

SCALDsystem™
UNIX™
PROGRAMMER'S
MANUAL

900-00038 Rev C



2820 Orchard Parkway
San Jose, CA 95134
408/945-9400
Telex 371 9004

Copyright 1979, Bell Telephone Laboratories,
Incorporated. Holders of a UNIXTM/32V software
license are permitted to copy this document,
or any portion thereof, as necessary for licensed
use of the software, provided this copyright
notice and statement of permission are included.

PREFACE

This manual describes the UNIX operating system as implemented in SCALDsystem. While a majority of the commands, system calls, subroutines, and maintenance facilities are based on the 4.2 Berkeley Software Distribution, features from Western Electric System 3 and Bell Laboratories Version 7 also are included. In addition, a number of unique "private" extensions have been developed by Valid Logic Systems Incorporated and embodied within SCALDsystem UNIX.

TABLE OF CONTENTS

VOLUME 1

1. Commands and Application Programs

intro	introduction to commands
adb	debugger
admin	create and administer SCCS files
ar	archive and library maintainer
at	execute commands at a later date
awk	pattern scanning and processing language
basename	strip filename affixes
bdiff	differential file comparator for large files
cal	print calendar
calendar	reminder service
cat	catenate and print
cb	C program beautifier
cc	C compiler
cd	change working directory
chgrp	change group
chmod	change mode
chsh	change default login shell
clear	clear terminal screen
cmp	compare two files
col	filter reverse line feeds
comb	combine SCCS deltas
comm	select or reject lines common to two sorted files
cp	copy
cpio	copy file archives in and out
cptree	directory tree file copy utility
crypt	encode/decode
csh	a shell (command interpreter) with C-like syntax
ctags	create a tags file
date	print and set the date
dd	convert and copy a file
delta	make a delta (change) to an SCCS file
deroff	remove nroff, troff, and eqn constructs
df	disk free
diff	differential file and directory comparator
diff3	3-way differential file comparator
du	summarize disk usage
echo	echo arguments
ed	text editor
ex	text editor
expand	expand tabs to spaces and vice versa
expr	evaluate arguments as an expression

Table of Contents

false	provide truth tables
file	determine file type
find	find files
get	get a version of an SCCS file
getline	get input line
grep	search a file for a pattern
groups	show group memberships
head	give first few lines
hostid	set or print identifier of current host system
hostname	set or print name of current host system
install	install binaries
iostat	report I/O statistics
join	relational database operator
kill	terminate a process with extreme prejudice
last	indicate last logins of users and teletypes
ld	link editor
lex	generator of lexical analysis programs
lint	a C program verifier
ln	make links
login	sign on
look	find files in a sorted list
lorder	find ordering relation for an object library
lpq	spool queue examination program
lpr	off line print
lprm	remove jobs from the line printer spooling queue
ls	list contents of directory
mail	send and receive mail
make	maintain program groups
man	find manual information by keywords; print out the manual
mesg	print or deny messages
mkdir	make a directory
more	file perusal filter for crt viewing
mv	move or rename files
newgrp	log in to a new group
nice	run a command at low priority (sh only)
nm	print name list
nroff	text formatting
od	octal, decimal, hex, ascii dump
pagesize	print system page size
passwd	change login password
plot	graphics filters
pr	pr to the line printer
printenv	print out the environment
prof	display profile data
ps	process status
pwd	working directory name
rev	reverse lines of a file
rlogin	remote login
rm	remove (unlink) files or directories
rmdel	remove a delta from an SCCS file
rmdir	remove (unlink) directories or files
rsh	remote shell
ruptime	show host status of local machines

rwho	who's logged in on local machines
sccshelp	ask for SCCS help
script	make typescript of terminal session
sed	stream editor
sh	command language
shownet	show Valid node status
size	size of an object file
sleep	suspend execution for an interval
sort	sort or merge files
spell	find spelling errors
split	split a file into pieces
strings	find printable strings in an object, or other binary, file
strip	remove symbols and relocation bits
stty	set terminal options
su	substitute user id temporarily
sum	sum and count blocks in a file
tail	deliver the last part of a file
tar	tape archiver
tee	pipe fitting
test	condition command
time	time a command
touch	update date last modified of a file
tp	manipulate tape archiver
tr	translate characters
troff	text formatting and typesetting
true	provide truth tables
tsort	topological sort
tty	get terminal name
unset	undo a previous get of an SCCS file
uniq	report repeated lines in a file
units	conversion program
vfontinfo	inspect and print out information about UNIX fonts
vi	screen oriented (visual) display editor based on ex
vpl	copy a spooled plot to raster printer/plotter
vpr	raster printer/plotter spooler
vtroff	troff to raster plotter
wall	write to all users
wc	word count
what	show what versions of object modules were used to construct a file
who	who is on the system
whoami	print effective current user id
write	write to another user
xstr	extract strings from C programs to implement shared strings
yacc	yet another compiler-compiler
yes	be repetitively affirmative

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 file
bind	bind a name to a socket
brk	change data segment size
chdir	change current working directory
chmod	change mode of file
chown	change owner and group of a file
chroot	change root directory
close	delete a descriptor
connect	initiate a connection on a socket
creat	create a new file
dup	duplicate a descriptor
execve	execute a file
exit	terminate a process
flock	apply or remove an advisory lock on an open file
fork	create a new process
fsync	synchronize a file's in-core state with that on a disk
getdtablesize	get descriptor table size
getgid	get group identity
getgroups	get group access list
gethostid	get/set unique identifier of current host
gethostname	get/set name of current host
getitimer	get/set value of interval timer
getpagesize	get system page size
getpgrp	get process group
getpid	get process identification
getpriority	get/set program scheduling priority
getrlimit	control maximum system resource consumption
getrusage	get information about resource utilization
getsockopt	get and set options on sockets
gettimeofday	get/set date and time
getuid	get user identity
ioctl	control device
kill	send signal to a process
killpg	send signal to a process group
link	make a hard link to a file
listen	listen for connections on a socket
lseek	move read/write pointer
mkdir	make a directory file
mknod	make a special file
mount	mount or remove file system
open	open a file for reading or writing, or create a new file
pipe	create an interprocess communication channel
profil	execution time profile
ptrace	process trace
read	read input
readlink	read value of symbolic link
reboot	reboot system or halt processor
recv	receive a message from a socket
rename	change the name of a file

Table of Contents

rmdir remove a directory file
send send a message from a socket
setgroups set group access list
setpgrp set process group
setregid set real and effective group ID
setreuid set real and effective user ID's
shutdown shut down part of a full-duplex connection
socket create an endpoint for communication
stat get file status
symlink make symbolic link to a file
sync update super-block
syscall indirect system call
truncate truncate a file to a specified length
umask set file creation mode mask
unlink remove directory entry
utimes set file times
vfork spawn new process in a virtual memory efficient way
vhangup virtually "hangup" the current control terminal
wait wait for process to terminate
write write on a file

Table of Contents

3. C Library Subroutines

intro	introduction to library functions
abort	generate a fault
abs	integer absolute value
atof	convert ASCII to numbers
bstring	bit and byte string operations
crypt	DES encryption
ctime	convert date and time to ASCII
ctype	character classification macros
directory	directory operations
ecvt	output conversion
end	last locations in program
execl	execute a file
exit	terminate a process after flushing any pending output
frexp	split into mantissa and exponent
getenv	value for environment name
getgrent	get group file entry
getlogin	get login name
getpass	read a password
getpwent	get password file entry
getwd	get current working directory pathname
malloc	memory allocator
mktemp	make a unique filename
monitor	prepare execution profile
nlist	get entries from name list
perror	system error messages
popen	initiate I/O to/from a process
psignal	system signal messages
qsort	quicker sort
random	better random number generator; routines for changing generators
regex	regular expression handler
scandir	scan a directory
setjmp	non-local goto
setuid	set user and group ID
sleep	suspend execution for interval
string	string operations
swab	swap bytes
system	issue a shell command
ttyname	find name of a terminal
varargs	variable argument list

3M. Math Library

intro	introduction to mathematical library functions
exp	exponential, logarithm, power, square root
floor	absolute value, floor, ceiling functions
gamma	log gamma function
hypot	Euclidean distance
j0	bessel functions
sin	trigonometric functions
sinh	hyperbolic functions

3N. Internet Network Library

byteorder convert values between host and network byte order
gethostent get network host entry
getnetent get network entry
getprotoent get protocol entry
getservent get service entry
inet Internet address manipulation routines

3S. C Standard I/O Library Subroutines

intro standard buffered input/output package
fclose close or flush a stream
ferror stream status inquires
fopen open a stream
fread buffered binary input/output
fseek reposition a stream
getc get character or word from stream
gets get a string from a stream
printf formatted output conversion
putc put character or word on a stream
puts put a string on a stream
scanf formatted input conversion
setbuf assign buffering to a stream
ungetc push character back into input stream

3X. Other Libraries

intro introduction to miscellaneous library functions
curses screen functions with "optimal" cursor motion
getfsent get file system descriptor file entry
initgroups initialize group access list
termcap terminal independent operation routines

3C. Compatibility Library Subroutines

intro introduction to compatibility library functions
alarm schedule signal after specified time
getpw get name from uid
nice set program priority
pause stop until signal
rand random number generator
signal simplified software signal facilities
stty set and get terminal state (defunct)
time get date and time
times get process times
utime set file times
vlimit control maximum system resource consumption
vtimes get information about resource utilization

Table of Contents

4. Special Files

intro	introduction to special files and hardware support
drum	paging device
ec	3Com 10Mb/s Ethernet interface
ip	internet protocol
lo	software loopback network interface
lp	line printer
mem	main memory
mtio	UNIX magtape interface
null	data sink
pty	pseudo terminal driver
tty	general terminal interface
va	Benson-Varian interface
vp	Versatec interface

5. File Formats

a.out	assembler and link editor output
ar	archive (library) file format
core	format of memory image file
dir	format of directories
disktab	disk description file
dump	incremental dump format
fs	format of file system volume
fstab	static information about the filesystems
group	group file
hosts	host name data base
mtab	mounted file system table
passwd	password file
protocols	protocol name data base
services	service name data base
stab	symbol table types
tar	tape archive file format
termcap	terminal capability data base
ttys	terminal initialization data
ttytype	data base of terminal types by port
types	primitive system data types
utmp	login records
vfont	font formats for the Benson-Varian or Versatec

7. Miscellaneous

intro	miscellaneous useful information pages
ascii	map of ASCII character set
environ	user environment
eqnchar	special character definitions for eqn
hier	file system hierarchy
mailaddr	mail addressing description
man	macros to typeset manual
me	macros for formatting papers

ms text formatting macros
term conventional names for terminals

8. System Maintenance

intro introduction to system maintenance and operation commands
ac login accounting
backup user-friendly backup procedure
chkhosts maintain network host tables
chown change owner
chuid change user/group ids on directory trees
clri clear i-node
conn administer extended file system
crash what happens when the system crashes
cron clock daemon
dcheck file system directory consistency check
dmesg collect system diagnostic messages to form error log
dump incremental file system dump
efsioctl EFS superuser access control
ether 3Com 3C400 driver control program
fsck file system consistency check and interactive repair
getfs list file systems
getty set terminal mode
halt stop the processor
icheck file system storage consistency check
init process control initialization
lpd line printer daemon
makekey generate encryption key
mkconfig maintain configuration file
mkfs construct a file system
mklost+found make a lost+found directory for fsck
mknod build special file
mkproto construct a prototype file system
mkusr procedure for adding new users
mount mount and dismount file system
ncheck generate names from i-numbers
newfs construct a new file system
pstat print system files
rc command script for auto-reboot and daemons
reboot reboot or halt system
restor incremental file system restore
rimioctl send ioctl commands to Rimfire controller
rlogind remote login server
rshd remote shell server
rwhod system status server
savecore save a core dump of the operating system
shutdown close down the system at a specific time
sync update the super block
tunefs tune up an existing file system
update periodically update the super block

Table of Contents

VOLUME 2

An Introduction to the C Shell

An Introduction to Display Editing with Vi

NROFF/TROFF User's Manual

INTRODUCTION TO VOLUME 1

This volume gives descriptions of the publicly available features of the UNIX/32V† system, as extended to provide a virtual memory environment and other enhancements at U. C. Berkeley. It does not attempt to provide perspective or tutorial information upon the UNIX operating system, its facilities, or its implementation. Various documents on those topics are contained in Volume 2. In particular, for an overview see 'The UNIX Time-Sharing System' by Ritchie and Thompson; for a tutorial see 'UNIX for Beginners' by Kernighan, and for an guide to the new features of this virtual version, see 'Getting started with Berkeley Software for UNIX on the VAX' in volume 2C.

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

The volume is divided into eight sections:

1. Commands
2. System calls
3. Subroutines
4. Special files
5. File formats and conventions
6. Games
7. Macro packages and language conventions
8. Maintenance commands and procedures

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

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

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

The special files section 4 discusses the characteristics of each system 'file' that actually refers to an I/O device. The names in this section refer to the DEC device names for the hardware, instead of the names of the special files themselves.

The file formats and conventions section 5 documents the structure of particular kinds of files; for example, the form of the output of the loader and assembler is given. Excluded are files used by only one command, for example the assembler's intermediate files.

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

† UNIX is a trademark of Bell Laboratories.

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

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

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

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

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

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

Boldface words are considered literals, and are typed just as they appear.

Square brackets [] around an argument indicate that the argument is optional. When an argument is given as 'name', it always refers to a file name.

Ellipses '...' are used to show that the previous argument-prototype may be repeated.

A final convention is used by the commands themselves. An argument beginning with a minus sign '-' is often taken to mean some sort of option-specifying argument even if it appears in a position where a file name could appear. Therefore, it is unwise to have files whose names begin with '-'.

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

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

A *see also* subsection gives pointers to related information.

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

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

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

HOW TO GET STARTED

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

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

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

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

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

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

The evidence that you have successfully logged in is that a shell program will type a prompt ('\$' or '%') to you. (The shells are described below under 'How to run a program.')

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

Logging out. There are three ways to log out:

By typing an end-of-file indication (EOT character, control-d) to the Shell. The Shell will terminate and the 'login:.' message will appear again.

You can log in directly as another user by giving a *login(1)* command.

If worse comes to worse, you can simply hang up the phone; but beware — some machines may lack the necessary hardware to detect that the phone has been hung up. Ask your system administrator if this is a problem on your machine.

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

UNIX terminal I/O is full-duplex. It has full read-ahead, which means that you can type at any time, even while a program is typing at you. Of course, if you type during output, the printed output will have the input characters interspersed. However, whatever you type will be saved up and interpreted in correct sequence. There is a limit to the amount of read-ahead, but it is generous and not likely to be exceeded unless the system is in trouble. When the read-ahead limit is exceeded, the system throws away all the saved characters (or beeps, if your prompt was a %).

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

The 'break' or 'interrupt' key causes an *interrupt signal*, as does the ASCII 'delete' (or 'rubout') character, which is not passed to programs. This signal generally causes whatever program you

are running to terminate. It is typically used to stop a long printout that you don't want. However, programs can arrange either to ignore this signal altogether, or to be notified when it happens (instead of being terminated). The editor, for example, catches interrupts and stops what it is doing, instead of terminating, so that an interrupt can be used to halt an editor printout without losing the file being edited. Many users change this interrupt character to be ^C (control-C) using *stty(1)*.

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

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

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

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

How to run a program; the shells. When you have successfully logged in, a program called a shell is listening to your terminal. The shell reads typed-in lines, splits them up into a command name and arguments, and executes the command. A command is simply an executable program. The Shell looks in several system directories to find the command. You can also place commands in your own directory and have the shell find them there. There is nothing special about system-provided commands except that they are kept in a directory where the shell can find them.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Text processing. Almost all text is entered through the editor *ex(1)* (often entered via *vi(1)*). The commands most often used to write text on a terminal are: *cat*, *pr*, *more* and *nroff*, all in section 1.

The *cat* command simply dumps ASCII text on the terminal, with no processing at all. The *pr* command paginates the text, supplies headings, and has a facility for multi-column output. *Nroff* is an elaborate text formatting program. Used naked, it requires careful forethought, but for ordinary documents it has been tamed; see *me(7)* and *ms(7)*.

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

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

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

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

files.

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

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

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

When you log in, a message-of-the-day may greet you before the first prompt.

CONVERTING FROM THE 6TH EDITION

There follows a catalogue of significant, mostly incompatible, changes that will affect old users converting from the sixth edition on a PDP-11. No attempt is made to list all new facilities, or even all minor, but easily spotted changes, just the bare essentials without which it will be almost impossible to do anything.

Addressing files. Byte addresses in files are now long (32-bit) integers. Accordingly *seek* has been replaced by *lseek(2)*. Every program that contains a *seek* must be modified. *Stat* and *fstat(2)* have been affected similarly, since file lengths are now 32- rather than 24-bit quantities.

Assembly language. This language is dead. Necromancy will be severely punished.

Stty and *gtty.* These system calls have been extensively altered, see *ioctl(2)* and *tty(4)*.

C language, lint. The syntax for initialization requires an equal sign = before an initializer, and brackets { } around compound initial values; arrays and structures are now initialized honestly. Assignment operators such as += and -= are now written in the reverse order: +=, -=. This removes the possibility of ambiguity in constructs such as x=-2, y=*p, and a=/*b. You will also certainly want to learn about

- long integers
- type definitions
- casts (for type conversion)
- unions (for more honest storage sharing)
- #include <filename> (which searches in standard places)

The program *lint(1)* checks for obsolete syntax and does strong type checking of C programs, singly or in groups that are expected to be loaded together. It is indispensable for conversion work.

Fortran. The old *fc* is replaced by *f77*, a true compiler for Fortran 77, compatible with C. There are substantial changes in the language; see 'A Portable Fortran 77 Compiler' in Volume 2.

Stream editor. The program *sed(1)* is adapted to massive, repetitive editing jobs of the sort encountered in converting to the new system. It is well worth learning.

Standard I/O. The old *fopen,getc,putc* complex and the old *-lp* package are both dead, and even *getchar* has changed. All have been replaced by the clean, highly efficient, *stdio* package, *intro(3S)*. The first things to know are that *getchar(3)* returns the integer EOF (-1) (which is not a possible byte value) on end of file, that 518-byte buffers are out, and that there is a defined FILE data type.

Make. The program *make(1)* handles the recompilation and loading of software in an orderly way from a 'makefile' recipe given for each piece of software. It remakes only as much as the modification dates of the input files show is necessary. The makefiles will guide you in building your new system.

Shell, chdir. F. L. Bauer once said Algol 68 is the Everest that must be climbed by every computer scientist because it is there. So it is with the shell for UNIX users. Everything beyond simple command invocation from a terminal is different. Even *chdir* is now spelled *cd*. If you wish to use *sh* (as opposed to *csh*) then you will want to study *sh(1)* long and hard.

C shell. *Csh(1)*, developed at Berkeley, has features comparable to *sh*. It includes a history mechanism that saves you from retyping all or part of previous commands, as well as an efficient aliasing (macro) mechanism. The job control facilities of the system, which make the system much more pleasant to use, are currently available only with *csh*. See *csh(1)* for a description. These features make *csh* pleasant to use interactively. *Csh* programs have a syntax reminiscent of *C*, while *sh* command programs have a syntax reminiscent of ALGOL-68.

Debugging. *Sdb* is a far more capable replacement for the debugger *cdb*, and debugs *C* and Fortran at the source level. For machine language debugging, *adb* replaces *db*. The first-time user should be especially careful about distinguishing */* and *?* in *adb* commands, and watching to make sure that the *x* whose value he asked for is the real *x*, and not just some absolute location equal to the stack offset of some automatic *x*. You can always use the 'true' name, *_x*, to pin down a *C* external variable.

Dsw. This little-known, but indispensable facility has been taken over by *rm -ri*.

Boot procedures. Needless to say, these are all different. See section 8 of this volume, and the other documentation you should have received with your tape.

CONVERTING FROM THE DECEMBER, 1979 BERKELEY DISTRIBUTION

There have been a number of significant changes and improvements in the system. This list just gives the bare essentials:

C language changes. The *C* compiler now accepts and checks essentially arbitrary length identifiers and preprocessor names. There is a new type available in type casts: *void* which signifies that a value is to be ignored. It is useful in keeping lint happy about values which are not used (especially values returned from procedures). Finally, the language has been changed so that field names need not be unique to structures; on the other hand, the compiler insists that you be more honest about types involved in pointer constructs or it will warn you.

Object file format. The object file format has been changed to include a string table, so that language compilers may have names longer than 8 characters in their resulting *a.out* files. Old *.o* files must be recreated. *A.out* files will still run on both this and the December 1979 version of the system; only the symbol tables are incompatible.

Archive format and table of contents. The archive format has been changed to one which is portable between the VAX and other machines (e.g. the PDP-11). Old VAX archives should be converted with *arcv(8)*; loader archives should just be recreated since the object files are also obsolete. Loader archives should have table-of-contents added by *ranlib(1)*; if they don't the loader will gripe when they are used.

New tty driver, job control facilities and csh. Hand in hand are new job control facilities, a new tty driver and a new version of the *C* shell which supports and uses all of this. See *tty(4)* and *csh(1)* for a quick introduction.

Pascal compiler. There is a true Pascal compiler, *pc(1)* which allows separate compilation as well as mixing in of FORTRAN and *C* code.

Error analyzer. There is an error analyzer program *error(1)*, which takes a set of error message and merges them back into the source files at the point of error. It can be used interactively to avoid inserting errors which are uninteresting. This program eliminates once and for all making lists of errors on small scraps of paper.

Mail forwarding. The system now provides mail forwarding and distribution facilities. Group and aliases are defined in the file *usr/lib/aliases* see *aliases(5)*. If you change this file you will have to rerun *newaliases(1)*. For any particular system a table in the source of the *delivermail* postman program may have to be changed so that it knows about the gateways on the local

machine.

System bootstrap procedures. These are totally changed; the system performs automatic reboots and preens the disks automatically at reboot. You should reread the appropriate pages in section 8 if you deal with system reboots.

CONVERTING FROM THE JUNE, 1981 BERKELEY DISTRIBUTION

Many many changes have been made. This list indicates those which are most visible to users.

Directory format. Directory entries are no longer fixed length. This forces user programs which read directories to be modified to use the *directory(3)* package.

Signals. A new signal package has replaced the previous signal mechanism as well as the "jobs library". When using the compatible *signal(3C)* interface routine, the two most important changes are: signal handlers are not reset to SIG_DFL when a process receives a signal, and while a signal handler is processing a signal, that signal is blocked until the handler returns. This has implications, in particular, for programs which process the suspend character typed at the terminal. Refer to *sigvec*, *sigblock*, *sigpause*, *sigstack*, and *sigsetmask(2)* for information about the new signal facilities.

File and path names. File names may now be up to 255 characters in length. Path names are restricted to be at most 1024 characters. These two constants are provided as MAXNAMLEN and MAXPATHLEN in *<sys/dir.h>* and *<sys/param.h>*, respectively.

System time. System time is provided in microsecond precision with 10 millisecond accuracy. The new system call *gettimeofday(2)* supplants the old *time(3)* call which is now a library routine. The major impact of this change is that programs are now written in a fashion which is independent of the line clock frequency.

Groups. A user may now be in many groups simultaneously. This has obviated the need for the *newgrp* command. See *getgroups(2)* for more information.

Stat and fstat return value. The structure returned by the *stat* and *fstat* system calls is now larger. This is due to inode numbers growing to 32-bits, time stamps expanding to 64-bits and other information being included in the return value. Consult *stat(2)* for more information.

Mail forwarding. The system now provides general internetwork mail forwarding and distribution facilities. The *sendmail(8)* program replaces the old *delivermail* facility.

Debuggers. The previous C source language debugger, *sdb*, has been replaced by a new one, *dbx(1)*. *Adb(1)* has been extended to simplify debugging of the operating system.

Networking support. Many new user programs provide access to the networking facilities. The *rlogin(1C)* and *rsh(1C)* programs are intended for communicating between UNIX systems. The *telnet(1C)* and *ftp(1C)* programs support the DARPA Internet standard protocols. The *netstat(1)* program is useful in watching network activity.

NAME

intro — introduction to commands

DESCRIPTION

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

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

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

SEE ALSO

Section (6) for computer games.

How to get started, in the Introduction.

DIAGNOSTICS

Upon termination each command returns two bytes of status, one supplied by the system giving the cause for termination, and (in the case of 'normal' termination) one supplied by the program, see *wait* and *exit(2)*. The former byte is 0 for normal termination, the latter is customarily 0 for successful execution, nonzero to indicate troubles such as erroneous parameters, bad or inaccessible data, or other inability to cope with the task at hand. It is called variously 'exit code', 'exit status' or 'return code', and is described only where special conventions are involved.

NAME

`adb` — debugger

SYNOPSIS

`adb [-w] [-k] [-ldir] [objfil [corfil]]`

DESCRIPTION

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

Objfil is normally an executable program file, preferably containing a symbol table; if not then the symbolic features of *adb* cannot be used although the file can still be examined. The default for *objfil* is `a.out`. *Corfil* is assumed to be a core image file produced after executing *objfil*, the default for *corfil* is `core`.

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

The `-k` option makes *adb* do UNIX kernel memory mapping; it should be used when *core* is a UNIX crash dump or `/dev/mem`.

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

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

In general requests to *adb* are of the form

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

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

The interpretation of an address depends on the context it is used in. If a subprocess is being debugged then addresses are interpreted in the usual way in the address space of the subprocess. If the operating system is being debugged either post-mortem or using the special file `/dev/mem` to interactive examine and/or modify memory the maps are set to map the kernel virtual addresses which start at `0x80000000` (on the VAX). ADDRESSES.

EXPRESSIONS

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

integer A number. The prefixes `0o` and `0O` ("zero oh") force interpretation in octal radix; the prefixes `0t` and `0T` force interpretation in decimal radix; the prefixes `0x` and `0X` force interpretation in hexadecimal radix. Thus `0o20 = 0t16 = 0x10 = sixteen`. If no prefix appears, then the *default radix* is used; see the `$d` command. The default radix is initially hexadecimal. The hexadecimal digits are `0123456789abcdefABCDEF` with the obvious values. Note that a hexadecimal number whose most significant digit would otherwise be an alphabetic character must have a `0x` (or `0X`) prefix (or a leading zero if the default radix is hexadecimal).

integer.fraction

A 32 bit floating point number.

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

< name

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

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

_ symbol

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

routine.name

The address of the variable *name* in the specified C routine. Both *routine* and *name* are *symbols*. If *name* is omitted the value is the address of the most recently activated C stack frame corresponding to *routine*. (This form is currently broken on the VAX; local variables can be examined only with *dbx(1)*.)

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

Monadic operators

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

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

-exp Integer negation.

~exp Bitwise complement.

#exp Logical negation.

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

e1 + e2 Integer addition.

e1 - e2 Integer subtraction.

e1 * e2 Integer multiplication.

e1 % e2 Integer division.

e1 & e2 Bitwise conjunction.

e1 | e2 Bitwise disjunction.

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

COMMANDS

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

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

/f Locations starting at *address* in *corfil* are printed according to the format *f* and *dot* is incremented as for '?'.
=f The value of *address* itself is printed in the styles indicated by the format *f*. (For i format '?' is printed for the parts of the instruction that reference subsequent words.)

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

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

newline

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

[?/]1 *value mask*

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

[?/]w *value ...*

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

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

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

> *name* *Dot* is assigned to the variable or register named.

