectable automatic margins, the *terminfo* file usually assumes that this is on; i.e., `am`. If the terminal has a command which moves to the first column of the next line, that command can be given as `nel` (newline). It does not matter if the command clears the remainder of the current line, so if the terminal has no `cr` or if it may still be possible to craft a working `nel` out of one or both of them.

These capabilities suffice to describe hardcopy and glass-tty terminals. Thus the model 33 teletype is described as

```
33|tty33|tty|model 33 teletype,
bel="G, cols#72, cr="M, cud1="J, hc, ind="J, os,
```

while the Lear Siegler ADM-3 is described as

```
adm3|3|lsi adm3,
am, bel="G, clear="Z, cols#80, cr="M, cub1="H, cud1="J,
ind="J, lines#24,
```

### Parameterized Strings

Cursor addressing and other strings requiring parameters in the terminal are described by a parameterized string capability, with *printf*(3S) like escapes `%x` in it. For example, to address the cursor, the `cup` capability is given, using two parameters: the row and column to address to. (Rows and columns are numbered from zero and refer to the physical screen visible to the user, not to any unseen memory.) If the terminal has memory relative cursor addressing, that can be indicated by `mrcup`.

The parameter mechanism uses a stack and special % codes to manipulate it. Typically a sequence will push one of the parameters onto the stack and then print it in some format. Often more complex operations are necessary.

The % encodings have the following meanings:

%%	outputs '%'
%d	print pop() as in printf
%2d	print pop() like %2d
%3d	print pop() like %3d
%02d	
%03d	as in printf
%c	print pop() gives %c
%s	print pop() gives %s
%p[1-9]	push ith parm
%P[a-z]	set variable [a-z] to pop()
%g[a-z]	get variable [a-z] and push it
%'c'	char constant c
%{nn}	integer constant nn
%+ %- %+ %/ %m	arithmetic (%m is mod): push(pop() op pop())
%& %  %^	bit operations: push(pop() op pop())
%= %> %<	logical operations: push(pop() op pop())
%! %~	unary operations push(op pop())
%i	add 1 to first two parms (for ANSI terminals)
%? expr %t thenpart %e elsepart %;	if-then-else, %e elsepart is optional. else-if's are possible ala Algol 68: %? c <sub>1</sub> %t b <sub>1</sub> %e c <sub>2</sub> %t b <sub>2</sub> %e c <sub>3</sub> %t b <sub>3</sub> %e c <sub>4</sub> %t b <sub>4</sub> %e %; c <sub>i</sub> are conditions, b <sub>i</sub> are bodies.

Binary operations are in postfix form with the operands in the usual order. That is, to get x-5 one would use "%gx%{5}%-".

Consider the HP2645, which, to get to row 3 and column 12, needs to be sent \E&a12c03Y padded for 6 milliseconds. Note that the order of the rows and columns is inverted here, and that the row and column are printed as two digits. Thus its cup capability is cup=6\E&p2%2dc%p1%2dY.

The Microterm ACT-IV needs the current row and column sent preceded by a `T`, with the row and column simply encoded in binary, `cup=T%p1%c%p2%c`. Terminals which use `%c` need to be able to backspace the cursor (`cub1`), and to move the cursor up one line on the screen (`cuu1`). This is necessary because it is not always safe to transmit `\n`, `\D` and `\r`, as the system may change or discard them. (The library routines dealing with terminfo set tty modes so that tabs are never expanded, so `\t` is safe to send. This turns out to be essential for the Ann Arbor 4080.)

A final example is the LSI ADM-3a, which uses row and column offset by a blank character, thus `cup=\E-%p1%'%+%c%p2%'%+%c`. After sending `\E-`, this pushes the first parameter, pushes the ASCII value for a space (32), adds them (pushing the sum on the stack in place of the two previous values) and outputs that value as a character. Then the same is done for the second parameter. More complex arithmetic is possible using the stack.

If the terminal has row or column absolute cursor addressing, these can be given as single parameter capabilities `hpa` (horizontal position absolute) and `vpa` (vertical position absolute). Sometimes these are shorter than the more general two parameter sequence (as with the `hp2645`) and can be used in preference to `cup`. If there are parameterized local motions (e.g., move *n* spaces to the right) these can be given as `cud`, `cub`, `cuf`, and `cuu` with a single parameter indicating how many spaces to move. These are primarily useful if the terminal does not have `cup`, such as the `TEKTRONIX 4025`.

#### Cursor Motions

If the terminal has a fast way to home the cursor (to very upper left corner of screen) then this can be given as `home`; similarly a fast way of getting to the lower left-hand corner can be given as `ll`; this may involve going up with `cuu1` from the home position, but a program should never do this itself (unless `ll` does) because it can make no assumption about the effect of moving up from the home position. Note that the home position is the same as addressing to (0,0): to the top left corner of the screen, not of memory. (Thus, the `\EH` sequence on HP terminals cannot be used for home.)

#### Area Clears

If the terminal can clear from the current position to the end of the line, leaving the cursor where it is, this should be given as `el`. If the terminal can clear from the current position to the end of the display, then this should be given as `ed`. `Ed` is only defined from the first column of a line. (Thus, it can be simulated by a request to delete a large number of lines, if a true `ed` is not available.)

### Insert/delete line

If the terminal can open a new blank line before the line where the cursor is, this should be given as *ill*; this is done only from the first position of a line. The cursor must then appear on the newly blank line. If the terminal can delete the line which the cursor is on, then this should be given as *dll*; this is done only from the first position on the line to be deleted. Versions of *ill* and *dll* which take a single parameter and insert or delete that many lines can be given as *il* and *dl*. If the terminal has a settable scrolling region (like the vt100) the command to set this can be described with the *csr* capability, which takes two parameters: the top and bottom lines of the scrolling region. The cursor position is, alas, undefined after using this command. It is possible to get the effect of insert or delete line using this command - the *sc* and *rc* (save and restore cursor) commands are also useful. Inserting lines at the top or bottom of the screen can also be done using *ri* or *ind* on many terminals without a true insert/delete line, and is often faster even on terminals with those features.

If the terminal has the ability to define a window as part of memory, which all commands affect, it should be given as the parameterized string *wind*. The four parameters are the starting and ending lines in memory and the starting and ending columns in memory, in that order.

If the terminal can retain display memory above, then the *da* capability should be given; if display memory can be retained below, then *db* should be given. These indicate that deleting a line or scrolling may bring non-blank lines up from below or that scrolling back with *ri* may bring down non-blank lines.

### Insert/Delete Character

There are two basic kinds of intelligent terminals with respect to insert/delete character which can be described using *terminfo*. The most common insert/delete character operations affect only the characters on the current line and shift characters off the end of the line rigidly. Other terminals, such as the Concept 100 and the Perkin Elmer Owl, make a distinction between typed and untyped blanks on the screen, shifting upon an insert or delete only to an untyped blank on the screen which is either eliminated, or expanded to two untyped blanks. You can determine the kind of terminal you have by clearing the screen and then typing text separated by cursor motions. Type *abc def* using local cursor motions (not spaces) between the *abc* and the *def*. Then position the cursor before the *abc* and put the terminal in insert mode. If typing characters causes the rest of the line to shift rigidly and characters to fall off the end, then your terminal does not distinguish between blanks and untyped



positions. If the abc shifts over to the def which then move together around the end of the current line and onto the next as you insert, you have the second type of terminal, and should give the capability `ln`, which stands for insert null. While these are two logically separate attributes (one line vs. multiline insert mode, and special treatment of untyped spaces) we have seen no terminals whose insert mode cannot be described with the single attribute.

Terminfo can describe both terminals which have an insert mode, and terminals which send a simple sequence to open a blank position on the current line. Give as `smir` the sequence to get into insert mode. Give as `rmir` the sequence to leave insert mode. Now give as `ich1` any sequence needed to be sent just before sending the character to be inserted. Most terminals with a true insert mode will not give `ich1`; terminals which send a sequence to open a screen position should give it here. (If your terminal has both, insert mode is usually preferable to `ich1`. Do not give both unless the terminal actually requires both to be used in combination.) If post insert padding is needed, give this as a number of milliseconds in `lp` (a string option). Any other sequence which may need to be sent after an insert of a single character may also be given in `ip`. If your terminal needs both to be placed into an 'insert mode' and a special code to precede each inserted character, then both `smir/rmir` and `ich1` can be given, and both will be used. The `ich` capability, with one parameter, `n`, will repeat the effects of `ich1` `n` times.

It is occasionally necessary to move around while in insert mode to delete characters on the same line (e.g., if there is a tab after the insertion position). If your terminal allows motion while in insert mode you can give the capability `mir` to speed up inserting in this case. Omitting `mir` will affect only speed. Some terminals (notably Datamedia's) must not have `mir` because of the way their insert mode works.

Finally, you can specify `dch1` to delete a single character, `dch` with one parameter, `n`, to delete `n` characters, and delete mode by giving `smdc` and `rmdc` to enter and exit delete mode (any mode the terminal needs to be placed in for `dch1` to work).

A command to erase `n` characters (equivalent to outputting `n` blanks without moving the cursor) can be given as `ech` with one parameter.

#### Highlighting, Underlining, and Visible Bells

If your terminal has one or more kinds of display attributes, these can be represented in a number of different ways. You should choose one display form as *standout mode*, representing a good, high contrast, easy-on-the-eyes.

format for highlighting error messages and other attention getters. (If you have a choice, reverse video plus half-bright is good, or reverse video alone.) The sequences to enter and exit standout mode are given as `sms0` and `rmso`, respectively. If the code to change into or out of standout mode leaves one or even two blank spaces on the screen, as the TVI 912 and Teleray 1061 do, then `xmc` should be given to tell how many spaces are left.

Codes to begin underlining and end underlining can be given as `smul` and `rmul` respectively. If the terminal has a code to underline the current character and move the cursor one space to the right, such as the Microterm Mime, this can be given as `uc`.

Other capabilities to enter various highlighting modes include `blink` (blinking) `bold` (bold or extra bright) `dim` (dim or half-bright) `invis` (blanking or invisible text) `prot` (protected) `rev` (reverse video) `sgr0` (turn off *all* attribute modes) `smacs` (enter alternate character set mode) and `rmacs` (exit alternate character set mode). Turning on any of these modes singly may or may not turn off other modes.

If there is a sequence to set arbitrary combinations of modes, this should be given as `sgr` (set attributes), taking 9 parameters. Each parameter is either 0 or 1, as the corresponding attribute is on or off. The 9 parameters are, in order: standout, underline, reverse, blink, dim, bold, blank, protect, alternate character set. Not all modes need be supported by `sgr`, only those for which corresponding separate attribute commands exist.

Terminals with the "magic cookie" glitch (`xmc`) deposit special "cookies" when they receive mode-setting sequences, which affect the display algorithm rather than having extra bits for each character. Some terminals, such as the HP 2621, automatically leave standout mode when they move to a new line or the cursor is addressed. Programs using standout mode should exit standout mode before moving the cursor or sending a newline, unless the `msgr` capability, asserting that it is safe to move in standout mode, is present.

If the terminal has a way of flashing the screen to indicate an error quietly (a bell replacement) then this can be given as `flash`; it must not move the cursor.

If the cursor needs to be made more visible than normal when it is not on the bottom line (to make, for example, a non-blinking underline into an easier to find block or blinking underline) give this sequence as `cvvis`. If there is a way to make the cursor completely invisible, give that as `cvis`. The capability `cnorm` should be given which undoes the effects of both of these modes.

If the terminal needs to be in a special mode when running a program that uses these capabilities, the codes to enter and exit this mode can be given as `smcup` and `rmcup`. This arises, for example, from terminals like the Concept with more than one page of memory. If the terminal has only memory relative cursor addressing and not screen relative cursor addressing, a one screen-sized window must be fixed into the terminal for cursor addressing to work properly. This is also used for the TEKTRONIX 4025, where `smcup` sets the command character to be the one used by `terminfo`.

If your terminal correctly generates underlined characters (with no special codes needed) even though it does not overstrike, then you should give the capability `ul`. If overstrikes are erasable with a blank, then this should be indicated by giving `eo`.

### Keypad

If the terminal has a keypad that transmits codes when the keys are pressed, this information can be given. Note that it is not possible to handle terminals where the keypad only works in local (this applies, for example, to the unshifted HP 2621 keys). If the keypad can be set to transmit or not transmit, give these codes as `smkx` and `rmkx`. Otherwise the keypad is assumed to always transmit. The codes sent by the left arrow, right arrow, up arrow, down arrow, and home keys can be given as `kcub1`, `kcuf1`, `kcuu1`, `kcud1`, and `khome` respectively. If there are function keys such as `f0`, `f1`, ..., `f10`, the codes they send can be given as `kf0`, `kf1`, ..., `kf10`. If these keys have labels other than the default `f0` through `f10`, the labels can be given as `lf0`, `lf1`, ..., `lf10`. The codes transmitted by certain other special keys can be given: `kll` (home down), `kbs` (backspace), `ktbc` (clear all tabs), `kctab` (clear the tab stop in this column), `kclr` (clear screen or erase key), `kdch1` (delete character), `kdll` (delete line), `krmir` (exit insert mode), `kel` (clear to end of line), `ked` (clear to end of screen), `kich1` (insert character or enter insert mode), `kill` (insert line), `knp` (next page), `kpp` (previous page), `kind` (scroll forward/down), `kri` (scroll backward/up), `khts` (set a tab stop in this column). In addition, if the keypad has a 3 by 3 array of keys including the four arrow keys, the other five keys can be given as `ka1`, `ka3`, `kb2`, `kc1`, and `kc3`. These keys are useful when the effects of a 3 by 3 directional pad are needed.

### Tabs and Initialization

If the terminal has hardware tabs, the command to advance to the next tab stop can be given as `ht` (usually control I). A "backtab" command which moves leftward to the next tab stop can be given as `cbt`. By convention, if the teletype

modes indicate that tabs are being expanded by the computer rather than being sent to the terminal, programs should not use `ht` or `cbt` even if they are present, since the user may not have the tab stops properly set. If the terminal has hardware tabs which are initially set every *n* spaces when the terminal is powered up, the numeric parameter it is given, showing the number of spaces the tabs are set to. This is normally used by the `tset` command to determine whether to set the mode for hardware tab expansion, and whether to set the tab stops. If the terminal has tab stops that can be saved in nonvolatile memory, the terminfo description can assume that they are properly set.

Other capabilities include `is1`, `is2`, and `is3`, initialization strings for the terminal, `iprog`, the path name of a program to be run to initialize the terminal, and `if`, the name of a file containing long initialization strings. These strings are expected to set the terminal into modes consistent with the rest of the terminfo description. They are normally sent to the terminal, by the `tset` program, each time the user logs in. They will be printed in the following order: `is1`; `is2`; setting tabs using `tbc` and `hts`; `if`; running the program `iprog`; and finally `is3`. Most initialization is done with `is2`. Special terminal modes can be set up without duplicating strings by putting the common sequences in `is2` and special cases in `is1` and `is3`. A pair of sequences that does a harder reset from a totally unknown state can be analogously given as `rs1`, `rs2`, `rf`, and `rs3`, analogous to `is2` and `if`. These strings are output by the `reset` program, which is used when the terminal gets into a wedged state. Commands are normally placed in `rs2` and `rf` only if they produce annoying effects on the screen and are not necessary when logging in. For example, the command to set the vt100 into 80-column mode would normally be part of `is2`, but it causes an annoying glitch of the screen and is not normally needed since the terminal is usually already in 80 column mode.

If there are commands to set and clear tab stops, they can be given as `tbc` (clear all tab stops) and `hts` (set a tab stop in the current column of every row). If a more complex sequence is needed to set the tabs than can be described by this, the sequence can be placed in `is2` or `if`.

#### Delays

Certain capabilities control padding in the teletype driver. These are primarily needed by hard copy terminals, and are used by the `tset` program to set teletype modes appropriately. Delays embedded in the capabilities `cr`, `ind`, `cub1`, `ff`, and `tab` will cause the appropriate delay bits to be set in the teletype driver. If `pb` (padding baud rate) is given, these values can be ignored at baud rates below the value of `pb`.

### Miscellaneous

If the terminal requires other than a null (zero) character as a pad, then this can be given as `pad`. Only the first character of the pad string is used.

If the terminal has an extra "status line" that is not normally used by software, this fact can be indicated. If the status line is viewed as an extra line below the bottom line, into which one can cursor address normally (such as the Heathkit h19's 25th line, or the 24th line of a vt100 which is set to a 23-line scrolling region), the capability `hs` should be given. Special strings to go to the beginning of the status line and to return from the status line can be given as `tsl` and `fsl`. (`fsl` must leave the cursor position in the same place it was before `tsl`. If necessary, the `sc` and `rc` strings can be included in `tsl` and `fsl` to get this effect.) The parameter `tsl` takes one parameter, which is the column number of the status line the cursor is to be moved to. If escape sequences and other special commands, such as `tab`, work while in the status line, the flag `eslok` can be given. A string which turns off the status line (or otherwise erases its contents) should be given as `dsl`. If the terminal has commands to save and restore the position of the cursor, give them as `sc` and `rc`. The status line is normally assumed to be the same width as the rest of the screen, e.g., cols. If the status line is a different width (possibly because the terminal does not allow an entire line to be loaded) the width, in columns, can be indicated with the numeric parameter `wsl`.