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

Miscellaneous commands. The available *modifiers* are:

- <*f* Read commands from the file *f*. If this command is executed in a file, further commands in the file are not seen. If *f* is omitted, the current input stream is terminated. If a *count* is given, and is zero, the command will be ignored. The value of the count will be placed in variable 9 before the first command in *f* is executed.
- <<*f* Similar to < except it can be used in a file of commands without causing the file to be closed. Variable 9 is saved during the execution of this command, and restored when it completes. There is a (small) finite limit to the number of << files that can be open at once.
- >*f* Append output to the file *f*, which is created if it does not exist. If *f* is omitted, output is returned to the terminal.
- ? Print process id, the signal which caused stoppage or termination, as well as the registers as \$r. This is the default if *modifier* is omitted.
- r Print the general registers and the instruction addressed by pc. *Dot* is set to pc.
- b Print all breakpoints and their associated counts and commands.
- c C stack backtrace. If *address* is given then it is taken as the address of the current frame instead of the contents of the frame-pointer register. If C is used then the names and (32 bit) values of all automatic and static variables are printed for each active function. (broken on the VAX). If *count* is given then only the first *count* frames are printed.
- d Set the default radix to *address* and report the new value. Note that *address* is interpreted in the (old) current radix. Thus "10\$d" never changes the default radix. To make decimal the default radix, use "0t10\$d".
- e The names and values of external variables are printed.
- w Set the page width for output to *address* (default 80).
- s Set the limit for symbol matches to *address* (default 255).
- o All integers input are regarded as octal.
- q Exit from *adb*.
- v Print all non zero variables in octal.
- m Print the address map.
- p (*Kernel debugging*) Change the current kernel memory mapping to map the designated **user structure** to the address given by the symbol *_u*. The *address* argument is the address of the user's user page table entries (on the VAX).

:*modifier*

Manage a subprocess. Available modifiers are:

- bc Set breakpoint at *address*. The breakpoint is executed *count*-1 times before causing a stop. Each time the breakpoint is encountered the command *c* is executed. If this command is omitted or sets *dot* to zero then the breakpoint

- causes a stop.
- d** Delete breakpoint at *address*.
 - r** Run *objfil* as a subprocess. If *address* is given explicitly then the program is entered at this point; otherwise the program is entered at its standard entry point. *count* specifies how many breakpoints are to be ignored before stopping. Arguments to the subprocess may be supplied on the same line as the command. An argument starting with < or > causes the standard input or output to be established for the command.
 - cs** The subprocess is continued with signal *s*, see *sigvec*(2). If *address* is given then the subprocess is continued at this address. If no signal is specified then the signal that caused the subprocess to stop is sent. Breakpoint skipping is the same as for **r**.
 - ss** As for **c** except that the subprocess is single stepped *count* times. If there is no current subprocess then *objfil* is run as a subprocess as for **r**. In this case no signal can be sent; the remainder of the line is treated as arguments to the subprocess.
 - k** The current subprocess, if any, is terminated.

VARIABLES

Adb provides a number of variables. Named variables are set initially by *adb* but are not used subsequently. Numbered variables are reserved for communication as follows.

- 0 The last value printed.
- 1 The last offset part of an instruction source.
- 2 The previous value of variable 1.
- 9 The count on the last \$< or \$<< command.

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

- b** The base address of the data segment.
- d** The data segment size.
- e** The entry point.
- m** The 'magic' number (0407, 0410 or 0413).
- s** The stack segment size.
- t** The text segment size.

ADDRESSES

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

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

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

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

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

FILES

a.out
core

SEE ALSO

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

DIAGNOSTICS

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

BUGS

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

NAME

`admin` — create and administer SCCS files

SYNOPSIS

```
admin [-n] [-i[name]] [-rrel] [-t[name]] [-fflag[flag-val]]
[-dflag[flag-val]] [-alogin] [-elogin] [-m[mrlist]] [-y[comment]]
[-h] [-z] files
```

DESCRIPTION

Admin is used to create new SCCS files and change parameters of existing ones. Arguments to *admin*, which may appear in any order, consist of keyletter arguments, which begin with `-`, and named files (note that SCCS file names must begin with the characters `s.`). If a named file doesn't exist, it is created, and its parameters are initialized according to the specified keyletter arguments. Parameters not initialized by a keyletter argument are assigned a default value. If a named file does exist, parameters corresponding to specified keyletter arguments are changed, and other parameters are left as is.

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

The keyletter arguments are as follows. Each is explained as though only one named file is to be processed since the effects of the arguments apply independently to each named file.

- `-n` This keyletter indicates that a new SCCS file is to be created.
- `-i[name]` The *name* of a file from which the text for a new SCCS file is to be taken. The text constitutes the first delta of the file (see `-r` keyletter for delta numbering scheme). If the `i` keyletter is used, but the file name is omitted, the text is obtained by reading the standard input until an end-of-file is encountered. If this keyletter is omitted, then the SCCS file is created empty. Only one SCCS file may be created by an *admin* command on which the `i` keyletter is supplied. Using a single *admin* to create two or more SCCS files require that they be created empty (no `-i` keyletter). Note that the `-i` keyletter implies the `-n` keyletter.
- `-rrel` The *release* into which the initial delta is inserted. This keyletter may be used only if the `-i` keyletter is also used. If the `-r` keyletter is not used, the initial delta is inserted into release 1. The level of the initial delta is always 1 (by default initial deltas are named 1.1).
- `-t[name]` The *name* of a file from which descriptive text for the SCCS file is to be taken. If the `-t` keyletter is used and *admin* is creating a new SCCS file (the `-n` and/or `-i` keyletters also used), the descriptive text file name must also be supplied. In the case of existing SCCS files: (1) a `-t` keyletter without a file name causes removal of descriptive text (if any) currently in the SCCS file, and (2) a `-t` keyletter with a file

- name causes text (if any) in the named file to replace the descriptive text (if any) currently in the SCCS file.
- f***flag* This keyletter specifies a *flag*, and, possibly, a value for the *flag*, to be placed in the SCCS file. Several f keyletters may be supplied on a single *admin* command line. The allowable *flags* and their values are:
- b** Allows use of the **-b** keyletter on a *get*(1) command to create branch deltas.
 - c***ceil* The highest release (i.e., "ceiling"), a number less than or equal to 9999, which may be retrieved by a *get*(1) command for editing. The default value for an unspecified c flag is 9999.
 - f***floor* The lowest release (i.e., "floor"), a number greater than 0 but less than 9999, which may be retrieved by a *get*(1) command for editing. The default value for an unspecified f flag is 1.
 - d***SID* The default delta number (SID) to be used by a *get*(1) command.
 - i** Causes the "No id keywords (ge6)" message issued by *get*(1) or *delta*(1) to be treated as a fatal error. In the absence of this flag, the message is only a warning. The message is issued if no SCCS identification keywords (see *get*(1)) are found in the text retrieved or stored in the SCCS file.
 - j** Allows concurrent *get*(1) commands for editing on the same SID of an SCCS file. This allows multiple concurrent updates to the same version of the SCCS file.
 - l***list* A *list* of releases to which deltas can no longer be made (*get* -e against one of these "locked" releases fails). The *list* has the following syntax:


```
<list> ::= <range> | <list> , <range>
<range> ::= RELEASE NUMBER | a
```

 The character **a** in the *list* is equivalent to specifying *all releases* for the named SCCS file.
 - n** Causes *delta*(1) to create a "null" delta in each of those releases (if any) being skipped when a delta is made in a *new* release (e.g., in making delta 5.1 after delta 2.7, releases 3 and 4 are skipped). These null deltas serve as "anchor points" so that branch deltas may later be created from them. The absence of this flag causes skipped releases to be non-existent in the SCCS file preventing branch deltas from being created from them in the future.
 - q***text* User definable text substituted for all occurrences of the %Q% keyword in SCCS file text retrieved by *get*(1).
 - m***mod* Module name of the SCCS file substituted for all occurrences of the %M% keyword in SCCS file text retrieved by *get*(1). If the **m** flag is not specified, the value assigned is the name of the SCCS file with the

- leading s. removed.
- ttype** Type of module in the SCCS file substituted for all occurrences of %Y% keyword in SCCS file text retrieved by *get*(1).
- v[pgm]** Causes *delta*(1) to prompt for Modification Request (MR) numbers as the reason for creating a delta. The optional value specifies the name of an MR number validity checking program (see *delta*(1)). (If this flag is set when creating an SCCS file, the *m* keyletter must also be used even if its value is null).
- dflag** Causes removal (deletion) of the specified *flag* from an SCCS file. The *-d* keyletter may be specified only when processing existing SCCS files. Several *-d* keyletters may be supplied on a single *admin* command. See the *-f* keyletter for allowable *flag* names.
- llist** A *list* of releases to be "unlocked". See the *-f* keyletter for a description of the *l* flag and the syntax of a *list*.
- alogin** A *login* name, or numerical UNIX group ID, to be added to the list of users which may make deltas (changes) to the SCCS file. A group ID is equivalent to specifying all *login* names common to that group ID. Several *a* keyletters may be used on a single *admin* command line. As many *logins*, or numerical group IDs, as desired may be on the list simultaneously. If the list of users is empty, then anyone may add deltas.
- eloin** A *login* name, or numerical group ID, to be erased from the list of users allowed to make deltas (changes) to the SCCS file. Specifying a group ID is equivalent to specifying all *login* names common to that group ID. Several *e* keyletters may be used on a single *admin* command line.
- y[comment]** The *comment* text is inserted into the SCCS file as a comment for the initial delta in a manner identical to that of *delta*(1). Omission of the *-y* keyletter results in a default comment line being inserted in the form:
 date and time created YY/MM/DD HH:MM:SS by *login*
 The *-y* keyletter is valid only if the *-i* and/or *-n* keyletters are specified (i.e., a new SCCS file is being created).
- m[mrlist]** The list of Modification Requests (MR) numbers is inserted into the SCCS file as the reason for creating the initial delta in a manner identical to *delta*(1). The *v* flag must be set and the MR numbers are validated if the *v* flag has a value (the name of an MR number validation program). Diagnostics will occur if the *v* flag is not set or MR validation fails.
- h** Causes *admin* to check the structure of the SCCS file (see *scsfile*(5)), and to compare a newly computed check-sum (the sum of all the characters in the SCCS file except those in the first line) with the check-sum

that is stored in the first line of the SCCS file. Appropriate error diagnostics are produced.

This keyletter inhibits writing on the file, so that it nullifies the effect of any other keyletters supplied, and is, therefore, only meaningful when processing existing files.

—z The SCCS file check-sum is recomputed and stored in the first line of the SCCS file (see —h, above).

Note that use of this keyletter on a truly corrupted file may prevent future detection of the corruption.

FILES

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

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

If it should be necessary to patch an SCCS file for any reason, the mode may be changed to 644 by the owner allowing use of *ed(1)*. *Care must be taken!* The edited file should *always* be processed by an *admin -h* to check for corruption followed by an *admin -z* to generate a proper check-sum. Another *admin -h* is recommended to ensure the SCCS file is valid.

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

SEE ALSO

delta(1), *ed(1)*, *get(1)*, *help(1)*, *prs(1)*, *what(1)*, *scsfile(5)*.

Source Code Control System User's Guide by L. E. Bonanni and C. A. Salemi.

DIAGNOSTICS

Use *help(1)* for explanations.

NAME

ar - archive and library maintainer

SYNOPSIS

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

DESCRIPTION

Ar maintains groups of files combined into a single archive file. Its main use is to create and update library files as used by the loader. It can be used, though, for any similar purpose.

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

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

FILES

*/tmp/v** temporaries

SEE ALSO

lorder(1), *ld*(1), *ar*(5)

BUGS

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

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

NAME

as - VAX-11 assembler

SYNOPSIS

as [-d124] [-L] [-W] [-V] [-J] [-R] [-t directory] [-o objfile] [name ...]

DESCRIPTION

As assembles the named files, or the standard input if no file name is specified. The available flags are:

- d Specifies the number of bytes to be assembled for offsets which involve forward or external references, and which have sizes unspecified in the assembly language. The default is -d4.
- L Save defined labels beginning with a 'L', which are normally discarded to save space in the resultant symbol table. The compilers generate such temporary labels.
- V Use virtual memory for some intermediate storage, rather than a temporary file.
- W Do not complain about errors.
- J Use long branches to resolve jumps when byte-displacement branches are insufficient. This must be used when a compiler-generated assembly contains branches of more than 32k bytes.
- R Make initialized data segments read-only, by concatenating them to the text segments. This obviates the need to run editor scripts on assembly code to make initialized data read-only and shared.
- t Specifies a directory to receive the temporary file, other than the default /tmp.

All undefined symbols in the assembly are treated as global.

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

FILES

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

SEE ALSO

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

AUTHORS

John F. Reiser
Robert R. Henry

BUGS

-J should be eliminated; the assembler should automatically choose among byte, word and long branches.

NAME

at — execute commands at a later time

SYNOPSIS

at time [day] [file]

DESCRIPTION

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

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

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

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

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

Standard output or error output is lost unless redirected.

FILES

/usr/lib/atrun executor (run by *cron*(8)).

in */usr/spool/at*:

```
yy.ddd.hhhh.*            activity for year yy, day dd, hour hhhh.
lasttimedone            last hhhh
past                    activities in progress
```

SEE ALSO

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

DIAGNOSTICS

Complains about various syntax errors and times out of range.

BUGS

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

NAME

awk — pattern scanning and processing language

SYNOPSIS

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

DESCRIPTION

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

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

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

A pattern-action statement has the form

```
pattern { action }
```

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

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

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

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

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

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

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

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

A relational expression is one of the following:

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

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

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

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

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

or by using the `-Fc` option.

Other variable names with special meanings include NF, the number of fields in the current record; NR, the ordinal number of the current record; FILENAME, the name of the current input file; OFS, the output field separator (default blank); ORS, the output record separator (default newline); and OFMT, the output format for numbers (default "%.6g").

EXAMPLES

Print lines longer than 72 characters:

```
length > 72
```

Print first two fields in opposite order:

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

Add up first column, print sum and average:

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

Print fields in reverse order:

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

Print all lines between start/stop pairs:

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

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

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

SEE ALSO

lex(1), sed(1)

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

BUGS

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

NAME

`basename` — strip filename affixes

SYNOPSIS

`basename` string [suffix]

DESCRIPTION

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

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

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

SEE ALSO

sh(1)

NAME

`bdiff` - big diff

SYNOPSIS

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

DESCRIPTION

Bdiff is used in a manner analogous to *diff*(1) to find which lines must be changed in two files to bring them into agreement. Its purpose is to allow processing of files which are too large for *diff*. *Bdiff* ignores lines common to the beginning of both files, splits the remainder of each file into *n*-line segments, and invokes *diff* upon corresponding segments. The value of *n* is 3500 by default. If the optional third argument is given, and it is numeric, it is used as the value for *n*. This is useful in those cases in which 3500-line segments are too large for *diff*, causing it to fail. If *file1* (*file2*) is -, the standard input is read. The optional -s (silent) argument specifies that no diagnostics are to be printed by *bdiff* (note, however, that this does not suppress possible exclamations by *diff*). If both optional arguments are specified, they must appear in the order indicated above.

The output of *bdiff* is exactly that of *diff*, with line numbers adjusted to account for the segmenting of the files (that is, to make it look as if the files had been processed whole). Note that because of the segmenting of the files, *bdiff* does not necessarily find a smallest sufficient set of file differences.

FILES

`/tmp/bd?????`

SEE ALSO

`diff`(1).

DIAGNOSTICS

Use *help*(1) for explanations.

NAME

cal — print calendar

SYNOPSIS

cal [month] year

DESCRIPTION

Cal prints a calendar for the specified year. If a month is also specified, a calendar just for that month is printed. *Year* can be between 1 and 9999. The *month* is a number between 1 and 12. The calendar produced is that for England and her colonies.

Try September 1752.

BUGS

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

NAME

calendar — reminder service

SYNOPSIS

calendar [-]

DESCRIPTION

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

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

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

FILES

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

SEE ALSO

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

BUGS

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

NAME

cat — catenate and print

SYNOPSIS

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

DESCRIPTION

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

```
cat file
```

displays the file on the standard output, and

```
cat file1 file2 >file3
```

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

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

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

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

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

SEE ALSO

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

BUGS

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

NAME

cb — C program beautifier

SYNOPSIS

cb

DESCRIPTION

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

NAME

`cc` — C compiler

SYNOPSIS

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

DESCRIPTION

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

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

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

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

- `-c` Suppress the loading phase of the compilation, and force an object file to be produced even if only one program is compiled.
- `-g` Have the compiler produce additional symbol table information for `dbx(1)`. Also pass the `-lg` flag to `ld(1)`.
- `-go` Have the compiler produce additional symbol table information for the obsolete debugger `sdb(1)`. Also pass the `-lg` flag to `ld(1)`.
- `-w` Suppress warning diagnostics.
- `-p` Arrange for the compiler to produce code which counts the number of times each routine is called. If loading takes place, replace the standard startup routine by one which automatically calls `monitor(3)` at the start and arranges to write out a `mon.out` file at normal termination of execution of the object program. An execution profile can then be generated by use of `prof(1)`.
- `-pg` Causes the compiler to produce counting code in the manner of `-p`, but invokes a run-time recording mechanism that keeps more extensive statistics and produces a `gmon.out` file at normal termination. Also, a profiling library is searched, in lieu of the standard C library. An execution profile can then be generated by use of `gprof(1)`.
- `-O` Invoke an object-code improver.
- `-R` Passed on to `as`, making initialized variables shared and read-only.
- `-S` Compile the named C programs, and leave the assembler-language output on corresponding files suffixed `.s`.
- `-E` Run only the macro preprocessor on the named C programs, and send the result to the standard output.
- `-C` prevent the macro preprocessor from eliding comments.
- `-o output`
Name the final output file `output`. If this option is used the file `a.out` will be left undisturbed.
- `-Dname=def`
`-Dname`
Define the `name` to the preprocessor, as if by `#define`. If no definition is given, the name is defined as `"1"`.
- `-Uname`
Remove any initial definition of `name`.

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

-B*string*

Find substitute compiler passes in the files named *string* with the suffixes *cpp*, *ccom* and *c2*. If *string* is empty, use a standard backup version.

-t[*p012*]

Find only the designated compiler passes in the files whose names are constructed by a **-B** option. In the absence of a **-B** option, the *string* is taken to be '/usr/c/'.

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

FILES

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

SEE ALSO

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

DIAGNOSTICS

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

BUGS

The compiler currently ignores advice to put **char**, **unsigned char**, **short** or **unsigned short** variables in registers. It previously produced poor, and in some cases incorrect, code for such declarations.

NAME

`cd` — change working directory

SYNOPSIS

`cd` *directory*

DESCRIPTION

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

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

SEE ALSO

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

NAME

`chgrp` — change group

SYNOPSIS

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

DESCRIPTION

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

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

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

FILES

`/etc/group`

SEE ALSO

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

NAME

chmod — change mode

SYNOPSIS

chmod mode file ...

DESCRIPTION

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

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

A symbolic *mode* has the form:

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

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

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

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

EXAMPLES

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

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

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

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

SEE ALSO

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

NAME

chsh — change default login shell

SYNOPSIS

chsh name [shell]

DESCRIPTION

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

An example use of this command would be

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

SEE ALSO

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

NAME

clear — clear terminal screen

SYNOPSIS

clear

DESCRIPTION

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

FILES

/etc/termcap terminal capability data base

NAME

`cmp` — compare two files

SYNOPSIS

`cmp [-l] [-s] file1 file2`

DESCRIPTION

The two files are compared. (If *file1* is '-', the standard input is used.) Under default options, *cmp* makes no comment if the files are the same; if they differ, it announces the byte and line number at which the difference occurred. If one file is an initial subsequence of the other, that fact is noted.

Options:

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

SEE ALSO

`diff(1)`, `comm(1)`

DIAGNOSTICS

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

NAME

`col` - filter reverse line feeds

SYNOPSIS

`col [-bfX]`

DESCRIPTION

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

Although `col` accepts half line motions in its input, it normally does not emit them on output. Instead, text that would appear between lines is moved to the next lower full line boundary. This treatment can be suppressed by the `-f` (fine) option; in this case the output from `col` may contain forward half line feeds (ESC-9), but will still never contain either kind of reverse line motion.

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

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

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

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

SEE ALSO

`troff(1)`, `tbl(1)`

BUGS

Can't back up more than 128 lines.

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

NAME

comb — combine SCCS deltas

SYNOPSIS

comb [-o] [-s] [-psid] [-clist] files

DESCRIPTION

Comb generates a shell procedure (see *sh(1)*) which, when run, will reconstruct the given SCCS files. The reconstructed files will, hopefully, be smaller than the original files. The arguments may be specified in any order, but all keyletter arguments apply to all named SCCS files. If a directory is named, *comb* behaves as though each file in the directory were specified as a named file, except that non-SCCS files (last component of the path name does not begin with s.) and unreadable files are silently ignored. If a name of - is given, the standard input is read; each line of the standard input is taken to be the name of an SCCS file to be processed; non-SCCS files and unreadable files are silently ignored.

The generated shell procedure is written on the standard output.

The keyletter arguments are as follows. Each is explained as though only one named file is to be processed, but the effects of any keyletter argument apply independently to each named file.

- p*SID* The SCCS *ID*entification string (*SID*) of the oldest delta to be preserved. All older deltas are discarded in the reconstructed file.
- clist* A *list* (see *get(1)* for the syntax of a *list*) of deltas to be preserved. All other deltas are discarded.
- o For each *get -e* generated, this argument causes the reconstructed file to be accessed at the release of the delta to be created, otherwise the reconstructed file would be accessed at the most recent ancestor. Use of the -o keyletter may decrease the size of the reconstructed SCCS file. It may also alter the shape of the delta tree of the original file.
- s This argument causes *comb* to generate a shell procedure which, when run, will produce a report giving, for each file: the file name, size (in blocks) after combining, original size (also in blocks), and percentage change computed by:

$$100 * (\text{original} - \text{combined}) / \text{original}$$

It is recommended that before any SCCS files are actually combined, one should use this option to determine exactly how much space is saved by the combining process.

If no keyletter arguments are specified, *comb* will preserve only leaf deltas and the minimal number of ancestors needed to preserve the tree.

FILES

s.COMB The name of the reconstructed SCCS file.
comb????? Temporary.

SEE ALSO

admin(1), delta(1), get(1), help(1), prs(1), sccsfile(5).
Source Code Control System User's Guide by L. E. Bonanni and C. A. Salemi.

DIAGNOSTICS

Use *help(1)* for explanations.

BUGS

Comb may rearrange the shape of the tree of deltas. It may not save any space; in fact, it is possible for the reconstructed file to actually be larger than the original.

NAME

comm - select or reject lines common to two sorted files

SYNOPSIS

comm [- [123]] file1 file2

DESCRIPTION

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

Flags 1, 2, or 3 suppress printing of the corresponding column. Thus **comm -12** prints only the lines common to the two files; **comm -23** prints only lines in the first file but not in the second; **comm -123** is a no-op.

SEE ALSO

cmp(1), **diff(1)**, **uniq(1)**

NAME

`cp` - copy

SYNOPSIS

`cp` [`-i`] [`-r`] *file1* *file2*

`cp` [`-i`] [`-r`] *file* ... *directory*

DESCRIPTION

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

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

Cp refuses to copy a file onto itself.

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

If the `-r` option is specified and any of the source files are directories, *cp* copies each subtree rooted at that name; in this case the destination must be a directory.

SEE ALSO

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

NAME

cpio - copy file archives in and out (to & from tape)

SYNOPSIS

cpio - o [**acBSv**]

cpio - i [**BcdfmMrtuv8q**] [*patterns*]

cpio - p [**adlmMSruvq**] *directory*

DESCRIPTION

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

Cpio - i (copy in) extracts from the standard input (which is assumed to be the product of a previous **cpio - o**) the names of files selected by zero or more *patterns* given in the name-generating notation of *sh* (1). In *patterns*, meta-characters **?**, *****, and **[...]** match the slash (/) character. A *pattern* can be omitted by using **!** *pattern*. The default for *patterns* is ***** (i.e., select all files).

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

Options:

- a** Reset *access* times of input files after they have been copied.
- B** Input/output is to be *blocked* 5,120 bytes (10 UNIX blocks) to the record. Note that this is different from the default that **tar** uses (20 UNIX blocks). This does not apply to the *pass* option; meaningful only with data directed to or from tape.
- c** Write header information in ASCII *character* form for portability.
- d** *Directories* are to be created as needed. This option is only necessary if the directories do not exist in the archive and do not exist in the destination directory. If the directories were placed in the archive, you do not need this option.
- f** Take as a parameter a *file* containing a list of file names to extract. For example:

```
cpio - iBf flist < /dev/rmt0
```

will extract from */dev/rmt0* the files whose names are listed, one per line, in the file "flist." No pattern metacharacters are recognized here. This option does not work with **cpio - p** or **cpio - o**.

- l** *Link* files rather than copying, where possible. Usable only with the **- p** option. **Cpio** always preserves links.
- m** Retain previous file *modification* time. This option is ineffective on directories that are being copied.
- M** Change mode and ownership of existing directories to match mode and ownership of corresponding directories on tape.
- q** Take the next argument as a filename. **Cpio** quits when the given filename is found.
- r** Interactively *rename* files. If the user types a null line, the file is skipped. Entering control **d** assumes a null line for the remaining files. This option is not available with **cpio - p**.
- s** *Swabs* the file bodies (but not the headers). Try it if file names come out scrambled.
- t** Print a *table of contents* of the input. No files are created. This list of files does not contain any "junk" and is suitable input to **cpio**.

- u** Copy *unconditionally* (normally, an older file will not replace a newer file with the same name).
- S** Causes symbolic links to be followed as if they are real files
- v** *Verbose*: causes a list of file names to be printed. When used with the **t** option, the table of contents looks like the output of an **ls - l** command (see *ls (1)*).

EXAMPLES

The first example below copies the contents of a directory into an archive (tape); the second duplicates a directory hierarchy:

```
find . - print cpio - oB >/dev/rmt0
```

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

The first example can be handled more efficiently by:

```
find . - cpio /dev/rmt0
```

To copy an archive (tape) in, use:

```
cpio - iBdmu < /dev/rmt0
```

NOTES

Cpio can archive special files (devices) if you are logged on as the super-user. **Tar** can not archive special files.

BUGS

There is no way, short of using **- r** interactively, of unrooting a **cpio** archive made with rooted file names (ones that start with '/').

Cpio changes modification dates by default; **tar** leaves them alone by default.

If you use pattern matching with the **- i** option, **cpio** always searches the whole archive (or tape) even if it has already found all the files listed. There is no way to use the rename (**- r**) option from a file instead of interactively.

With the **- o** option, if you have a directory file as input, it adds the directory to the tape but does not recursively add the directory's files (unlike **tar**).

NAME

`cptree` - copy directory tree

SYNOPSIS

`cptree` *fromdir* *todir*

DESCRIPTION

cptree recursively copies a directory hierarchy to another existing directory. *fromdir* is the top of the hierarchy to be copied and *todir* represents the top of the resulting directory tree.

EXAMPLE

```
mkdir /tmp/newdir
cptree /usr/adm /tmp/newdir
```

SEE ALSO

`tar(1)`

NAME

`crypt` — encode/decode

SYNOPSIS

`crypt` [*password*]

DESCRIPTION

Crypt reads from the standard input and writes on the standard output. The *password* is a key that selects a particular transformation. If no *password* is given, *crypt* demands a key from the terminal and turns off printing while the key is being typed in. *Crypt* encrypts and decrypts with the same key:

```
crypt key <clear >cypher
crypt key <cypher | pr
```

will print the clear.

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

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

Crypt implements a one-rotor machine designed along the lines of the German Enigma, but with a 256-element rotor. Methods of attack on such machines are known, but not widely; moreover the amount of work required is likely to be large.

The transformation of a key into the internal settings of the machine is deliberately designed to be expensive, i.e. to take a substantial fraction of a second to compute. However, if keys are restricted to (say) three lower-case letters, then encrypted files can be read by expending only a substantial fraction of five minutes of machine time.

Since the key is an argument to the *crypt* command, it is potentially visible to users executing *ps*(1) or a derivative. To minimize this possibility, *crypt* takes care to destroy any record of the key immediately upon entry. No doubt the choice of keys and key security are the most vulnerable aspect of *crypt*.

FILES

`/dev/tty` for typed key

SEE ALSO

`ed`(1), `makekey`(8)

BUGS

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

NAME

`cs`h — a shell (command interpreter) with C-like syntax

SYNOPSIS

`cs`h [`-cefinstvVxX`] [`arg ...`]

DESCRIPTION

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

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

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

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

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

Lexical structure

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

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

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

Commands

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

Any of the above may be placed in `'('` `')'` to form a simple command (which may be a component of a pipeline, etc.) It is also possible to separate pipelines with `'|'` or `'&&'` indicating, as in the C language, that the second is to be executed only if the first fails or succeeds respectively. (See *Expressions*.)

Jobs

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

```
[1] 1234
```

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

If you are running a job and wish to do something else you may hit the key ^Z (control-Z) which sends a STOP signal to the current job. The shell will then normally indicate that the job has been 'Stopped', and print another prompt. You can then manipulate the state of this job, putting it in the background with the *bg* command, or run some other commands and then eventually bring the job back into the foreground with the foreground command *fg*. A ^Z takes effect immediately and is like an interrupt in that pending output and unread input are discarded when it is typed. There is another special key ^Y which does not generate a STOP signal until a program attempts to *read(2)* it. This can usefully be typed ahead when you have prepared some commands for a job which you wish to stop after it has read them.

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

There are several ways to refer to jobs in the shell. The character '%' introduces a job name. If you wish to refer to job number 1, you can name it as '%1'. Just naming a job brings it to the foreground; thus '%1' is a synonym for 'fg %1', bringing job 1 back into the foreground. Similarly saying '%1 &' resumes job 1 in the background. Jobs can also be named by prefixes of the string typed in to start them, if these prefixes are unambiguous, thus '%ex' would normally restart a suspended *ex(1)* job, if there were only one suspended job whose name began with the string 'ex'. It is also possible to say '%?string' which specifies a job whose text contains *string*, if there is only one such job.

The shell maintains a notion of the current and previous jobs. In output pertaining to jobs, the current job is marked with a '+' and the previous job with a '-'. The abbreviation '%+' refers to the current job and '%-' refers to the previous job. For close analogy with the syntax of the *history* mechanism (described below), '%%' is also a synonym for the current job.

Status reporting

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

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

Substitutions

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

History substitutions

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

Commands input from the terminal which consist of one or more words are saved on the history list. The history substitutions reintroduce sequences of words from these saved commands into the input stream. The size of which is controlled by the *history* variable; the previous command is always retained, regardless of its value. Commands are numbered sequentially from 1.

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

```

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

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

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

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

```

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

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

h	Remove a trailing pathname component, leaving the head.
r	Remove a trailing '.xxx' component, leaving the root name.
e	Remove all but the extension '.xxx' part.
s//r/	Substitute /for r
t	Remove all leading pathname components, leaving the tail.
&	Repeat the previous substitution.
g	Apply the change globally, prefixing the above, e.g. 'g&'.
p	Print the new command but do not execute it.
q	Quote the substituted words, preventing further substitutions.
x	Like q, but break into words at blanks, tabs and newlines.

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

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

A history reference may be given without an event specification, e.g. '\$'. In this case the reference is to the previous command unless a previous history reference occurred on the same line in which case this form repeats the previous reference. Thus '!foo?| '\$' gives the first and last arguments from the command matching '?foo?'.

A special abbreviation of a history reference occurs when the first non-blank character of an input line is a '|'. This is equivalent to '!s|' providing a convenient shorthand for substitutions on the text of the previous line. Thus '|b|lib' fixes the spelling of 'lib' in the previous command. Finally, a history substitution may be surrounded with '{' and '}' if necessary to insulate it from the characters which follow. Thus, after 'ls -ld ~paul' we might do '!{l}a' to do 'ls -ld ~paula', while '!a' would look for a command starting 'la'.

Quotations with ' and "

The quotation of strings by "" and "" can be used to prevent all or some of the remaining substitutions. Strings enclosed in "" are prevented any further interpretation. Strings enclosed in "" may be expanded as described below.

In both cases the resulting text becomes (all or part of) a single word; only in one special case (see *Command Substitution* below) does a "" quoted string yield parts of more than one word; "" quoted strings never do.

Alias substitution

The shell maintains a list of aliases which can be established, displayed and modified by the *alias* and *unalias* commands. After a command line is scanned, it is parsed into distinct commands and the first word of each command, left-to-right, is checked to see if it has an alias. If it does, then the text which is the alias for that command is reread with the history mechanism available as though that command were the previous input line. The resulting words replace the command and argument list. If no reference is made to the history list, then the argument list is left unchanged.

Thus if the alias for 'ls' is 'ls -l' the command 'ls /usr' would map to 'ls -l /usr', the argument list here being undisturbed. Similarly if the alias for 'lookup' was 'grep !f /etc/passwd' then 'lookup bill' would map to 'grep bill /etc/passwd'.

If an alias is found, the word transformation of the input text is performed and the aliasing process begins again on the reformed input line. Looping is prevented if the first word of the new text is the same as the old by flagging it to prevent further aliasing. Other loops are detected and cause an error.

Note that the mechanism allows aliases to introduce parser metasyntax. Thus we can 'alias print 'pr \!* | lpr'' to make a command which *pr*'s its arguments to the line printer.

Variable substitution

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

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

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

After the input line is aliased and parsed, and before each command is executed, variable substitution is performed keyed by '\$' characters. This expansion can be prevented by preceding the '\$' with a '\' except within ""'s where it **always** occurs, and within '''s where it **never** occurs. Strings quoted by '' are interpreted later (see *Command substitution* below) so '\$' substitution does not occur there until later, if at all. A '\$' is passed unchanged if followed by a blank, tab, or end-of-line.

Input/output redirections are recognized before variable expansion, and are variable expanded separately. Otherwise, the command name and entire argument list are expanded together. It is thus possible for the first (command) word to this point to generate more than one word, the first of which becomes the command name, and the rest of which become arguments.

Unless enclosed in "" or given the ':q' modifier the results of variable substitution may eventually be command and filename substituted. Within "" a variable whose value consists of multiple words expands to a (portion of) a single word, with the words of the variables value separated by blanks. When the ':q' modifier is applied to a substitution the variable will expand to multiple words with each word separated by a blank and quoted to prevent later command or filename substitution.

The following metasequences are provided for introducing variable values into the shell input. Except as noted, it is an error to reference a variable which is not set.

\$name
\${name}

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

If *name* is not a shell variable, but is set in the environment, then that value is returned

(but `:` modifiers and the other forms given below are not available in this case).

`$name[selector]`
`${name[selector]}`

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

`$#name`
`${#name}`

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

`$0`

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

`$number`
`${number}`

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

`$*`

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

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

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

`$?name`
`${?name}`

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

`$?0`

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

`$$`

Substitute the (decimal) process number of the (parent) shell.

`$<`

Substitutes a line from the standard input, with no further interpretation thereafter. It can be used to read from the keyboard in a shell script.

Command and filename substitution

The remaining substitutions, command and filename substitution, are applied selectively to the arguments of builtin commands. This means that portions of expressions which are not evaluated are not subjected to these expansions. For commands which are not internal to the shell, the command name is substituted separately from the argument list. This occurs very late, after input-output redirection is performed, and in a child of the main shell.

Command substitution

Command substitution is indicated by a command enclosed in `'`'`. The output from such a command is normally broken into separate words at blanks, tabs and newlines, with null words being discarded, this text then replacing the original string. Within `'`'`s, only newlines force new words; blanks and tabs are preserved.

In any case, the single final newline does not force a new word. Note that it is thus possible for a command substitution to yield only part of a word, even if the command outputs a complete line.

Filename substitution

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

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

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

The metanotation 'a{b,c,d}e' is a shorthand for 'abe ace ade'. Left to right order is preserved, with results of matches being sorted separately at a low level to preserve this order. This construct may be nested. Thus '~source/s1/{oldls,ls}.c' expands to '/usr/source/s1/oldls.c /usr/source/s1/ls.c' whether or not these files exist without any chance of error if the home directory for 'source' is '/usr/source'. Similarly './{memo,*box}' might expand to './memo ../box ../mbox'. (Note that 'memo' was not sorted with the results of matching '*box'.) As a special case '{', '}' and '{} are passed undisturbed.

Input/output

The standard input and standard output of a command may be redirected with the following syntax:

< name

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

<< word

Read the shell input up to a line which is identical to *word*. *Word* is not subjected to variable, filename or command substitution, and each input line is compared to *word* before any substitutions are done on this input line. Unless a quoting '\', '"', '' or '' appears in *word* variable and command substitution is performed on the intervening lines, allowing '\ ' to quote '\$', '\ ' and '''. Commands which are substituted have all blanks, tabs, and newlines preserved, except for the final newline which is dropped. The resultant text is placed in an anonymous temporary file which is given to the command as standard input.

> name

>! name

>& name

>&! name

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

If the variable *noclobber* is set, then the file must not exist or be a character special file (e.g. a terminal or */dev/null*) or an error results. This helps prevent accidental destruction of files. In this case the *!* forms can be used and suppress this check.

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

>> name

>>& name

>>! name

>>&! name

Uses file *name* as standard output like *>* but places output at the end of the file. If the variable *noclobber* is set, then it is an error for the file not to exist unless one of the *!* forms is given. Otherwise similar to *>*.

A command receives the environment in which the shell was invoked as modified by the input-output parameters and the presence of the command in a pipeline. Thus, unlike some previous shells, commands run from a file of shell commands have no access to the text of the commands by default; rather they receive the original standard input of the shell. The *<<* mechanism should be used to present inline data. This permits shell command scripts to function as components of pipelines and allows the shell to block read its input. Note that the default standard input for a command run detached is **not** modified to be the empty file */dev/null*; rather the standard input remains as the original standard input of the shell. If this is a terminal and if the process attempts to read from the terminal, then the process will block and the user will be notified (see **Jobs** above.)

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

Expressions

A number of the builtin commands (to be described subsequently) take expressions, in which the operators are similar to those of C, with the same precedence. These expressions appear in the *@*, *exit*, *if*, and *while* commands. The following operators are available:

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

Here the precedence increases to the right, *'=='* *'!='* *'=~'* and *'!~'*, *'<='* *'>='* *'<'* and *'>'*, *'<<'* and *'>>'*, *'+'* and *'-'*, *'*'* *'/'* and *'%'* being, in groups, at the same level. The *'=='* *'!='* *'=~'* and *'!~'* operators compare their arguments as strings; all others operate on numbers. The operators *'=~'* and *'!~'* are like *'!='* and *'=='* except that the right hand side is a *pattern* (containing, e.g. *'*'*s, *'?'*s and instances of *'[...]'*) against which the left hand operand is matched. This reduces the need for use of the *switch* statement in shell scripts when all that is really needed is pattern matching.

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

Also available in expressions as primitive operands are command executions enclosed in *'{'* and *'}'* and file enquiries of the form *'- / name'* where */* is one of:

r	read access
w	write access
x	execute access
e	existence
o	ownership
z	zero size
f	plain file
d	directory

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

Control flow

The shell contains a number of commands which can be used to regulate the flow of control in command files (shell scripts) and (in limited but useful ways) from terminal input. These commands all operate by forcing the shell to reread or skip in its input and, due to the implementation, restrict the placement of some of the commands.

The *foreach*, *switch*, and *while* statements, as well as the *if-then-else* form of the *if* statement require that the major keywords appear in a single simple command on an input line as shown below.

If the shell's input is not seekable, the shell buffers up input whenever a loop is being read and performs seeks in this internal buffer to accomplish the rereading implied by the loop. (To the extent that this allows, backward goto's will succeed on non-seekable inputs.)

Builtin commands

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

alias

alias name

alias name wordlist

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

alloc

Shows the amount of dynamic core in use, broken down into used and free core, and address of the last location in the heap. With an argument shows each used and free block on the internal dynamic memory chain indicating its address, size, and whether it is used or free. This is a debugging command and may not work in production versions of the shell; it requires a modified version of the system memory allocator.

bg

bg %job ...

Puts the current or specified jobs into the background, continuing them if they were stopped.

break

Causes execution to resume after the *end* of the nearest enclosing *foreach* or *while*. The remaining commands on the current line are executed. Multi-level breaks are thus possible by writing them all on one line.

breaksw

Causes a break from a *switch*, resuming after the *endsw*.

case label:

A label in a *switch* statement as discussed below.

cd**cd name****chdir****chdir name**

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

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

continue

Continue execution of the nearest enclosing *while* or *foreach*. The rest of the commands on the current line are executed.

default:

Labels the default case in a *switch* statement. The default should come after all *case* labels.

dirs

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

echo wordlist**echo -n wordlist**

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

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

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

eval arg ...

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

exec command

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

exit**exit(expr)**

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

fg**fg %job ...**

Brings the current or specified jobs into the foreground, continuing them if they were stopped.

foreach name (wordlist)

...
end

The variable *name* is successively set to each member of *wordlist* and the sequence of commands between this command and the matching *end* are executed. (Both *foreach* and *end* must appear alone on separate lines.)

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

glob wordlist

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

goto word

The specified *word* is filename and command expanded to yield a string of the form 'label'. The shell rewinds its input as much as possible and searches for a line of the form 'label:' possibly preceded by blanks or tabs. Execution continues after the specified line.

hashstat

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

history**history** *n***history** -r *n***history** -h *n*

Displays the history event list; if *n* is given only the *n* most recent events are printed. The -r option reverses the order of printout to be most recent first rather than oldest first. The -h option causes the history list to be printed without leading numbers. This is used to produce files suitable for sourcing using the -h option to *source*.

if (expr) command

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

if (expr) then

...
else if (expr2) then

...
else

...
endif

If the specified *expr* is true then the commands to the first *else* are executed; else if *expr2* is true then the commands to the second *else* are executed, etc. Any number of *else-if* pairs are possible; only one *endif* is needed. The *else* part is likewise optional. (The words *else* and *endif* must appear at the beginning of input lines; the *if* must appear alone on its input line or after an *else*.)

jobs**jobs -l**

Lists the active jobs; given the **-l** options lists process id's in addition to the normal information.

kill %job**kill -sig %job ...****kill pid****kill -sig pid ...****kill -l**

Sends either the TERM (terminate) signal or the specified signal to the specified jobs or processes. Signals are either given by number or by names (as given in */usr/include/signal.h*, stripped of the prefix "SIG"). The signal names are listed by "kill -l". There is no default, saying just 'kill' does not send a signal to the current job. If the signal being sent is TERM (terminate) or HUP (hangup), then the job or process will be sent a CONT (continue) signal as well.

limit**limit resource****limit resource maximum-use**

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

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

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

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

login

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

logout

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

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

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

nohup

nohup command

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

notify

notify %job ...

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

onintr

onintr -

onintr label

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

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

popd

popd +n

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

pushd

pushd name

pushd +n

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

rehash

Causes the internal hash table of the contents of the directories in the *path* variable to be recomputed. This is needed if new commands are added to directories in the *path* while you are logged in. This should only be necessary if you add commands to one of your own directories, or if a systems programmer changes the contents of one of the system directories.

repeat count command

The specified *command* which is subject to the same restrictions as the *command* in the one line *if* statement above, is executed *count* times. I/O redirections occur exactly once, even if *count* is 0.

set

set name

set name=word

set name[index]=word

set name=(wordlist)

The first form of the command shows the value of all shell variables. Variables which

have other than a single word as value print as a parenthesized word list. The second form sets *name* to the null string. The third form sets *name* to the single *word*. The fourth form sets the *index*'th component of *name* to *word*; this component must already exist. The final form sets *name* to the list of words in *wordlist*. In all cases the value is command and filename expanded.

These arguments may be repeated to set multiple values in a single set command. Note however, that variable expansion happens for all arguments before any setting occurs.

setenv name value

Sets the value of environment variable *name* to be *value*, a single string. The most commonly used environment variable USER, TERM, and PATH are automatically imported to and exported from the *cs*h variables *user*, *term*, and *path*; there is no need to use *setenv* for these.

shift

shift variable

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

source name

source -h name

The shell reads commands from *name*. *Source* commands may be nested; if they are nested too deeply the shell may run out of file descriptors. An error in a *source* at any level terminates all nested *source* commands. Normally input during *source* commands is not placed on the history list; the -h option causes the commands to be placed in the history list without being executed.

stop

stop %job ...

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

suspend

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

switch (string)

case str1:

...

breaksw

...

default:

...

breaksw

endsw

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

time

time command

With no argument, a summary of time used by this shell and its children is printed. If

arguments are given the specified simple command is timed and a time summary as described under the *time* variable is printed. If necessary, an extra shell is created to print the time statistic when the command completes.

umask**umask value**

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

unalias pattern

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

unhash

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

unlimit resource**unlimit**

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

unset pattern

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

unsetenv pattern

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

wait

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

while (expr)**end**

While the specified expression evaluates non-zero, the commands between the *while* and the matching *end* are evaluated. *Break* and *continue* may be used to terminate or continue the loop prematurely. (The *while* and *end* must appear alone on their input lines.) Prompting occurs here the first time through the loop as for the *foreach* statement if the input is a terminal.

%job

Brings the specified job into the foreground.

%job &

Continues the specified job in the background.

@**@ name = expr****@ name[index] = expr**

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

must already exist.

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

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

Pre-defined and environment variables

The following variables have special meaning to the shell. Of these, *argv*, *cwd*, *home*, *path*, *prompt*, *shell* and *status* are always set by the shell. Except for *cwd* and *status* this setting occurs only at initialization; these variables will not then be modified unless this is done explicitly by the user.

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

argv	Set to the arguments to the shell, it is from this variable that positional parameters are substituted, i.e. <code>\$1</code> is replaced by <code>\$argv[1]</code> , etc.
cdpath	Gives a list of alternate directories searched to find subdirectories in <i>chdir</i> commands.
cwd	The full pathname of the current directory.
echo	Set when the <code>-x</code> command line option is given. Causes each command and its arguments to be echoed just before it is executed. For non-builtin commands all expansions occur before echoing. Builtin commands are echoed before command and filename substitution, since these substitutions are then done selectively.
histchars	Can be given a string value to change the characters used in history substitution. The first character of its value is used as the history substitution character, replacing the default character <code>!</code> . The second character of its value replaces the character <code>↑</code> in quick substitutions.
history	Can be given a numeric value to control the size of the history list. Any command which has been referenced in this many events will not be discarded. Too large values of <i>history</i> may run the shell out of memory. The last executed command is always saved on the history list.
home	The home directory of the invoker, initialized from the environment. The filename expansion of <code>~</code> refers to this variable.
ignoreeof	If set the shell ignores end-of-file from input devices which are terminals. This prevents shells from accidentally being killed by control-D's.
mail	The files where the shell checks for mail. This is done after each command completion which will result in a prompt, if a specified interval has elapsed. The shell says 'You have new mail.' if the file exists with an access time not greater than its modify time. If the first word of the value of <i>mail</i> is numeric it specifies a different mail checking interval, in seconds, than the default, which is 10 minutes. If multiple mail files are specified, then the shell says 'New mail in <i>name</i> ' when there is mail in the file <i>name</i> .

noclobber	As described in the section on <i>Input/output</i> , restrictions are placed on output redirection to insure that files are not accidentally destroyed, and that '>>' redirections refer to existing files.
noglob	If set, filename expansion is inhibited. This is most useful in shell scripts which are not dealing with filenames, or after a list of filenames has been obtained and further expansions are not desirable.
nonomatch	If set, it is not an error for a filename expansion to not match any existing files; rather the primitive pattern is returned. It is still an error for the primitive pattern to be malformed, i.e. 'echo [' still gives an error.
notify	If set, the shell notifies asynchronously of job completions. The default is to rather present job completions just before printing a prompt.
path	Each word of the path variable specifies a directory in which commands are to be sought for execution. A null word specifies the current directory. If there is no <i>path</i> variable then only full path names will execute. The usual search path is '.', '/bin' and '/usr/bin', but this may vary from system to system. For the super-user the default search path is '/etc', '/bin' and '/usr/bin'. A shell which is given neither the <i>-c</i> nor the <i>-t</i> option will normally hash the contents of the directories in the <i>path</i> variable after reading <i>.cshrc</i> , and each time the <i>path</i> variable is reset. If new commands are added to these directories while the shell is active, it may be necessary to give the <i>rehash</i> or the commands may not be found.
prompt	The string which is printed before each command is read from an interactive terminal input. If a '!' appears in the string it will be replaced by the current event number unless a preceding '\ ' is given. Default is '% ', or '# ' for the super-user.
savehist	is given a numeric value to control the number of entries of the history list that are saved in <i>~/history</i> when the user logs out. Any command which has been referenced in this many events will be saved. During start up the shell sources <i>~/history</i> into the history list enabling history to be saved across logins. Too large values of <i>savehist</i> will slow down the shell during start up.
shell	The file in which the shell resides. This is used in forking shells to interpret files which have execute bits set, but which are not executable by the system. (See the description of <i>Non-builtin Command Execution</i> below.) Initialized to the (system-dependent) home of the shell.
status	The status returned by the last command. If it terminated abnormally, then 0200 is added to the status. Builtin commands which fail return exit status '1', all other builtin commands set status '0'.
time	Controls automatic timing of commands. If set, then any command which takes more than this many cpu seconds will cause a line giving user, system, and real times and a utilization percentage which is the ratio of user plus system times to real time to be printed when it terminates.
verbose	Set by the <i>-v</i> command line option, causes the words of each command to be printed after history substitution.

Non-builtin command execution

When a command to be executed is found to not be a builtin command the shell attempts to execute the command via *execve(2)*. Each word in the variable *path* names a directory from which the shell will attempt to execute the command. If it is given neither a *-c* nor a *-t* option, the shell will hash the names in these directories into an internal table so that it will

only try an *exec* in a directory if there is a possibility that the command resides there. This greatly speeds command location when a large number of directories are present in the search path. If this mechanism has been turned off (via *unhash*), or if the shell was given a *-c* or *-t* argument, and in any case for each directory component of *path* which does not begin with a '/', the shell concatenates with the given command name to form a path name of a file which it then attempts to execute.

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

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

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

Argument list processing

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

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

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

Signal handling

The shell normally ignores *quit* signals. Jobs running detached (either by '&' or the *bg* or %... & commands) are immune to signals generated from the keyboard, including hangups. Other signals have the values which the shell inherited from its parent. The shells handling of interrupts and terminate signals in shell scripts can be controlled by *onintr*. Login shells catch the *terminate* signal; otherwise this signal is passed on to children from the state in the shell's parent. In no case are interrupts allowed when a login shell is reading the file *logout*.

AUTHOR

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

FILES

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

LIMITATIONS

Words can be no longer than 1024 characters. The system limits argument lists to 10240 characters. The number of arguments to a command which involves filename expansion is limited to 1/6'th the number of characters allowed in an argument list. Command substitutions may substitute no more characters than are allowed in an argument list. To detect looping, the shell restricts the number of *alias* substitutions on a single line to 20.

SEE ALSO

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

BUGS

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

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

Control over *tty* output after processes are started is primitive; perhaps this will inspire someone to work on a good virtual terminal interface. In a virtual terminal interface much more interesting things could be done with output control.

Alias substitution is most often used to clumsily simulate shell procedures; shell procedures should be provided rather than aliases.

Commands within loops, prompted for by '?', are not placed in the *history* list. Control structure should be parsed rather than being recognized as built-in commands. This would allow control commands to be placed anywhere, to be combined with '|', and to be used with '&' and ';' metasyntax.

It should be possible to use the ':' modifiers on the output of command substitutions. All and more than one ':' modifier should be allowed on '\$' substitutions.

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

NAME

`ctags` — create a tags file

SYNOPSIS

`ctags [-BFatuwvx] name ...`

DESCRIPTION

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

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

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

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

Files whose name ends in `.c` or `.h` are assumed to be C source files and are searched for C routine and macro definitions. Others are first examined to see if they contain any Pascal or Fortran routine definitions; if not, they are processed again looking for C definitions.

Other options are:

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

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

FILES

`tags` output tags file

SEE ALSO

ex(1), *vi*(1)

AUTHOR

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

BUGS

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

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

Does not know about #ifdefs.

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

NAME

date — print and set the date

SYNOPSIS

date [-u] [yymmddhhmm [.ss]]

DESCRIPTION

If no arguments are given, the current date and time are printed. If a date is specified, the current date is set. The *-u* flag is used to display the date in GMT (universal) time. This flag may also be used to set GMT time. *yy* is the last two digits of the year; the first *mm* is the month number; *dd* is the day number in the month; *hh* is the hour number (24 hour system); the second *mm* is the minute number; *.ss* is optional and is the seconds. For example:

```
date 10080045
```

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

FILES

/usr/adm/wtmp to record time-setting

SEE ALSO

utmp(5)

DIAGNOSTICS

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

BUGS

The system attempts to keep the date in a format closely compatible with VMS. VMS, however, uses local time (rather than GMT) and does not understand daylight savings time. Thus if you use both UNIX and VMS, VMS will be running on GMT.

NAME

dd — convert and copy a file

SYNOPSIS

dd [option=value] ...

DESCRIPTION

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

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

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

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

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

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

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

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

SEE ALSO

cp(1), tr(1)

DIAGNOSTICS

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

BUGS

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

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

NAME

delta — make a delta (change) to an SCCS file

SYNOPSIS

delta [**-r**SID] [**-s**] [**-n**] [**-g**list] [**-m**[mrlist]] [**-y**[comment]] [**-p**]
files

DESCRIPTION

Delta is used to permanently introduce into the named SCCS file changes that were made to the file retrieved by *get*(1) (called the *g-file*, or generated file).

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

Delta may issue prompts on the standard output depending upon certain keyletters specified and flags (see *admin*(1)) that may be present in the SCCS file (see **-m** and **-y** keyletters below).

Keyletter arguments apply independently to each named file.

- r**SID Uniquely identifies which delta is to be made to the SCCS file. The use of this keyletter is necessary only if two or more outstanding *gets* for editing (**get -e**) on the same SCCS file were done by the same person (login name). The SID value specified with the **-r** keyletter can be either the SID specified on the *get* command line or the SID to be made as reported by the *get* command (see *get*(1)). A diagnostic results if the specified SID is ambiguous, or, if necessary and omitted on the command line.
- s** Suppresses the issue, on the standard output, of the created delta's SID, as well as the number of lines inserted, deleted and unchanged in the SCCS file.
- n** Specifies retention of the edited *g-file* (normally removed at completion of delta processing).
- g**list Specifies a *list* (see *get*(1) for the definition of *list*) of deltas which are to be *ignored* when the file is accessed at the change level (SID) created by this delta.
- m**[mrlist] If the SCCS file has the *v* flag set (see *admin*(1)) then a Modification Request (MR) number *must* be supplied as the reason for creating the new delta.