If the terminal can move up or down half a line, this can be indicated with `hu` (half-line up) and `hd` (half-line down). This is primarily useful for superscripts and subscripts on hardcopy terminals. If a hardcopy terminal can eject to the next page (form feed), give this as `ff` (usually control L).

If there is a command to repeat a given character a given number of times (to save time transmitting a large number of identical characters) this can be indicated with the parameterized string `rep`. The first parameter is the character to be repeated and the second is the number of times to repeat it. Thus, `tparam(repeat_char, 'x', 10)` is the same as `'xxxxxxxxxx'`.

If the terminal has a settable command character, such as the TEKTRONIX 4025, this can be indicated with `cmdch`. A prototype command character is chosen which is used in all capabilities. This character is given in the `cmdch` capability to identify it. The following convention is supported on some UNIX systems: The environment is to be searched for a `CC` variable, and if found, all occurrences of the prototype character are replaced with the character in the environment variable.

Terminal descriptions that do not represent a specific kind of known terminal, such as *switch*, *dialup*, *patch*, and *network*, should include the **gn** (generic) capability so that programs can complain that they do not know how to talk to the terminal. (This capability does not apply to *virtual* terminal descriptions for which the escape sequences are known.)

If the terminal uses xon/xoff handshaking for flow control, give **xon**. Padding information should still be included so that routines can make better decisions about costs, but actual pad characters will not be transmitted.

If the terminal has a "meta key" which acts as a shift key, setting the 8th bit of any character transmitted, this fact can be indicated with **km**. Otherwise, software will assume that the 8th bit is parity and it will usually be cleared. If strings exist to turn this "meta mode" on and off, they can be given as **smm** and **rmm**.

If the terminal has more lines of memory than will fit on the screen at once, the number of lines of memory can be indicated with **lm**. A value of **lm#0** indicates that the number of lines is not fixed, but that there is still more memory than fits on the screen.

If the terminal is one of those supported by the UNIX virtual terminal protocol, the terminal number can be given as **vt**.

Media copy strings which control an auxiliary printer connected to the terminal can be given as **mc0**: print the contents of the screen, **mc4**: turn off the printer, and **mc5**: turn on the printer. When the printer is on, all text sent to the terminal will be sent to the printer. It is undefined whether the text is also displayed on the terminal screen when the printer is on. A variation **mc5p** takes one parameter, and leaves the printer on for as many characters as the value of the parameter, then turns the printer off. The parameter should not exceed 255. All text, including **mc4**, is transparently passed to the printer while an **mc5p** is in effect.

Strings to program function keys can be given as **pfkey**, **pfloc**, and **pfx**. Each of these strings takes two parameters: the function key number to program (from 0 to 10) and the string to program it with. Function key numbers out of this range may program undefined keys in a terminal dependent manner. The difference between the capabilities is that **pfkey** causes pressing the given key to be the same as the user typing the given string; **pfloc** causes the string to be executed by the terminal in local; and **pfx** causes the string to be transmitted to the computer.

**Glitches and Braindamage**

Hazeltine terminals, which do not allow "" characters to be displayed should indicate `hz`.

Terminals which ignore a linefeed immediately after an am wrap, such as the Concept and vt100, should indicate `xenl`.

If `el` is required to get rid of standout (instead of merely writing normal text on top of it), `xhp` should be given.

Teleray terminals, where tabs turn all characters moved over to blanks, should indicate `xt` (destructive tabs). This glitch is also taken to mean that it is not possible to position the cursor on top of a "magic cookie", that to erase standout mode it is instead necessary to use delete and insert line.

The Beehive Superbee, which is unable to correctly transmit the escape or control C characters, has `xsb`, indicating that the `f1` key is used for escape and `f2` for control C. (Only certain Superbees have this problem, depending on the ROM.)

Other specific terminal problems may be corrected by adding more capabilities of the form `xx`.

**Similar Terminals**

If there are two very similar terminals, one can be defined as being just like the other with certain exceptions. The string capability `use` can be given with the name of the similar terminal. The capabilities given before `use` override those in the terminal type invoked by `use`. A capability can be cancelled by placing `xx@` to the left of the capability definition, where `xx` is the capability. For example, the entry

```
2621-nl, smkx@, rmkx@, use=2621,
```

defines a `2621-nl` that does not have the `smkx` or `rmkx` capabilities, and hence does not turn on the function key labels when in visual mode. This is useful for different modes for a terminal, or for different user preferences.

**FILES**

`/usr/lib/terminfo/??*` files containing terminal descriptions

**SEE ALSO**

`curses(3X)`, `printf(3S)`, `term(5)`,  
`tic(1M)` in the *System Administrator Reference Manual*.

**NAME**

ttytype — data base of terminal types by port

**DESCRIPTION**

*Ttytype* is a database containing, for each tty port on the system, the kind of terminal that is attached to it. There is one line per port, containing the terminal kind (as a name listed in *termcap(5)*), a space, and the name of the tty, minus */dev/*.

This information is read by *tset(1)* and by *login(1)* to initialize the TERM environment variable at login time.

**EXAMPLE**

```
dw console
3a tty0
h19 tty1
h19 tty2
du ttyd0
```

**FILES**

*/etc/ttytype*

**SEE ALSO**

*tset(1)*, *login(1)*.



**NAME**

utmp, wtmp — utmp and wtmp entry formats

**SYNOPSIS**

```
#include <sys/types.h>
#include <utmp.h>
```

**DESCRIPTION**

These files, which hold user and accounting information for such commands as *who(1)*, *write(1)*, and *login(1)*, have the following structure as defined by *<utmp.h>*:

```
#define  UTMP_FILE  "/etc/utmp"
#define  WTMP_FILE  "/etc/wtmp"
#define  ut_name    ut_user

struct utmp {
    char    ut_user[8];        /* User login name */
    char    ut_id[4];         /* /etc/inittab id (usually line #) */
    char    ut_line[12];     /* device name (console, lxxx) */
    short   ut_pid;          /* process id */
    short   ut_type;         /* type of entry */
    struct  exit_status {
        short  e_termination; /* Process termination status */
        short  e_exit;        /* Process exit status */
    } ut_exit;               /* The exit status of a process
                             * marked as DEAD_PROCESS. */
    time_t  ut_time;        /* time entry was made */
    char    ut_host[16];    /* host name if remote */
};

/* Definitions for ut_type */
#define  EMPTY          0
#define  RUN_LVL        1
#define  BOOT_TIME     2
#define  OLD_TIME       3
#define  NEW_TIME       4
#define  INIT_PROCESS   5          /* Process spawned by "init" */
#define  LOGIN_PROCESS  6          /* A "getty" process waiting for login */
#define  USER_PROCESS   7          /* A user process */
#define  DEAD_PROCESS   8
#define  ACCOUNTING     9
#define  UTMXTYPE      ACCOUNTING /* Largest legal value of ut_type */

/* Special strings or formats used in the "ut_line" field when */
/* accounting for something other than a process. */
/* No string for the ut_line field can be more than 11 chars + */
/* a NULL in length. */

#define  RUN_LVL_MSG  "run—level %c"
#define  BOOT_MSG    "system boot"
#define  OLD_TIME_MSG "old time"
#define  NEW_TIME_MSG "new time"
```

**FILES**

/usr/include/utmp.h

UTMP(4)

UTMP(4)

*/etc/utmp*

*/etc/wtmp*

**SEE ALSO**

login(1), who(1), write(1), getut(3C).

**INTRO(5)**

**INTRO(5)**

**NAME**

intro -- introduction to miscellany

**DESCRIPTION**

This section describes miscellaneous facilities such as macro packages, character set tables, etc.

**NAME**

networking -- introduction to networking facilities

**SYNOPSIS**

```
#include <sys/socket.h>
#include <net/route.h>
#include <net/lf.h>
```

**DESCRIPTION**

This section briefly describes the networking facilities available in the system. Documentation in this part of section 5 is broken up into three areas: *protocol-families*, *protocols*, and *network interfaces*. Entries describing a protocol-family are marked "5F", while entries describing protocol use are marked "5P". Hardware support for network interfaces are found among the standard "5" entries.

All network protocols are associated with a specific *protocol-family*. A protocol-family provides basic services to the protocol implementation to allow it to function within a specific network environment. These services may include packet fragmentation and reassembly, routing, addressing, and basic transport. A protocol-family may support multiple methods of addressing, though the current protocol implementations do not. A protocol-family is normally comprised of a number of protocols, one per *socket(2N)* type. It is not required that a protocol-family support all socket types. A protocol-family may contain multiple protocols supporting the same socket abstraction.

A protocol supports one of the socket abstractions detailed in *socket(2N)*. A specific protocol may be accessed either by creating a socket of the appropriate type and protocol-family, or by requesting the protocol explicitly when creating a socket. Protocols normally accept only one type of address format, usually determined by the addressing structure inherent in the design of the protocol-family/network architecture. Certain semantics of the basic socket abstractions are protocol specific. All protocols are expected to support the basic model for their particular socket type, but may, in addition, provide non-standard facilities or extensions to a mechanism. For example, a protocol supporting the SOCK\_STREAM abstraction may allow more than one byte of out-of-band data to be transmitted per out-of-band message.

A network interface is similar to a device interface. Network interfaces comprise the lowest layer of the networking subsystem, interacting with the actual transport hardware. An interface may support one or more protocol families, and/or address formats.

**PROTOCOLS**

The system currently supports only the DARPA Internet protocols fully. Raw socket interfaces are provided to IP protocol layer of the DARPA Internet, to the IMP link layer (1822), and to Xerox PUP-I layer operating on top of 3Mb/s Ethernet interfaces. Consult the appropriate manual pages in this section for more information regarding the support for each protocol family.

**ADDRESSING**

Associated with each protocol family is an address format. The following address formats are used by the system:

```

#define AF_UNIX          1      /* local to host (pipes, portals) */
#define AF_INET          2      /* internetwork: UDP, TCP, etc. */
#define AF_IMPLINK      3      /* arpanet imp addresses */
#define AF_PUP           4      /* pup protocols: e.g. BSP */

```

## ROUTING

The network facilities provided limited packet routing. A simple set of data structures comprise a "routing table" used in selecting the appropriate network interface when transmitting packets. This table contains a single entry for each route to a specific network or host. A user process, the routing daemon, maintains this data base with the aid of two socket specific *ioctl(2)* commands, SIOCADDRT and SIOCDELRT. The commands allow the addition and deletion of a single routing table entry, respectively. Routing table manipulations may only be carried out by super-user.

A routing table entry has the following form, as defined in `<neth/route.h>`;

```

struct rentry {
    u_long      rt_hash;
    struct      sockaddr rt_dst;
    struct      sockaddr rt_gateway;
    short      rt_flags;
    short      rt_refcnt;
    u_long      rt_use;
    struct      ifnet *rt_ifp;
};

```

with *rt\_flags* defined from,

```

#define RTF_UP          0x1      /* route usable */
#define RTF_GATEWAY     0x2      /* destination is a gateway */
#define RTF_HOST        0x4      /* host entry (not otherwise) */

```

Routing table entries come in three flavors: for a specific host, for all hosts on a specific network, for any destination not matched by entries of the first two types (a wildcard route). When the system is booted, each network interface autoconfigured installs a routing table entry when it wishes to have packets sent through it. Normally the interface specifies the route through it is a "direct" connection to the destination host or network. If the route is direct, the transport layer of a protocol family usually requests the packet be sent to the same host specified in the packet. Otherwise, the interface may be requested to address the packet to an entity different from the eventual recipient (i.e. the packet is forwarded).

Routing table entries installed by a user process may not specify the hash, reference count, use, or interface fields; these are filled in by the routing routines. If a route is in use when it is deleted (*rt\_refcnt* is non-zero), the resources associated with it will not be reclaimed until further references to it are released.

The routing code returns EEXIST if requested to duplicate an existing entry, ESRCH if requested to delete a non-existent entry, or ENOBUFS if insufficient resources were available to install a new route.

User processes read the routing tables through the `/dev/kmem` device.

The *rt\_use* field contains the number of packets sent along the route. This value is used to select among multiple routes to the same destination. When multiple routes to the same destination exist, the least used route is selected.

A wildcard routing entry is specified with a zero destination address value. Wildcard routes are used only when the system fails to find a route to the destination host and network. The combination of wildcard routes and routing redirects can provide an economical mechanism for routing traffic.

## INTERFACES

Each network interface in a system corresponds to a path through which messages may be sent and received. A network interface usually has a hardware device associated with it, though certain interfaces such as the loopback interface, *lo(5)*, do not.

At boot time each interface which has underlying hardware support makes itself known to the system during the autoconfiguration process. Once the interface has acquired its address it is expected to install a routing table entry so that messages may be routed through it. Most interfaces require some part of their address specified with an *SIOCSIFADDR* *ioctl* before they will allow traffic to flow through them. On interfaces where the network-link layer address mapping is static, only the network number is taken from the *ioctl*; the remainder is found in a hardware specific manner. On interfaces which provide dynamic network-link layer address mapping facilities (e.g. 10Mb/s Ethernets), the entire address specified in the *ioctl* is used.

The following *ioctl* calls may be used to manipulate network interfaces. Unless specified otherwise, the request takes an *ifreq* structure as its parameter. This structure has the form

```
struct   ifreq {
    char   ifr_name[16];           /* name of interface (e.g. "cc0") */
    union {
        struct   sockaddr ifr_addr;
        struct   sockaddr ifr_dstaddr;
        short    ifr_flags;
    } ifr_ifru;
#define   ifr_addr   ifr_ifru.ifru_addr   /* address */
#define   ifr_dstaddr   ifr_ifru.ifru_dstaddr /* other end of p-to-p link */
#define   ifr_flags   ifr_ifru.ifru_flags /* flags */
};
```

### SIOCSIFADDR

Set interface address. Following the address assignment, the "initialization" routine for the interface is called.

### SIOCGIFADDR

Get interface address.

### SIOCSIFDSTADDR

Set point to point address for interface.

### SIOCGIFDSTADDR

Get point to point address for interface.

### SIOCSIFFLAGS

Set interface flags field. If the interface is marked down, any

processes currently routing packets through the interface are notified.

### SIOCGIFFLAGS

Get interface flags.

### SIOCGIFCONF

Get interface configuration list. This request takes an *ifconf* structure (see below) as a value-result parameter. The *ifc\_len* field should be initially set to the size of the buffer pointed to by *ifc\_buf*. On return it will contain the length, in bytes, of the configuration list.

```
/*
 * Structure used in SIOCGIFCONF request.
 * Used to retrieve interface configuration
 * for machine (useful for programs which
 * must know all networks accessible).
 */
struct ifconf {
    int ifc_len; /* size of associated buffer */
    union {
        caddr_t ifcu_buf;
        struct ifreq *ifcu_req;
    } ifc_ifcu;
#define ifc_buf ifc_ifcu.ifcu_buf /* buffer address */
#define ifc_req ifc_ifcu.ifcu_req /* array of structures returned */
};
```

### SEE ALSO

socket(2N), ioctl(2), routed(8N)

**NAME**

ascii — map of ASCII character set

**SYNOPSIS**

cat /usr/pub/ascii

**DESCRIPTION**

*Ascii* is a map of the ASCII character set, giving both octal and hexadecimal equivalents of each character, to be printed as needed. It contains:

000 nul	001 soh	002 stx	003 etx	004 eot	005 enq	006 ack	007 bel
010 bs	011 ht	012 nl	013 vt	014 np	015 cr	016 so	017 si
020 dle	021 dc1	022 dc2	023 dc3	024 dc4	025 nak	026 syn	027 etb
030 can	031 em	032 sub	033 esc	034 fs	035 gs	036 rs	037 us
040 sp	041 !	042 "	043 #	044 \$	045 %	046 &	047 '
050 (	051 )	052 *	053 +	054 ,	055 -	056 .	057 /
060 0	061 1	062 2	063 3	064 4	065 5	066 6	067 7
070 8	071 9	072 :	073 ;	074 <	075 =	076 >	077 ?
100 @	101 A	102 B	103 C	104 D	105 E	106 F	107 G
110 H	111 I	112 J	113 K	114 L	115 M	116 N	117 O
120 P	121 Q	122 R	123 S	124 T	125 U	126 V	127 W
130 X	131 Y	132 Z	133 [	134 \	135 ]	136 ^	137 _
140 `	141 a	142 b	143 c	144 d	145 e	146 f	147 g
150 h	151 i	152 j	153 k	154 l	155 m	156 n	157 o
160 p	161 q	162 r	163 s	164 t	165 u	166 v	167 w
170 x	171 y	172 z	173 {	174	175 }	176 ~	177 del

00 nul	01 soh	02 stx	03 etx	04 eot	05 enq	06 ack	07 bel
08 bs	09 ht	0a nl	0b vt	0c np	0d cr	0e so	0f si
10 dle	11 dc1	12 dc2	13 dc3	14 dc4	15 nak	16 syn	17 etb
18 can	19 em	1a sub	1b esc	1c fs	1d gs	1e rs	1f us
20 sp	21 !	22 "	23 #	24 \$	25 %	26 &	27 '
28 (	29 )	2a *	2b +	2c ,	2d -	2e .	2f /
30 0	31 1	32 2	33 3	34 4	35 5	36 6	37 7
38 8	39 9	3a :	3b ;	3c <	3d =	3e >	3f ?
40 @	41 A	42 B	43 C	44 D	45 E	46 F	47 G
48 H	49 I	4a J	4b K	4c L	4d M	4e N	4f O
50 P	51 Q	52 R	53 S	54 T	55 U	56 V	57 W
58 X	59 Y	5a Z	5b [	5c \	5d ]	5e ^	5f _
60 `	61 a	62 b	63 c	64 d	65 e	66 f	67 g
68 h	69 i	6a j	6b k	6c l	6d m	6e n	6f o
70 p	71 q	72 r	73 s	74 t	75 u	76 v	77 w
78 x	79 y	7a z	7b {	7c	7d }	7e ~	7f del

**FILES**

/usr/pub/ascii



**NAME**

arp -- Address Resolution Protocol

**DESCRIPTION**

ARP is a protocol used to dynamically map between DARPA Internet and 10Mb/s Ethernet addresses on a local area network. It is used by all the 10Mb/s Ethernet interface drivers and is not directly accessible to users.

ARP caches Internet-Ethernet address mappings. When an interface requests a mapping for an address not in the cache, ARP queues the message which requires the mapping and broadcasts a message on the associated network requesting the address mapping. If a response is provided, the new mapping is cached and any pending messages are transmitted. ARP itself is not Internet or Ethernet specific; this implementation, however, is. ARP will queue at most one packet while waiting for a mapping request to be responded to; only the most recently "transmitted" packet is kept.

ARP watches passively for hosts impersonating the local host (i.e. a host which responds to an ARP mapping request for the local host's address) and will, optionally, periodically probe a network looking for impostors.

**DIAGNOSTICS**

duplicate IP address!! sent from ethernet address: %x %x %x %x %x %x .  
ARP has discovered another host on the local network which responds to mapping requests for its own Internet address.

**NAME**

environ — user environment

**SYNOPSIS**

```
extern char **environ;
```

**DESCRIPTION**

An array of strings called the 'environment' is made available by *exec(2)* when a process begins. By convention these strings have the form '*name=value*'. The following names are used by various commands:

- PATH** The sequence of directory prefixes that *sh*, *time*, *nice(1)*, etc., apply in searching for a file known by an incomplete path name. The prefixes are separated by ':'. *Login(1)* sets:  
PATH=:/bin:/usr/bin.
- HOME** A user's login directory, set by *login(1)* from the password file *passwd(4)*.
- TERM** The kind of terminal for which output is to be prepared. This information is used by commands, such as *nroff*, *more*, or *vi*, which may exploit special terminal capabilities. See *techemcap* or *termcap(5)* for a list of terminal types.
- SHELL** The file name of the users login shell.
- TERMCAP** The string describing the terminal in **TERM**, or the name of the *termcap* file, see *termcap(5)*.
- EXINIT** A startup list of commands read by *ex(1)*, *edit(1)*, and *vi(1)*.
- LOGNAME** The login name of the user.
- TZ** Time zone information. The format is *xxx#zzz* where *xxx* is standard local time zone abbreviation, *n* is the difference in hours from GMT, and *zzz* is the abbreviation for the daylight-saving local time zone, if any; for example, **EST5EDT**.

Further names may be placed in the environment by the *export* command and '*name=value*' arguments in *sh(1)*, or by the *setenv* command if you use *csh(1)*. Arguments may also be placed in the environment at the point of an *exec(2)*. It is unwise to conflict with certain *sh(1)* variables that are frequently exported by ".profile" files: MAIL, PS1, PS2, IFS.

**SEE ALSO**

*csh(1)*, *ex(1)*, *login(1)*, *sh(1)*, *exec(2)*, *system(3S)*, *termcap(5)*, *tty(7)*.

**NAME**

*eqnchar* — special character definitions for *eqn* and *neqn*

**SYNOPSIS**

*eqn* /usr/pub/*eqnchar* [ files ] | *troff* [ options ]

*neqn* /usr/pub/*eqnchar* [ files ] | *nroff* [ options ]

**DESCRIPTION**

*Eqnchar* contains *troff* and *nroff* character definitions for constructing characters that are not available on the Wang Laboratories, Inc. C/A/T phototypesetter. These definitions are primarily intended for use with *eqn* and *neqn*; *eqnchar* contains definitions for the following characters:

<i>ciplus</i>	⊕			<i>square</i>	□
<i>citimes</i>	⊗	<i>langle</i>	{	<i>circle</i>	○
<i>wig</i>	~	<i>rangle</i>	}	<i>blot</i>	◻
<i>-wig</i>	≅	<i>hbar</i>	ℏ	<i>bullet</i>	•
<i>&gt;wig</i>	≳	<i>ppd</i>	⊥	<i>prop</i>	∝
<i>&lt;wig</i>	≲	<i>&lt;-&gt;</i>	↔	<i>empty</i>	∅
<i>=wig</i>	≡	<i>&lt;=&gt;</i>	↔	<i>member</i>	∈
<i>star</i>	*	<	⊲	<i>nomem</i>	∉
<i>bigstar</i>	*	>	⊳	<i>cup</i>	∪
<i>=dot</i>	⋮	<i>ang</i>	∟	<i>cap</i>	∩
<i>orsign</i>	∨	<i>rang</i>	└	<i>incl</i>	⊆
<i>andsign</i>	∧	<i>3dot</i>	⋮	<i>subset</i>	⊂
<i>=del</i>	⊖	<i>thf</i>	∴	<i>supset</i>	⊃
<i>oppA</i>	∇	<i>quarter</i>	¼	<i>lsubset</i>	⊆
<i>oppE</i>	∃	<i>3quarter</i>	¾	<i>lsupset</i>	⊇
<i>angstrom</i>	Å	<i>degree</i>	°	<i>scrL</i>	<i>scrL</i>
<i>==&lt;</i>	--<	<i>--&gt;</i>	-->		

**FILES**

/usr/pub/*eqnchar*

**SEE ALSO**

*eqn*(1), *nroff*(1), *troff*(1).

**NAME**

fcntl — file control options

**SYNOPSIS**

#include &lt;fcntl.h&gt;

**DESCRIPTION**

The *fcntl(2)* function provides for control over open files. The *include* file describes *requests* and *arguments* to *fcntl* and *open(2)*.

```

/* Flag values accessible to open(2) and fcntl(2) */
/* (The first three can only be set by open) */
#define O_RDONLY 0
#define O_WRONLY 1
#define O_RDWR 2
#define O_NDELAY 04 /* Non-blocking I/O */
#define O_APPEND 010 /* append (writes guaranteed at the end) */

/* Flag values accessible only to open(2) */
#define O_CREAT 00400 /* open with file create (uses third open arg) */
#define O_TRUNC 01000 /* open with truncation */
#define O_EXCL 02000 /* exclusive open */

/* fcntl(2) requests */
#define F_DUPFD 0 /* Duplicate files */
#define F_GETFD 1 /* Get files flags */
#define F_SETFD 2 /* Set files flags */
#define F_GETFL 3 /* Get file flags */
#define F_SETFL 4 /* Set file flags */
#define F_GETLK 5 /* Get blocking file locks */
#define F_SETLK 6 /* Set or clear file locks and fail on busy */
#define F_SETLKW 7 /* Set or clear file locks and wait on busy */
#define F_GETOWN 8 /* Get owner */
#define F_SETOWN 9 /* Set owner */

/* file segment locking control structure */
struct flock {
short l_type;
short l_whence;
long l_start;
long l_len; /* if 0 then until EOF */
int l_pid; /* returned with F_GETLK */

/* file segment locking types */
#define F_RDLCK 01 /* Read lock */
#define F_WRLCK 02 /* Write lock */
#define F_UNLCK 03 /* Remove locks */

```

**SEE ALSO**

fcntl(2), open(2).

**NAME**

`greek` — graphics for the extended TTY-37 type-box

**SYNOPSIS**

`cat /usr/pub/greek [ | greek -Tterminal ]`

**DESCRIPTION**

`Greek` gives the mapping from ASCII to the "shift-out" graphics in effect between SO and SI on TELETYPE® Model 37 terminals equipped with a 128-character type-box. These are the default greek characters produced by `nroff`. The filters of `greek(1)` attempt to print them on various other terminals. The file contains:

alpha	$\alpha$	A	beta	$\beta$	B	gamma	$\gamma$	\
GAMMA	$\Gamma$	G	delta	$\delta$	D	DELTA	$\Delta$	W
epsilon	$\epsilon$	S	zeta	$\zeta$	Q	eta	$\eta$	N
THETA	$\Theta$	T	theta	$\theta$	O	lambda	$\lambda$	L
LAMBDA	$\Lambda$	E	mu	$\mu$	M	nu	$\nu$	@
xi	$\xi$	X	pi	$\pi$	J	PI	$\Pi$	P
rho	$\rho$	K	sigma	$\sigma$	Y	SIGMA	$\Sigma$	R
tau	$\tau$	I	phi	$\phi$	U	PHI	$\Phi$	F
psi	$\psi$	V	PSI	$\Psi$	H	omega	$\omega$	C
OMEGA	$\Omega$	Z	nabla	$\nabla$	[	not	$\neg$	-
partial	$\partial$	]	integral	$\int$	^			

**FILES**

`/usr/pub/greek`

**SEE ALSO**

`300(1)`, `4014(1)`, `450(1)`, `greek(1)`, `tc(1)`, `nroff(1)`.

**NAME**

inet — Internet protocol family

**SYNOPSIS**

```
#include <sys/types.h>
#include <netinet/in.h>
```

**DESCRIPTION**

The Internet protocol family is a collection of protocols layered atop the *Internet Protocol* (IP) transport layer, and utilizing the Internet address format. The Internet family provides protocol support for the SOCK\_STREAM, SOCK\_DGRAM, and SOCK\_RAW socket types; the SOCK\_RAW interface provides access to the IP protocol.

**ADDRESSING**

Internet addresses are four byte quantities, stored in network standard format (on the VAX these are word and byte reversed). The include file `<netinet/in.h>` defines this address as a discriminated union.

Sockets bound to the Internet protocol family utilize the following addressing structure,

```
struct sockaddr_in {
    short      sin_family;
    u_short    sin_port;
    struct     in_addr sin_addr;
    char       sin_zero[8];
};
```

Sockets may be created with the address INADDR\_ANY to effect wildcard matching on incoming messages.

**PROTOCOLS**

The Internet protocol family is comprised of the IP transport protocol, Internet Control Message Protocol (ICMP), Transmission Control Protocol (TCP), and User Datagram Protocol (UDP). TCP is used to support the SOCK\_STREAM abstraction while UDP is used to support the SOCK\_DGRAM abstraction. A raw interface to IP is available by creating an Internet socket of type SOCK\_RAW. The ICMP message protocol is not directly accessible.

**SEE ALSO**

tcp(5P), udp(5P), ip(5P)

**CAVEAT**

The Internet protocol support is subject to change as the Internet protocols develop. Users should not depend on details of the current implementation, but rather the services exported.

**NAME**

ip — Internet Protocol

**SYNOPSIS**

```
#include <sys/socket.h>
#include <netinet/in.h>

s = socket(AF_INET, SOCK_RAW, 0);
```

**DESCRIPTION**

IP is the transport layer protocol used by the Internet protocol family. It may be accessed through a raw socket when developing new protocols, or special purpose applications. IP sockets are connectionless, and are normally used with the *sendto* and *recvfrom* calls, though the *connect(2N)* call may also be used to fix the destination for future packets (in which case the *read(2)* or *recv(2N)* and *write(3)* or *send(2N)* system calls may be used).

Outgoing packets automatically have an IP header prepended to them (based on the destination address and the protocol number the socket is created with). Likewise, incoming packets have their IP header stripped before being sent to the user.

**DIAGNOSTICS**

A socket operation may fail with one of the following errors returned:

- [EISCONN] when trying to establish a connection on a socket which already has one, or when trying to send a datagram with the destination address specified and the socket is already connected;
- [ENOTCONN] when trying to send a datagram, but no destination address is specified, and the socket hasn't been connected;
- [ENOBUFS] when the system runs out of memory for an internal data structure;
- [EADDRNOTAVAIL] when an attempt is made to create a socket with a network address for which no network interface exists.

**SEE ALSO**

*send(2N)*, *recv(2N)*, *intro(5N)*, *inet(5F)*

**BUGS**

- One should be able to send and receive ip options.
- The protocol should be settable after socket creation.

**NAME**

lo — software loopback network interface

**SYNOPSIS**

**pseudo-device loop**

**DESCRIPTION**

The *loop* interface is a software loopback mechanism which may be used for performance analysis, software testing, and/or local communication. By default, the loopback interface is accessible at address 127.0.0.1 (non-standard); this address may be changed with the SIOCSIFADDR ioctl.

**DIAGNOSTICS**

**lo%d: can't handle af%d.** The interface was handed a message with addresses formatted in an unsuitable address family; the packet was dropped.

**SEE ALSO**

intro(5N), inet(5F)

**BUGS**

It should handle all address and protocol families. An approved network address should be reserved for this interface.



## NAME

`man` - macros for formatting entries in this manual

## SYNOPSIS

`nroff -man files`

`troff -man [ -rs1 ] files`

## DESCRIPTION

These *troff*(1) macros are used to lay out the format of the entries of this manual. These macros are used by the *man*(1) command.

The default page size is 8.5"×11", with a 6.5"×10" text area; the `-rs1` option reduces these dimensions to 6"×9" and 4.75"×8.375", respectively; this option (which is *not* effective in *nroff*) also reduces the default type size from 10-point to 9-point, and the vertical line spacing from 12-point to 10-point. The `-rV2` option may be used to set certain parameters to values appropriate for certain Versatec printers: it sets the line length to 82 characters, the page length to 84 lines, and it inhibits underlining; this option should not be confused with the `-Tvp` option of the *man*(1) command, which is available at some UNIX System sites.

Any *text* argument below may be one to six "words". Double quotes (") may be used to include blanks in a "word". If *text* is empty, the special treatment is applied to the next line that contains text to be printed. For example, `I` may be used to italicize a whole line, or `.SM` followed by `.B` to make small bold text. By default, hyphenation is turned off for *nroff*, but remains on for *troff*.

Type font and size are reset to default values before each paragraph and after processing font- and size-setting macros, e.g., `.I`, `.RB`, `.SM`. Tab stops are neither used nor set by any macro except `.DT` and `.TH`.

Default units for indents *in* are ens. When *in* is omitted, the previous indent is used. This remembered indent is set to its default value (7.2 ens in *troff*, 5 ens in *nroff*—this corresponds to 0.5" in the default page size) by `.TH`, `.P`, and `.RS`, and restored by `.RE`.

`.TH t s c n` Set the title and entry heading; *t* is the title, *s* is the section number, *c* is extra commentary, e.g., "local", *n* is new manual name. Invokes `.DT` (see below).

`.SH text` Place subhead *text*, e.g., SYNOPSIS, here.

`.SS text` Place sub-subhead *text*, e.g., Options, here.

`.B text` Make *text* bold.

**.I** *text*      Make *text* italic.  
**.SM** *text*      Make *text* 1 point smaller than default point size.  
**.RI** *a b*        Concatenate roman *a* with italic *b*, and alternate these two fonts for up to six arguments. Similar macros alternate between any two of roman, italic, and bold:

**.IR .RB .BR .IB .BI**

**.P**

Begin a paragraph with normal font, point size, and indent. **.PP** is a synonym for **.P**.

**.HP** *in*

Begin paragraph with hanging indent.

**.TP** *in*

Begin indented paragraph with hanging tag. The next line that contains text to be printed is taken as the tag. If the tag does not fit, it is printed on a separate line.

**.JP** *t in*

Same as **.TP** *in* with tag *t*; often used to get an indented paragraph without a tag.

**.RS** *in*

Increase relative indent (initially zero). Indent all output an extra *in* units from the current left margin.

**.RE** *k*

Return to the *k*th relative indent level (initially, *k*=1; *k*=0 is equivalent to *k*=1); if *k* is omitted, return to the most recent lower indent level.

**.PM** *m*

Produces proprietary markings; where *m* may be P for PRIVATE, N for NOTICE, BP for BELL LABORATORIES PROPRIETARY, or BR for BELL LABORATORIES RESTRICTED.