If **-m** is not used and the standard input is a terminal, the prompt MRs? is issued on the standard output before the standard input is read; if the standard input is not a terminal, no prompt is issued. The MRs? prompt always precedes the **comments?** prompt (see **-y** keyletter).

MRs in a list are separated by blanks and/or tab characters. An unescaped new-line character terminates the MR list.

Note that if the *v* flag has a value (see *admin(1)*), it is taken to be the name of a program (or shell procedure) which will validate the correctness of the MR numbers. If a non-zero exit status is returned from MR number validation program, *delta* terminates (it is assumed that the MR numbers were not all valid).

- y**[*comment*] Arbitrary text used to describe the reason for making the delta. A null string is considered a valid *comment*.
If **-y** is not specified and the standard input is a terminal, the prompt **comments?** is issued on the standard output before the standard input is read; if the standard input is not a terminal, no prompt is issued. An unescaped new-line character terminates the comment text.
- p** Causes *delta* to print (on the standard output) the SCCS file differences before and after the delta is applied in a *diff(1)* format.

FILES

All files of the form *?-file* are explained in the *Source Code Control System User's Guide*. The naming convention for these files is also described there.

- g-file** Existed before the execution of *delta*; removed after completion of *delta*.
- p-file** Existed before the execution of *delta*; may exist after completion of *delta*.
- q-file** Created during the execution of *delta*; removed after completion of *delta*.
- x-file** Created during the execution of *delta*; renamed to SCCS file after completion of *delta*.
- z-file** Created during the execution of *delta*; removed during the execution of *delta*.
- d-file** Created during the execution of *delta*; removed after completion of *delta*.
- /usr/bin/bdiff** Program to compute differences between the "gotten" file and the *g-file*.

WARNINGS

Lines beginning with an SOH ASCII character (binary 001) cannot be placed in the SCCS file unless the SOH is escaped. This character has special meaning to SCCS (see *sccsfile(5)*) and will cause an error.

A *get* of many SCCS files, followed by a *delta* of those files, should be avoided when the *get* generates a large amount of data. Instead, multiple *get/delta* sequences should be used.

If the standard input (**-**) is specified on the *delta* command line, the **-m** (if necessary) and **-y** keyletters *must* also be present. Omission of these keyletters causes an error to occur.

SEE ALSO

admin(1), *bdiff(1)*, *get(1)*, *help(1)*, *prs(1)*, *sccsfile(5)*.
Source Code Control System User's Guide by L. E. Bonanni and C. A. Salemi.

DIAGNOSTICS

Use *help(1)* for explanations.

NAME

`deroff` — remove *nroff*, *troff*, *tbl* and *eqn* constructs

SYNOPSIS

`deroff [-w] file ...`

DESCRIPTION

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

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

SEE ALSO

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

BUGS

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

NAME

df — disk free

SYNOPSIS

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

DESCRIPTION

Df prints out the amount of free disk space available on the specified *filesystem*, e.g. “/dev/rp0a”, or on the filesystem in which the specified *file*, e.g. “\$HOME”, is contained. If no file system is specified, the free space on all of the normally mounted file systems is printed. The reported numbers are in kilobytes.

Other options are:

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

FILES

/etc/fstab list of normally mounted filesystems

SEE ALSO

fstab(5), icheck(8), quot(8)

NAME

diff — differential file and directory comparator

SYNOPSIS

diff [**-l**] [**-r**] [**-s**] [**-cefh**] [**-b**] dir1 dir2

diff [**-cefh**] [**-b**] file1 file2

diff [**-Dstring**] [**-b**] file1 file2

DESCRIPTION

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

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

-Sname

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

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

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

```
n1 a n3,n4
n1,n2 d n3
n1,n2 c n3,n4
```

These lines resemble *ed* commands to convert *file1* into *file2*. The numbers after the letters pertain to *file2*. In fact, by exchanging 'a' for 'd' and reading backward one may ascertain equally how to convert *file2* into *file1*. As in *ed*, identical pairs where $n1 = n2$ or $n3 = n4$ are abbreviated as a single number.

Following each of these lines come all the lines that are affected in the first file flagged by '<', then all the lines that are affected in the second file flagged by '>'.

Except for **-b**, which may be given with any of the others, the following options are mutually exclusive:

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

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

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

- f** produces a script similar to that of **-e**, not useful with *ed*, and in the opposite order.

- c** produces a diff with lines of context. The default is to present 3 lines of context and may be changed, e.g to 10, by **-c10**. With **-c** the output format is modified slightly: the output beginning with identification of the files involved and their creation dates and then each change is separated by a line with a dozen *'s. The lines removed from *file1* are marked with '-'; those added to *file2* are marked '+'. Lines which are changed from one file to the other are marked in both files with '!'.
-Dstring causes *diff* to create a merged version of *file1* and *file2* on the standard output, with C preprocessor controls included so that a compilation of the result without defining *string* is equivalent to compiling *file1*, while defining *string* will yield *file2*.
- b** causes trailing blanks (spaces and tabs) to be ignored, and other strings of blanks to compare equal.

FILES

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

SEE ALSO

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

DIAGNOSTICS

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

BUGS

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

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

NAME

diff3 — 3-way differential file comparison

SYNOPSIS

diff3 [-ex3] file1 file2 file3

DESCRIPTION

Diff3 compares three versions of a file, and publishes disagreeing ranges of text flagged with these codes:

```

=====      all three files differ
=====1     file1 is different
=====2     file2 is different
=====3     file3 is different

```

The type of change suffered in converting a given range of a given file to some other is indicated in one of these ways:

```

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

```

The original contents of the range follows immediately after a c indication. When the contents of two files are identical, the contents of the lower-numbered file is suppressed.

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

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

FILES

```

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

```

SEE ALSO

diff(1)

BUGS

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

NAME

`du` — summarize disk usage

SYNOPSIS

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

DESCRIPTION

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

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

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

SEE ALSO

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

BUGS

Non-directories given as arguments (not under `-a` option) are not listed. If there are too many distinct linked files, *du* counts the excess files multiply.

NAME

echo — echo arguments

SYNOPSIS

echo [-n] [arg] ...

DESCRIPTION

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

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

NAME

ed — text editor

SYNOPSIS

ed [-] [-x] [name]

DESCRIPTION

Ed is the standard text editor.

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

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

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

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

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

1. Any character except a special character matches itself. Special characters are the regular expression delimiter plus \[, and sometimes ^*\$.
2. A . matches any character.
3. A \ followed by any character except a digit or () matches that character.
4. A nonempty string *s* bracketed [*s*] (or [^*s*]) matches any character in (or not in) *s*. In *s*, \ has no special meaning, and] may only appear as the first letter. A substring *a-b*, with *a* and *b* in ascending ASCII order, stands for the inclusive range of ASCII characters.
5. A regular expression of form 1-4 followed by * matches a sequence of 0 or more matches of the regular expression.
6. A regular expression, *x*, of form 1-8, bracketed \(*x*\) matches what *x* matches.
7. A \ followed by a digit *n* matches a copy of the string that the bracketed regular expression beginning with the *n*th \(\(matched.
8. A regular expression of form 1-8, *x*, followed by a regular expression of form 1-7, *y* matches a match for *x* followed by a match for *y*, with the *x* match being as long as possible while still permitting a *y* match.
9. A regular expression of form 1-8 preceded by ^ (or followed by \$), is constrained to matches that begin at the left (or end at the right) end of a line.
10. A regular expression of form 1-9 picks out the longest among the leftmost matches in a line.
11. An empty regular expression stands for a copy of the last regular expression encountered.

Regular expressions are used in addresses to specify lines and in one command (see *s* below) to specify a portion of a line which is to be replaced. If it is desired to use one of the regular expression metacharacters as an ordinary character, that character may be preceded by '\'. This also applies to the character bounding the regular expression (often '/') and to '\' itself.

To understand addressing in *ed* it is necessary to know that at any time there is a *current line*. Generally speaking, the current line is the last line affected by a command; however, the exact effect on the current line is discussed under the description of the command. Addresses are constructed as follows.

1. The character '.' addresses the current line.
2. The character '\$' addresses the last line of the buffer.
3. A decimal number *n* addresses the *n*-th line of the buffer.
4. 'x' addresses the line marked with the name *x*, which must be a lower-case letter. Lines are marked with the *k* command described below.
5. A regular expression enclosed in slashes '/' addresses the line found by searching forward from the current line and stopping at the first line containing a string that matches the regular expression. If necessary the search wraps around to the beginning of the buffer.
6. A regular expression enclosed in queries '?' addresses the line found by searching backward from the current line and stopping at the first line containing a string that matches the regular expression. If necessary the search wraps around to the end of the buffer.
7. An address followed by a plus sign '+' or a minus sign '-' followed by a decimal number specifies that address plus (resp. minus) the indicated number of lines. The plus sign may be omitted.
8. If an address begins with '+' or '-' the addition or subtraction is taken with respect to the current line; e.g. '-5' is understood to mean '.-5'.
9. If an address ends with '+' or '-', then 1 is added (resp. subtracted). As a consequence of this rule and rule 8, the address '-' refers to the line before the current line. Moreover, trailing '+' and '-' characters have cumulative effect, so '--' refers to the current line less 2.
10. To maintain compatibility with earlier versions of the editor, the character '^' in addresses is equivalent to '-'.

Commands may require zero, one, or two addresses. Commands which require no addresses regard the presence of an address as an error. Commands which accept one or two addresses assume default addresses when insufficient are given. If more addresses are given than such a command requires, the last one or two (depending on what is accepted) are used.

Addresses are separated from each other typically by a comma ','. They may also be separated by a semicolon ';'. In this case the current line '.' is set to the previous address before the next address is interpreted. This feature can be used to determine the starting line for forward and backward searches ('/', '?'). The second address of any two-address sequence must correspond to a line following the line corresponding to the first address. The special form '%' is an abbreviation for the address pair '1,\$'.

In the following list of *ed* commands, the default addresses are shown in parentheses. The parentheses are not part of the address, but are used to show that the given addresses are the default.

As mentioned, it is generally illegal for more than one command to appear on a line. However, most commands may be suffixed by 'p' or by 'l', in which case the current line is either printed or listed respectively in the way discussed below. Commands may also be suffixed by 'n',

meaning the output of the command is to be line numbered. These suffixes may be combined in any order.

(.)a
<text>

The append command reads the given text and appends it after the addressed line. '.' is left on the last line input, if there were any, otherwise at the addressed line. Address '0' is legal for this command; text is placed at the beginning of the buffer.

(.,.)c
<text>

The change command deletes the addressed lines, then accepts input text which replaces these lines. '.' is left at the last line input; if there were none, it is left at the line preceding the deleted lines.

(.,.)d

The delete command deletes the addressed lines from the buffer. The line originally after the last line deleted becomes the current line; if the lines deleted were originally at the end, the new last line becomes the current line.

e filename

The edit command causes the entire contents of the buffer to be deleted, and then the named file to be read in. '.' is set to the last line of the buffer. The number of characters read is typed. 'filename' is remembered for possible use as a default file name in a subsequent *r* or *w* command. If 'filename' is missing, the remembered name is used.

E filename

This command is the same as *e*, except that no diagnostic results when no *w* has been given since the last buffer alteration.

f filename

The filename command prints the currently remembered file name. If 'filename' is given, the currently remembered file name is changed to 'filename'.

(1,\$)g/regular expression/command list

In the global command, the first step is to mark every line which matches the given regular expression. Then for every such line, the given command list is executed with '.' initially set to that line. A single command or the first of multiple commands appears on the same line with the global command. All lines of a multi-line list except the last line must be ended with '\'. *A*, *i*, and *c* commands and associated input are permitted; the '.' terminating input mode may be omitted if it would be on the last line of the command list. The commands *g* and *v* are not permitted in the command list.

(.)i
<text>

This command inserts the given text before the addressed line. '.' is left at the last line input, or, if there were none, at the line before the addressed line. This command differs from the *a* command only in the placement of the text.

(.,.+1)j

This command joins the addressed lines into a single line; intermediate newlines simply disappear. '.' is left at the resulting line.

(.)kx

The mark command marks the addressed line with name *x*, which must be a lower-case

letter. The address form '*x*' then addresses this line.

(.,.)l

The list command prints the addressed lines in an unambiguous way: non-graphic characters are printed in two-digit octal, and long lines are folded. The *l* command may be placed on the same line after any non-i/o command.

(.,.)ma

The move command repositions the addressed lines after the line addressed by *a*. The last of the moved lines becomes the current line.

(.,.)n

The number command prints the addressed lines with line numbers and a tab at the left.

(.,.)p

The print command prints the addressed lines. '.' is left at the last line printed. The *p* command may be placed on the same line after any non-i/o command.

(.,.)P

This command is a synonym for *p*.

q The quit command causes *ed* to exit. No automatic write of a file is done.

Q This command is the same as *q*, except that no diagnostic results when no *w* has been given since the last buffer alteration.

(\$)r filename

The read command reads in the given file after the addressed line. If no file name is given, the remembered file name, if any, is used (see *e* and *f* commands). The file name is remembered if there was no remembered file name already. Address '0' is legal for *r* and causes the file to be read at the beginning of the buffer. If the read is successful, the number of characters read is typed. '.' is left at the last line read in from the file.

(.,.)s/regular expression/replacement/ or,

(.,.)s/regular expression/replacement/g

The substitute command searches each addressed line for an occurrence of the specified regular expression. On each line in which a match is found, all matched strings are replaced by the replacement specified, if the global replacement indicator 'g' appears after the command. If the global indicator does not appear, only the first occurrence of the matched string is replaced. It is an error for the substitution to fail on all addressed lines. Any punctuation character may be used instead of '/' to delimit the regular expression and the replacement. '.' is left at the last line substituted.

An ampersand '&' appearing in the replacement is replaced by the string matching the regular expression. The special meaning of '&' in this context may be suppressed by preceding it by '\'. The characters '\n' where *n* is a digit, are replaced by the text matched by the *n*-th regular subexpression enclosed between '\(' and '\)'. When nested, parenthesized subexpressions are present, *n* is determined by counting occurrences of '\(' starting from the left.

Lines may be split by substituting new-line characters into them. The new-line in the replacement string must be escaped by preceding it by '\'.

One or two trailing delimiters may be omitted, implying the 'p' suffix. The special form 's' followed by *no* delimiters repeats the most recent substitute command on the addressed lines. The 's' may be followed by the letters *r* (use the most recent regular expression for the left hand side, instead of the most recent left hand side of a substitute command), *p* (complement the setting of the *p* suffix from the previous substitution), or *g* (complement the setting of the *g* suffix). These letters may be combined in any order.

(.,.)ta

This command acts just like the *m* command, except that a copy of the addressed lines is placed after address *a* (which may be 0). '.' is left on the last line of the copy.

(.,.)u

The undo command restores the buffer to its state before the most recent buffer modifying command. The current line is also restored. Buffer modifying commands are *a*, *c*, *d*, *g*, *i*, *k*, and *v*. For purposes of undo, *g* and *v* are considered to be a single buffer modifying command. Undo is its own inverse.

When *ed* runs out of memory (at about 8000 lines on any 16 bit mini-computer such as the PDP-11) This full undo is not possible, and *u* can only undo the effect of the most recent substitute on the current line. This restricted undo also applies to editor scripts when *ed* is invoked with the - option.

(1, \$) v/regular expression/command list

This command is the same as the global command *g* except that the command list is executed *g* with '.' initially set to every line *except* those matching the regular expression.

(1, \$) w filename

The write command writes the addressed lines onto the given file. If the file does not exist, it is created. The file name is remembered if there was no remembered file name already. If no file name is given, the remembered file name, if any, is used (see *e* and *f* commands). '.' is unchanged. If the command is successful, the number of characters written is printed.

(1, \$) W filename

This command is the same as *w*, except that the addressed lines are appended to the file.

(1, \$) wq filename

This command is the same as *w* except that afterwards a *q* command is done, exiting the editor after the file is written.

x A key string is demanded from the standard input. Later *r*, *e* and *w* commands will encrypt and decrypt the text with this key by the algorithm of *crypt*(1). An explicitly empty key turns off encryption. (.+1)z or,

(.+1)zn

This command scrolls through the buffer starting at the addressed line. 22 (or *n*, if given) lines are printed. The last line printed becomes the current line. The value *n* is sticky, in that it becomes the default for future *z* commands.

(\$) =

The line number of the addressed line is typed. '.' is unchanged by this command.

!<shell command>

The remainder of the line after the '!' is sent to *sh*(1) to be interpreted as a command. '.' is unchanged.

(.+1, .+1) <newline>

An address alone on a line causes the addressed line to be printed. A blank line alone is equivalent to '.+1p'; it is useful for stepping through text. If two addresses are present with no intervening semicolon, *ed* prints the range of lines. If they are separated by a semicolon, the second line is printed.

If an interrupt signal (ASCII DEL) is sent, *ed* prints '?interrupted' and returns to its command level.

Some size limitations: 512 characters per line, 256 characters per global command list, 64 characters per file name, and, on mini computers, 128K characters in the temporary file. The limit on the number of lines depends on the amount of core: each line takes 2 words.

When reading a file, *ed* discards ASCII NUL characters and all characters after the last newline. It refuses to read files containing non-ASCII characters.

FILES

/tmp/e*

edhup: work is saved here if terminal hangs up

SEE ALSO

B. W. Kernighan, *A Tutorial Introduction to the ED Text Editor*

B. W. Kernighan, *Advanced editing on UNIX*

ex(1), sed(1), crypt(1)

DIAGNOSTICS

'?name' for inaccessible file; '?self-explanatory message' for other errors.

To protect against throwing away valuable work, a *q* or *e* command is considered to be in error, unless a *w* has occurred since the last buffer change. A second *q* or *e* will be obeyed regardless.

BUGS

The *l* command mishandles DEL.

The *undo* command causes marks to be lost on affected lines.

The *x* command, *-x* option, and special treatment of hangups only work on UNIX.

NAME

ex, edit — text editor

SYNOPSIS

ex [-] [-v] [-t tag] [-r] [+command] [-l] name ...
edit [ex options]

DESCRIPTION

Ex is the root of a family of editors: *edit*, *ex* and *vi*. *Ex* is a superset of *ed*, with the most notable extension being a display editing facility. Display based editing is the focus of *vi*.

If you have not used *ed*, or are a casual user, you will find that the editor *edit* is convenient for you. It avoids some of the complexities of *ex* used mostly by systems programmers and persons very familiar with *ed*.

If you have a CRT terminal, you may wish to use a display based editor; in this case see *vi*(1), which is a command which focuses on the display editing portion of *ex*.

DOCUMENTATION

The document *Edit: A tutorial* provides a comprehensive introduction to *edit* assuming no previous knowledge of computers or the UNIX system.

The *Ex Reference Manual — Version 3.5* is a comprehensive and complete manual for the command mode features of *ex*, but you cannot learn to use the editor by reading it. For an introduction to more advanced forms of editing using the command mode of *ex* see the editing documents written by Brian Kernighan for the editor *ed*; the material in the introductory and advanced documents works also with *ex*.

An Introduction to Display Editing with Vi introduces the display editor *vi* and provides reference material on *vi*. All of these documents can be found in volume 2c of the Programmer's Manual. In addition, the *Vi Quick Reference* card summarizes the commands of *vi* in a useful, functional way, and is useful with the *Introduction*.

FILES

/usr/lib/ex?.?strings	error messages
/usr/lib/ex?.?recover	recover command
/usr/lib/ex?.?preserve	preserve command
/etc/termcap	describes capabilities of terminals
~/exrc	editor startup file
/tmp/Exnnnnn	editor temporary
/tmp/Rxnnnnn	named buffer temporary
/usr/preserve	preservation directory

SEE ALSO

awk(1), ed(1), grep(1), sed(1), grep(1), vi(1), termcap(5), environ(7)

AUTHOR

Originally written by William Joy

Mark Horton has maintained the editor since version 2.7, adding macros, support for many unusual terminals, and other features such as word abbreviation mode.

BUGS

The *undo* command causes all marks to be lost on lines changed and then restored if the marked lines were changed.

Undo never clears the buffer modified condition.

The *z* command prints a number of logical rather than physical lines. More than a screen full of output may result if long lines are present.

File input/output errors don't print a name if the command line '-' option is used.

There is no easy way to do a single scan ignoring case.

The editor does not warn if text is placed in named buffers and not used before exiting the editor.

Null characters are discarded in input files, and cannot appear in resultant files.

NAME

`expand`, `unexpand` - expand tabs to spaces, and vice versa

SYNOPSIS

```
expand [ - tabstop ] [ - tab1,tab2,...,tabn ] [ file ... ]  
unexpand [ - a ] [ file ... ]
```

DESCRIPTION

Expand processes the named files or the standard input writing the standard output with tabs changed into blanks. Backspace characters are preserved into the output and decrement the column count for tab calculations. *Expand* is useful for pre-processing character files (before sorting, looking at specific columns, etc.) that contain tabs.

If a single *tabstop* argument is given then tabs are set *tabstop* spaces apart instead of the default 8. If multiple tabstops are given then the tabs are set at those specific columns.

Unexpand puts tabs back into the data from the standard input or the named files and writes the result on the standard output. By default only leading blanks and tabs are reconverted to maximal strings of tabs. If the `- a` option is given, then tabs are inserted whenever they would compress the resultant file by replacing two or more characters.

NAME

expr — evaluate arguments as an expression

SYNOPSIS

expr *arg* ...

DESCRIPTION

The arguments are taken as an expression. After evaluation, the result is written on the standard output. Each token of the expression is a separate argument.

The operators and keywords are listed below. The list is in order of increasing precedence, with equal precedence operators grouped.

expr | *expr*

yields the first *expr* if it is neither null nor '0', otherwise yields the second *expr*.

expr & *expr*

yields the first *expr* if neither *expr* is null or '0', otherwise yields '0'.

expr *relop* *expr*

where *relop* is one of < <= = != >= >, yields '1' if the indicated comparison is true, '0' if false. The comparison is numeric if both *expr* are integers, otherwise lexicographic.

expr + *expr*

expr - *expr*

addition or subtraction of the arguments.

expr * *expr*

expr / *expr*

expr % *expr*

multiplication, division, or remainder of the arguments.

expr : *expr*

The matching operator compares the string first argument with the regular expression second argument; regular expression syntax is the same as that of *ed*(1). The `\(...\)` pattern symbols can be used to select a portion of the first argument. Otherwise, the matching operator yields the number of characters matched ('0' on failure).

(*expr*)

parentheses for grouping.

Examples:

To add 1 to the Shell variable *a*:

```
a=`expr $a + 1`
```

To find the filename part (least significant part) of the pathname stored in variable *a*, which may or may not contain '/':

```
expr $a : '.*\(.*)' ↑ $a
```

Note the quoted Shell metacharacters.

SEE ALSO

sh(1), *test*(1)

DIAGNOSTICS

Expr returns the following exit codes:

- | | |
|---|--|
| 0 | if the expression is neither null nor '0', |
| 1 | if the expression is null or '0', |
| 2 | for invalid expressions. |

NAME

false, true — provide truth values

SYNOPSIS

true

false

DESCRIPTION

True and *false* are usually used in a Bourne shell script. They test for the appropriate status "true" or "false" before running (or failing to run) a list of commands.

EXAMPLE

```
while false
do
    command list
done
```

SEE ALSO

csh(1), sh(1), true(1)

DIAGNOSTICS

False has exit status nonzero.

NAME

file — determine file type

SYNOPSIS

file file ...

DESCRIPTION

File performs a series of tests on each argument in an attempt to classify it. If an argument appears to be ascii, *file* examines the first 512 bytes and tries to guess its language.

BUGS

It often makes mistakes. In particular it often suggests that command files are C programs. Does not recognize Pascal or LISP.

NAME

find — find files

SYNOPSIS

find pathname-list expression

DESCRIPTION

Find recursively descends the directory hierarchy for each pathname in the *pathname-list* (i.e., one or more pathnames) seeking files that match a boolean *expression* written in the primaries given below. In the descriptions, the argument *n* is used as a decimal integer where *+n* means more than *n*, *-n* means less than *n* and *n* means exactly *n*.

-name filename

True if the *filename* argument matches the current file name. Normal Shell argument syntax may be used if escaped (watch out for '[', '?' and '*').

-perm onum

True if the file permission flags exactly match the octal number *onum* (see *chmod(1)*). If *onum* is prefixed by a minus sign, more flag bits (017777, see *stat(2)*) become significant and the flags are compared: $(\text{flags}\&\text{onum}) = \text{onum}$.

-type c True if the type of the file is *c*, where *c* is **b**, **c**, **d**, **f** or **l** for block special file, character special file, directory, plain file, or symbolic link.

-links n True if the file has *n* links.

-user uname

True if the file belongs to the user *uname* (login name or numeric user ID).

-group gname

True if the file belongs to group *gname* (group name or numeric group ID).

-size n True if the file is *n* blocks long (512 bytes per block).

-inum n True if the file has inode number *n*.

-atime n True if the file has been accessed in *n* days.

-mtime n

True if the file has been modified in *n* days.

-exec command

True if the executed command returns a zero value as exit status. The end of the command must be punctuated by an escaped semicolon. A command argument '{}' is replaced by the current pathname.

-ok command

Like **-exec** except that the generated command is written on the standard output, then the standard input is read and the command executed only upon response *y*.

-print Always true; causes the current pathname to be printed.

-newer file

True if the current file has been modified more recently than the argument *file*.

The primaries may be combined using the following operators (in order of decreasing precedence):

- 1) A parenthesized group of primaries and operators (parentheses are special to the Shell and must be escaped).
- 2) The negation of a primary ('!' is the unary *not* operator).
- 3) Concatenation of primaries (the *and* operation is implied by the juxtaposition of two primaries).

- 4) Alternation of primaries ('-o' is the *or* operator).

EXAMPLE

To remove all files named 'a.out' or '*.o' that have not been accessed for a week:

```
find / \( -name a.out -o -name '*.o' \) -atime +7 -exec rm {} \;
```

FILES

/etc/passwd
/etc/group

SEE ALSO

sh(1), test(1), fs(5)

BUGS

The syntax is painful.

NAME

get — get a version of an SCCS file

SYNOPSIS

```
get [-rSID] [-ccutoff] [-ilist] [-xlist] [-aseq-no.] [-k] [-e]
[-l{p}] [-p] [-m] [-n] [-s] [-b] [-g] [-t] file ...
```

DESCRIPTION

Get generates an ASCII text file from each named SCCS file according to the specifications given by its keyletter arguments, which begin with *-*. The arguments may be specified in any order, but all keyletter arguments apply to all named SCCS files. If a directory is named, *get* behaves as though each file in the directory were specified as a named file, except that non-SCCS files (last component of the path name does not begin with *s.*) and unreadable files are silently ignored. If a name of *-* is given, the standard input is read; each line of the standard input is taken to be the name of an SCCS file to be processed. Again, non-SCCS files and unreadable files are silently ignored.

The generated text is normally written into a file called the *g-file* whose name is derived from the SCCS file name by simply removing the leading *s.*; (see also *FILES*, below).

Each of the keyletter arguments is explained below as though only one SCCS file is to be processed, but the effects of any keyletter argument applies independently to each named file.

-rSID The SCCS *ID*entification string (SID) of the version (*delta*) of an SCCS file to be retrieved. Table 1 below shows, for the most useful cases, what version of an SCCS file is retrieved (as well as the SID of the version to be eventually created by *delta*(1) if the *-e* keyletter is also used), as a function of the SID specified.

-ccutoff *Cutoff* date-time, in the form:

```
YY[MM[DD[HH[MM[SS]]]]]
```

No changes (*deltas*) to the SCCS file which were created after the specified *cutoff* date-time are included in the generated ASCII text file. Units omitted from the date-time default to their maximum possible values; that is, *-c7502* is equivalent to *-c750228235959*. Any number of non-numeric characters may separate the various 2 digit pieces of the *cutoff* date-time. This feature allows one to specify a *cutoff* date in the form: *"-c77/2/2 9:22:25"*. Note that this implies that one may use the *%E%* and *%U%* identification keywords (see below) for nested *gets* within, say the input to a *send*(1C) command:

```
^!get "-c%E% %U%" s.file
```

-e Indicates that the *get* is for the purpose of editing or making a change (*delta*) to the SCCS file via a subsequent use of *delta*(1). The *-e* keyletter used in a *get* for a particular version (SID) of the SCCS file prevents further *gets* for editing on the same SID until *delta* is executed or the *j* (joint edit) flag is set in the SCCS file (see *admin*(1)). Concurrent use of *get -e* for different SIDs is always allowed.

If the *g-file* generated by *get* with an *-e* keyletter is accidentally ruined in the process of editing it, it may be regenerated by re-executing the *get* command with the *-k* keyletter in place of the *-e* keyletter.

SCCS file protection specified via the ceiling, floor, and authorized user list stored in the SCCS file (see *admin(1)*) are enforced when the *-e* keyletter is used.

- b** Used with the *-e* keyletter to indicate that the new delta should have an SID in a new branch as shown in Table 1. This keyletter is ignored if the *b* flag is not present in the file (see *admin(1)*) or if the retrieved *delta* is not a leaf *delta*. (A leaf *delta* is one that has no successors on the SCCS file tree.)
Note: A branch *delta* may always be created from a non-leaf *delta*.
- ilist** A *list* of deltas to be included (forced to be applied) in the creation of the generated file. The *list* has the following syntax:

```

<list> ::= <range> | <list> , <range>
<range> ::= SID | SID - SID

```

 SID, the SCCS Identification of a delta, may be in any form shown in the "SID Specified" column of Table 1. Partial SIDs are interpreted as shown in the "SID Retrieved" column of Table 1.
- xlist** A *list* of deltas to be excluded (forced not to be applied) in the creation of the generated file. See the *-i* keyletter for the *list* format.
- k** Suppresses replacement of identification keywords (see below) in the retrieved text by their value. The *-k* keyletter is implied by the *-e* keyletter.
- l[p]** Causes a delta summary to be written into an *l-file*. If *-lp* is used then an *l-file* is not created; the delta summary is written on the standard output instead. See *FILES* for the format of the *l-file*.
- p** Causes the text retrieved from the SCCS file to be written on the standard output. No *g-file* is created. All output which normally goes to the standard output goes to file descriptor 2 instead, unless the *-s* keyletter is used, in which case it disappears.
- s** Suppresses all output normally written on the standard output. However, fatal error messages (which always go to file descriptor 2) remain unaffected.
- m** Causes each text line retrieved from the SCCS file to be preceded by the SID of the delta that inserted the text line in the SCCS file. The format is: SID, followed by a horizontal tab, followed by the text line.
- n** Causes each generated text line to be preceded with the %M% identification keyword value (see below). The format is: %M% value, followed by a horizontal tab, followed by the text line. When both the *-m* and *-n* keyletters are used, the format is: %M% value, followed by a horizontal tab, followed by the *-m* keyletter generated format.
- g** Suppresses the actual retrieval of text from the SCCS file. It is primarily used to generate an *l-file*, or to verify the existence of a particular SID.
- t** Used to access the most recently created ("top") delta in a given release (e.g., *-r1*), or release and level (e.g., *-r1.2*).

-aseq-no. The delta sequence number of the SCCS file delta (version) to be retrieved (see *sccsfile(5)*). This keyletter is used by the *comb(1)* command; it is not a generally useful keyletter, and users should not use it. If both the **-r** and **-a** keyletters are specified, the **-a** keyletter is used. Care should be taken when using the **-a** keyletter in conjunction with the **-e** keyletter, as the SID of the delta to be created may not be what one expects. The **-r** keyletter can be used with the **-a** and **-e** keyletters to control the naming of the SID of the delta to be created.

For each file processed, *get* responds (on the standard output) with the SID being accessed and with the number of lines retrieved from the SCCS file.

If the **-e** keyletter is used, the SID of the delta to be made appears after the SID accessed and before the number of lines generated. If there is more than one named file or if a directory or standard input is named, each file name is printed (preceded by a new-line) before it is processed. If the **-i** keyletter is used included deltas are listed following the notation "Included"; if the **-x** keyletter is used, excluded deltas are listed following the notation "Excluded".

TABLE 1. Determination of SCCS Identification String

SID* Specified	-b Keyletter Used†	Other Conditions	SID Retrieved	SID of Delta to be Created
none‡	no	R defaults to mR	mR.mL	mR.(mL+1)
none‡	yes	R defaults to mR	mR.mL	mR.mL.(mB+1).1
R	no	R > mR	mR.mL	R.1***
R	no	R = mR	mR.mL	mR.(mL+1)
R	yes	R > mR	mR.mL	mR.mL.(mB+1).1
R	yes	R = mR	mR.mL	mR.mL.(mB+1).1
R	-	R < mR and R does <i>not</i> exist	hR.mL**	hR.mL.(mB+1).1
R	-	Trunk succ.# in release > R and R exists	R.mL	R.mL.(mB+1).1
R.L	no	No trunk succ.	R.L	R.(L+1)
R.L	yes	No trunk succ.	R.L	R.L.(mB+1).1
R.L	-	Trunk succ. in release ≥ R	R.L	R.L.(mB+1).1
R.L.B	no	No branch succ.	R.L.B.mS	R.L.B.(mS+1)
R.L.B	yes	No branch succ.	R.L.B.mS	R.L.(mB+1).1
R.L.B.S	no	No branch succ.	R.L.B.S	R.L.B.(S+1)
R.L.B.S	yes	No branch succ.	R.L.B.S	R.L.(mB+1).1
R.L.B.S	-	Branch succ.	R.L.B.S	R.L.(mB+1).1

* "R", "L", "B", and "S" are the "release", "level", "branch", and "sequence" components of the SID, respectively; "m" means "maximum". Thus, for example, "R.mL" means "the maximum level number within release R"; "R.L.(mB+1).1" means "the first sequence number on the *new* branch (i.e., maximum branch number plus one) of level L within release R". Note that if the SID specified is of the form "R.L", "R.L.B", or "R.L.B.S", each of the specified components *must* exist.

** "hR" is the highest *existing* release that is lower than the specified, *nonexistent*, release R.

- *** This is used to force creation of the *first* delta in a *new* release.
- # Successor.
- † The *-b* keyletter is effective only if the *b* flag (see *admin(1)*) is present in the file. An entry of *-* means “irrelevant”.
- ‡ This case applies if the *d* (default SID) flag is *not* present in the file. If the *d* flag is present in the file, then the SID obtained from the *d* flag is interpreted as if it had been specified on the command line. Thus, one of the other cases in this table applies.

IDENTIFICATION KEYWORDS

Identifying information is inserted into the text retrieved from the SCCS file by replacing *identification keywords* with their value wherever they occur. The following keywords may be used in the text stored in an SCCS file:

<i>Keyword</i>	<i>Value</i>
%M%	Module name: either the value of the <i>m</i> flag in the file (see <i>admin(1)</i>), or if absent, the name of the SCCS file with the leading <i>s.</i> removed.
%I%	SCCS identification (SID) (%R%.%L%.%B%.%S%) of the retrieved text.
%R%	Release.
%L%	Level.
%B%	Branch.
%S%	Sequence.
%D%	Current date (YY/MM/DD).
%H%	Current date (MM/DD/YY).
%T%	Current time (HH:MM:SS).
%E%	Date newest applied delta was created (YY/MM/DD).
%G%	Date newest applied delta was created (MM/DD/YY).
%U%	Time newest applied delta was created (HH:MM:SS).
%Y%	Module type: value of the <i>t</i> flag in the SCCS file (see <i>admin(1)</i>).
%F%	SCCS file name.
%P%	Fully qualified SCCS file name.
%Q%	The value of the <i>q</i> flag in the file (see <i>admin(1)</i>).
%C%	Current line number. This keyword is intended for identifying messages output by the program such as “this shouldn’t have happened” type errors. It is <i>not</i> intended to be used on every line to provide sequence numbers.
%Z%	The 4-character string <i>@(#)</i> recognizable by <i>what(1)</i> .
%W%	A shorthand notation for constructing <i>what(1)</i> strings for UNIX program files. %W% = %Z%%M%<horizontal-tab>%I%
%A%	Another shorthand notation for constructing <i>what(1)</i> strings for non-UNIX program files. %A% = %Z%%Y% %M% %I%%Z%

FILES

Several auxiliary files may be created by *get*. These files are known generically as the *g-file*, *l-file*, *p-file*, and *z-file*. The letter before the hyphen is called the tag. An auxiliary file name is formed from the SCCS file name: the last component of all SCCS file names must be of the form *s.module-name*, the auxiliary files are named by replacing the leading *s* with the tag. The *g-file* is an exception to this scheme: the *g-file* is named by removing the *s.* prefix. For example, *s.xyz.c*, the auxiliary file names would be *xyz.c*, *l.xyz.c*, *p.xyz.c*, and *z.xyz.c*, respectively.

The *g-file*, which contains the generated text, is created in the current directory (unless the *-p* keyletter is used). A *g-file* is created in all cases, whether or not any lines of text were generated by the *get*. It is owned by the real user. If the *-k* keyletter is used or implied its mode is 644; otherwise its mode is 444. Only the real user need have write permission in

the current directory.

The *l-file* contains a table showing which deltas were applied in generating the retrieved text. The *l-file* is created in the current directory if the `-l` keyletter is used; its mode is 444 and it is owned by the real user. Only the real user need have write permission in the current directory.

Lines in the *l-file* have the following format:

- a. A blank character if the delta was applied;
* otherwise.
- b. A blank character if the delta was applied or wasn't applied and ignored;
* if the delta wasn't applied and wasn't ignored.
- c. A code indicating a "special" reason why the delta was or was not applied:
"I": Included.
"X": Excluded.
"C": Cut off (by a `-c` keyletter).
- d. Blank.
- e. SCCS identification (SID).
- f. Tab character.
- g. Date and time (in the form YY/MM/DD HH:MM:SS) of creation.
- h. Blank.
- i. Login name of person who created *delta*.

The comments and MR data follow on subsequent lines, indented one horizontal tab character. A blank line terminates each entry.

The *p-file* is used to pass information resulting from a *get* with an `-e` keyletter along to *delta*. Its contents are also used to prevent a subsequent execution of *get* with an `-e` keyletter for the same SID until *delta* is executed or the joint edit flag, `j`, (see *admin*(1)) is set in the SCCS file. The *p-file* is created in the directory containing the SCCS file and the effective user must have write permission in that directory. Its mode is 644 and it is owned by the effective user. The format of the *p-file* is: the gotten SID, followed by a blank, followed by the SID that the new delta will have when it is made, followed by a blank, followed by the login name of the real user, followed by a blank, followed by the date-time the *get* was executed, followed by a blank and the `-i` keyletter argument if it was present, followed by a blank and the `-x` keyletter argument if it was present, followed by a new-line. There can be an arbitrary number of lines in the *p-file* at any time; no two lines can have the same new delta SID.

The *z-file* serves as a *lock-out* mechanism against simultaneous updates. Its contents are the binary (2 bytes) process ID of the command (i.e., *get*) that created it. The *z-file* is created in the directory containing the SCCS file for the duration of *get*. The same protection restrictions as those for the *p-file* apply for the *z-file*. The *z-file* is created mode 444.

SEE ALSO

admin(1), *delta*(1), *help*(1), *prs*(1), *what*(1), *sccsfile*(5).
Source Code Control System User's Guide by L. E. Bonanni and C. A. Salemi.

DIAGNOSTICS

Use *help*(1) for explanations.

BUGS

If the effective user has write permission (either explicitly or implicitly) in the directory containing the SCCS files, but the real user doesn't, then only one file may be named when the `-e` keyletter is used.

NAME

`getline` - get a line from stdin

SYNOPSIS

`getline`

DESCRIPTION

getline retrieves a line of text from the standard input device (normally a terminal), waiting for a carriage return or newline to signal the end of input. It is useful for handling user input to shell scripts.

EXAMPLE

a shell script to retrieve a user's name and acknowledge it:

```
echo -n "enter your name: "  
username = `getline`  
echo Hello $username
```

SEE ALSO

`sh(1)`, `csh(1)`

NAME

grep, **egrep**, **fgrep** — search a file for a pattern

SYNOPSIS

grep [option] ... expression [file] ...

egrep [option] ... [expression] [file] ...

fgrep [option] ... [strings] [file]

DESCRIPTION

Commands of the *grep* family search the input *files* (standard input default) for lines matching a pattern. Normally, each line found is copied to the standard output. *Grep* patterns are limited regular expressions in the style of *ex(1)*; it uses a compact nondeterministic algorithm. *Egrep* patterns are full regular expressions; it uses a fast deterministic algorithm that sometimes needs exponential space. *Fgrep* patterns are fixed strings; it is fast and compact. The following options are recognized.

- v** All lines but those matching are printed.
- x** (Exact) only lines matched in their entirety are printed (*fgrep* only).
- c** Only a count of matching lines is printed.
- l** The names of files with matching lines are listed (once) separated by newlines.
- n** Each line is preceded by its relative line number in the file.
- b** Each line is preceded by the block number on which it was found. This is sometimes useful in locating disk block numbers by context.
- i** The case of letters is ignored in making comparisons — that is, upper and lower case are considered identical. This applies to *grep* and *fgrep* only.
- s** Silent mode. Nothing is printed (except error messages). This is useful for checking the error status.
- w** The expression is searched for as a word (as if surrounded by ‘\<’ and ‘\>’, see *ex(1)*.) (*grep* only)
- e expression**
Same as a simple *expression* argument, but useful when the *expression* begins with a **-**.
- f file** The regular expression (*egrep*) or string list (*fgrep*) is taken from the *file*.

In all cases the file name is shown if there is more than one input file. Care should be taken when using the characters **\$**, *****, **[**, **^**, **|**, **(**, **)** and **** in the *expression* as they are also meaningful to the Shell. It is safest to enclose the entire *expression* argument in single quotes ‘ ’.

Fgrep searches for lines that contain one of the (newline-separated) *strings*.

Egrep accepts extended regular expressions. In the following description ‘character’ excludes newline:

- A **** followed by a single character other than newline matches that character.
- The character **^** matches the beginning of a line.
- The character **\$** matches the end of a line.
- A **.** (period) matches any character.
- A single character not otherwise endowed with special meaning matches that character.
- A string enclosed in brackets **[]** matches any single character from the string. Ranges of ASCII character codes may be abbreviated as in ‘a-z0-9’. A **]** may occur only as the first character of the string. A literal **-** must be placed where it can’t be mistaken

as a range indicator.

A regular expression followed by an * (asterisk) matches a sequence of 0 or more matches of the regular expression. A regular expression followed by a + (plus) matches a sequence of 1 or more matches of the regular expression. A regular expression followed by a ? (question mark) matches a sequence of 0 or 1 matches of the regular expression.

Two regular expressions concatenated match a match of the first followed by a match of the second.

Two regular expressions separated by | or newline match either a match for the first or a match for the second.

A regular expression enclosed in parentheses matches a match for the regular expression.

The order of precedence of operators at the same parenthesis level is [] then *+? then concatenation then | and newline.

Ideally there should be only one *grep*, but we don't know a single algorithm that spans a wide enough range of space-time tradeoffs.

SEE ALSO

ex(1), sed(1), sh(1)

DIAGNOSTICS

Exit status is 0 if any matches are found, 1 if none, 2 for syntax errors or inaccessible files.

BUGS

Lines are limited to 256 characters; longer lines are truncated.

NAME

groups — show group memberships

SYNOPSIS

groups [user]

DESCRIPTION

The *groups* command shows the groups to which you or the optionally specified user belong. Each user belongs to a group specified in the password file */etc/passwd* and possibly to other groups as specified in the file */etc/group*. If you do not own a file but belong to the group which it is owned by then you are granted group access to the file.

When a new file is created it is given the group of the containing directory.

SEE ALSO

setgroups(2)

FILES

/etc/passwd, */etc/group*

BUGS

More groups should be allowed.

NAME

head - give first few lines

SYNOPSIS

head [-count] [file ...]

DESCRIPTION

This filter gives the first *count* lines of each of the specified files, or of the standard input. If *count* is omitted it defaults to 10.

SEE ALSO

tail(1)

NAME

`hostid` — set or print identifier of current host system

SYNOPSIS

`hostid` [identifier]

DESCRIPTION

The *hostid* command prints the identifier of the current host. This numeric value is expected to be unique across all hosts and is normally set to the host's Internet address. The super-user can set the *hostid* by giving an argument; this is usually done in the startup script `/etc/rc.local`.

SEE ALSO

`gethostid(2)`, `sethostid(2)`

NAME

`hostname` — set or print name of current host system

SYNOPSIS

`hostname [nameofhost]`

DESCRIPTION

The *hostname* command prints the name of the current host, as given before the “login” prompt. The super-user can set the hostname by giving an argument; this is usually done in the startup script `/etc/rc.local`.

SEE ALSO

`gethostname(2)`, `sethostname(2)`

NAME

install — install binaries

SYNOPSIS

install [**-c**] [**-m mode**] [**-o owner**] [**-g group**] [**-s**] *binary* *destination*

DESCRIPTION

Binary is moved (or copied if **-c** is specified) to *destination*. If *destination* already exists, it is removed before *binary* is moved. If the destination is a directory then *binary* is moved into the *destination* directory with its original file-name.

The mode for *Destination* is set to 755; the **-m mode** option may be used to specify a different mode.

Destination is changed to owner root; the **-o owner** option may be used to specify a different owner.

Destination is changed to group staff; the **-g group** option may be used to specify a different group.

If the **-s** option is specified the binary is stripped after being installed.

Install refuses to move a file onto itself.

SEE ALSO

chgrp(1), chmod(1), cp(1), mv(1), strip(1), chown(8)

NAME

`iostat` — report I/O statistics

SYNOPSIS

`iostat` [*interval* [*count*]]

DESCRIPTION

Iostat iteratively reports the number of characters read and written to terminals, and, for each disk, the number of seeks transfers per second, kilobytes transferred per second, and the milliseconds per average seek. It also gives the percentage of time the system has spent in user mode, in user mode running low priority (niced) processes, in system mode, and idling.

To compute this information, for each disk, seeks and data transfer completions and number of words transferred are counted; for terminals collectively, the number of input and output characters are counted. Also, each sixtieth of a second, the state of each disk is examined and a tally is made if the disk is active. From these numbers and given the transfer rates of the devices it is possible to determine average seek times for each device.

The optional *interval* argument causes *iostat* to report once each *interval* seconds. The first report is for all time since a reboot and each subsequent report is for the last interval only.

The optional *count* argument restricts the number of reports.

FILES

`/dev/kmem`
`/vmunix`

SEE ALSO

`vmstat(1)`

NAME

join — relational database operator

SYNOPSIS

join [options] file1 file2

DESCRIPTION

Join forms, on the standard output, a join of the two relations specified by the lines of *file1* and *file2*. If *file1* is '-', the standard input is used.

File1 and *file2* must be sorted in increasing ASCII collating sequence on the fields on which they are to be joined, normally the first in each line.

There is one line in the output for each pair of lines in *file1* and *file2* that have identical join fields. The output line normally consists of the common field, then the rest of the line from *file1*, then the rest of the line from *file2*.

Fields are normally separated by blank, tab or newline. In this case, multiple separators count as one, and leading separators are discarded.

These options are recognized:

- a *n* In addition to the normal output, produce a line for each unpairable line in file *n*, where *n* is 1 or 2.
- e *s* Replace empty output fields by string *s*.
- j *n m* Join on the *m*th field of file *n*. If *n* is missing, use the *m*th field in each file.
- o *list* Each output line comprises the fields specified in *list*, each element of which has the form *n.m*, where *n* is a file number and *m* is a field number.
- t *c* Use character *c* as a separator (tab character). Every appearance of *c* in a line is significant.

SEE ALSO

sort(1), comm(1), awk(1)

BUGS

With default field separation, the collating sequence is that of *sort -b*; with *-t*, the sequence is that of a plain sort.

The conventions of *join*, *sort*, *comm*, *uniq*, *look* and *awk(1)* are wildly incongruous.

NAME

kill — terminate a process with extreme prejudice

SYNOPSIS

kill [**-sig**] processid ...

kill -l

DESCRIPTION

Kill sends the TERM (terminate, 15) signal to the specified processes. If a signal name or number preceded by '-' is given as first argument, that signal is sent instead of terminate (see *sigvec(2)*). The signal names are listed by 'kill -l', and are as given in */usr/include/signal.h*, stripped of the common SIG prefix.

The terminate signal will kill processes that do not catch the signal; 'kill -9 ...' is a sure kill, as the KILL (9) signal cannot be caught. By convention, if process number 0 is specified, all members in the process group (i.e. processes resulting from the current login) are signaled (but beware: this works only if you use *sh(1)*; not if you use *cs(1)*.) The killed processes must belong to the current user unless he is the super-user.

The process number of an asynchronous process started with '&' is reported by the shell. Process numbers can also be found by using *Kill* is a built-in to *cs(1)*; it allows job specifiers "%..." so process id's are not as often used as *kill* arguments. See *cs(1)* for details.

SEE ALSO

cs(1), *ps(1)*, *kill(2)*, *sigvec(2)*

BUGS

An option to kill process groups ala *killpg(2)* should be provided; a replacement for "kill 0" for *cs(1)* users should be provided.

NAME

`last` — indicate last logins of users and teletypes

SYNOPSIS

`last [-N] [name ...] [tty ...]`

DESCRIPTION

Last will look back in the *wtmp* file which records all logins and logouts for information about a user, a teletype or any group of users and teletypes. Arguments specify names of users or teletypes of interest. Names of teletypes may be given fully or abbreviated. For example 'last 0' is the same as 'last tty0'. If multiple arguments are given, the information which applies to any of the arguments is printed. For example 'last root console' would list all of "root's" sessions as well as all sessions on the console terminal. *Last* will print the sessions of the specified users and teletypes, most recent first, indicating the times at which the session began, the duration of the session, and the teletype which the session took place on. If the session is still continuing or was cut short by a reboot, *last* so indicates.

The pseudo-user **reboot** logs in at reboots of the system, thus

```
last reboot
```

will give an indication of mean time between reboot.

Last with no arguments prints a record of all logins and logouts, in reverse order. The `-N` option limits the report to N lines.

If *last* is interrupted, it indicates how far the search has progressed in *wtmp*. If interrupted with a quit signal (generated by a control-\) *last* indicates how far the search has progressed so far, and the search continues.

FILES

`/usr/adm/wtmp` login data base
`/usr/adm/shutdownlog` which records shutdowns and reasons for same

SEE ALSO

`wtmp(5)`, `ac(8)`, `lastcomm(1)`

AUTHOR

Howard Katseff

NAME

`ld` — link editor

SYNOPSIS

`ld [option] ... file ...`

DESCRIPTION

Ld combines several object programs into one, resolves external references, and searches libraries. In the simplest case several object *files* are given, and *ld* combines them, producing an object module which can be either executed or become the input for a further *ld* run. (In the latter case, the `-r` option must be given to preserve the relocation bits.) The output of *ld* is left on `a.out`. This file is made executable only if no errors occurred during the load.

The argument routines are concatenated in the order specified. The entry point of the output is the beginning of the first routine (unless the `-e` option is specified).

If any argument is a library, it is searched exactly once at the point it is encountered in the argument list. Only those routines defining an unresolved external reference are loaded. If a routine from a library references another routine in the library, and the library has not been processed by *ranlib*(1), the referenced routine must appear after the referencing routine in the library. Thus the order of programs within libraries may be important. The first member of a library should be a file named `__SYMDEF`, which is understood to be a dictionary for the library as produced by *ranlib*(1); the dictionary is searched iteratively to satisfy as many references as possible.

The symbols `__etext`, `__edata` and `__end` (`etext`, `edata` and `end` in C) are reserved, and if referred to, are set to the first location above the program, the first location above initialized data, and the first location above all data respectively. It is erroneous to define these symbols.

Ld understands several options. Except for `-l`, they should appear before the file names.

- `-A` This option specifies incremental loading, i.e. linking is to be done in a manner so that the resulting object may be read into an already executing program. The next argument is the name of a file whose symbol table will be taken as a basis on which to define additional symbols. Only newly linked material will be entered into the text and data portions of `a.out`, but the new symbol table will reflect every symbol defined before and after the incremental load. This argument must appear before any other object file in the argument list. The `-T` option may be used as well, and will be taken to mean that the newly linked segment will commence at the corresponding address (which must be a multiple of 1024). The default value is the old value of `__end`.
- `-D` Take the next argument as a hexadecimal number and pad the data segment with zero bytes to the indicated length.
- `-d` Force definition of common storage even if the `-r` flag is present.
- `-e` The following argument is taken to be the name of the entry point of the loaded program; location 0 is the default.
- `-lx` This option is an abbreviation for the library name `‘/lib/libx.a’`, where *x* is a string. If that does not exist, *ld* tries `‘/usr/lib/libx.a’`. A library is searched when its name is encountered, so the placement of a `-l` is significant.
- `-M` produce a primitive load map, listing the names of the files which will be loaded.
- `-N` Do not make the text portion read only or sharable. (Use "magic number" 0407.)
- `-n` Arrange (by giving the output file a 0410 "magic number") that when the output file is executed, the text portion will be read-only and shared among all users executing the file. This involves moving the data areas up to the first possible 1024 byte boundary following the end of the text.

- o** The *name* argument after **-o** is used as the name of the *ld* output file, instead of **a.out**.
- r** Generate relocation bits in the output file so that it can be the subject of another *ld* run. This flag also prevents final definitions from being given to common symbols, and suppresses the 'undefined symbol' diagnostics.
- S** 'Strip' the output by removing all symbols except locals and globals.
- s** 'Strip' the output, that is, remove the symbol table and relocation bits to save space (but impair the usefulness of the debuggers). This information can also be removed by *strip(1)*.
- T** The next argument is a hexadecimal number which sets the text segment origin. The default origin is 0.
- t** ("trace") Print the name of each file as it is processed.
- u** Take the following argument as a symbol and enter it as undefined in the symbol table. This is useful for loading wholly from a library, since initially the symbol table is empty and an unresolved reference is needed to force the loading of the first routine.
- X** Save local symbols except for those whose names begin with 'L'. This option is used by *cc(1)* to discard internally-generated labels while retaining symbols local to routines.
- x** Do not preserve local (non-globl) symbols in the output symbol table; only enter external symbols. This option saves some space in the output file.
- y***sym* Indicate each file in which *sym* appears, its type and whether the file defines or references it. Many such options may be given to trace many symbols. (It is usually necessary to begin *sym* with an '_', as external C, FORTRAN and Pascal variables begin with underscores.)
- z** Arrange for the process to be loaded on demand from the resulting executable file (413 format) rather than preloaded. This is the default. Results in a 1024 byte header on the output file followed by a text and data segment each of which have size a multiple of 1024 bytes (being padded out with nulls in the file if necessary). With this format the first few BSS segment symbols may actually appear (from the output of *size(1)*) to live in the data segment; this to avoid wasting the space resulting from data segment size roundup.