**.DT**

Restore default tab settings (every 7.2 ens in *troff*, 5 ens in *nroff*).

**.PD** *v*

Set the interparagraph distance to *v* vertical spaces. If *v* is omitted, set the interparagraph distance to the default value (0.4*v* in *troff*, 1*v* in *nroff*).

The following *strings* are defined:

**\\*R**            © in *troff*, (Reg.) in *nroff*.  
**\\*S**            Change to default type size.  
**\\*(Tm**         Trademark indicator.

The following *number registers* are given default values by **.TH**:

**IN** Left margin indent relative to subheads (default is 7.2 ens in *troff*, 5 ens in *nroff*).

**LL** Line length including **IN**.

**PD** Current interparagraph distance.

**CAVEATS**

In addition to the macros, strings, and number registers mentioned above, there are defined a number of *internal* macros, strings, and number registers. Except for names predefined by *troff* and number registers *d*, *m*, and *y*, all such internal names are of the form *XA*, where *X* is one of *(*, *)*, *]*, and *}*, and *A* stands for any alphanumeric character.

If a manual entry needs to be preprocessed by *cw*(1), *eqn*(1) (or *neqn*), and/or *tbl*(1), it must begin with a special line (described in *man*(1)), causing the *man* command to invoke the appropriate preprocessor(s).

The programs that prepare the Table of Contents and the Permuted Index for this Manual assume the *NAME* section of each entry consists of a single line of input that has the following format:

```
name[, name, name ...] \- explanatory text
```

The macro package increases the inter-word spaces (to eliminate ambiguity) in the *SYNOPSIS* section of each entry.

The macro package itself uses only the roman font (so that one can replace, for example, the bold font by the constant-width font—see *cw*(1)). Of course, if the input text of an entry contains requests for other fonts (e.g., *I*, *.RB*, *VI*), the corresponding fonts must be mounted.

**EXAMPLE**

```
nroff -man man.5
```

to *nroff* this manual section.

**FILES**

```
/usr/lib/tmac/tmac.an
/usr/lib/macros/cmp.[nt].[dt].an
/usr/lib/macros/ucmp.[nt].an
```

**SEE ALSO**

man(1), mroff(1), troff(1).

**BUGS**

If the argument to .TH contains *any* blanks and is *not* enclosed by double quotes (""), there will be bird-dropping-like things on the output.

**NAME**

math — math functions and constants

**SYNOPSIS**

#include &lt;math.h&gt;

**DESCRIPTION**

This file contains declarations of all the functions in the Math Library (described in Section 3M), as well as various functions in the C Library (Section 3C) that return floating-point values.

It defines the structure and constants used by the *matherr*(3M) error-handling mechanisms, including the following constant used as an error-return value:

**HUGE**                   The maximum value of a single-precision floating-point number.

The following mathematical constants are defined for user convenience:

**M\_E**                    The base of natural logarithms (*e*).

**M\_LOG2E**               The base-2 logarithm of *e*.

**M\_LOG10E**              The base-10 logarithm of *e*.

**M\_LN2**                  The natural logarithm of 2.

**M\_LN10**                 The natural logarithm of 10.

**M\_PI**                   The ratio of the circumference of a circle to its diameter. (There are also several fractions of its reciprocal and its square root.)

**M\_SQRT2**                The positive square root of 2.

**M\_SQRT1\_2**             The positive square root of 1/2.

For the definitions of various machine-dependent "constants," see the description of the <*values.h*> header file.

**FILES**

/usr/include/math.h

**SEE ALSO**intro(3), *matherr*(3M), *values*(5).

**NAME**

**mm** — the MM macro package for formatting documents

**SYNOPSIS**

```

mm [ options ] [ files ]
nroff -mm [ options ] [ files ]
nroff -cm [ options ] [ files ]

mmt [ options ] [ files ]
troff -mm [ options ] [ files ]
troff -cm [ options ] [ files ]

```

**DESCRIPTION**

This package provides a formatting capability for a very wide variety of documents. It is the standard package used by the BTL typing pools and documentation centers. The manner in which a document is typed in and edited is essentially independent of whether the document is to be eventually formatted at a terminal or is to be phototypeset. See the references below for further details.

The **-mm** option causes *nroff* and *troff*(1) to use the non-compacted version of the macro package, while the **-cm** option results in the use of the compacted version, thus speeding up the process of loading the macro package.

**FILES**

<code>/usr/lib/tmac/tmac.m</code>	pointer to the non-compacted version of the package
<code>/usr/lib/macros/mm[nt]</code>	non-compacted version of the package
<code>/usr/lib/macros/cmp.[nt].[dt].m</code>	compacted version of the package
<code>/usr/lib/macros/ucmp.[nt].m</code>	initializers for the compacted version of the package

**SEE ALSO**

`mm`(1), `mmt`(1), `nroff`(1), `troff`(1).  
**MM** in the *Document Processing Guide*.

**NAME**

mosd — the OSDD adapter macro package for formatting documents

**SYNOPSIS**

```
osdd [ options ] [ files ]
mm -mosd [ options ] [ files ]
nroff -mm -mosd [ options ] [ files ]
nroff -cm -mosd [ options ] [ files ]

mmt -mosd [ options ] [ files ]
troff -mm -mosd [ options ] [ files ]
troff -cm -mosd [ options ] [ files ]
```

**DESCRIPTION**

The OSDD adapter macro package is a tool used in conjunction with the MM macro package to prepare Operations Systems Deliverable Documentation. Many of the OSDD Standards are different than the default format provided by MM. The OSDD adapter package sets the appropriate MM options for automatic production of the OSDD Standards. The OSDD adapter package also generates the correct OSDD page headers and footers, heading styles, Table of Contents format, etc.

OSDD document (input) files are prepared with the MM macros. Additional information which must be given at the beginning of the document file is specified by the following string definitions:

```
.ds H1 document-number
.ds H2 section-number
.ds H3 issue-number
.ds H4 date
.ds H5 rating
```

The *document-number* should be of the standard 10 character format. The words "Section" and "Issue" should not be included in the string definitions; they will be supplied automatically when the document is printed. For example:

```
.ds H1 OPA-1P135-01
.ds H2 4
.ds H3 2
```

automatically produces  
OPA-1P135-01  
Section 4  
Issue 2

as the document page header. Quotation marks are not used in string definitions.

If certain information is not to be included in a page header, then the string is defined as null; e.g.,

```
.ds H2
```

means that there is no *section-number*.

The OSDD Standards require that the *Table of Contents* be numbered beginning with *Page 1*. By default, the first page of text will be numbered *Page 2*. If the *Table of Contents* has more than one page, for example *n*, then either `-rPn+1` must be included as a command line option or `.nr P n` must be included in the document file. For example, if the *Table of*

*Contents* is four pages then use `-rP5` on the command line or `.nr P 4` in the document file.

The OSDD Standards require that certain information such as the document rating appear on the *Document Index* or on the *Table of Contents* page if there is no index. By default, it is assumed that an index has been prepared separately. If there is no index, the following must be included in the document file:

```
.nr Di 0
```

This will ensure that the necessary information is included on the *Table of Contents* page.

The OSDD Standards require that all numbered figures be placed at the end of the document. The `.Fg` macro is used to produce full page figures. This macro produces a blank page with the appropriate header, footer, and figure caption. Insertion of the actual figure on the page is a manual operation. The macro usage is

```
.Fg page-count "figure caption"
```

where *page-count* is the number of pages required for a multi-page figure (default 1 page).

Figure captions are produced by the `.Fg` macro using the `.BS/.BE` macros. Thus the `.BS/.BE` macros are also not available for users. The `.Fg` macro cannot be used within the document unless the final `.Fg` in a series of figures is followed by a `.SK` macro to force out the last figure page.

The *Table of Contents* for OSDD documents (see Figure 4 in Section 4.1 of the OSDD Standards) is produced with:

```
.Tc
System Type
System Name
Document Type
.Td
```

The `.Tc/.Td` macros are used instead of the `.TC` macro from MM.

By default, the adapter package causes the NOTICE disclosure statement to be printed. The `.PM` macro may be used to suppress the NOTICE or to replace it with the PRIVATE disclosure statement as follows:

```
.PM none printed
.PM P PRIVATE printed
.PM N NOTICE printed (default)
```

The `.P` macro is used for paragraphs. The `Np` register is set automatically to indicate the paragraph numbering style. It is very important that the `.P` macro be used correctly. All paragraphs (including those immediately following a `.H` macro) must use a `.P` macro. Unless there is a `.P` macro, there will not be a number generated for the paragraph. Similarly, the `.P` macro should not be used for text which is not a paragraph. The `.SP` macro may be appropriate for these cases, e.g., for "paragraphs" within a list item.

The page header format is produced automatically in accordance with the OSDD Standards. The OSDD Adapter macro package uses the `.TP` macro for this purpose. Therefore the `.TP` macro normally available in MM is not available for users.

## FILES

`/usr/lib/tmac/tmac.osd`



**SEE ALSO**

**mm(1), mmt(1), nroff(1), troff(1), mm(5).**  
**MM** in the *Document Processing Guide*.

**NAME**

mptx — the macro package for formatting a permuted index

**SYNOPSIS**

**nroff** -mptx [ options ] [ files ]

**troff** -mptx [ options ] [ files ]

**DESCRIPTION**

This package provides a definition for the .xx macro used for formatting a permuted index as produced by *ptx*(1). This package does not provide any other formatting capabilities such as headers and footers. If these or other capabilities are required, the *mptx* macro package may be used in conjunction with the *MM* macro package. In this case, the -mptx option must be invoked *after* the -mm call. For example:

nroff -cm -mptx file

or

mm -mptx file

**FILES**

/usr/lib/tmac/tmac.ptx pointer to the non-compacted version of the package

/usr/lib/macros/ptx non-compacted version of the package

**SEE ALSO**

mm(1), nroff(1), ptx(1), troff(1), mm(5).

**NAME**

**mv** — a troff macro package for typesetting view graphs and slides

**SYNOPSIS**

**mv** [ -a ] [ options ] [ files ]

**troff** [ -a ] [ -rX1 ] -mv { options } [ files ]

**DESCRIPTION**

This package makes it easy to typeset view graphs and projection slides in a variety of sizes. A few macros (briefly described below) accomplish most of the formatting tasks needed in making transparencies. All of the facilities of *troff*(1), *cw*(1), *eqn*(1), and *tbl*(1) are available for more difficult tasks.

The output can be previewed on most terminals, and, in particular, on the Tektronix 4014, as well as on the Versatec printer. For these two devices, specify the -rX1 option (this option is automatically specified by the *mv* command -q.v.- when that command is invoked with the -T4014 or -Tvp options). To preview output on other terminals, specify the -a option.

The available macros are:

**.VS** [n] [i] [d] Foil-start macro; foil size is to be 7"×7"; n is the foil number, i is the foil identification, d is the date; the foil-start macro resets all parameters (indent, point size, etc.) to initial default values, except for the values of i and d arguments inherited from a previous foil-start macro; it also invokes the .A macro (see below).

The naming convention for this and the following eight macros is that the first character of the name (V or S) distinguishes between view graphs and slides, respectively, while the second character indicates whether the foil is square (S), small wide (w), small high (h), big wide (W), or big high (H). Slides are "skinnier" than the corresponding view graphs: the ratio of the longer dimension to the shorter one is larger for slides than for view graphs. As a result, slide foils can be used for view graphs, but not vice versa; on the other hand, view graphs can accommodate a bit more text.

**.Vw** [n] [i] [d] Same as .VS, except that foil size is 7" wide × 5" high.  
**.Vh** [n] [i] [d] Same as .VS, except that foil size is 5"×7".  
**.VW** [n] [i] [d] Same as .VS, except that foil size is 7"×5.4".  
**.VH** [n] [i] [d] Same as .VS, except that foil size is 7"×9".  
**.Sw** [n] [i] [d] Same as .VS, except that foil size is 7"×5".  
**.Sh** [n] [i] [d] Same as .VS, except that foil size is 5"×7".  
**.SW** [n] [i] [d] Same as .VS, except that foil size is 7"×5.4".  
**.SH** [n] [i] [d] Same as .VS, except that foil size is 7"×9".  
**.A** [x] Place text that follows at the first indentation level (left margin); the presence of x suppresses the ½ line spacing from the preceding text.

**.B** [m [s] ] Place text that follows at the second indentation level; text is preceded by a mark; m is the mark (default is a large bullet); s is the increment or decrement to the point size of the mark with respect to the prevailing

- point size (default is 0); if *s* is 100, it causes the point size of the mark to be the same as that of the *default* mark.
- .C** [*m* [*s*]] Same as **.B**, but for the third indentation level; default mark is a dash.
- .D** [*m* [*s*]] Same as **.B**, but for the fourth indentation level; default mark is a small bullet.
- .T** *string* *String* is printed as an over-size, centered title.
- .I** [*in*] [*a* [*x*]] Change the current text indent (does not affect titles); *in* is the indent (in inches unless dimensioned, default is 0); if *in* is signed, it is an increment or decrement; the presence of *a* invokes the **.A** macro (see below) and passes *x* (if any) to it.
- .S** [*p*] [*l*] Set the point size and line length; *p* is the point size (default is "previous"); if *p* is 100, the point size reverts to the *initial* default for the current foil-start macro; if *p* is signed, it is an increment or decrement (default is 18 for **.VS**, **.VH**, and **.SH**, and 14 for the other foil-start macros); *l* is the line length (in inches unless dimensioned; default is 4.2" for **.Vh**, 3.8" for **.Sh**, 5" for **.SH**, and 6" for the other foil-start macros).
- .DF** *n f* [*n f*...] Define font positions; may not appear within a foil's input text (i.e., it may only appear after all the input text for a foil, but before the next foil-start macro); *n* is the position of font *f*; up to four "*n f*" pairs may be specified; the first font named becomes the *prevailing* font; the initial setting is (H is a synonym for G):  
**.DF** 1 H 2 I 3 B 4 S
- .DV** [*a*] [*b*] [*c*] [*d*] Alter the vertical spacing between indentation levels; *a* is the spacing for **.A**, *b* is for **.B**, *c* is for **.C**, and *d* is for **.D**; all non-null arguments must be dimensioned; null arguments leave the corresponding spacing unaffected; initial setting is:  
**.DV** .5v .5v .5v 0v
- .U** *str1* [*str2*] Underline *str1* and concatenate *str2* (if any) to it.

The last four macros in the above list do not cause a break; the **.I** macro causes a break only if it is invoked with more than one argument; all the other macros cause a break.

The macro package also recognizes the following upper-case synonyms for the corresponding lower-case *traff* requests:

**.AD .BR .CE .FI .HY .NA .NF .NH .NX .SO .SP .TA .TI**

The **Tm** string produces the trademark symbol.

The input tilde (~) character is translated into a blank on output.

## FILES

/usr/lib/tmac/tmac.v  
 /usr/lib/macros/vmca

## SEE ALSO

cw(1), eqn(1), mmt(1), tbl(1), troff(1).

## BUGS

The **.VW** and **.SW** foils are meant to be 9" wide by 7" high, but because

the typesetter paper is generally only 8" wide, they are printed 7" wide by 5.4" high and have to be enlarged by a factor of 9/7 before use as view graphs; this makes them less than totally useful.

**NAME**

prof — profile within a function

**SYNOPSIS**

```
#define MARK
#include <prof.h>
void MARK (name)
```

**DESCRIPTION**

*MARK* will introduce a mark called *name* that will be treated the same as a function entry point. Execution of the mark will add to a counter for that mark, and program-counter time spent will be accounted to the immediately preceding mark or to the function if there are no preceding marks within the active function.

*Name* may be any combination of up to six letters, numbers or underscores. Each *name* in a single compilation must be unique, but may be the same as any ordinary program symbol.

For marks to be effective, the symbol *MARK* must be defined before the header file *<prof.h>* is included. This may be defined by a preprocessor directive as in the synopsis, or by a command line argument, i.e:

```
cc -p -DMARK foo.c
```

If *MARK* is not defined, the *MARK(name)* statements may be left in the source files containing them and will be ignored.

**EXAMPLE**

In this example, marks can be used to determine how much time is spent in each loop. Unless this example is compiled with *MARK* defined on the command line, the marks are ignored.

```
#include <prof.h>

foo()
{
    int i, j;
    .
    .
    .
    MARK(loop1);
    for (i = 0; i < 2000; i++) {
        . . .
    }
    MARK(loop2);
    for (j = 0; j < 2000; j++) {
        . . .
    }
}
```

**SEE ALSO**

prof(1), profil(2), monitor(3C).