FILES

/lib/lib*.a	libraries
/usr/lib/lib*.a	more libraries
/usr/local/lib/lib*.a	still more libraries
a.out	output file

SEE ALSO

as(1), ar(1), cc(1), ranlib(1)

BUGS

There is no way to force data to be page aligned. *Ld* pads images which are to be demand loaded from the file system to the next page boundary to avoid a bug in the system.

NAME

`lex` — generator of lexical analysis programs

SYNOPSIS

`lex [-tvfn] [file] ...`

DESCRIPTION

Lex generates programs to be used in simple lexical analysis of text. The input *files* (standard input default) contain regular expressions to be searched for, and actions written in C to be executed when expressions are found.

A C source program, 'lex.yy.c' is generated, to be compiled thus:

```
cc lex.yy.c -ll
```

This program, when run, copies unrecognized portions of the input to the output, and executes the associated C action for each regular expression that is recognized.

The options have the following meanings.

- `-t` Place the result on the standard output instead of in file "lex.yy.c".
- `-v` Print a one-line summary of statistics of the generated analyzer.
- `-n` Opposite of `-v`; `-n` is default.
- `-f` "Faster" compilation: don't bother to pack the resulting tables; limited to small programs.

EXAMPLE

```
lex lexcommands
```

would draw *lex* instructions from the file *lexcommands*, and place the output in *lex.yy.c*

```
%%
[A-Z] putchar(yytext[0]+'a'-'A');
[ ]+$
[ ]+  putchar(' ');
```

is an example of a *lex* program that would be put into a *lex* command file. This program converts upper case to lower, removes blanks at the end of lines, and replaces multiple blanks by single blanks.

SEE ALSO

yacc(1), sed(1)

M. E. Lesk and E. Schmidt, *LEX — Lexical Analyzer Generator*

NAME

lint — a C program verifier

SYNOPSIS

lint [-abchnpux] file ...

DESCRIPTION

Lint attempts to detect features of the C program *files* which are likely to be bugs, or non-portable, or wasteful. It also checks the type usage of the program more strictly than the compilers. Among the things which are currently found are unreachable statements, loops not entered at the top, automatic variables declared and not used, and logical expressions whose value is constant. Moreover, the usage of functions is checked to find functions which return values in some places and not in others, functions called with varying numbers of arguments, and functions whose values are not used.

By default, it is assumed that all the *files* are to be loaded together; they are checked for mutual compatibility. Function definitions for certain libraries are available to *lint*; these libraries are referred to by a conventional name, such as '-lm', in the style of *ld*(1). Arguments ending in *.ln* are also treated as library files. To create lint libraries, use the -C option:

```
lint -Cfoo files . . .
```

where *files* are the C sources of library *foo*. The result is a file *llib-foo.ln* in the correct library format suitable for linting programs using *foo*.

Any number of the options in the following list may be used. The -D, -U, and -I options of *cc*(1) are also recognized as separate arguments.

- p** Attempt to check portability to the *IBM* and *GCOS* dialects of C.
- h** Apply a number of heuristic tests to attempt to intuit bugs, improve style, and reduce waste.
- b** Report *break* statements that cannot be reached. (This is not the default because, unfortunately, most *lex* and many *yacc* outputs produce dozens of such comments.)
- v** Suppress complaints about unused arguments in functions.
- x** Report variables referred to by extern declarations, but never used.
- a** Report assignments of long values to int variables.
- c** Complain about casts which have questionable portability.
- u** Do not complain about functions and variables used and not defined, or defined and not used (this is suitable for running *lint* on a subset of files out of a larger program).
- n** Do not check compatibility against the standard library.
- z** Do not complain about structures that are never defined (e.g. using a structure pointer without knowing its contents.).

Exit(2) and other functions which do not return are not understood; this causes various lies.

Certain conventional comments in the C source will change the behavior of *lint*:

```
/*NOTREACHED*/
```

at appropriate points stops comments about unreachable code.

```
/*VARARGSn*/
```

suppresses the usual checking for variable numbers of arguments in the following function declaration. The data types of the first *n* arguments are checked; a missing *n* is taken to be 0.

```
/*NOSTRICT*/
```

shuts off strict type checking in the next expression.

/*ARGSUSED*/

turns on the **-v** option for the next function.

/*LINTLIBRARY*/

at the beginning of a file shuts off complaints about unused functions in this file.

AUTHOR

S.C. Johnson. Lint library construction implemented by Edward Wang.

FILES

<code>/usr/lib/lint/lint[12]</code>	programs
<code>/usr/lib/lint/lilib-ic.in</code>	declarations for standard functions
<code>/usr/lib/lint/lilib-ic</code>	human readable version of above
<code>/usr/lib/lint/lilib-port.in</code>	declarations for portable functions
<code>/usr/lib/lint/lilib-port</code>	human readable . . .
<code>lilib-l*.in</code>	library created with -C

SEE ALSO

`cc(1)`

S. C. Johnson, *Lint, a C Program Checker*

BUGS

There are some things you just **can't** get lint to shut up about.

NAME

ln — make links

SYNOPSIS

ln [*-s*] *name1* [*name2*]
ln *name* ... *directory*

DESCRIPTION

A link is a directory entry referring to a file; the same file (together with its size, all its protection information, etc.) may have several links to it. There are two kinds of links: hard links and symbolic links.

By default *ln* makes hard links. A hard link to a file is indistinguishable from the original directory entry; any changes to a file are effective independent of the name used to reference the file. Hard links may not span file systems and may not refer to directories.

The *-s* option causes *ln* to create symbolic links. A symbolic link contains the name of the file to which it is linked. The referenced file is used when an *open(2)* operation is performed on the link. A *stat(2)* on a symbolic link will return the linked-to file; an *lstat(2)* must be done to obtain information about the link. The *readlink(2)* call may be used to read the contents of a symbolic link. Symbolic links may span file systems and may refer to directories.

Given one or two arguments, *ln* creates a link to an existing file *name1*. If *name2* is given, the link has that name; *name2* may also be a directory in which to place the link; otherwise it is placed in the current directory. If only the directory is specified, the link will be made to the last component of *name1*.

Given more than two arguments, *ln* makes links to all the named files in the named directory. The links made will have the same name as the files being linked to.

SEE ALSO

rm(1), *cp(1)*, *mv(1)*, *link(2)*, *readlink(2)*, *stat(2)*, *symlink(2)*

NAME

login — sign on

SYNOPSIS

login [username]

DESCRIPTION

The *login* command is used when a user initially signs on, or it may be used at any time to change from one user to another. The latter case is the one summarized above and described here. See “How to Get Started” for how to dial up initially.

If *login* is invoked without an argument, it asks for a user name, and, if appropriate, a password. Echoing is turned off (if possible) during the typing of the password, so it will not appear on the written record of the session.

After a successful login, accounting files are updated and the user is informed of the existence of mail, and the message of the day is printed, as is the time he last logged in (unless he has a “.hushlogin” file in his home directory — this is mostly used to make life easier for non-human users, such as *uucp*).

Login initializes the user and group IDs and the working directory, then executes a command interpreter (usually *sh*(1)) according to specifications found in a password file. Argument 0 of the command interpreter is “-sh”, or more generally the name of the command interpreter with a leading dash (“-”) prepended.

Login also initializes the environment *environ*(7) with information specifying home directory, command interpreter, terminal type (if available) and user name.

If the file */etc/nologin* exists *login* prints its contents on the user's terminal and exits. This is used by *shutdown*(8) to stop users logging in when the system is about to go down.

Login is recognized by *sh*(1) and *csh*(1) and executed directly (without forking).

FILES

<i>/etc/utmp</i>	accounting
<i>/usr/adm/wtmp</i>	accounting
<i>/usr/spool/mail/*</i>	mail
<i>/etc/motd</i>	message-of-the-day
<i>/etc/passwd</i>	password file
<i>/etc/nologin</i>	stops logins
<i>.hushlogin</i>	makes login quieter
<i>/etc/securetty</i>	lists ttys that root may log in on

SEE ALSO

init(8), *getty*(8), *mail*(1), *passwd*(1), *passwd*(5), *environ*(7), *shutdown*(8)

DIAGNOSTICS

“Login incorrect,” if the name or the password is bad.

“No Shell”, “cannot open password file”, “no directory”: consult a programming counselor.

BUGS

An undocumented option, *-r* is used by the remote login server, *rlogind*(8C) to force *login* to enter into an initial connection protocol.

NAME

look — find lines in a sorted list

SYNOPSIS

look [**-df**] *string* [*file*]

DESCRIPTION

Look consults a sorted *file* and prints all lines that begin with *string*. It uses binary search.

The options **d** and **f** affect comparisons as in *sort*(1):

d 'Dictionary' order: only letters, digits, tabs and blanks participate in comparisons.

f Fold. Upper case letters compare equal to lower case.

If no *file* is specified, */usr/dict/words* is assumed with collating sequence **-df**.

FILES

/usr/dict/words

SEE ALSO

sort(1), *grep*(1)

NAME

lorder — find ordering relation for an object library

SYNOPSIS

lorder file ...

DESCRIPTION

The input is one or more object or library archive (see *ar(1)*) files. The standard output is a list of pairs of object file names, meaning that the first file of the pair refers to external identifiers defined in the second. The output may be processed by *tsort(1)* to find an ordering of a library suitable for one-pass access by *ld(1)*.

This brash one-liner intends to build a new library from existing '.o' files.

```
ar cr library `lorder *.o | tsort`
```

The need for *lorder* may be vitiated by use of *ranlib(1)*, which converts an ordered archive into a randomly accessed library.

FILES

*symref, *symdef
nm(1), sed(1), sort(1), join(1)

SEE ALSO

tsort(1), *ld(1)*, *ar(1)*, *ranlib(1)*

BUGS

The names of object files, in and out of libraries, must end with '.o'; nonsense results otherwise.

NAME

lpq - spool queue examination program

SYNOPSIS

lpq [+[*n*]] [-l] [-Pprinter] [job # ...] [user ...]

DESCRIPTION

lpq examines the spooling area used by *lpd*(8) for printing files on the line printer, and reports the status of the specified jobs or all jobs associated with a user. *lpq* invoked without any arguments reports on any jobs currently in the queue. A **-P** flag may be used to specify a particular printer, otherwise the default line printer is used (or the value of the **PRINTER** variable in the environment). If a **+** argument is supplied, *lpq* displays the spool queue until it empties. Supplying a number immediately after the **+** sign indicates that *lpq* should sleep *n* seconds in between scans of the queue. All other arguments supplied are interpreted as user names or job numbers to filter out only those jobs of interest.

For each job submitted (i.e. invocation of *lpr*(1)) *lpq* reports the user's name, current rank in the queue, the names of files comprising the job, the job identifier (a number which may be supplied to *lprm*(1) for removing a specific job), and the total size in bytes. The **-l** option causes information about each of the files comprising the job to be printed. Normally, only as much information as will fit on one line is displayed. Job ordering is dependent on the algorithm used to scan the spooling directory and is supposed to be FIFO (First in First Out). File names comprising a job may be unavailable (when *lpr*(1) is used as a sink in a pipeline) in which case the file is indicated as "(standard input)".

If *lpq* warns that there is no daemon present (i.e. due to some malfunction), the *lpc*(8) command can be used to restart the printer daemon.

FILES

<i>/etc/termcap</i>	for manipulating the screen for repeated display
<i>/etc/printcap</i>	to determine printer characteristics
<i>/usr/spool/*</i>	the spooling directory, as determined from <i>printcap</i>
<i>/usr/spool/*/cf*</i>	control files specifying jobs
<i>/usr/spool/*/lock</i>	the lock file to obtain the currently active job

SEE ALSO

lpr(1), *lprm*(1), *lpc*(8), *lpd*(8)

BUGS

Due to the dynamic nature of the information in the spooling directory *lpq* may report unreliably. Output formatting is sensitive to the line length of the terminal; this can result in widely spaced columns.

DIAGNOSTICS

Unable to open various files. The lock file being malformed. Garbage files when there is no daemon active, but files in the spooling directory.

NAME

`lpr` — off line print

SYNOPSIS

```
lpr [ -Pprinter ] [ -#num ] [ -C class ] [ -J job ] [ -T title ] [ -l [ numcols ] ] [ -1234 font ] [ -wnum ] [ -pltdgvcfrmhs ] [ name ... ]
```

DESCRIPTION

`lpr` uses a spooling daemon to print the named files when facilities become available. If no names appear, the standard input is assumed. The `-P` option may be used to force output to a specific printer. Normally, the default printer is used (site dependent), or the value of the environment variable `PRINTER` is used.

The following single letter options are used to notify the line printer spooler that the files are not standard text files. The spooling daemon will use the appropriate filters to print the data accordingly.

- `-p` Use `pr(1)` to format the files (equivalent to `print`).
- `-l` Use a filter which allows control characters to be printed and suppresses page breaks.
- `-t` The files are assumed to contain data from `troff(1)` (cat phototypesetter commands).
- `-n` The files are assumed to contain data from `ditroff` (device independent troff).
- `-d` The files are assumed to contain data from `tex(1)` (DVI format from Stanford).
- `-g` The files are assumed to contain standard plot data as produced by the `plot(3X)` routines (see also `plot(1G)` for the filters used by the printer spooler).
- `-v` The files are assumed to contain a raster image for devices like the Benson Varian.
- `-c` The files are assumed to contain data produced by `cifplot(1)`.
- `-f` Use a filter which interprets the first character of each line as a standard FORTRAN carriage control character.

The remaining single letter options have the following meaning.

- `-r` Remove the file upon completion of spooling or upon completion of printing (with the `-s` option).
- `-m` Send mail upon completion.
- `-h` Suppress the printing of the burst page.
- `-s` Use symbolic links. Usually files are copied to the spool directory.

The `-C` option takes the following argument as a job classification for use on the burst page. For example,

```
lpr -C EECS foo.c
```

causes the system name (the name returned by `hostname(1)`) to be replaced on the burst page by `EECS`, and the file `foo.c` to be printed.

The `-J` option takes the following argument as the job name to print on the burst page. Normally, the first file's name is used.

The `-T` option uses the next argument as the title used by `pr(1)` instead of the file name.

To get multiple copies of output, use the `-#num` option, where `num` is the number of copies desired of each file named. For example,

```
lpr -#3 foo.c bar.c more.c
```

would result in 3 copies of the file `foo.c`, followed by 3 copies of the file `bar.c`, etc. On the other hand,

```
cat foo.c bar.c more.c | lpr -#3
```

will give three copies of the concatenation of the files.

The `-i` option causes the output to be indented. If the next argument is numeric, it is used as the number of blanks to be printed before each line; otherwise, 8 characters are printed.

The `-w` option takes the immediately following number to be the page width for `pr`.

The `-s` option will use `symlink(2)` to link data files rather than trying to copy them so large files can be printed. This means the files should not be modified or removed until they have been printed.

The option `-1234` Specifies a font to be mounted on font position `i`. The daemon will construct a `.railmag` file referencing `/usr/lib/vfont/name.size`.

FILES

<code>/etc/passwd</code>	personal identification
<code>/etc/printcap</code>	printer capabilities data base
<code>/usr/lib/lpd*</code>	line printer daemons
<code>/usr/spool/*</code>	directories used for spooling
<code>/usr/spool/*/cf*</code>	daemon control files
<code>/usr/spool/*/df*</code>	data files specified in "cf" files
<code>/usr/spool/*/tf*</code>	temporary copies of "cf" files

SEE ALSO

`lpq(1)`, `lprm(1)`, `pr(1)`, `symlink(2)`, `printcap(5)`, `lpc(8)`, `lpd(8)`

DIAGNOSTICS

If you try to spool too large a file, it will be truncated. `Lpr` will object to printing binary files. If a user other than root prints a file and spooling is disabled, `lpr` will print a message saying so and will not put jobs in the queue. If a connection to `lpd` on the local machine cannot be made, `lpr` will say that the daemon cannot be started. Diagnostics may be printed in the daemon's log file regarding missing spool files by `lpd`.

BUGS

Fonts for `troff` and `tex` reside on the host with the printer. It is currently not possible to use local font libraries.

NAME

lprm — remove jobs from the line printer spooling queue

SYNOPSIS

lprm [**-P***printer*] [**-**] [job # ...] [user ...]

DESCRIPTION

Lprm will remove a job, or jobs, from a printer's spool queue. Since the spooling directory is protected from users, using *lprm* is normally the only method by which a user may remove a job.

Lprm without any arguments will delete the currently active job if it is owned by the user who invoked *lprm*.

If the **-** flag is specified, *lprm* will remove all jobs which a user owns. If the super-user employs this flag, the spool queue will be emptied entirely. The owner is determined by the user's login name and host name on the machine where the *lpr* command was invoked.

Specifying a user's name, or list of user names, will cause *lprm* to attempt to remove any jobs queued belonging to that user (or users). This form of invoking *lprm* is useful only to the super-user.

A user may dequeue an individual job by specifying its job number. This number may be obtained from the *lpq*(1) program, e.g.

```
% lpq -l
```

```
1st: ken                [job #013ucbarpa]
      (standard input)  100 bytes
```

```
% lprm 13
```

Lprm will announce the names of any files it removes and is silent if there are no jobs in the queue which match the request list.

Lprm will kill off an active daemon, if necessary, before removing any spooling files. If a daemon is killed, a new one is automatically restarted upon completion of file removals.

The **-P** option may be used to specify the queue associated with a specific printer (otherwise the default printer, or the value of the **PRINTER** variable in the environment is used).

FILES

```
/etc/printcap      printer characteristics file
/usr/spool/*       spooling directories
/usr/spool/*/lock  lock file used to obtain the pid of the current
                  daemon and the job number of the currently active job
```

SEE ALSO

lpr(1), *lpq*(1), *lpd*(8)

DIAGNOSTICS

"Permission denied" if the user tries to remove files other than his own.

BUGS

Since there are race conditions possible in the update of the lock file, the currently active job may be incorrectly identified.

NAME

ls — list contents of directory

SYNOPSIS

ls [**-acdfgilqrstu1ACLFR**] name ...

DESCRIPTION

For each directory argument, *ls* lists the contents of the directory; for each file argument, *ls* repeats its name and any other information requested. By default, the output is sorted alphabetically. When no argument is given, the current directory is listed. When several arguments are given, the arguments are first sorted appropriately, but file arguments are processed before directories and their contents.

There are a large number of options:

- l** List in long format, giving mode, number of links, owner, size in bytes, and time of last modification for each file. (See below.) If the file is a special file the size field will instead contain the major and minor device numbers. If the file is a symbolic link the pathname of the linked-to file is printed preceded by “->”.
- g** Include the group ownership of the file in a long output.
- t** Sort by time modified (latest first) instead of by name.
- a** List all entries; in the absence of this option, entries whose names begin with a period (.) are *not* listed.
- s** Give size in kilobytes of each file.
- d** If argument is a directory, list only its name; often used with **-l** to get the status of a directory.
- L** If argument is a symbolic link, list the file or directory the link references rather than the link itself.
- r** Reverse the order of sort to get reverse alphabetic or oldest first as appropriate.
- u** Use time of last access instead of last modification for sorting (with the **-t** option) and/or printing (with the **-l** option).
- c** Use time of file creation for sorting or printing.
- i** For each file, print the i-number in the first column of the report.
- f** Force each argument to be interpreted as a directory and list the name found in each slot. This option turns off **-l**, **-t**, **-s**, and **-r**, and turns on **-a**; the order is the order in which entries appear in the directory.
- F** cause directories to be marked with a trailing ‘/’, sockets with a trailing ‘=’, symbolic links with a trailing ‘@’, and executable files with a trailing ‘*’.
- R** recursively list subdirectories encountered.
- 1** force one entry per line output format; this is the default when output is not to a terminal.
- C** force multi-column output; this is the default when output is to a terminal.
- q** force printing of non-graphic characters in file names as the character ‘?’; this is the default when output is to a terminal.

The mode printed under the **-l** option contains 11 characters which are interpreted as follows: the first character is

- d** if the entry is a directory;
- b** if the entry is a block-type special file;

- c** if the entry is a character-type special file;
- l** if the entry is a symbolic link;
- s** if the entry is a socket, or
- if the entry is a plain file.

The next 9 characters are interpreted as three sets of three bits each. The first set refers to owner permissions; the next to permissions to others in the same user-group; and the last to all others. Within each set the three characters indicate permission respectively to read, to write, or to execute the file as a program. For a directory, 'execute' permission is interpreted to mean permission to search the directory. The permissions are indicated as follows:

- r** if the file is readable;
- w** if the file is writable;
- x** if the file is executable;
- if the indicated permission is not granted.

The group-execute permission character is given as **s** if the file has the set-group-id bit set; likewise the user-execute permission character is given as **s** if the file has the set-user-id bit set.

The last character of the mode (normally 'x' or '-') is **t** if the 1000 bit of the mode is on. See *chmod(1)* for the meaning of this mode.

When the sizes of the files in a directory are listed, a total count of blocks, including indirect blocks is printed.

FILES

/etc/passwd to get user id's for 'ls -l'.
/etc/group to get group id's for 'ls -g'.

BUGS

Newline and tab are considered printing characters in file names.

The output device is assumed to be 80 columns wide.

The option setting based on whether the output is a teletype is undesirable as "ls -s" is much different than "ls -s | lpr". On the other hand, not doing this setting would make old shell scripts which used *ls* almost certain losers.

NAME

mail - send or receive mail among users

SYNOPSIS

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

DESCRIPTION

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

newline

Go on to next message.

d Delete message and go on to the next.

p Print message again.

- Go back to previous message.

s [*file*] ...

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

w [*file*] ...

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

m [*person*] ...

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

EOT (control-D)

Put unexamined mail back in the mailbox and stop.

q Same as EOT.

!command

Escape to the Shell to do command.

* Print a command summary.

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

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

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

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

FILES

/etc/passwd	to identify sender and locate persons
/u0/spool/mail	post office for incoming mail
mbox	saved mail
/tmp/ma*	temp file
/usr/mail/*.lock	lock for mail directory
dead.letter	unmailable text

SEE ALSO

write(1)

BUGS

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

Normally anybody can read your mail. An installation can overcome this by making *mail* a set-user-id command that owns the mail directory.

NAME

make - maintain program groups

SYNOPSIS

make [-f *makefile*] [option] ... file ...

DESCRIPTION

Make executes commands in *makefile* to update one or more target *names*. *Name* is typically a program. If no -f option is present, 'makefile' and 'Makefile' are tried in order. If *makefile* is '-', the standard input is taken. More than one -f option may appear

Make updates a target if it depends on prerequisite files that have been modified since the target was last modified, if the target does not exist, or if the keyword **ALWAYS** is specified in the dependency list.

Makefile contains a sequence of entries that specify dependencies. The first line of an entry is a blank-separated list of targets, then a colon, then a list of prerequisite files. Text following a semicolon, and all following lines that begin with a tab, are shell commands to be executed to update the target. If a name appears on the left of more than one 'colon' line, then it depends on all of the names on the right of the colon on those lines, but only one command sequence may be specified for it. If a name appears on a line with a double colon :: then the command sequence following that line is performed only if the name is out of date with respect to the names to the right of the double colon, and is not affected by other double colon lines on which that name may appear.

Two special forms of a name are recognized. A name like *a(b)* means the file named *b* stored in the archive named *a*. A name like *a((b))* means the file stored in archive *a* containing the entry point *b*.

Sharp and newline surround comments.

The following makefile says that 'pgm' depends on two files 'a.o' and 'b.o', and that they in turn depend on '.c' files and a common file 'incl'.

```
pgm: a.o b.o
    cc a.o b.o - lm - o pgm
a.o: incl a.c
    cc - c a.c
b.o: incl b.c
    cc - c b.c
```

Makefile entries of the form

```
string1 = string2
```

are macro definitions. Subsequent appearances of $\$(string1)$ or $\${string1}$ are replaced by *string2*. If *string1* is a single character, the parentheses or braces are optional.

Make infers prerequisites for files for which *makefile* gives no construction commands. For example, a '.c' file may be inferred as prerequisite for a '.o' file and be compiled to produce the '.o' file. Thus the preceding example can be done more briefly:

```
pgm: a.o b.o
    cc a.o b.o - lm - o pgm
a.o b.o: incl
```

Prerequisites are inferred according to selected suffixes listed as the 'prerequisites' for the special name '.SUFFIXES'; multiple lists accumulate; an empty list clears what came before. Order is significant; the first possible name for which both a file and a rule as described in the next paragraph exist is inferred. The default list is

```
.SUFFIXES: .out .o .c .e .r .f .y .l .s .p
```

The rule to create a file with suffix *s2* that depends on a similarly named file with suffix *s1* is specified as an entry for the 'target' *s1s2*. In such an entry, the special macro *\$** stands for the target name with suffix deleted, *\$@* for the full target name, *\$<* for the complete list of prerequisites, and *\$?* for the list of prerequisites that are out of date. For example, a rule for making optimized '.o' files from '.c' files is

```
.c.o: ; cc - c - O - o $@ $*.c
```

Certain macros are used by the default inference rules to communicate optional arguments to any resulting compilations. In particular, 'CFLAGS' is used for *cc(1)* options, 'FFLAGS' for *f77(1)* options, 'PFLAGS' for *pc(1)* options, and 'LFLAGS' and 'YFLAGS' for *lex* and *yacc(1)* options. In addition, the macro 'MFLAGS' is filled in with the initial command line options supplied to *make*. This simplifies maintaining a hierarchy of makefiles as one may then invoke *make* on makefiles in subdirectories and pass along useful options such as *-k*.

Command lines are executed one at a time, each by its own shell. A line is printed when it is executed unless the special target '.SILENT' is in *makefile*, or the first character of the command is '@'.

Commands returning nonzero status (see *intro(1)*) cause *make* to terminate unless the special target '.IGNORE' is in *makefile* or the command begins with <tab><hyphen>.

Interrupt and quit cause the target to be deleted unless the target is a directory or depends on the special name '.PRECIOUS'.

Other options:

- **i** Equivalent to the special entry '.IGNORE:'.
- **k** When a command returns nonzero status, abandon work on the current entry, but continue on branches that do not depend on the current entry.
- **n** Trace and print, but do not execute the commands needed to update the targets.
- **t** Touch, i.e. update the modified date of targets, without executing any commands.
- **r** Equivalent to an initial special entry '.SUFFIXES:' with no list.
- **s** Equivalent to the special entry '.SILENT:'.

FILES

makefile, Makefile

SEE ALSO

sh(1), *touch(1)*, *f77(1)*, *pc(1)*

S. I. Feldman *Make - A Program for Maintaining Computer Programs*

BUGS

Some commands return nonzero status inappropriately. Use *-i* to overcome the difficulty.

Commands that are directly executed by the shell, notably *cd(1)*, are ineffectual across newlines in *make*.

NAME

man — find manual information by keywords; print out the manual

SYNOPSIS

man **-k** keyword ...
man **-f** file ...
man [**-**] [**-t**] [section] title ...

DESCRIPTION

Man is a program which gives information from the programmers manual. It can be asked for one line descriptions of commands specified by name, or for all commands whose description contains any of a set of keywords. It can also provide on-line access to the sections of the printed manual.