**NAME**

pty — pseudo terminal driver

**DESCRIPTION**

The *pty* driver provides support for a device-pair termed a *pseudo terminal*. A pseudo terminal is a pair of character devices, a *master* device and a *slave* device. The slave device provides processes an interface identical to that described in *termio(7)*. However, whereas all other devices which provide the interface described in *termio(7)* have a hardware device of some sort behind them, the slave device has, instead, another process manipulating it through the master half of the pseudo terminal. That is, anything written on the master device is given to the slave device as input and anything written on the slave device is presented as input on the master device.

The following *ioctl* calls apply only to pseudo terminals:

**TIOCSTOP**

Stops output to a terminal (e.g. like typing ^S). Takes no parameter.

**TIOCSTART**

Restarts output (stopped by TIOCSTOP or by typing ^S). Takes no parameter.

**TIOCPKT**

Enable/disable *packet* mode. Packet mode is enabled by specifying (by reference) a nonzero parameter and disabled by specifying (by reference) a zero parameter. When applied to the master side of a pseudo terminal, each subsequent *read* from the terminal will return data written on the slave part of the pseudo terminal preceded by a zero byte (symbolically defined as TIOCPKT\_DATA), or a single byte reflecting control status information. In the latter case, the byte is an inclusive-or of zero or more of the bits:

**TIOCPKT\_FLUSHREAD**

whenever the read queue for the terminal is flushed.

**TIOCPKT\_FLUSHWRITE**

whenever the write queue for the terminal is flushed.

**TIOCPKT\_STOP**

whenever output to the terminal is stopped a la ^S.

**TIOCPKT\_START**

whenever output to the terminal is restarted.

**TIOCPKT\_DOSTOP**

whenever *t\_stop* is ^S and *t\_startc* is ^Q.

**TIOCPKT\_NOSTOP**

whenever the start and stop characters are not ^S/^Q.

This mode is used by *rlogin(1N)* and *rlogind(8N)* to implement a remote-echoed, locally ^S/^Q flow-controlled remote login with proper back-flushing of output; it can be used by other similar programs.

**TIOCREMOTE**

A mode for the master half of a pseudo terminal, independent of TIOCPKT. This mode causes input to the pseudo terminal to be flow controlled and not input edited (regardless of the terminal

mode). Each write to the control terminal produces a record boundary for the process reading the terminal. In normal usage, a write of data is like the data typed as a line on the terminal; a write of 0 bytes is like typing an end-of-file character. TIOCREMOTE can be used when doing remote line editing in a window manager, or whenever flow controlled input is required.

**FILES**

/dev/pty[p-r][0-9a-f] master pseudo terminals  
/dev/tty[p-r][0-9a-f] slave pseudo terminals

**DIAGNOSTICS**

None.

**BUGS**

It is not possible to send an EOT.



## NAME

regexp — regular expression compile and match routines

## SYNOPSIS

```
#define INIT <declarations>
#define GETC() <getc code>
#define PEEKC() <peekc code>
#define UNGETC(c) <ungetc code>
#define RETURN(pointer) <return code>
#define ERROR(val) <error code>
#include <regexp.h>

char *compile (instring, expbuf, endbuf, eof)
char *instring, *expbuf, *endbuf;
int eof;

int step (string, expbuf)
char *string, *expbuf;

extern char *loc1, *loc2, *locs;
extern int circf, sed, nbra;
```

## DESCRIPTION

This page describes general-purpose regular expression matching routines in the form of *ed(1)*, defined in */usr/include/regexp.h*. Programs such as *ed(1)*, *sed(1)*, *grep(1)*, *bs(1)*, *expr(1)*, etc., which perform regular expression matching use this source file. In this way, only this file need be changed to maintain regular expression compatibility.

The interface to this file is unpleasantly complex. Programs that include this file must have the following five macros declared before the "#include <regexp.h>" statement. These macros are used by the *compile* routine.

- |                          |   |
|--------------------------|---|
| GETC()                   | Return the value of the next character in the regular expression pattern. Successive calls to GETC() should return successive characters of the regular expression.   |
| PEEKC()                  | Return the next character in the regular expression. Successive calls to PEEKC() should return the same character (which should also be the next character returned by GETC()).   |
| UNGETC(c)                | Cause the argument <i>c</i> to be returned by the next call to GETC() (and PEEKC()). No more than one character of pushback is ever needed and this character is guaranteed to be the last character read by GETC(). The value of the macro UNGETC( <i>c</i> ) is always ignored. |
| RETURN( <i>pointer</i> ) | This macro is used on normal exit of the <i>compile</i> routine. The value of the argument <i>pointer</i> is a pointer to the character after the last character of the compiled regular expression. This is useful to programs which have memory allocation to manage.           |
| ERROR( <i>val</i> )      | This is the abnormal return from the <i>compile</i> routine. The argument <i>val</i> is an error number (see table  |

below for meanings). This call should never return.

ERROR	MEANING
11	Range endpoint too large.
16	Bad number.
25	"\digit" out of range.
36	Illegal or missing delimiter.
41	No remembered search string.
42	{ ( ) imbalance.
43	Too many \(.
44	More than 2 numbers given in \{ \}.
45	} expected after \.
46	First number exceeds second in \{ \}.
49	[ ] imbalance.
50	Regular expression overflow.

The syntax of the *compile* routine is as follows:

```
compile(instring, expbuf, endbuf, eof)
```

The first parameter *instring* is never used explicitly by the *compile* routine but is useful for programs that pass down different pointers to input characters. It is sometimes used in the INIT declaration (see below). Programs which call functions to input characters or have characters in an external array can pass down a value of ((char \*) 0) for this parameter.

The next parameter *expbuf* is a character pointer. It points to the place where the compiled regular expression will be placed.

The parameter *endbuf* is one more than the highest address where the compiled regular expression may be placed. If the compiled expression cannot fit in (*endbuf* - *expbuf*) bytes, a call to ERROR(50) is made.

The parameter *eof* is the character which marks the end of the regular expression. For example, in *ed(1)*, this character is usually a /.

Each program that includes this file must have a #define statement for INIT. This definition will be placed right after the declaration for the function *compile* and the opening curly brace {}. It is used for dependent declarations and initializations. Most often it is used to set a register variable to point the beginning of the regular expression so that this register variable can be used in the declarations for GETC(), PEEKC() and UNGETC(). Otherwise it can be used to declare external variables that might be used by GETC(), PEEKC() and UNGETC(). See the example below of the declarations taken from *grep(1)*.

There are other functions in this file which perform actual regular expression matching, one of which is the function *step*. The call to *step* is as follows:

```
step(string, expbuf)
```

The first parameter to *step* is a pointer to a string of characters to be checked for a match. This string should be null terminated.

The second parameter *expbuf* is the compiled regular expression which was obtained by a call of the function *compile*.

The function *step* returns non-zero if the given string matches the regular expression, and zero if the expressions do not match. If there is a match, two external character pointers are set as a side effect to the call to *step*.

The variable set in *step* is *loc1*. This is a pointer to the first character that matched the regular expression. The variable *loc2*, which is set by the function *advance*, points to the character after the last character that matches the regular expression. Thus if the regular expression matches the entire line, *loc1* will point to the first character of *string* and *loc2* will point to the null at the end of *string*.

*Step* uses the external variable *circf* which is set by *compile* if the regular expression begins with `^`. If this is set then *step* will try to match the regular expression to the beginning of the string only. If more than one regular expression is to be compiled before the first is executed the value of *circf* should be saved for each compiled expression and *circf* should be set to that saved value before each call to *step*.

The function *advance* is called from *step* with the same arguments as *step*. The purpose of *step* is to step through the *string* argument and call *advance* until *advance* returns non-zero indicating a match or until the end of *string* is reached. If one wants to constrain *string* to the beginning of the line in all cases, *step* need not be called; simply call *advance*.

When *advance* encounters a `*` or `{ \ }` sequence in the regular expression, it will advance its pointer to the string to be matched as far as possible and will recursively call itself trying to match the rest of the string to the rest of the regular expression. As long as there is no match, *advance* will back up along the string until it finds a match or reaches the point in the string that initially matched the `*` or `{ \ }`. It is sometimes desirable to stop this backing up before the initial point in the string is reached. If the external character pointer *locs* is equal to the point in the string at sometime during the backing up process, *advance* will break out of the loop that backs up and will return zero. This is used by *ed(1)* and *sed(1)* for substitutions done globally (not just the first occurrence, but the whole line) so, for example, expressions like *s/y\*/g* do not loop forever.

The additional external variables *sed* and *nbra* are used for special purposes.

#### EXAMPLES

The following is an example of how the regular expression macros and calls look from *grep(1)*:

```
#define INIT          register char *sp = instring;
#define GETC()        (*sp++)
#define PEEKC()       (*sp)
#define UNGETC(c)     (--sp)
#define RETURN(c)     return;
#define ERROR(c)      regerr()
#include <regexp.h>
...
                (void) compile(*argv, expbuf, &expbuf[ESIZE], '^0');
...
                if (step(linebuf, expbuf))
                    succeed();
```

#### FILES

/usr/include/regexp.h

#### SEE ALSO

bs(1), ed(1), expr(1), grep(1), sed(1).

**BUGS**

The handling of *circf* is kludgy.

The actual code is probably easier to understand than this manual page.

**NAME**

stat — data returned by stat system call

**SYNOPSIS**

```
#include <sys/types.h>
#include <sys/stat.h>
```

**DESCRIPTION**

The system calls *stat* and *fstat* return data whose structure is defined by this include file. The encoding of the field *st\_mode* is defined in this file also.

```
/*
 * Structure of the result of stat
 */
struct stat {
    dev_t    st_dev;
    ino_t    st_ino;
    ushort  st_mode;
    short    st_nlink;
    ushort  st_uid;
    ushort  st_gid;
    dev_t    st_rdev;
    off_t    st_size;
    time_t   st_atime;
    time_t   st_mtime;
    time_t   st_ctime;
};
```

```
#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_IFIFO   0010000 /* fifo */
#define S_ISUID   04000 /* set user id on execution */
#define S_ISGID   02000 /* set group id on execution */
#define S_ISVTX   01000 /* save swapped text even after use */
#define S_IRREAD  00400 /* read permission, owner */
#define S_IWRITE  00200 /* write permission, owner */
#define S_IXEXEC  00100 /* execute/search permission, owner */
```

**FILES**

```
/usr/include/sys/types.h
/usr/include/sys/stat.h
```

**SEE ALSO**

stat(2), types(5).

**NAME**

tcp - Internet Transmission Control Protocol

**SYNOPSIS**

```
#include <sys/socket.h>
#include <netinet/in.h>

s = socket(AF_INET, SOCK_STREAM, 0);
```

**DESCRIPTION**

The TCP protocol provides reliable, flow-controlled, two-way transmission of data. It is a byte-stream protocol used to support the SOCK\_STREAM abstraction. TCP uses the standard Internet address format and, in addition, provides a per-host collection of port addresses. Thus, each address is composed of an Internet address specifying the host and network, with a specific TCP port on the host identifying the peer entity.

Sockets utilizing the tcp protocol are either active or passive. Active sockets initiate connections to passive sockets. By default TCP sockets are created active; to create a passive socket the *listen(2N)* system call must be used after binding the socket with the *bind(2N)* system call. Only passive sockets may use the *accept(2N)* call to accept incoming connections. Only active sockets may use the *connect(2N)* call to initiate connections.

Passive sockets may underspecify their location to match incoming connection requests from multiple networks. This technique, termed wildcard addressing, allows a single server to provide service to clients on multiple networks. To create a socket which listens on all networks, the Internet address INADDR\_ANY must be bound. The TCP port may still be specified at this time; if the port is not specified the system will assign one. Once a connection has been established the socket's address is fixed by the peer entity's location. The address assigned the socket is the address associated with the network interface through which packets are being transmitted and received. Normally this address corresponds to the peer entity's network.

**DIAGNOSTICS**

A socket operation may fail with one of the following errors returned:

- |                 |   |
|-----------------|---|
| [EISCONN]       | when trying to establish a connection on a socket which already has one;  |
| [ENOBUFS]       | when the system runs out of memory for an internal data structure;  |
| [ETIMEDOUT]     | when a connection was dropped due to excessive retransmissions;   |
| [ECONNRESET]    | when the remote peer forces the connection to be closed;  |
| [ECONNREFUSED]  | when the remote peer actively refuses connection establishment (usually because no process is listening to the port); |
| [EADDRINUSE]    | when an attempt is made to create a socket with a port which has already been allocated;                              |
| [EADDRNOTAVAIL] | when an attempt is made to create a socket with a network address for which no network interface                      |

exists.

**SEE ALSO**

intro(5N), inet(5F)

**BUGS**

It should be possible to send and receive TCP options. The system always tries to negotiate the maximum TCP segment size to be 1024 bytes. This can result in poor performance if an intervening network performs excessive fragmentation.

## NAME

term — conventional names for terminals

## DESCRIPTION

These names are used by certain commands (e.g., *nroff*, *mm(1)*, *man(1)*, *tabs(1)*) and are maintained as part of the shell environment (see *sh(1)*, *profile(4)*, and *environ(5)*) in the variable \$TERM:

1520	Datamedia 1520
1620	Diablo 1620 and others using the HyType II printer
1620-12	same, in 12-pitch mode
2621	Hewlett-Packard HP2621 series
2631	Hewlett-Packard 2631 line printer
2631-c	Hewlett-Packard 2631 line printer - compressed mode
2631-e	Hewlett-Packard 2631 line printer - expanded mode
2640	Hewlett-Packard HP2640 series
2645	Hewlett-Packard HP264n series (other than the 2640 series)
300	DASI/DTC/GSI 300 and others using the HyType I printer
300-12	same, in 12-pitch mode
300s	DASI/DTC/GSI 300s
382	DTC 382
300s-12	same, in 12-pitch mode
3045	Datamedia 3045
33	TELETYPE® Terminal Model 33 KSR
37	TELETYPE Terminal Model 37 KSR
40-2	TELETYPE Terminal Model 40/2
40-4	TELETYPE Terminal Model 40/4
4540	TELETYPE Terminal Model 4540
3270	IBM Model 3270
4000a	Trendata 4000a
4014	Tektronix 4014
43	TELETYPE Model 43 KSR
450	DASI 450 (same as Diablo 1620)
450-12	same, in 12-pitch mode
735	Texas Instruments TI735 and TI725
745	Texas Instruments TI745
dumb	generic name for terminals that lack reverse line-feed and other special escape sequences
sync	generic name for synchronous TELETYPE 4540-compatible terminals
hp	Hewlett-Packard (same as 2645)
lp	generic name for a line printer
tn1200	General Electric TermiNet 1200
tn300	General Electric TermiNet 300

Up to 8 characters, chosen from [-a-z0-9], make up a basic terminal name. Terminal sub-models and operational modes are distinguished by suffixes beginning with a -. Names should generally be based on original vendors, rather than local distributors. A terminal acquired from one vendor should not have more than one distinct basic name.

Commands whose behavior depends on the type of terminal should accept arguments of the form *-Tterm* where *term* is one of the names given above; if no such argument is present, such commands should obtain the terminal type from the environment variable \$TERM, which, in turn,



should contain *term*.

See */etc/termcap* on your system for a complete list.

**SEE ALSO**

*mm(1)*, *nroff(1)*, *sh(1)*, *stty(1)*, *tabs(1)*, *tplot(1G)*, *profile(4)*, *environ(5)*.

**BUGS**

This is a small candle trying to illuminate a large, dark problem. Programs that ought to adhere to this nomenclature do so somewhat fitfully.

**NAME**

*termcap* – terminal capability data base

**SYNOPSIS**

*/etc/termcap*

**DESCRIPTION**

*Termcap* is a data base describing terminals used, e.g., by *vi*(1). Terminals are described in *termcap* by giving a set of capabilities which they have, and by describing how operations are performed. Padding requirements and initialization sequences are included in *termcap*.

Entries in *termcap* consist of a number of ':' separated fields. The first entry for each terminal gives the names which are known for the terminal, separated by '|' characters. The first name is always 2 characters long and is used by older version 6 systems which store the terminal type in a 16 bit word in a system-wide data base. The second name given is the most common abbreviation for the terminal, and the last name given should be a long name fully identifying the terminal. The second name should contain no blanks; the last name may well contain blanks for readability.

**CAPABILITIES**

(P) indicates padding may be specified

(P\*) indicates that padding may be based on no. lines affected

**Name Type Pad? Description**

ae	str	(P)	End alternate character set
al	str	(P*)	Add new blank line
am	bool		Terminal has automatic margins
as	str	(P)	Start alternate character set
bc	str		Backspace if not ^H
bs	bool		Terminal can backspace with ^H
bt	str	(P)	Back tab
bw	bool		Backspace wraps from column 0 to last column
CC	str		Command character in prototype if terminal settable
cd	str	(P*)	Clear to end of display
ce	str	(P)	Clear to end of line
ch	str	(P)	Like cm but horizontal motion only, line stays same
cl	str	(P*)	Clear screen
cm	str	(P)	Cursor motion
co	num		Number of columns in a line
cr	str	(P*)	Carriage return, (default ^M)

cs	str	(P)	Change scrolling region (vt100), like cm
cv	str	(P)	Like ch but vertical only.
da	bool		Display may be retained above
dB	num		Number of millisecc of bs delay needed
db	bool		Display may be retained below
dC	num		Number of millisecc of cr delay needed
dc	str	(P*)	Delete character
dF	num		Number of millisecc of ff delay needed
dl	str	(P*)	Delete line
dm	str		Delete mode (enter)
dN	num		Number of millisecc of nl delay needed
do	str		Down one line
dT	num		Number of millisecc of tab delay needed
ed	str		End delete mode
ei	str		End insert mode; give :ei=: if ic
eo	str		Can erase overstrikes with a blank
ff	str	(P*)	Hardcopy terminal page eject (default "L")
hc	bool		Hardcopy terminal
hd	str		Half-line down (forward 1/2 linefeed)
ho	str		Home cursor (if no cm)
hu	str		Half-line up (reverse 1/2 linefeed)
hz	str		Hazeltine; can't print ""s
ic	str	(P)	Insert character
if	str		Name of file containing is
im	str		Insert mode (enter); give :im=: if ic
in	bool		Insert mode distinguishes nulls on display
ip	str	(P*)	Insert pad after character inserted
is	str		Terminal initialization string
k0-k9	str		Sent by other function keys 0-9
kb	str		Sent by backspace key
kd	str		Sent by terminal down arrow key
ke	str		Out of keypad transmit mode
kh	str		Sent by home key
kl	str		Sent by terminal left arrow key
kn	num		Number of other keys
ko	str		Termcap entries for other non-function keys
kr	str		Sent by terminal right arrow key
ks	str		Put terminal in keypad transmit mode
ku	str		Sent by terminal up arrow key

l0-19	str	Labels on other function keys
li	num	Number of lines on screen or page
ll	str	Last line, first column (if no cm)
ma	str	Arrow key map, used by vi version 2 only
mi	bool	Safe to move while in insert mode
ml	str	Memory lock on above cursor.
ms	bool	Safe to move while in standout and underline mode
mu	str	Memory unlock (turn off memory lock).
nc	bool	No correctly working carriage return (DM2500,H2000)
nd	str	Non-destructive space (cursor right)
nl	str (P*)	Newline character (default \n)
ns	bool	Terminal is a CRT but doesn't scroll.
os	bool	Terminal overstrikes
pc	str	Pad character (rather than null)
pt	bool	Has hardware tabs (may need to be set with is)
se	str	End stand out mode
sf	str (P)	Scroll forwards
sg	num	Number of blank chars left by so or se
so	str	Begin stand out mode
sr	str (P)	Scroll reverse (backwards)
ta	str (P)	Tab (other than ^I or with padding)
tc	str	Entry of similar terminal - must be last
te	str	String to end programs that use cm
ti	str	String to begin programs that use cm
uc	str	Underscore one char and move past it
ue	str	End underscore mode
ug	num	Number of blank chars left by us or ue
ul	bool	Terminal underlines even though it doesn't overstrike
up	str	Upline (cursor up)
us	str	Start underscore mode
vb	str	Visible bell (may not move cursor)
ve	str	Sequence to end open/visual mode
vs	str	Sequence to start open/visual mode
xb	bool	Beehive (f1=escape, f2=ctrl C)
xn	bool	A newline is ignored after a wrap (Concept)
xr	bool	Return acts like ce \r \n (Delta Data)
xs	bool	Standout not erased by writing over it (HP 2647)
xt	bool	Tabs are destructive, magic so char (Telera y 1061)

### A Sample Entry

The following entry, which describes the Concept-100, is among the more complex entries in the *termcap* file as of this writing. (This particular concept entry is outdated and is used as an example only.)

```
c1|c100|concept100:is=AEUAE7AE5EBERENHAEKAE200Ae&A200A
:al=3*E'R:amb:cd=16*E'C:ce=16E'S:cl=2*L:cm=AE%+ %+ :co#80A
:dc=16E'A:dl=3*E'B:el=AE200:eo:im=AE*Pinip=16*li#24:mind=AE-:
:se=AEdEe:so=AEAEeta=8*ul:up=AE;vb=AEkEK:xn:
```

Entries may continue onto multiple lines by giving a \ as the last character of a line, and that empty fields may be included for readability (here between the last field on a line and the first field on the next). Capabilities in *termcap* are of three types: Boolean capabilities which indicate that the terminal has some particular feature, numeric capabilities giving the size of the terminal or the size of particular delays, and string capabilities, which give a sequence which can be used to perform particular terminal operations.

### Types of Capabilities

All capabilities have two letter codes. For instance, the fact that the Concept has automatic margins (i.e. an automatic return and linefeed when the end of a line is reached) is indicated by the capability *am*. Hence the description of the Concept includes *am*. Numeric capabilities are followed by the character '#' and then the value. Thus *co* which indicates the number of columns the terminal has gives the value '80' for the Concept.

Finally, string valued capabilities, such as *ce* (clear to end of line sequence) are given by the two character code, an '=', and then a string ending at the next following ':'. A delay in milliseconds may appear after the '=' in such a capability, and padding characters are supplied by the editor after the remainder of the string is sent to provide this delay. The delay can be either an integer, e.g. '20', or an integer followed by an '\*', i.e. '3\*'. A '\*' indicates that the padding required is proportional to the number of lines affected by the operation, and the amount given is the per-affected-unit padding required. When a '\*' is specified, it is sometimes useful to give a delay of the form '3.5' specify a delay per unit to tenths of milliseconds.

A number of escape sequences are provided in the string valued capabilities for easy encoding of characters there. A *\E* maps to an ESCAPE character, *\x* maps to a control-x for any appropriate x, and the sequences *\n* *\r* *\t* *\b* *\f* give a new-line, return, tab, backspace and formfeed. Finally, characters may be given as

three octal digits after a \, and the characters ^ and \ may be given as \^ and \\ . If it is necessary to place a : in a capability it must be escaped in octal as \072. If it is necessary to place a null character in a string capability it must be encoded as \200. The routines which deal with *termcap* use C strings, and strip the high bits of the output very late so that a \200 comes out as a \000 would.

### Preparing Descriptions

We now outline how to prepare descriptions of terminals. The most effective way to prepare a terminal description is by imitating the description of a similar terminal in *termcap* and to build up a description gradually, using partial descriptions with *ex* to check that they are correct. Be aware that a very unusual terminal may expose deficiencies in the ability of the *termcap* file to describe it or bugs in *ex*. To easily test a new terminal description you can set the environment variable *TERMCAP* to a pathname of a file containing the description you are working on and the editor will look there rather than in */etc/termcap*. *TERMCAP* can also be set to the *termcap* entry itself to avoid reading the file when starting up the editor. (This only works on version 7 systems.)

### Basic capabilities

The number of columns on each line for the terminal is given by the *co* numeric capability. If the terminal is a CRT, then the number of lines on the screen is given by the *li* capability. If the terminal wraps around to the beginning of the next line when it reaches the right margin, then it should have the *am* capability. If the terminal can clear its screen, then this is given by the *cl* string capability. If the terminal can backspace, then it should have the *bs* capability, unless a backspace is accomplished by a character other than ^H (ugh) in which case you should give this character as the *bc* string capability. If it overstrikes (rather than clearing a position when a character is struck over) then it should have the *os* capability.

A very important point here is that the local cursor motions encoded in *termcap* are undefined at the left and top edges of a CRT terminal. The editor will never attempt to backspace around the left edge, nor will it attempt to go up locally off the top. The editor assumes that feeding off the bottom of the screen will cause the screen to scroll up, and the *am* capability tells whether the cursor sticks at the right edge of the screen. If the terminal has switch selectable automatic margins, the *termcap* file usually assumes that this is on, i.e. *am*.

These capabilities suffice to describe hardcopy and glass-tty terminals. Thus the model 33 teletype is described as

```
␣|33|tty33:co#72:os
```

while the Lear Siegler ADM-3 is described as

```
cl|adm3|3|lsi adm3:am:bs:cl=~Z:li#24:co#80
```

### Cursor addressing

Cursor addressing in the terminal is described by a *cm* string capability, with *printf(3s)* like escapes *%x* in it. These substitute to encodings of the current line or column position, while other characters are passed through unchanged. If the *cm* string is thought of as being a function, then its arguments are the line and then the column to which motion is desired, and the *%* encodings have the following meanings:

```
%d  as in printf, 0 origin
%2   like %2d
%3   like %3d
%.   like %c
%+x  adds x to value, then %.
%>xy if value > x adds y, no output.
%r   reverses order of line and column, no output
%i   increments line/column (for 1 origin)
%%   gives a single %
%n   exclusive or row and column with 0140 (DM2500)
%B   BCD (16*(x/10)) + (x%10), no output.
%D   Reverse coding (x-2*(x%16)), no output. (Delta Data).
```

Consider the HP2645, which, to get to row 3 and column 12, needs to be sent `\E&a12c03Y` padded for 6 milliseconds. Note that the order of the rows and columns is inverted here, and that the row and column are printed as two digits. Thus its *cm* capability is `cm=6\E&%r%2c%2Y`. The Microterm ACT-IV needs the current row and column sent preceded by a `T`, with the row and column simply encoded in binary, `cm=~T%.%.` Terminals which use `%.` need to be able to backspace the cursor (`bs` or `bc`), and to move the cursor up one line on the screen (`up` introduced below). This is necessary because it is not always safe to transmit `\t`, `\n` `^D` and `\r`, as the system may change or discard them.

A final example is the LSI ADM-3a, which uses row and column offset by a blank character, thus `cm=\E= %+ %+`.

### Cursor motions

If the terminal can move the cursor one position to the right, leaving the character at the current position unchanged, then this sequence should be given as `nd` (non-destructive space). If it can move the cursor up a line on the screen in the same column, this should be given as `up`. If the terminal has no cursor addressing capability, but can home the cursor (to very upper left corner of screen) then this can be given as `ho`; similarly a fast way of getting to the lower left hand corner can be given as `ll`; this may involve going up with `up` from the home position, but the editor will never do this itself (unless `ll` does) because it makes no assumption about the effect of moving up from the home position.

#### Area clears

If the terminal can clear from the current position to the end of the line, leaving the cursor where it is, this should be given as `ce`. If the terminal can clear from the current position to the end of the display, then this should be given as `cd`. The editor only uses `cd` from the first column of a line.

#### Insert/delete line

If the terminal can open a new blank line before the line where the cursor is, this should be given as `al`; this is done only from the first position of a line. The cursor must then appear on the newly blank line. If the terminal can delete the line which the cursor is on, then this should be given as `dl`; this is done only from the first position on the line to be deleted. If the terminal can scroll the screen backwards, then this can be given as `sb`, but just `al` suffices. If the terminal can retain display memory above then the `da` capability should be given; if display memory can be retained below then `db` should be given. These let the editor understand that deleting a line on the screen may bring non-blank lines up from below or that scrolling back with `sb` may bring down non-blank lines.

#### Insert/delete character

There are two basic kinds of intelligent terminals with respect to insert/delete character which can be described using *termcap*. The most common insert/delete character operations affect only the characters on the current line and shift characters off the end of the line rigidly. Other terminals, such as the Concept 100 and the Perkin Elmer Owl, make a distinction between typed and untyped blanks on the screen, shifting upon an insert or delete only to an untyped blank on the screen which is either eliminated, or expanded to two untyped blanks. You can find out which kind of terminal you have by clearing the screen and then typing text separated by cursor motions. Type `abc def` using local cursor motions (not spaces) between the `abc` and the `def`. Then position the cursor before the `abc` and put the terminal in insert mode. If typing



characters causes the rest of the line to shift rigidly and characters to fall off the end, then your terminal does not distinguish between blanks and untyped positions. If the abc shifts over to the def which then move together around the end of the current line and onto the next as you insert, you have the second type of terminal, and should give the capability `in`, which stands for insert null. If your terminal does something different and unusual then you may have to modify the editor to get it to use the insert mode your terminal defines. We have seen no terminals which have an insert mode not falling into one of these two classes.

The editor can handle both terminals which have an insert mode, and terminals which send a simple sequence to open a blank position on the current line. Give as `im` the sequence to get into insert mode, or give it an empty value if your terminal uses a sequence to insert a blank position. Give as `ei` the sequence to leave insert mode (give this, with an empty value also if you gave `im` so). Now give as `ic` any sequence needed to be sent just before sending the character to be inserted. Most terminals with a true insert mode will not give `ic`, terminals which send a sequence to open a screen position should give it here. (Insert mode is preferable to the sequence to open a position on the screen if your terminal has both.) If post insert padding is needed, give this as a number of milliseconds in `ip` (a string option). Any other sequence which may need to be sent after an insert of a single character may also be given in `ip`.

It is occasionally necessary to move around while in insert mode to delete characters on the same line (e.g. if there is a tab after the insertion position). If your terminal allows motion while in insert mode you can give the capability `mi` to speed up inserting in this case. Omitting `mi` will affect only speed. Some terminals (notably Datamedia's) must not have `mi` because of the way their insert mode works.

Finally, you can specify delete mode by giving `dm` and `ed` to enter and exit delete mode, and `dc` to delete a single character while in delete mode.

#### **Highlighting, underlining, and visible bells**

If your terminal has sequences to enter and exit standout mode these can be given as `so` and `se` respectively. If there are several flavors of standout mode (such as inverse video, blinking, or underlining – half bright is not usually an acceptable standout mode unless the terminal is in inverse video mode constantly) the preferred mode is inverse video by itself. If the code to change into or out of standout mode leaves one or even two blank spaces on the screen, as the TVI 912 and Teleray 1061 do, then `ug` should be given to tell how many spaces are left.

Codes to begin underlining and end underlining can be given as *us* and *ue* respectively. If the terminal has a code to underline the current character and move the cursor one space to the right, such as the Microterm Mime, this can be given as *uc*. (If the underline code does not move the cursor to the right, give the code followed by a nondestructive space.)

Many terminals, such as the HP 2621, automatically leave standout mode when they move to a new line or the cursor is addressed. Programs using standout mode should exit standout mode before moving the cursor or sending a newline.

If the terminal has a way of flashing the screen to indicate an error quietly (a bell replacement) then this can be given as *vb*; it must not move the cursor. If the terminal should be placed in a different mode during open and visual modes of *ex*, this can be given as *vs* and *ve*, sent at the start and end of these modes respectively. These can be used to change, e.g., from a underline to a block cursor and back.

If the terminal needs to be in a special mode when running a program that addresses the cursor, the codes to enter and exit this mode can be given as *ti* and *te*. This arises, for example, from terminals like the Concept with more than one page of memory. If the terminal has only memory relative cursor addressing and not screen relative cursor addressing, a one screen-sized window must be fixed into the terminal for cursor addressing to work properly.

If your terminal correctly generates underlined characters (with no special codes needed) even though it does not overstrike, then you should give the capability *ul*. If overstrikes are erasable with a blank, then this should be indicated by giving *eo*.

### Keypad

If the terminal has a keypad that transmits codes when the keys are pressed, this information can be given. Note that it is not possible to handle terminals where the keypad only works in local (this applies, for example, to the unshifted HP 2621 keys). If the keypad can be set to transmit or not transmit, give these codes as *ks* and *ke*. Otherwise the keypad is assumed to always transmit. The codes sent by the left arrow, right arrow, up arrow, down arrow, and home keys can be given as *kl*, *kr*, *ku*, *kd*, and *kh* respectively. If there are function keys such as *f0*, *f1*, ..., *f9*, the codes they send can be given as *k0*, *k1*, ..., *k9*. If these keys have labels other than the default *f0* through *f9*, the labels can be given as *l0*, *l1*, ..., *l9*. If there are other keys that transmit the same code as the terminal expects for the corresponding function, such as clear screen, the *termcap* 2 letter codes can be given in the *ko* capability, for example, *:ko=c,l,lf,fb,*

which says that the terminal has clear, home down, scroll down, and scroll up keys that transmit the same thing as the cl, ll, sf, and sb entries.

The ma entry is also used to indicate arrow keys on terminals which have single character arrow keys. It is obsolete but still in use in version 2 of vi, which must be run on some minicomputers due to memory limitations. This field is redundant with kl, kr, ku, kd, and kh. It consists of groups of two characters. In each group, the first character is what an arrow key sends, the second character is the corresponding vi command. These commands are h for kl, j for kd, k for ku, l for kr, and H for kh. For example, the mime would be :ma=^KJ^Zk^Xl: indicating arrow keys left (^H), down (^K), up (^Z), and right (^X). (There is no home key on the mime.)

### Miscellaneous

If the terminal requires other than a null (zero) character as a pad, then this can be given as pc.

If tabs on the terminal require padding, or if the terminal uses a character other than ^I to tab, then this can be given as ta.