When given the option **-k** and a set of keywords, *man* prints out a one line synopsis of each manual sections whose listing in the table of contents contains that keyword.

When given the option **-f** and a list of file names, *man* attempts to locate manual sections related to those files, printing out the table of contents lines for those sections.

When neither **-k** nor **-f** is specified, *man* formats a specified set of manual pages. If a section specifier is given *man* looks in the that section of the manual for the given *titles*. *Section* is an Arabic section number (3 for instance). The number may followed by a single letter classifier (**lg** for instance) indicating a graphics program in section 1. If *section* is omitted, *man* searches all sections of the manual, giving preference to commands over subroutines in system libraries, and printing the first section it finds, if any.

If the standard output is a teletype, or if the flag **-** is given, *man* pipes its output through *cat*(1) with the option **-s** to crush out useless blank lines, *ul*(1) to create proper underlines for different terminals, and through *more*(1) to stop after each page on the screen. Hit a space to continue, a control-D to scroll 11 more lines when the output stops.

The **-t** flag causes *man* to arrange for the specified section to be *troff*ed to a suitable raster output device; see *vtroff*(1).

FILES

/usr/man/man?/*
/usr/man/cat?/*

SEE ALSO

more(1), *ul*(1), *whereis*(1), *catman*(8)

BUGS

The manual is supposed to be reproducible either on the phototypesetter or on a typewriter. However, on a typewriter some information is necessarily lost.

NAME

mesg — permit or deny messages

SYNOPSIS

mesg [n] [y]

DESCRIPTION

Mesg with argument *n* forbids messages via *write* and *talk*(1) by revoking non-user write permission on the user's terminal. *Mesg* with argument *y* reinstates permission. All by itself, *mesg* reports the current state without changing it.

FILES

/dev/tty*

SEE ALSO

write(1), talk(1)

DIAGNOSTICS

Exit status is 0 if messages are receivable, 1 if not, 2 on error.

NAME

mkdir — make a directory

SYNOPSIS

mkdir dirname ...

DESCRIPTION

Mkdir creates specified directories in mode 777. Standard entries, `.`, for the directory itself, and `..`, for its parent, are made automatically.

Mkdir requires write permission in the parent directory.

SEE ALSO

`rm(1)`.

DIAGNOSTICS

Mkdir returns exit code 0 if all directories were successfully made; otherwise, it prints a diagnostic and returns non-zero.

NAME

more, *page* — file perusal filter for crt viewing

SYNOPSIS

more [**-cdfisu**] [**-n**] [**+linenumber**] [**+/pattern**] [name ...]

page *more options*

DESCRIPTION

More is a filter which allows examination of a continuous text one screenful at a time on a soft-copy terminal. It normally pauses after each screenful, printing --More-- at the bottom of the screen. If the user then types a carriage return, one more line is displayed. If the user hits a space, another screenful is displayed. Other possibilities are enumerated later.

The command line options are:

- n** An integer which is the size (in lines) of the window which *more* will use instead of the default.
- c** *More* will draw each page by beginning at the top of the screen and erasing each line just before it draws on it. This avoids scrolling the screen, making it easier to read while *more* is writing. This option will be ignored if the terminal does not have the ability to clear to the end of a line.
- d** *More* will prompt the user with the message "Hit space to continue, Rubout to abort" at the end of each screenful. This is useful if *more* is being used as a filter in some setting, such as a class, where many users may be unsophisticated.
- f** This causes *more* to count logical, rather than screen lines. That is, long lines are not folded. This option is recommended if *nroff* output is being piped through *ul*, since the latter may generate escape sequences. These escape sequences contain characters which would ordinarily occupy screen positions, but which do not print when they are sent to the terminal as part of an escape sequence. Thus *more* may think that lines are longer than they actually are, and fold lines erroneously.
- l** Do not treat ^L (form feed) specially. If this option is not given, *more* will pause after any line that contains a ^L , as if the end of a screenful had been reached. Also, if a file begins with a form feed, the screen will be cleared before the file is printed.
- s** Squeeze multiple blank lines from the output, producing only one blank line. Especially helpful when viewing *nroff* output, this option maximizes the useful information present on the screen.
- u** Normally, *more* will handle underlining such as produced by *nroff* in a manner appropriate to the particular terminal: if the terminal can perform underlining or has a stand-out mode, *more* will output appropriate escape sequences to enable underlining or stand-out mode for underlined information in the source file. The **-u** option suppresses this processing.

+ linenumber

Start up at *linenumber*.

+ /pattern

Start up two lines before the line containing the regular expression *pattern*.

If the program is invoked as *page*, then the screen is cleared before each screenful is printed (but only if a full screenful is being printed), and $k - 1$ rather than $k - 2$ lines are printed in each screenful, where k is the number of lines the terminal can display.

More looks in the file */etc/termcap* to determine terminal characteristics, and to determine the default window size. On a terminal capable of displaying 24 lines, the default window size is 22 lines.

More looks in the environment variable *MORE* to pre-set any flags desired. For example, if you prefer to view files using the *-c* mode of operation, the *cs*h command *setenv MORE -c* or the *sh* command sequence *MORE='-c' ; export MORE* would cause all invocations of *more*, including invocations by programs such as *man* and *msgs*, to use this mode. Normally, the user will place the command sequence which sets up the *MORE* environment variable in the *.cshrc* or *.profile* file.

If *more* is reading from a file, rather than a pipe, then a percentage is displayed along with the *--More--* prompt. This gives the fraction of the file (in characters, not lines) that has been read so far.

Other sequences which may be typed when *more* pauses, and their effects, are as follows (*i* is an optional integer argument, defaulting to 1) :

- i* <space>
display *i* more lines, (or another screenful if no argument is given)
- ^D display 11 more lines (a "scroll"). If *i* is given, then the scroll size is set to *i*.
- d same as ^D (control-D)
- iz* same as typing a space except that *i*, if present, becomes the new window size.
- is* skip *i* lines and print a screenful of lines
- if* skip *i* screenfuls and print a screenful of lines
- q or Q Exit from *more*.
- = Display the current line number.
- v Start up the editor *vi* at the current line.
- h Help command; give a description of all the *more* commands.
- i*/*expr* search for the *i*-th occurrence of the regular expression *expr*. If there are less than *i* occurrences of *expr*, and the input is a file (rather than a pipe), then the position in the file remains unchanged. Otherwise, a screenful is displayed, starting two lines before the place where the expression was found. The user's erase and kill characters may be used to edit the regular expression. Erasing back past the first column cancels the search command.
- in* search for the *i*-th occurrence of the last regular expression entered.
- ' (single quote) Go to the point from which the last search started. If no search has been performed in the current file, this command goes back to the beginning of the file.
- !*command*
invoke a shell with *command*. The characters '%' and '!' in "*command*" are replaced with the current file name and the previous shell command respectively. If there is no current file name, '%' is not expanded. The sequences "\%" and "\!" are replaced by "%" and "!" respectively.
- i*:*n* skip to the *i*-th next file given in the command line (skips to last file if *n* doesn't make sense)
- i*:*p* skip to the *i*-th previous file given in the command line. If this command is given in the middle of printing out a file, then *more* goes back to the beginning of the file. If *i* doesn't make sense, *more* skips back to the first file. If *more* is not reading from a file, the bell is rung and nothing else happens.
- :*f* display the current file name and line number.

:q or :Q

exit from *more* (same as q or Q).

(dot) repeat the previous command.

The commands take effect immediately, i.e., it is not necessary to type a carriage return. Up to the time when the command character itself is given, the user may hit the line kill character to cancel the numerical argument being formed. In addition, the user may hit the erase character to redisplay the --More--(xx%) message.

At any time when output is being sent to the terminal, the user can hit the quit key (normally control- \backslash). *More* will stop sending output, and will display the usual --More-- prompt. The user may then enter one of the above commands in the normal manner. Unfortunately, some output is lost when this is done, due to the fact that any characters waiting in the terminal's output queue are flushed when the quit signal occurs.

The terminal is set to *noecho* mode by this program so that the output can be continuous. What you type will thus not show on your terminal, except for the / and ! commands.

If the standard output is not a teletype, then *more* acts just like *cat*, except that a header is printed before each file (if there is more than one).

A sample usage of *more* in previewing *nroff* output would be

```
nroff -ms +2 doc.n | more -s
```

AUTHOR

Eric Shienbrood, minor revisions by John Foderaro and Geoffrey Peck

FILES

/etc/termcap	Terminal data base
/usr/lib/more.help	Help file

SEE ALSO

csh(1), man(1), msgs(1), script(1), sh(1), environ(7)

NAME

mv — move or rename files

SYNOPSIS

mv [**-i**] [**-f**] [**-**] file1 file2

mv [**-i**] [**-f**] [**-**] file ... directory

DESCRIPTION

Mv moves (changes the name of) *file1* to *file2*.

If *file2* already exists, it is removed before *file1* is moved. If *file2* has a mode which forbids writing, *mv* prints the mode (see *chmod(2)*) and reads the standard input to obtain a line; if the line begins with *y*, the move takes place; if not, *mv* exits.

In the second form, one or more *files* (plain files or directories) are moved to the *directory* with their original file-names.

Mv refuses to move a file onto itself.

Options:

- i** stands for interactive mode. Whenever a move is to supercede an existing file, the user is prompted by the name of the file followed by a question mark. If he answers with a line starting with 'y', the move continues. Any other reply prevents the move from occurring.
- f** stands for force. This option overrides any mode restrictions or the **-i** switch.
- means interpret all the following arguments to *mv* as file names. This allows file names starting with minus.

SEE ALSO

cp(1), *ln(1)*

BUGS

If *file1* and *file2* lie on different file systems, *mv* must copy the file and delete the original. In this case the owner name becomes that of the copying process and any linking relationship with other files is lost.

NAME

newgrp — log in to a new group

SYNOPSIS

newgrp [group]

DESCRIPTION

Newgrp changes the group identification of its caller, analogously to *login*(1). The same person remains logged in, and the current directory is unchanged, but calculations of access permissions to files are performed with respect to the new group ID.

Newgrp without an argument changes the group identification to the group in the password file; in effect it changes the group identification back to the caller's original group.

A password is demanded if the group has a password and the user himself does not, or if the group has a password and the user is not listed in */etc/group* as being a member of that group.

When most users log in, they are members of the group named **other**.

FILES

/etc/group
/etc/passwd

SEE ALSO

login(1), *group*(5).

BUGS

There is no convenient way to enter a password into */etc/group*. Use of group passwords is not encouraged, because, by their very nature, they encourage poor security practices. Group passwords may disappear in the future.

NAME

nice, *nohup* — run a command at low priority (*sh* only)

SYNOPSIS

nice [*-number*] *command* [*arguments*]

nohup *command* [*arguments*]

DESCRIPTION

Nice executes *command* with low scheduling priority. If the *number* argument is present, the priority is incremented (higher numbers mean lower priorities) by that amount up to a limit of 20. The default *number* is 10.

The super-user may run commands with priority higher than normal by using a negative priority, e.g. '*--10*'.

Nohup executes *command* immune to hangup and terminate signals from the controlling terminal. The priority is incremented by 5. *Nohup* should be invoked from the shell with '&' in order to prevent it from responding to interrupts by or stealing the input from the next person who logs in on the same terminal. The syntax of *nice* is also different.

FILES

nohup.out standard output and standard error file under *nohup*

SEE ALSO

csh(1), *setpriority*(2), *renice*(8)

DIAGNOSTICS

Nice returns the exit status of the subject command.

BUGS

Nice and *nohup* are particular to *sh*(1). If you use *csh*(1), then commands executed with "&" are automatically immune to hangup signals while in the background. There is a builtin command *nohup* which provides immunity from terminate, but it does not redirect output to *nohup.out*.

Nice is built into *csh*(1) with a slightly different syntax than described here. The form "*nice +10*" nices to positive nice, and "*nice -10*" can be used by the super-user to give a process more of the processor.

NAME

nm — print name list

SYNOPSIS

nm [**-gnopru**] [file ...]

DESCRIPTION

Nm prints the name list (symbol table) of each object *file* in the argument list. If an argument is an archive, a listing for each object file in the archive will be produced. If no *file* is given, the symbols in "a.out" are listed.

Each symbol name is preceded by its value (blanks if undefined) and one of the letters **U** (undefined), **A** (absolute), **T** (text segment symbol), **D** (data segment symbol), **B** (bss segment symbol), **C** (common symbol), **f** file name, or **-** for sdb symbol table entries (see **-a** below). If the symbol is local (non-external) the type letter is in lower case. The output is sorted alphabetically.

Options are:

- g** Print only global (external) symbols.
- n** Sort numerically rather than alphabetically.
- o** Prepend file or archive element name to each output line rather than only once.
- p** Don't sort; print in symbol-table order.
- r** Sort in reverse order.
- u** Print only undefined symbols.

SEE ALSO

ar(1), ar(5), a.out(5), stab(5)

NAME

nroff — text formatting

SYNOPSIS

nroff [option] ... [file] ...

DESCRIPTION

Nroff formats text in the named *files* for typewriter-like devices. See also *troff(1)*. The full capabilities of *nroff* are described in the *Nroff/Troff User's Manual*.

If no *file* argument is present, the standard input is read. An argument consisting of a single minus (–) is taken to be a file name corresponding to the standard input.

The options, which may appear in any order so long as they appear *before* the files, are:

- olist** Print only pages whose page numbers appear in the comma-separated *list* of numbers and ranges. A range *N–M* means pages *N* through *M*; an initial *–N* means from the beginning to page *N*; and a final *N–* means from *N* to the end.
- nN** Number first generated page *N*.
- sN** Stop every *N* pages. *Nroff* will halt prior to every *N* pages (default *N*=1) to allow paper loading or changing, and will resume upon receipt of a newline.
- mname** Prepend the macro file */usr/lib/tmac/tmac.name* to the input *files*.
- raN** Set register *a* (one-character) to *N*.
- i** Read standard input after the input files are exhausted.
- q** Invoke the simultaneous input-output mode of the **rd** request.
- Tname** Prepare output for specified terminal. Known *names* are **37** for the (default) Teletype Corporation Model 37 terminal, **tn300** for the GE TermiNet 300 (or any terminal without half-line capability), **300S** for the DASI-300S, **300** for the DASI-300, and **450** for the DASI-450 (Diablo Hyterm).
- e** Produce equally-spaced words in adjusted lines, using full terminal resolution.
- h** Use output tabs during horizontal spacing to speed output and reduce output character count. Tab settings are assumed to be every 8 nominal character widths.

FILES

*/tmp/ta** temporary file
*/usr/lib/tmac/tmac.** standard macro files
*/usr/lib/term/** terminal driving tables for *nroff*

SEE ALSO

J. F. Ossanna, *Nroff/Troff user's manual*
 B. W. Kernighan, *A TROFF Tutorial*
troff(1), *eqn(1)*, *tbl(1)*, *ms(7)*, *me(7)*, *man(7)*, *col(1)*

NAME

od — octal, decimal, hex, ascii dump

SYNOPSIS

od [*-format*] [*file*] [[+]*offset*[.][*b*] [*label*]]

DESCRIPTION

Od displays *file*, or it's standard input, in one or more dump formats as selected by the first argument. If the first argument is missing, *-o* is the default. Dumping continues until end-of-file.

The meanings of the format argument characters are:

- a** Interpret bytes as characters and display them with their ASCII names. If the **p** character is given also, then bytes with even parity are underlined. The **P** character causes bytes with odd parity to be underlined. Otherwise the parity bit is ignored.
- b** Interpret bytes as unsigned octal.
- c** Interpret bytes as ASCII characters. Certain non-graphic characters appear as C escapes: null=`\0`, backspace=`\b`, formfeed=`\f`, newline=`\n`, return=`\r`, tab=`\t`; others appear as 3-digit octal numbers. Bytes with the parity bit set are displayed in octal.
- d** Interpret (short) words as unsigned decimal.
- f** Interpret long words as floating point.
- h** Interpret (short) words as unsigned hexadecimal.
- i** Interpret (short) words as signed decimal.
- l** Interpret long words as signed decimal.
- o** Interpret (short) words as unsigned octal.
- s**[*n*] Look for strings of ascii graphic characters, terminated with a null byte. *N* specifies the minimum length string to be recognized. By default, the minimum length is 3 characters.
- v** Show all data. By default, display lines that are identical to the last line shown are not output, but are indicated with an "*" in column 1.
- w**[*n*] Specifies the number of input bytes to be interpreted and displayed on each output line. If **w** is not specified, 16 bytes are read for each display line. If *n* is not specified, it defaults to 32.
- x** Interpret (short) words as hexadecimal.

An upper case format character implies the long or double precision form of the object.

The *offset* argument specifies the byte offset into the file where dumping is to commence. By default this argument is interpreted in octal. A different radix can be specified; If "." is appended to the argument, then *offset* is interpreted in decimal. If *offset* begins with "x" or "0x", it is interpreted in hexadecimal. If "b" ("B") is appended, the offset is interpreted as a block count, where a block is 512 (1024) bytes. If the *file* argument is omitted, an *offset* argument must be preceded by "+".

The radix of the displayed address will be the same as the radix of the *offset*, if specified; otherwise it will be octal.

Label will be interpreted as a pseudo-address for the first byte displayed. It will be shown in "()" following the file offset. It is intended to be used with core images to indicate the real memory address. The syntax for *label* is identical to that for *offset*.

SEE ALSO

adb(1)

BUGS

A file name argument can't start with "+". A hexadecimal offset can't be a block count. Only one file name argument can be given.

It is an historical botch to require specification of object, radix, and sign representation in a single character argument.

NAME

pagesize — print system page size

SYNOPSIS

pagesize

DESCRIPTION

Pagesize prints the size of a page of memory in bytes, as returned by *getpagesize(2)*. This program is useful in constructing portable shell scripts.

SEE ALSO

getpagesize(2)

NAME

passwd — change login password

SYNOPSIS

passwd [name]

DESCRIPTION

This command changes (or installs) a password associated with the user *name* (your own name by default).

The program prompts for the old password and then for the new one. The caller must supply both. The new password must be typed twice, to forestall mistakes.

New passwords must be at least four characters long if they use a sufficiently rich alphabet and at least six characters long if monospace. These rules are relaxed if you are insistent enough.

Only the owner of the name or the super-user may change a password; the owner must prove he knows the old password.

FILES

/etc/passwd

SEE ALSO

login(1), passwd(5), crypt(3)

Robert Morris and Ken Thompson, *UNIX password security*

BUGS

The password file information should be kept in a different data structure allowing indexed access; *dbm(3X)* would probably be suitable.

NAME

plot — graphics filters

SYNOPSIS

plot [-Tterminal [raster]]

DESCRIPTION

These commands read plotting instructions (see *plot(5)*) from the standard input, and in general produce plotting instructions suitable for a particular *terminal* on the standard output.

If no *terminal* type is specified, the environment parameter \$TERM (see *environ(7)*) is used. Known *terminals* are:

4014 Tektronix 4014 storage scope.

450 DASI Hyterm 450 terminal (Diablo mechanism).

300 DASI 300 or GSI terminal (Diablo mechanism).

300S DASI 300S terminal (Diablo mechanism).

ver Versatec D1200A printer-plotter. This version of *plot* places a scan-converted image in '/usr/tmp/raster' and sends the result directly to the plotter device rather than to the standard output. The optional argument causes a previously scan-converted file *raster* to be sent to the plotter.

FILES

/usr/bin/tek
/usr/bin/t450
/usr/bin/t300
/usr/bin/t300s
/usr/bin/vplot
/usr/tmp/raster

SEE ALSO

plot(3X), plot(5)

BUGS

There is no lockout protection for /usr/tmp/raster.

NAME

pr — print file

SYNOPSIS

pr [option] ... [file] ...

DESCRIPTION

Pr produces a printed listing of one or more *files*. The output is separated into pages headed by a date, the name of the file or a specified header, and the page number. If there are no file arguments, *pr* prints its standard input.

Options apply to all following files but may be reset between files:

- n** Produce *n*-column output.
- +n** Begin printing with page *n*.
- h** Take the next argument as a page header.
- wn** For purposes of multi-column output, take the width of the page to be *n* characters instead of the default 72.
- f** Use formfeeds instead of newlines to separate pages. A formfeed is assumed to use up two blank lines at the top of a page. (Thus this option does not affect the effective page length.)
- ln** Take the length of the page to be *n* lines instead of the default 66.
- t** Do not print the 5-line header or the 5-line trailer normally supplied for each page.
- sc** Separate columns by the single character *c* instead of by the appropriate amount of white space. A missing *c* is taken to be a tab.
- m** Print all *files* simultaneously, each in one column,

Inter-terminal messages via *write*(1) are forbidden during a *pr*.

FILES

/dev/tty? to suspend messages.

SEE ALSO

cat(1)

DIAGNOSTICS

There are no diagnostics when *pr* is printing on a terminal.

NAME

`printenv` — print out the environment

SYNOPSIS

`printenv` [*name*]

DESCRIPTION

Printenv prints out the values of the variables in the environment. If a *name* is specified, only its value is printed.

If a *name* is specified and it is not defined in the environment, *printenv* returns exit status 1, else it returns status 0.

SEE ALSO

`sh(1)`, `environ(7)`, `csh(1)`

NAME

prof — display profile data

SYNOPSIS

prof [**-a**] [**-l**] [**-n**] [**-z**] [**-s**] [**-v** [**-low** [**-high**]]] [a.out [mon.out ...]]

DESCRIPTION

Prof interprets the file produced by the *monitor* subroutine. Under default modes, the symbol table in the named object file (*a.out* default) is read and correlated with the profile file (*mon.out* default). For each external symbol, the percentage of time spent executing between that symbol and the next is printed (in decreasing order), together with the number of times that routine was called and the number of milliseconds per call. If more than one profile file is specified, the output represents the sum of the profiles.

In order for the number of calls to a routine to be tallied, the **-p** option of *cc*, *f77* or *pc* must have been given when the file containing the routine was compiled. This option also arranges for the profile file to be produced automatically.

Options are:

- a** all symbols are reported rather than just external symbols.
- l** the output is sorted by symbol value.
- n** the output is sorted by number of calls
- s** a summary profile file is produced in *mon.sum*. This is really only useful when more than one profile file is specified.
- v** all printing is suppressed and a graphic version of the profile is produced on the standard output for display by the *plot(1)* filters. When plotting, the numbers *low* and *high*, by default 0 and 100, may be given to cause a selected percentage of the profile to be plotted with accordingly higher resolution.
- z** routines which have zero usage (as indicated by call counts and accumulated time) are nevertheless printed in the output.

FILES

mon.out for profile
 a.out for namelist
 mon.sum for summary profile

SEE ALSO

monitor(3), *profil(2)*, *cc(1)*, *plot(1G)*

BUGS

Beware of quantization errors.

Is confused by *f77* which puts the entry points at the bottom of subroutines and functions.

NAME

ps — process status

SYNOPSIS

ps [**acegklstuvwx#**]

DESCRIPTION

Ps prints information about processes. Normally, only your processes are candidates to be printed by *ps*; specifying **a** causes other users processes to be candidates to be printed; specifying **x** includes processes without control terminals in the candidate pool.

All output formats include, for each process, the process id **PID**, control terminal of the process **TT**, cpu time used by the process **TIME** (this includes both user and system time), the state **STAT** of the process, and an indication of the **COMMAND** which is running. The state is given by a sequence of four letters, e.g. "RWNA". The first letter indicates the runnability of the process: **R** for runnable processes, **T** for stopped processes, **P** for processes in page wait, **D** for those in disk (or other short term) waits, **S** for those sleeping for less than about 20 seconds, and **I** for idle (sleeping longer than about 20 seconds) processes. The second letter indicates whether a process is swapped out, showing **W** if it is, or a blank if it is loaded (in-core); a process which has specified a soft limit on memory requirements and which is exceeding that limit shows **>**; such a process is (necessarily) not swapped. The third letter indicates whether a process is running with altered CPU scheduling priority (nice); if the process priority is reduced, an **N** is shown, if the process priority has been artificially raised then a '**<**' is shown; processes running without special treatment have just a blank. The final letter indicates any special treatment of the process for virtual memory replacement; the letters correspond to options to the *vadvise*(2) call; currently the possibilities are **A** standing for **VA_ANOM**, **S** for **VA_SEQL** and blank for **VA_NORM**; an **A** typically represents a *lisp*(1) in garbage collection, **S** is typical of large image processing programs which are using virtual memory to sequentially address voluminous data.

Here are the options:

- a** asks for information about all processes with terminals (ordinarily only one's own processes are displayed).
- c** prints the command name, as stored internally in the system for purposes of accounting, rather than the command arguments, which are kept in the process' address space. This is more reliable, if less informative, since the process is free to destroy the latter information.
- e** Asks for the environment to be printed as well as the arguments to the command.
- g** Asks for all processes. Without this option, *ps* only prints "interesting" processes. Processes are deemed to be uninteresting if they are process group leaders. This normally eliminates top-level command interpreters and processes waiting for users to login on free terminals.
- k** causes the file */vmcore* is used in place of */dev/kmem* and */dev/mem*. This is used for post-mortem system debugging.
- l** asks for a long listing, with fields **PPID**, **CP**, **PRI**, **NI**, **ADDR**, **SIZE**, **RSS** and **WCHAN** as described below.
- s** Adds the size **SSIZ** of the kernel stack of each process (for use by system maintainers) to the basic output format.
- tx** restricts output to processes whose controlling tty is *x* (which should be specified as printed by *ps*, e.g. *t3* for tty3, *tco* for console, *td0* for ttyd0, *t?* for processes with no tty, *t* for processes at the current tty, etc). This option must be the last one given.
- u** A user oriented output is produced. This includes fields **USER**, **%CPU**, **NICE**, **SIZE**, and

RSS as described below.

- v A version of the output containing virtual memory statistics is output. This includes fields RE, SL, PAGEIN, SIZE, RSS, LIM, TSIZ, TRS, %CPU and %MEM, described below.
- w Use a wide output format (132 columns rather than 80); if repeated, e.g. ww, use arbitrarily wide output. This information is used to decide how much of long commands to print.
- x asks even about processes with no terminal.
- # A process number may be given, (indicated here by #), in which case the output is restricted to that process. This option must also be last.

A second argument tells *ps* where to look for *core* if the *k* option is given, instead of */vmcore*. A third argument is the name of a swap file to use instead of the default */dev/drum*. If a fourth argument is given, it is taken to be the file containing the system's namelist. Otherwise, */vmunix* is used.

Fields which are not common to all output formats:

USER	name of the owner of the process
%CPU	cpu utilization of the process; this is a decaying average over up to a minute of previous (real) time. Since the time base over which this is computed varies (since processes may be very young) it is possible for the sum of all %CPU fields to exceed 100%.
NICE	(or NI) process scheduling increment (see <i>setpriority(2)</i>)
SIZE	virtual size of the process (in 1024 byte units)
RSS	real memory (resident set) size of the process (in 1024 byte units)
LIM	soft limit on memory used, specified via a call to <i>setrlimit(2)</i> ; if no limit has been specified then shown as <i>xx</i>
TSIZ	size of text (shared program) image
TRS	size of resident (real memory) set of text
%MEM	percentage of real memory used by this process.
RE	residency time of the process (seconds in core)
SL	sleep time of the process (seconds blocked)
PAGEIN	number of disk i/o's resulting from references by the process to pages not loaded in core.
UID	numerical user-id of process owner
PPID	numerical id of parent of process
CP	short-term cpu utilization factor (used in scheduling)
PRI	process priority (non-positive when in non-interruptible wait)
ADDR	swap address of the process
WCHAN	event on which process is waiting (an address in the system), with the initial part of the address trimmed off e.g. 80004000 prints as 4000.