Hazeltine terminals, which don't allow '^' characters to be printed should indicate hz. Datamedia terminals, which echo carriage-return linefeed for carriage return and then ignore a following linefeed should indicate nc. Early Concept terminals, which ignore a linefeed immediately after an am wrap, should indicate xn. If an erase-eol is required to get rid of standout (instead of merely writing on top of it), xs should be given. Teleray terminals, where tabs turn all characters moved over to blanks, should indicate xt. Other specific terminal problems may be corrected by adding more capabilities of the form xx.

Other capabilities include is, an initialization string for the terminal, and if, the name of a file containing long initialization strings. These strings are expected to properly clear and then set the tabs on the terminal, if the terminal has settable tabs. If both are given, is will be printed before if. This is useful where if is /usr/lib/tabsel/std but is clears the tabs first.

### Similar Terminals

If there are two very similar terminals, one can be defined as being just like the other with certain exceptions. The string capability tc can be given with the name of the similar terminal. This capability must be last and the combined length of the two entries must not exceed 1024. Since *termlib* routines search the entry from left to right, and since the tc capability is replaced by the corresponding entry, the capabilities given at the left override the ones in the

similar terminal. A capability can be cancelled with `xx@` where `xx` is the capability. For example, the entry

```
hn|2621nl:ks@:ke@:tc-2621:
```

defines a 2621nl that does not have the `ks` or `ke` capabilities, and hence does not turn on the function key labels when in visual mode. This is useful for different modes for a terminal, or for different user preferences.

#### FILES

`/etc/termcap` file containing terminal descriptions

#### SEE ALSO

`ex(1)`, `tset(1)`, `ul(1)`, `vi(1)`, `termcap(3X)`.

#### BUGS

*Ex* allows only 256 characters for string capabilities, and the routines in *termcap(3X)* do not check for overflow of this buffer. The total length of a single entry (excluding only escaped newlines) may not exceed 1024.

The `ma`, `vs`, and `ve` entries are specific to the *vi* program.

Not all programs support all entries. There are entries that are not supported by any program.

#### AUTHOR

William Joy

Mark Horton added underlining and keypad support

**NAME**

types — primitive system data types

**SYNOPSIS**

#include &lt;sys/types.h&gt;

**DESCRIPTION**

The data types defined in the include file are used in UNIX System code; some data of these types are accessible to user code:

```
typedef struct { int r[1]; } * physadr;
typedef long      daddr_t;
typedef char *    caddr_t;
typedef unsigned int  uint;
typedef unsigned short ushort;
typedef ushort     ino_t;
typedef short      cnt_t;
typedef long       time_t;
typedef int        label_t[10];
typedef short      dev_t;
typedef long       off_t;
typedef long       paddr_t;
typedef long       key_t;
```

The form *daddr\_t* is used for disk addresses except in an i-node on disk, see *fs(4)*. Times are encoded in seconds since 00:00:00 GMT, January 1, 1970. The major and minor parts of a device code specify kind and unit number of a device and are installation-dependent. Offsets are measured in bytes from the beginning of a file. The *label\_t* variables are used to save the processor state while another process is running.

**SEE ALSO***fs(4)*.

**NAME**

udp — Internet User Datagram Protocol

**SYNOPSIS**

```
#include <sys/socket.h>
#include <netinet/in.h>

s = socket(AF_INET, SOCK_DGRAM, 0);
```

**DESCRIPTION**

UDP is a simple, unreliable datagram protocol which is used to support the SOCK\_DGRAM abstraction for the Internet protocol family. UDP sockets are connectionless, and are normally used with the *sendto* and *recvfrom* calls, though the *connect(2N)* call may also be used to fix the destination for future packets (in which case the *recv(2N)* or *send(2N)* system calls may be used).

UDP address formats are identical to those used by TCP. In particular UDP provides a port identifier in addition to the normal Internet address format. Note that the UDP port space is separate from the TCP port space (i.e. a UDP port may not be connected to a TCP port). In addition broadcast packets may be sent (assuming the underlying network supports this) by using a reserved broadcast address; this address is network interface dependent.

**DIAGNOSTICS**

A socket operation may fail with one of the following errors returned:

- [EISCONN] when trying to establish a connection on a socket which already has one, or when trying to send a datagram with the destination address specified and the socket is already connected;
- [ENOTCONN] when trying to send a datagram, but no destination address is specified, and the socket hasn't been connected;
- [ENOBUFS] when the system runs out of memory for an internal data structure;
- [EADDRINUSE] when an attempt is made to create a socket with a port which has already been allocated;
- [EADDRNOTAVAIL] when an attempt is made to create a socket with a network address for which no network interface exists.

**SEE ALSO**

*send(2N)*, *recv(2N)*, *intro(5N)*, *inet(5F)*

**NAME**

values — machine-dependent values

**SYNOPSIS**

#include &lt;values.h&gt;

**DESCRIPTION**

This file contains a set of manifest constants, conditionally defined for particular processor architectures.

The model assumed for integers is binary representation (one's or two's complement), where the sign is represented by the value of the high-order bit.

<b>BITS</b> ( <i>type</i> )	The number of bits in a specified type (e.g., int).
<b>HIBITS</b>	The value of a short integer with only the high-order bit set (in most implementations, 0x8000).
<b>HIBITL</b>	The value of a long integer with only the high-order bit set (in most implementations, 0x80000000).
<b>HIBITI</b>	The value of a regular integer with only the high-order bit set (usually the same as HIBITS or HIBITL).
<b>MAXSHORT</b>	The maximum value of a signed short integer (in most implementations, 0x7FFF $\equiv$ 32767).
<b>MAXLONG</b>	The maximum value of a signed long integer (in most implementations, 0x7FFFFFFF $\equiv$ 2147483647).
<b>MAXINT</b>	The maximum value of a signed regular integer (usually the same as MAXSHORT or MAXLONG).
<b>MAXFLOAT, LN_MAXFLOAT</b>	The maximum value of a single-precision floating-point number, and its natural logarithm.
<b>MAXDOUBLE, LN_MAXDOUBLE</b>	The maximum value of a double-precision floating-point number, and its natural logarithm.
<b>MINFLOAT, LN_MINFLOAT</b>	The minimum positive value of a single-precision floating-point number, and its natural logarithm.
<b>MINDOUBLE, LN_MINDOUBLE</b>	The minimum positive value of a double-precision floating-point number, and its natural logarithm.
<b>FSIGNIF</b>	The number of significant bits in the mantissa of a single-precision floating-point number.
<b>DSIGNIF</b>	The number of significant bits in the mantissa of a double-precision floating-point number.

**FILES**

/usr/include/values.h

**SEE ALSO**

intro(3), math(5).

## NAME

`varargs` — handle variable argument list

## SYNOPSIS

```
#include <varargs.h>
va_alist
va_dcl
void va_start(pvar)
va_list pvar;
type va_arg(pvar, type)
va_list pvar;
void va_end(pvar)
va_list pvar;
```

## DESCRIPTION

This set of macros allows portable procedures that accept variable argument lists to be written. Routines that have variable argument lists (such as *printf(3S)*) but do not use *varargs* are inherently nonportable, as different machines use different argument-passing conventions.

`va_alist` is used as the parameter list in a function header.

`va_dcl` is a declaration for `va_alist`. No semicolon should follow `va_dcl`.

`va_list` is a type defined for the variable used to traverse the list.

`va_start` is called to initialize `pvar` to the beginning of the list.

`va_arg` will return the next argument in the list pointed to by `pvar`. *Type* is the type the argument is expected to be. Different types can be mixed, but it is up to the routine to know what type of argument is expected, as it cannot be determined at runtime.

`va_end` is used to clean up.

Multiple traversals, each bracketed by `va_start ... va_end`, are possible.

## EXAMPLE

This example is a possible implementation of *execl(2)*.

```
#include <varargs.h>
#define MAXARGS 100

/* execl is called by
   execl(file, arg1, arg2, ..., (char *)0);
*/
execl(va_alist)
va_dcl
{
    va_list ap;
    char *file;
    char *args[MAXARGS];
    int argno = 0;

    va_start(ap);
    file = va_arg(ap, char *);
    while ((args[argno++] = va_arg(ap, char *)) != (char *)0)
        ;
}
```



```
        va_end(ap);  
    return execv(file, args);  
}
```

**SEE ALSO**

exec(2), printf(3S).

**BUGS**

It is up to the calling routine to specify how many arguments there are, since it is not always possible to determine this from the stack frame. For example, *execl* is passed a zero pointer to signal the end of the list. *Printf* can tell how many arguments are there by the format.

It is non-portable to specify a second argument of *char*, *short*, or *float* to *va\_arg*, since arguments seen by the called function are not *char*, *short*, or *float*. C converts *char* and *short* arguments to *int* and converts *float* arguments to *double* before passing them to a function.



**NAME**

intro - introduction to games

**DESCRIPTION**

This section describes the recreational and educational programs found in the directory /usr/games. The availability of these programs may vary from system to system.

**NAME**

adventure — an exploration game

**SYNOPSIS**

/usr/games/adventure

**DESCRIPTION**

The object of the game is to locate and explore Colossal Cave, find the treasures hidden there, and bring them back to the building with you. The program is self-describing to a point, but part of the game is to discover its rules.

To terminate a game, type "quit"; to save a game for later resumption, type "suspend".

**BUGS**

Saving a game creates a large executable file instead of just the information needed to resume the game.

**NAME**

aliens - The alien invaders attack the earth

**SYNOPSIS**

`/usr/games/aliens`

**DESCRIPTION**

This is a UNIX version of Space Invaders. The program is pretty much self documenting.

**FILES**

`/usr/games/lib/aliens.log` Score file

**BUGS**

The program is a CPU hog. It needs to be re-written. It doesn't do well on terminals that run slower than 9600 baud.

**NAME**

arithmetic — provide drill in number facts

**SYNOPSIS**

/usr/games/arithmetic [ + - x / ] [ range ]

**DESCRIPTION**

*Arithmetic* types out simple arithmetic problems, and waits for an answer to be typed in. If the answer is correct, it types back "Right!", and a new problem. If the answer is wrong, it replies "What?", and waits for another answer. Every twenty problems, it publishes statistics on correctness and the time required to answer.

To quit the program, type an interrupt (delete).

The first optional argument determines the kind of problem to be generated; +, -, x, and / respectively cause addition, subtraction, multiplication, and division problems to be generated. One or more characters can be given; if more than one is given, the different types of problems will be mixed in random order; default is + - .

*Range* is a decimal number; all addends, subtrahends, differences, multipliers, divisors, and quotients will be less than or equal to the value of *range*. Default *range* is 10.

At the start, all numbers less than or equal to *range* are equally likely to appear. If the respondent makes a mistake, the numbers in the problem which was missed become more likely to reappear.

As a matter of educational philosophy, the program will not give correct answers, since the learner should, in principle, be able to calculate them. Thus the program is intended to provide drill for someone just past the first learning stage, not to teach number facts *de novo*. For almost all users, the relevant statistic should be time per problem, not percent correct.

**NAME**

autorobots — Escape from the automatic robots

**SYNOPSIS**

/usr/games/autorobots

**DESCRIPTION**

The object of the game *autorobots* is to move around inside of the box on the screen without getting eaten by the robots chasing you and without running into any robots or junk heaps. The robots move continuously.

If a robot runs into another robot or junk heap while chasing you, they crash and leave a junk heap.

You start out with 10 robots worth 10 points each. If you defeat all of them, you get 20 robots worth 20 points each. Then 30, etc. Until you get eaten!

The game keeps track of the top ten scores and prints them at the end of the game.

The valid commands are described on the screen.

**NAME**

back — the game of backgammon

**SYNOPSIS**

`/usr/games/back`

**DESCRIPTION**

*Back* is a program which provides a partner for the game of backgammon. It is designed to play at three different levels of skill, one of which you must select. In addition to selecting the opponent's level, you may also indicate that you would like to roll your own dice during your turns (for the superstitious players). You will also be given the opportunity to move first. The practice of each player rolling one die for the first move is not incorporated.

The points are numbered 1–24, with 1 being white's extreme inner table, 24 being brown's inner table, 0 being the bar for removed white pieces and 25 the bar for brown. For details on how moves are expressed, type *y* when *back* asks "Instructions?" at the beginning of the game. When *back* first asks "Move?", type ? to see a list of move options other than entering your numerical move.

When the game is finished, *back* will ask you if you want postmortem statistics. If you respond with *y*, *back* will attempt to append to or create a file `.backlog` in your HOME directory.

**FILES**

`/usr/games/lib/backrules` rules file  
`$HOME/.backlog` log file

**BUGS**

The only level really worth playing is "expert", and it only plays the forward game.

Doubling is not implemented.



**NAME**

`bcd` — convert to antique media

**SYNOPSIS**

`/usr/games/bcd text`

**DESCRIPTION**

*Bcd* converts the literal *text* into a form familiar to old-timers.

This program works best on hard copy terminals.

**NAME**

bj — the game of black jack

**SYNOPSIS**

/usr/games/bj

**DESCRIPTION**

*Bj* is a serious attempt at simulating the dealer in the game of black jack (or twenty-one) as might be found in Reno. The following rules apply:

The bet is \$2 every hand.

A player "natural" (black jack) pays \$3. A dealer natural loses \$2. Both dealer and player naturals is a "push" (no money exchange).

If the dealer has an ace up, the player is allowed to make an "insurance" bet against the chance of a dealer natural. If this bet is not taken, play resumes as normal. If the bet is taken, it is a side bet where the player wins \$2 if the dealer has a natural and loses \$1 if the dealer does not.

If the player is dealt two cards of the same value, he is allowed to "double". He is allowed to play two hands, each with one of these cards. (The bet is doubled also; \$2 on each hand.)

If a dealt hand has a total of ten or eleven, the player may "double down". He may double the bet (\$2 to \$4) and receive exactly one more card on that hand.

Under normal play, the player may "hit" (draw a card) as long as his total is not over twenty-one. If the player "busts" (goes over twenty-one), the dealer wins the bet.

When the player "stands" (decides not to hit), the dealer hits until he attains a total of seventeen or more. If the dealer busts, the player wins the bet.

If both player and dealer stand, the one with the largest total wins. A tie is a push.

The machine deals and keeps score. The following questions will be asked at appropriate times. Each question is answered by y followed by a new-line for "yes", or just new-line for "no".

? (means, "do you want a hit?")

Insurance?

Double down?

Every time the deck is shuffled, the dealer so states and the "action" (total bet) and "standing" (total won or lost) is printed. To exit, hit the interrupt key (DEL) and the action and standing will be printed.

**NAME**

chase — Try to escape the killer robots

**SYNOPSIS**

`/usr/games/chase [ nrobots ] [ nfences ]`

**DESCRIPTION**

The object of the game *chase* is to move around inside of the box on the screen without getting eaten by the robots chasing you and without running into anything.

If a robot runs into another robot while chasing you, they crash and leave a junk heap. If a robot runs into a fence, it is destroyed.

If you can survive until all the robots are destroyed, you have won!

If you do not specify either *nrobots* or *nfences*, chase will prompt you for them.

The valid commands are described on the screen.

**NAME**

craps — the game of craps

**SYNOPSIS**

/usr/games/craps

**DESCRIPTION**

*Craps* is a form of the game of craps that is played in Las Vegas. The program simulates the *roller*, while the user (the *player*) places bets. The player may choose, at any time, to bet with the roller or with the *House*. A bet of a negative amount is taken as a bet with the House, any other bet is a bet with the roller.

The player starts off with a "bankroll" of \$2,000.

The program prompts with:

bet?

The bet can be all or part of the player's bankroll. Any bet over the total bankroll is rejected and the program prompts with *bet?* until a proper bet is made.

Once the bet is accepted, the roller throws the dice. The following rules apply (the player wins or loses depending on whether the bet is placed with the roller or with the House; the odds are even). The *first* roll is the roll immediately following a bet:

1. On the first roll:

7 or 11	wins for the roller;
2, 3, or 12	wins for the House;
any other number	is the <i>point</i> , roll again (Rule 2 applies).

2. On subsequent rolls:

point	roller wins;
7	House wins;
any other number	roll again.

If a player loses the entire bankroll, the House will offer to lend the player an additional \$2,000. The program will prompt:

marker?

A *yes* (or *y*) consummates the loan. Any other reply terminates the game.

If a player owes the House money, the House reminds the player, before a bet is placed, how many markers are outstanding.

If, at any time, the bankroll of a player who has outstanding markers exceeds \$2,000, the House asks:

Repay marker?

A reply of *yes* (or *y*) indicates the player's willingness to repay the loan. If only 1 marker is outstanding, it is immediately repaid. However, if more than 1 marker are outstanding, the House asks:

How many?

markers the player would like to repay. If an invalid number is entered (or just a carriage return), an appropriate message is printed and the program will prompt with *How many?* until a valid number is entered.

If a player accumulates 10 markers (a total of \$20,000 borrowed from the House), the program informs the player of the situation and exits.

Should the bankroll of a player who has outstanding markers exceed \$50,000, the *total* amount of money borrowed will be *automatically* repaid to the House.

Any player who accumulates \$100,000 or more breaks the bank. The program then prompts:

New game?

to give the House a chance to win back its money.

Any reply other than *yes* is considered to be a *no* (except in the case of *bet?* or *How many?*). To exit, send an interrupt (break), DEL, or control-D. The program will indicate whether the player won, lost, or broke even.

#### MISCELLANEOUS

The random number generator for the die numbers uses the seconds from the time of day. Depending on system usage, these numbers, at times, may seem strange but occurrences of this type in a real dice situation are not uncommon.

**NAME**

cribbage — the card game cribbage

**SYNOPSIS**

/usr/games/cribbage [ -[r][e][q] ] name ...

**DESCRIPTION**

*Cribbage* plays the card game cribbage, with the program playing one hand and the user the other. The program will initially ask the user if the rules of the game are needed -- if so, it will print out the appropriate section from *According to Hoyle with more (1)*.

*Cribbage* options include:

- e When the player makes a mistake scoring his hand or crib, provide an explanation of the correct score. (This is especially useful for beginning players.)
- q Print a shorter form of all messages -- this is only recommended for users who have played the game without specifying this option.
- r Instead of asking the player to cut the deck, the program will randomly cut the deck.

*Cribbage* first asks the player whether he wishes to play a short game (once around, to 61) or a long game (twice around, to 121). A response of 's' will result in a short game, any other response will play a long game.

At the start of the first game, the program asks the player to cut the deck to determine who gets the first crib. The user should respond with a number between 0 and 51, indicating how many cards down the deck is to be cut. The player who cuts the lower ranked card gets the first crib. If more than one game is played, the loser of the previous game gets the first crib in the current game.

For each hand, the program first prints the player's hand, whose crib it is, and then asks the player to discard two cards into the crib. The cards are prompted for one per line, and are typed as explained below.

After discarding, the program cuts the deck (if it is the player's crib) or asks the player to cut the deck (if it's its crib); in the later case, the appropriate response is a number from 0 to 39 indicating how far down the remaining 40 cards are to be cut.

After cutting the deck, play starts with the non-dealer (the person who doesn't have the crib) leading the first card. Play continues, as per cribbage, until all cards are exhausted. The program keeps track of the scoring of all points and the total of the cards on the table.

After play, the hands are scored. The program requests the player to score his hand (and the crib, if it is his) by printing out the appropriate cards (and the cut card enclosed in brackets). Play continues until one player reaches the game limit (61 or 121).

A carriage return when a numeric input is expected is equivalent to typing the lowest legal value; when cutting the deck this is equivalent to choosing the top card.

Cards are specified as rank followed by suit. The ranks may be specified as one of: 'a', '2', '3', '4', '5', '6', '7', '8', '9', 't', 'j', 'q', and 'k', or alternatively, one of: ace, two, three, four, five, six, seven, eight, nine, ten, jack, queen, and king. Suits may be specified as: 's', 'h', 'd', and 'c', or alternatively as: spades, hearts, diamonds, and clubs. A card may be specified as: <rank> <suit>, or: <rank> of <suit>. If the single letter rank and suit designations are used, the space separating the suit and rank may be left out. Also, if only one card of the desired rank is playable, typing the rank is sufficient. For example, if your hand was 2H, 4D, 5C, 6H, JC, KD and it was desired to discard the king of diamonds, any of the following could be typed: k, king, kd, k d, k of d, king d, king of d, k diamonds, k of diamonds, king diamonds, or king of diamonds.

**FILES**

/usr/games/cribbage

**AUTHOR**

Earl T. Cohen

**CUBIC (6)**

**SEE TTT**

**CUBIC (6)**



**NAME**

fish — play "Go Fish"

**SYNOPSIS**

`/usr/games/fish`

**DESCRIPTION**

*Fish* plays the game of Go Fish, a childrens' card game. The Object is to accumulate 'books' of 4 cards with the same face value. The players alternate turns; each turn begins with one player selecting a card from his hand, and asking the other player for all cards of that face value. If the other player has one or more cards of that face value in his hand, he gives them to the first player, and the first player makes another request. Eventually, the first player asks for a card which is not in the second player's hand: he replies 'GO FISH!' The first player then draws a card from the 'pool' of undealt cards. If this is the card he had last requested, he draws again. When a book is made, either through drawing or requesting, the cards are laid down and no further action takes place with that face value.

To play the computer, simply make guesses by typing a, 2, 3, 4, 5, 6, 7, 8, 9, 10, j, q, or k when asked. Hitting return gives you information about the size of my hand and the pool, and tells you about my books. Saying 'p' as a first guess puts you into 'pro' level; the default is pretty dumb.

**NAME**

fortune — print a random, hopefully interesting, adage

**SYNOPSIS**

fortune

**DESCRIPTION**

*Fortune* prints out a random adage.

**FILES**

*/usr/games/lib/fortunes*

**NAME**

hangman — guess the word

**SYNOPSIS**

/usr/games/hangman [ arg ]

**DESCRIPTION**

*Hangman* chooses a word at least seven letters long from a dictionary. The user is to guess letters one at a time.

The optional argument *arg* names an alternate dictionary.

**FILES**

/usr/lib/w2006

**BUGS**

Hyphenated compounds are run together.

**NAME**

life — play the game of life

**SYNOPSIS**

life [-r]

**DESCRIPTION**

Life is a pattern generating game set up for interactive use on a video terminal. The way it operates is: You use a series of commands to set up a pattern on the screen then let it generate further patterns from that pattern.

The algorithm used is: For each square in the matrix, look at it and its eight adjacent neighbors. If the present square is not occupied and exactly three of its neighbor squares are occupied, then that square will be occupied in the next pattern. If the present square is occupied and two or three of its neighbor squares are occupied, then that square will be occupied in the next pattern. Otherwise, the present square will not be occupied in the next pattern.

The edges of the screen are normally treated as an unoccupied void. If you specify the `-r` option on the command line, the screen is treated as a sphere; that is, the top and bottom lines are considered adjacent and the left and right columns are considered adjacent.

The pattern generation number and the number of occupied squares are displayed in the lower left hand corner.

Below is a list of commands available to the user. A # stands for any number. A ^ followed by a capital letter represents a control character.

- #,#a     Add a block of elements. The first number specifies the horizontal width. The second number specifies the vertical width. If a number is not specified, the default is 1.
- #c       Step through the next # patterns. If no number is specified, step forever. The operation can be aborted by typing rubout (delete).
- #,#d     Delete a block of elements. The first number specifies the horizontal width. The second number specifies the vertical width. If a number is not specified, the default is 1.
- #f       Generate a little flier at the present location. The number (modulo 8) determines the direction.
- #,#g     Move to absolute screen location. The first number specifies the horizontal location. The second number specifies the vertical location. If a number is not specified, the default is 0.
- #h       Move left # steps. If no number is specified, the default is 1.
- #j       Move down # steps. The default is 1.
- #k       Move up # steps. The default is 1.
- #l       Move right # steps. The default is 1.
- #n       Step through the next # patterns. If no number is specified, generate the next pattern. The operation can be aborted by typing rubout (delete).
- p        Put the last yanked or deleted block at the present location.

- q** Quit.
- #, #y** Yank a block of elements. The first number specifies the horizontal width. The second number specifies the vertical width. If a number is not specified, the default is 1.
- C** Clear the pattern.
- #F** Generate a big flier at the present location. The number (modulo 8) determines the direction.
- #H** Move to the left margin.
- #J** Move to the bottom margin.
- #K** Move to the top margin.
- #L** Move to the right margin.
- ^H** Move left # steps. If no number is specified, the default is 1.
- ^J** Move down # steps. The default is 1.
- ^K** Move up # steps. The default is 1.
- ^L** Move right # steps. The default is 1.
- ^R** Redraw the screen. This is used for those occasions when the terminal screws up.
- .** Repeat the last add (a) or delete (d) operation.
- ;** Repeat the last move (h, j, k, l) operation.

**BUGS**

The following features are planned but not implemented:

- #, #S** Save the selected area in a file.
- R** Restore from a file.
- m** Generate a macro command.
- !** Shell escape.
- e** Edit a file.
- i** Input commands from a file.

**AUTHOR**

Asa Romberger

**MAZE(6)**

**MAZE(6)**

**NAME**

maze — generate a maze

**SYNOPSIS**

/usr/games/maze

**DESCRIPTION**

*Maze* asks a few questions and then prints a maze.

**BUGS**

Some mazes (especially small ones) have no solutions.

**NAME**

moo — guessing game

**SYNOPSIS**

/usr/games/moo

**DESCRIPTION**

*Moo* is a guessing game imported from England. The computer picks a number consisting of four distinct decimal digits. The player guesses four distinct digits being scored on each guess. A "cow" is a correct digit in an incorrect position. A "bull" is a correct digit in a correct position. The game continues until the player guesses the number (a score of four bulls).

**NAME**

number — convert Arabic numerals to English

**SYNOPSIS**

*/usr/games/number*

**DESCRIPTION**

*Number* copies the standard input to the standard output, changing each decimal number to a fully spelled out version.



**NAME**

quiz — test your knowledge

**SYNOPSIS**

/usr/games/quiz [ -i file ] [ -t ] [ category1 category2 ]

**DESCRIPTION**

*Quiz* gives associative knowledge tests on various subjects. It asks items chosen from *category1* and expects answers from *category2*, or vice versa. If no categories are specified, *quiz* gives instructions and lists the available categories.

*Quiz* tells a correct answer whenever you type a bare new-line. At the end of input, upon interrupt, or when questions run out, *quiz* reports a score and terminates.

The `-t` flag specifies “tutorial” mode, where missed questions are repeated later, and material is gradually introduced as you learn.

The `-i` flag causes the named file to be substituted for the default index file. The lines of these files have the syntax:

```
line      = category new-line | category : line
category  = alternate | category | alternate
alternate = empty | alternate primary
primary   = character | [ category ] | option
option    = { category }
```

The first category on each line of an index file names an information file. The remaining categories specify the order and contents of the data in each line of the information file. Information files have the same syntax. Backslash `\` is used as with *sh*(1) to quote syntactically significant characters or to insert transparent new-lines into a line. When either a question or its answer is empty, *quiz* will refrain from asking it.

**FILES**

/usr/games/lib/quiz/index  
/usr/games/lib/quiz/\*

**BUGS**

The construct “`a|ab`” doesn’t work in an information file. Use “`a{b}`”.

**NAME**

rain — animated raindrops display

**SYNOPSIS**

rain

**DESCRIPTION**

*Rain's* display is modeled after the VAX/VMS program of the same name. The terminal has to be set for 9600 baud to obtain the proper effect.

As with all programs that use *termcap*, the TERM environment variable must be set (and exported) to the type of the terminal being used.

**FILES**

/etc/termcap

**AUTHOR**

Eric P. Scott

**NAME**

robots — Escape from the robots

**SYNOPSIS**

`/usr/games/robots`

**DESCRIPTION**

The object of the game *robots* is to move around inside of the box on the screen without getting eaten by the robots chasing you and without running into anything.

If a robot runs into another robot while chasing you, they crash and leave a junk heap.

You start out with 10 robots worth 10 points each. If you defeat all of them, you get 20 robots worth 20 points each. Then 30, etc. Until you get eaten!

The game keeps track of the top ten scores and prints them at the end of the game.

The valid commands are described on the screen.

**NAME**

trek - trekkie game

**SYNOPSIS**

/usr/games/trek [ [ -a ] file ]

**DESCRIPTION**

*Trek* is a game of space glory and war. Below is a summary of commands. For complete documentation, see *Trek* by Eric Allman.

If a filename is given, a log of the game is written onto that file. If the `--a` flag is given before the filename, that file is appended to, not truncated.

The game will ask you what length game you would like. Valid responses are short, medium, and long. You may also type restart, which restarts a previously saved game. You will then be prompted for the skill, to which you must respond novice, fair, good, expert, commadore, or impossible. You should normally start out with a novice and work up.

In general, throughout the game, if you forget what is appropriate the game will tell you what it expects if you just type in a question mark.

**COMMAND SUMMARY**

abandon	capture
cloak up/down	
computer request; ...	damages
destruct	dock
help	impulse course distance
lrscan	move course distance
phasers automatic amount	
phasers manual amt1 course1 spread1 ...	
torpedo course [yes] angle/no	
ram course distance	rest time
shell	shields up/down
srscan [yes/no]	
status	terminate yes/no
undock	visual course
warp warp_factor	

**AUTHOR**

Eric Allman

**NAME**

ttt, cubic - tic-tac-toe

**SYNOPSIS**

*/usr/games/ttt*

*/usr/games/cubic*

**DESCRIPTION**

*Ttt* is the X and O game popular in the first grade. This is a learning program that never makes the same mistake twice.

Although it learns, it learns slowly. It must lose nearly 80 games to completely know the game.

*Cubic* plays three-dimensional tic-tac-toe on a 4x4x4 board. Moves are specified as a sequence of three coordinate numbers in the range 1-4.

**FILES**

*/usr/games/ttt.k* learning file

**NAME**

twinkle — twinkle stars on the screen

**SYNOPSIS**

/usr/games/twinkle [-+ [ssave]] [density1] [density2]

**DESCRIPTION**

*Twinkle* causes a specified density of 'stars' to twinkle on the screen. The following options are available;

- print out the present screen density (the percentage of the screen that will be filled with stars) in the lower left hand corner of the screen. This number will change as stars go on and off.
- + do not 'randomize' before starting. The screen starts out completely blank and stars are added, bit by bit. In this case the density rises beyond the specified density, then falls to the required percentage.
- s save binary density on file 'save', in case you want to see the density curve that a particular density specification produced during the life of the show.

**density** If no density is specified, density is .5 (50% of the screen will be filled with stars).

If only *density1* is given, density is  $1/\text{density1}$

If both *density1* and *density2* are given, *density* is the resultant of  $\text{density1}/(\text{density1} + \text{density2})$ .

**EXAMPLE**

twinkle -+ 2 6

would start from a blank screen and twinkle stars to a final density of 2/8, or 25%. The densities would be shown in the lower left hand corner, as a three-place decimal.

**AUTHOR**

Asa Romberger

**NAME**

worm — Play the growing worm game

**SYNOPSIS**

worm [ *size* ]

**DESCRIPTION**

In *worm*, you are a little worm, your body is the "o"s on the screen and your head is the "@". You move with the hjkl keys (as in the game snake). If you don't press any keys, you continue in the direction you last moved. The upper case HJKL keys move you as if you had pressed several (9 for HL and 5 for JK) of the corresponding lower case key (unless you run into a digit, then it stops).

On the screen you will see a digit; if your worm eats the digit, it will grow longer. The actual amount by which the worm will grow longer depends upon which digit was eaten. The object of the game is to see how long you can make the worm grow.

The game ends when the worm runs into either the sides of the screen, or itself. The current score (how much the worm has grown) is kept in the upper left corner of the screen.

The optional argument, if present, is the initial length of the worm.

**BUGS**

If the initial length of the worm is set to less than one or more than 75, various strange things happen.

**NAME**

worms -- animate worms on a display terminal

**SYNOPSIS**

worms [ **-field** ] [ **-length #** ] [ **-number #** ] [ **-trail** ]

**DESCRIPTION**

**-field** makes a "field" for the worm(s) to eat; **-trail** causes each worm to leave a trail behind it. You can figure out the rest by yourself.

**FILES**

/etc/termcap

**DIAGNOSTICS**

Invalid length

Value not in range  $2 \leq \text{length} \leq 1024$

Invalid number of worms

Value not in range  $1 \leq \text{number} \leq 40$

**TERM**: parameter not set

The **TERM** environment variable is not defined. Do

**TERM**=terminal type

**export TERM**

Unknown terminal type

Your terminal type (as determined from the **TERM** environment variable) is not defined in /etc/termcap.

Terminal not capable of cursor motion

Your terminal is too stupid to run this program.

Out of memory

This should never happen.

**BUGS**

The lower-right-hand character position will not be updated properly on a terminal that wraps at the right margin.

Terminal initialization is not performed.

**AUTHOR**

Eric P. Scott



**NAME**

wump - the game of hunt-the-wumpus

**SYNOPSIS**

/usr/games/wump

**DESCRIPTION**

*Wump* plays the game of "Hunt the Wumpus." A Wumpus is a creature that lives in a cave with several rooms connected by tunnels. You wander among the rooms, trying to shoot the Wumpus with an arrow, meanwhile avoiding being eaten by the Wumpus and falling into Bottomless Pits. There are also Super Bats which are likely to pick you up and drop you in some random room.

The program asks various questions which you answer one per line; it will give a more detailed description if you want.

This program is based on one described in *People's Computer Company*, 2, 2 (November 1973).

**BUGS**

It will never replace Adventure.



## **Colophon**

**Composed at UniSoft Systems Inc.  
on the UniPlus<sup>+</sup> Operating System  
Designed by the Documentation Department  
Printed in Times Roman on Sequoia Matt**