F flags associated with process as in *<sys/proc.h>*:

SLOAD	000001	in core
SSYS	000002	swapper or pager process
SLOCK	000004	process being swapped out
SSWAP	000008	save area flag
STRC	000010	process is being traced
SWTED	000020	another tracing flag
SULOCK	000040	user settable lock in core
SPAGE	000080	process in page wait state
SKEEP	000100	another flag to prevent swap out

SDLYU	000200	delayed unlock of pages
SWEXIT	000400	working on exiting
SPHYSIO	000800	doing physical i/o (bio.c)
SVFORK	001000	process resulted from vfork()
SVFDONE	002000	another vfork flag
SNOVM	004000	no vm, parent in a vfork()
SPAGI	008000	init data space on demand from inode
SANOM	010000	system detected anomalous vm behavior
SUANOM	020000	user warned of anomalous vm behavior
STIMO	040000	timing out during sleep
SDETACH	080000	detached inherited by init
SOUSIG	100000	using old signal mechanism

A process that has exited and has a parent, but has not yet been waited for by the parent is marked <defunct>; a process which is blocked trying to exit is marked <exiting>; *Ps* makes an educated guess as to the file name and arguments given when the process was created by examining memory or the swap area. The method is inherently somewhat unreliable and in any event a process is entitled to destroy this information, so the names cannot be counted on too much.

FILES

/vmunix system namelist
/dev/kmem kernel memory
/dev/drum swap device
/vmcore core file
/dev searched to find swap device and tty names

SEE ALSO

kill(1), w(1)

BUGS

Things can change while *ps* is running; the picture it gives is only a close approximation to reality.

NAME

`pwd` — working directory name

SYNOPSIS

`pwd`

DESCRIPTION

Pwd prints the pathname of the working (current) directory.

SEE ALSO

`cd(1)`, `cd(1)`, `getwd(3)`

BUGS

In *cd(1)* the command *dirs* is always faster (although it can give a different answer in the rare case that the current directory or a containing directory was moved after the shell descended into it).

NAME

rev — reverse lines of a file

SYNOPSIS

rev [file] ...

DESCRIPTION

Rev copies the named files to the standard output, reversing the order of characters in every line. If no file is specified, the standard input is copied.

NAME

rlogin - remote login

SYNOPSIS

```
rlogin rhost [ - e c ] [ - l username ]  
rhost [ - e c ] [ - l username ]
```

DESCRIPTION

Rlogin connects your terminal on the current local host system *lhost* to the remote host system *rhost*.

Each host has a file */etc/hosts.equiv* which contains a list of *rhost*'s which which it shares account names. (The host names must be the standard names as described in *rsh(1c)* and printed by *login(1)*.) When you *rlogin* as the same user on an equivalent host, you don't need to give a password. Each user may also have a private equivalence list in a file *.rhosts* in his login directory. Each line in this file should contain a *rhost* and a *username* separated by a space, giving additional cases where logins without passwords are to be permitted. If the originating user is not equivalent to the remote user, then a login and password will be prompted for on the remote machine as in *login(1)*.

Your remote terminal type is the same as your local terminal type (as given in your environment *TERM* variable). All echoing takes place at the remote site, so that (except for delays) the *rlogin* is transparent. Flow control via *^S* and *^Q* and flushing of input and output on interrupts are handled properly. A line of the form "*~.*" disconnects from the remote host, where "*~*" is the escape character. A different escape character may be specified by the *- e* option.

SEE ALSO

rsh(1c), *rlogind(8c)*

FILES

*/usr/hosts/** for *rhost* version of the command

BUGS

More terminal characteristics should be propagated.

NAME

rm, rmdir - remove (unlink) files or directories

SYNOPSIS

rm [**-f**] [**-r**] [**-i**] [**-**] file ...

rmdir dir ...

DESCRIPTION

Rm removes the entries for one or more files from a directory. If an entry was the last link to the file, the file is destroyed. Removal of a file requires write permission in its directory, but neither read nor write permission on the file itself.

If a file has no write permission and the standard input is a terminal, its permissions are printed and a line is read from the standard input. If that line begins with 'y' the file is deleted, otherwise the file remains. No questions are asked and no errors are reported when the **-f** (force) option is given.

If a designated file is a directory, an error comment is printed unless the optional argument **-r** has been used. In that case, *rm* recursively deletes the entire contents of the specified directory, and the directory itself.

If the **-i** (interactive) option is in effect, *rm* asks whether to delete each file, and, under **-r**, whether to examine each directory.

The null option **-** indicates that all the arguments following it are to be treated as file names. This allows the specification of file names starting with a minus.

Rmdir removes entries for the named directories, which must be empty.

SEE ALSO

rm(1), unlink(2), rmdir(2)

NAME

`rmdel` - remove a delta from an SCCS file

SYNOPSIS

`rmdel` -rSID files

DESCRIPTION

Rmdel removes the delta specified by the *SID* from each named SCCS file. The delta to be removed must be the newest (most recent) delta in its branch in the delta chain of each named SCCS file. In addition, the specified must *not* be that of a version being edited for the purpose of making a delta (i. e., if a *p-file* (see *get(1)*) exists for the named SCCS file, the specified must *not* appear in any entry of the *p-file*).

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

The exact permissions necessary to remove a delta are documented in the *Source Code Control System User's Guide*. Simply stated, they are either (1) if you make a delta you can remove it; or (2) if you own the file and directory you can remove a delta.

FILES

x-file (see *delta(1)*)
z-file (see *delta(1)*)

SEE ALSO

delta(1), *get(1)*, *help(1)*, *prs(1)*, *sccsfile(5)*.
Source Code Control System User's Guide by L. E. Bonanni and C. A. Salemi.

DIAGNOSTICS

Use *help(1)* for explanations.

NAME

`rmdir`, `rm` - remove (unlink) directories or files

SYNOPSIS

`rmdir` *dir* ...

`rm` [`-f`] [`-r`] [`-i`] [`-`] *file* ...

DESCRIPTION

Rmdir removes entries for the named directories, which must be empty.

Rm removes the entries for one or more files from a directory. If an entry was the last link to the file, the file is destroyed. Removal of a file requires write permission in its directory, but neither read nor write permission on the file itself.

If a file has no write permission and the standard input is a terminal, its permissions are printed and a line is read from the standard input. If that line begins with 'y' the file is deleted, otherwise the file remains. No questions are asked and no errors are reported when the `-f` (force) option is given.

If a designated file is a directory, an error comment is printed unless the optional argument `-r` has been used. In that case, *rm* recursively deletes the entire contents of the specified directory, and the directory itself.

If the `-i` (interactive) option is in effect, *rm* asks whether to delete each file, and, under `-r`, whether to examine each directory.

The null option `-` indicates that all the arguments following it are to be treated as file names. This allows the specification of file names starting with a minus.

SEE ALSO

`rm(1)`, `unlink(2)`, `rmdir(2)`

NAME

rsh - remote shell

SYNOPSIS

```
rsh host [ - l username ] [ - n ] command
host [ - l username ] [ - n ] command
```

DESCRIPTION

Rsh connects to the specified *host*, and executes the specified *command*. *Rsh* copies its standard input to the remote command, the standard output of the remote command to its standard output, and the standard error of the remote command to its standard error. Interrupt, quit and terminate signals are propagated to the remote command; *rsh* normally terminates when the remote command does.

The remote username used is the same as your local username, unless you specify a different remote name with the *-l* option. This remote name must be equivalent (in the sense of *rlogin(1c)*) to the originating account; no provision is made for specifying a password with a command.

If you omit *command*, then instead of executing a single command, you will be logged in on the remote host using *rlogin(1c)*.

Shell metacharacters which are not quoted are interpreted on local machine, while quoted metacharacters are interpreted on the remote machine. Thus the command

```
rsh otherhost cat remotefile >> localfile
```

appends the remote file *remotefile* to the localfile *localfile*, while

```
rsh otherhost cat remotefile ">>" otherremotefile
```

appends *remotefile* to *otherremotefile*.

Host names are given in the file */etc/hosts*. Each host has one standard name (the first name given in the file), which is rather long and unambiguous, and optionally one or more nicknames. The host names for local machines are also commands in the directory */usr/hosts*; if you put this directory in your search path then the **rsh** can be omitted.

FILES

```
/etc/hosts
/usr/hosts/*
```

SEE ALSO

rlogin(1c), *rshd(8c)*

BUGS

If you are using *cs(1)* and put a *rsh(1c)* in the background without redirecting its input away from the terminal, it will block even if no reads are posted by the remote command. If no input is desired you should redirect the input of *rsh* to */dev/null* using the *-n* option.

You cannot run an interactive command (like *rogue(6)* or *vi(1)*); use *rlogin(1c)*.

Stop signals stop the local *rsh* process only; this is arguably wrong, but currently hard to fix for reasons too complicated to explain here.

NAME

ruptime — show host status of local machines

SYNOPSIS

ruptime [**-a**] [**-l**] [**-t**] [**-u**]

DESCRIPTION

Ruptime gives a status line like *uptime* for each machine on the local network; these are formed from packets broadcast by each host on the network once a minute.

Machines for which no status report has been received for 5 minutes are shown as being down.

Users idle an hour or more are not counted unless the **-a** flag is given.

Normally, the listing is sorted by host name. The **-l**, **-t**, and **-u** flags specify sorting by load average, uptime, and number of users, respectively.

FILES

/usr/spool/rwho/whod.* data files

SEE ALSO

rwho(1C)

NAME

`shownet` - show VALID node status

SYNOPSIS

`shownet`

DESCRIPTION

shownet displays the set of currently reachable VALID nodes on the local Ethernet. Also displayed is the set of nodes that have been reachable in the past but are no longer active on the net. This information can also be extracted from the *conn(8V)* show command. The advantage of the *shownet* program is that the display is denser and only shows reachability, since that is what most users need.

DIAGNOSTICS

`shownet`: failure reading node list (5, I/O error)

Usually caused by version mismatch between *shownet* and the kernel.

`shownet`: cannot open (13, Permission denied) /net

shownet was not installed by root with the set user and set group on execution permissions.

NAME

rwho — who's logged in on local machines

SYNOPSIS

rwho [**-a**]

DESCRIPTION

The *rwho* command produces output similar to *who*, but for all machines on the local network. If no report has been received from a machine for 5 minutes then *rwho* assumes the machine is down, and does not report users last known to be logged into that machine.

If a users hasn't typed to the system for a minute or more, then *rwho* reports this idle time. If a user hasn't typed to the system for an hour or more, then the user will be omitted from the output of *rwho* unless the **-a** flag is given.

FILES

/usr/spool/rwho/whod.* information about other machines

SEE ALSO

ruptime(1C), rwhod(8C)

BUGS

This is unwieldy when the number of machines on the local net is large.

NAME

sccshelp - ask for help with Source Code Control System

SYNOPSIS

sccshelp [args]

DESCRIPTION

Sccshelp finds information to explain a message from a command or to explain the use of a command. Zero or more arguments may be supplied. If no arguments are given, sccshelp will prompt for an argument.

The arguments may be either message numbers (which normally appear in parentheses following messages) or command names of one of the following types:

- type 1 Begins with non-numeric, ends in numerics. The non-numeric prefix is usually an abbreviation for the program or set of routines that produced the message (e.g., ge6, for message 6 from the get command).
- type 2 Does not contain numerics (as a command such as get).
- type 3 Is all numeric (e.g., 212).

The response of the program will be the explanatory information related to the argument, if there is any.

When all else fails, try "sccshelp stuck."

FILES

/usr/lib/sccshelp directory containing files of message text.

DIAGNOSTICS

Use sccshelp(1) for explanations.

NAME

script — make typescript of terminal session

SYNOPSIS

script [**-a**] [*file*]

DESCRIPTION

Script makes a typescript of everything printed on your terminal. The typescript is written to *file*, or appended to *file* if the **-a** option is given. It can be sent to the line printer later with *lpr*. If no file name is given, the typescript is saved in the file *typescript*.

The script ends when the forked shell exits.

This program is useful when using a crt and a hard-copy record of the dialog is desired, as for a student handing in a program that was developed on a crt when hard-copy terminals are in short supply.

BUGS

Script places **everything** in the log file. This is not what the naive user expects.

NAME

sed — stream editor

SYNOPSIS

sed [**-n**] [**-e** script] [**-f** sfile] [file] ...

DESCRIPTION

Sed copies the named *files* (standard input default) to the standard output, edited according to a script of commands. The **-f** option causes the script to be taken from file *sfile*; these options accumulate. If there is just one **-e** option and no **-f**'s, the flag **-e** may be omitted. The **-n** option suppresses the default output.

A script consists of editing commands, one per line, of the following form:

[address [, address]] function [arguments]

In normal operation *sed* cyclically copies a line of input into a *pattern space* (unless there is something left after a 'D' command), applies in sequence all commands whose *addresses* select that pattern space, and at the end of the script copies the pattern space to the standard output (except under **-n**) and deletes the pattern space.

An *address* is either a decimal number that counts input lines cumulatively across files, a '\$' that addresses the last line of input, or a context address, '/regular expression/', in the style of *ed*(1) modified thus:

The escape sequence '\n' matches a newline embedded in the pattern space.

A command line with no addresses selects every pattern space.

A command line with one address selects each pattern space that matches the address.

A command line with two addresses selects the inclusive range from the first pattern space that matches the first address through the next pattern space that matches the second. (If the second address is a number less than or equal to the line number first selected, only one line is selected.) Thereafter the process is repeated, looking again for the first address.

Editing commands can be applied only to non-selected pattern spaces by use of the negation function '!' (below).

In the following list of functions the maximum number of permissible addresses for each function is indicated in parentheses.

An argument denoted *text* consists of one or more lines, all but the last of which end with '\ ' to hide the newline. Backslashes in *text* are treated like backslashes in the replacement string of an 's' command, and may be used to protect initial blanks and tabs against the stripping that is done on every script line.

An argument denoted *rfile* or *wfile* must terminate the command line and must be preceded by exactly one blank. Each *wfile* is created before processing begins. There can be at most 10 distinct *wfile* arguments.

(1) a\
text

Append. Place *text* on the output before reading the next input line.

(2) b *label*

Branch to the ':' command bearing the *label*. If *label* is empty, branch to the end of the script.

(2) c\
text

Change. Delete the pattern space. With 0 or 1 address or at the end of a 2-address range, place *text* on the output. Start the next cycle.

- (2) d Delete the pattern space. Start the next cycle.
- (2) D Delete the initial segment of the pattern space through the first newline. Start the next cycle.
- (2) g Replace the contents of the pattern space by the contents of the hold space.
- (2) G Append the contents of the hold space to the pattern space.
- (2) h Replace the contents of the hold space by the contents of the pattern space.
- (2) H Append the contents of the pattern space to the hold space.
- (1) i\
text
Insert. Place *text* on the standard output.
- (2) n Copy the pattern space to the standard output. Replace the pattern space with the next line of input.
- (2) N Append the next line of input to the pattern space with an embedded newline. (The current line number changes.)
- (2) p Print. Copy the pattern space to the standard output.
- (2) P Copy the initial segment of the pattern space through the first newline to the standard output.
- (1) q Quit. Branch to the end of the script. Do not start a new cycle.
- (2) r *rfile*
Read the contents of *rfile*. Place them on the output before reading the next input line.
- (2) s/*regular expression*/*replacement*/*flags*
Substitute the *replacement* string for instances of the *regular expression* in the pattern space. Any character may be used instead of '/'. For a fuller description see *ed(1)*. *Flags* is zero or more of
 - g Global. Substitute for all nonoverlapping instances of the *regular expression* rather than just the first one.
 - p Print the pattern space if a replacement was made.
 - w *wfile* Write. Append the pattern space to *wfile* if a replacement was made.
- (2) t *label*
Test. Branch to the ':' command bearing the *label* if any substitutions have been made since the most recent reading of an input line or execution of a 't'. If *label* is empty, branch to the end of the script.
- (2) w *wfile*
Write. Append the pattern space to *wfile*.
- (2) x Exchange the contents of the pattern and hold spaces.
- (2) y/*string1*/*string2*/
Transform. Replace all occurrences of characters in *string1* with the corresponding character in *string2*. The lengths of *string1* and *string2* must be equal.
- (2)! *function*
Don't. Apply the *function* (or group, if *function* is '{') only to lines *not* selected by the address(es).
- (0) : *label*
This command does nothing; it bears a *label* for 'b' and 't' commands to branch to.
- (1) = Place the current line number on the standard output as a line.

- (2) { Execute the following commands through a matching '}' only when the pattern space is selected.
- (0) An empty command is ignored.

SEE ALSO

ed(1), grep(1), awk(1), lex(1)

NAME

sh, for, case, if, while, :, ., break, continue, cd, eval, exec, exit, export, login, read, readonly, set, shift, times, trap, umask, wait — command language

SYNOPSIS

sh [**-ceiknrstuvx**] [**arg**] ...

DESCRIPTION

Sh is a command programming language that executes commands read from a terminal or a file. See **invocation** for the meaning of arguments to the shell.

Commands.

A *simple-command* is a sequence of non blank *words* separated by blanks (a blank is a **tab** or a **space**). The first word specifies the name of the command to be executed. Except as specified below the remaining words are passed as arguments to the invoked command. The command name is passed as argument 0 (see *execve(2)*). The *value* of a simple-command is its exit status if it terminates normally or 200+*status* if it terminates abnormally (see *sigvec(2)* for a list of status values).

A *pipeline* is a sequence of one or more *commands* separated by |. The standard output of each command but the last is connected by a *pipe(2)* to the standard input of the next command. Each command is run as a separate process; the shell waits for the last command to terminate.

A *list* is a sequence of one or more *pipelines* separated by ;, &, && or || and optionally terminated by ; or &. ; and & have equal precedence which is lower than that of && and ||, && and || also have equal precedence. A semicolon causes sequential execution; an ampersand causes the preceding *pipeline* to be executed without waiting for it to finish. The symbol && (||) causes the *list* following to be executed only if the preceding *pipeline* returns a zero (non zero) value. Newlines may appear in a *list*, instead of semicolons, to delimit commands.

A *command* is either a simple-command or one of the following. The value returned by a command is that of the last simple-command executed in the command.

for name [in word ...] do list done

Each time a **for** command is executed *name* is set to the next word in the **for** word list. If **in word ...** is omitted, **in "\$@"** is assumed. Execution ends when there are no more words in the list.

case word in [pattern [| pattern] ...) list ;;] ... esac

A **case** command executes the *list* associated with the first pattern that matches *word*. The form of the patterns is the same as that used for file name generation.

if list then list [elif list then list] ... [else list] fi

The *list* following **if** is executed and if it returns zero the *list* following **then** is executed. Otherwise, the *list* following **elif** is executed and if its value is zero the *list* following **then** is executed. Failing that the **else list** is executed.

while list [do list] done

A **while** command repeatedly executes the **while list** and if its value is zero executes the **do list**; otherwise the loop terminates. The value returned by a **while** command is that of the last executed command in the **do list**. **until** may be used in place of **while** to negate the loop termination test.

(*list*) Execute *list* in a subshell.

{ *list* } *list* is simply executed.

The following words are only recognized as the first word of a command and when not quoted.

if then else elif fi case in esac for while until do done { }

Command substitution.

The standard output from a command enclosed in a pair of back quotes (``) may be used as part or all of a word; trailing newlines are removed.

Parameter substitution.

The character \$ is used to introduce substitutable parameters. Positional parameters may be assigned values by set. Variables may be set by writing

```
name=value [ name=value ] ...
```

\${parameter}

A *parameter* is a sequence of letters, digits or underscores (a *name*), a digit, or any of the characters * @ # ? - \$!. The value, if any, of the parameter is substituted. The braces are required only when *parameter* is followed by a letter, digit, or underscore that is not to be interpreted as part of its name. If *parameter* is a digit, it is a positional parameter. If *parameter* is * or @ then all the positional parameters, starting with \$1, are substituted separated by spaces. \$0 is set from argument zero when the shell is invoked.

\${parameter - word}

If *parameter* is set, substitute its value; otherwise substitute *word*.

\${parameter = word}

If *parameter* is not set, set it to *word*; the value of the parameter is then substituted. Positional parameters may not be assigned to in this way.

\${parameter ? word}

If *parameter* is set, substitute its value; otherwise, print *word* and exit from the shell. If *word* is omitted, a standard message is printed.

\${parameter + word}

If *parameter* is set, substitute *word*; otherwise substitute nothing.

In the above *word* is not evaluated unless it is to be used as the substituted string. (So that, for example, echo \${d-'pwd'} will only execute *pwd* if *d* is unset.)

The following *parameters* are automatically set by the shell.

#	The number of positional parameters in decimal.
-	Options supplied to the shell on invocation or by set.
?	The value returned by the last executed command in decimal.
\$	The process number of this shell.
!	The process number of the last background command invoked.

The following *parameters* are used but not set by the shell.

HOME	The default argument (home directory) for the <i>cd</i> command.
PATH	The search path for commands (see execution).
MAIL	If this variable is set to the name of a mail file, the shell informs the user of the arrival of mail in the specified file.
PS1	Primary prompt string, by default '\$ '.
PS2	Secondary prompt string, by default '> '.
IFS	Internal field separators, normally space , tab , and newline .

Blank interpretation.

After parameter and command substitution, any results of substitution are scanned for internal field separator characters (those found in \$IFS) and split into distinct arguments where such characters are found. Explicit null arguments (" or '') are retained. Implicit null arguments (those resulting from *parameters* that have no values) are removed.

File name generation.

Following substitution, each command word is scanned for the characters *, ? and [. If one of these characters appears, the word is regarded as a pattern. The word is replaced with alphabetically sorted file names that match the pattern. If no file name is found that matches the pattern, the word is left unchanged. The character . at the start of a file name or immediately following a /, and the character /, must be matched explicitly.

- * Matches any string, including the null string.
- ? Matches any single character.
- [...] Matches any one of the characters enclosed. A pair of characters separated by - matches any character lexically between the pair.

Quoting.

The following characters have a special meaning to the shell and cause termination of a word unless quoted.

; & () | < > newline space tab

A character may be *quoted* by preceding it with a \. \newline is ignored. All characters enclosed between a pair of quote marks (' '), except a single quote, are quoted. Inside double quotes (" ") parameter and command substitution occurs and \ quotes the characters \ ' " and \$.

"\$*" is equivalent to "\$1 \$2 ..." whereas

"\$@" is equivalent to "\$1" "\$2"

Prompting.

When used interactively, the shell prompts with the value of PS1 before reading a command. If at any time a newline is typed and further input is needed to complete a command, the secondary prompt (\$PS2) is issued.

Input output.

Before a command is executed its input and output may be redirected using a special notation interpreted by the shell. The following may appear anywhere in a simple-command or may precede or follow a *command* and are not passed on to the invoked command. Substitution occurs before *word* or *digit* is used.

< *word* Use file *word* as standard input (file descriptor 0).

> *word* Use file *word* as standard output (file descriptor 1). If the file does not exist, it is created; otherwise it is truncated to zero length.

>> *word*

Use file *word* as standard output. If the file exists, output is appended (by seeking to the end); otherwise the file is created.

<< *word*

The shell input is read up to a line the same as *word*, or end of file. The resulting document becomes the standard input. If any character of *word* is quoted, no interpretation is placed upon the characters of the document; otherwise, parameter and command substitution occurs, \newline is ignored, and \ is used to quote the characters \ \$ ' and the first character of *word*.

< & *digit*

The standard input is duplicated from file descriptor *digit*; see *dup(2)*. Similarly for the standard output using > .

< & - The standard input is closed. Similarly for the standard output using > .

If one of the above is preceded by a digit, the file descriptor created is that specified by the digit (instead of the default 0 or 1). For example,

```
... 2>&1
```

creates file descriptor 2 to be a duplicate of file descriptor 1.

If a command is followed by **&** then the default standard input for the command is the empty file (`/dev/null`). Otherwise, the environment for the execution of a command contains the file descriptors of the invoking shell as modified by input output specifications.

Environment.

The environment is a list of name-value pairs that is passed to an executed program in the same way as a normal argument list; see *execve(2)* and *environ(7)*. The shell interacts with the environment in several ways. On invocation, the shell scans the environment and creates a *parameter* for each name found, giving it the corresponding value. Executed commands inherit the same environment. If the user modifies the values of these *parameters* or creates new ones, none of these affects the environment unless the **export** command is used to bind the shell's *parameter* to the environment. The environment seen by any executed command is thus composed of any unmodified name-value pairs originally inherited by the shell, plus any modifications or additions, all of which must be noted in **export** commands.

The environment for any *simple-command* may be augmented by prefixing it with one or more assignments to *parameters*. Thus these two lines are equivalent

```
TERM=450 cmd args
(export TERM; TERM=450; cmd args)
```

If the **-k** flag is set, *all* keyword arguments are placed in the environment, even if they occur after the command name. The following prints 'a=b c' and 'c':

```
echo a=b c
set -k
echo a=b c
```

Signals.

The **INTERRUPT** and **QUIT** signals for an invoked command are ignored if the command is followed by **&**; otherwise signals have the values inherited by the shell from its parent. (But see also **trap**.)

Execution.

Each time a command is executed the above substitutions are carried out. Except for the 'special commands' listed below a new process is created and an attempt is made to execute the command via an *execve(2)*.

The shell parameter **\$PATH** defines the search path for the directory containing the command. Each alternative directory name is separated by a colon (:). The default path is **:/bin:/usr/bin**. If the command name contains a */*, the search path is not used. Otherwise, each directory in the path is searched for an executable file. If the file has execute permission but is not an *a.out* file, it is assumed to be a file containing shell commands. A subshell (i.e., a separate process) is spawned to read it. A parenthesized command is also executed in a subshell.

Special commands.

The following commands are executed in the shell process and except where specified no input output redirection is permitted for such commands.

- :** No effect; the command does nothing.
- . file** Read and execute commands from *file* and return. The search path **\$PATH** is used to find the directory containing *file*.
- break [n]**
Exit from the enclosing **for** or **while** loop, if any. If *n* is specified, break *n* levels.
- continue [n]**
Resume the next iteration of the enclosing **for** or **while** loop. If *n* is specified, resume

at the *n*-th enclosing loop.

cd [*arg*]

Change the current directory to *arg*. The shell parameter **\$HOME** is the default *arg*.

eval [*arg ...*]

The arguments are read as input to the shell and the resulting command(s) executed.

exec [*arg ...*]

The command specified by the arguments is executed in place of this shell without creating a new process. Input output arguments may appear and if no other arguments are given cause the shell input output to be modified.

exit [*n*]

Causes a non interactive shell to exit with the exit status specified by *n*. If *n* is omitted, the exit status is that of the last command executed. (An end of file will also exit from the shell.)

export [*name ...*]

The given names are marked for automatic export to the *environment* of subsequently-executed commands. If no arguments are given, a list of exportable names is printed.

login [*arg ...*]

Equivalent to 'exec login arg ...'.

read *name ...*

One line is read from the standard input; successive words of the input are assigned to the variables *name* in order, with leftover words to the last variable. The return code is 0 unless the end-of-file is encountered.

readonly [*name ...*]

The given names are marked readonly and the values of the these names may not be changed by subsequent assignment. If no arguments are given, a list of all readonly names is printed.

set [**-eknptuvx** [*arg ...*]]

-e If non interactive, exit immediately if a command fails.

-k All keyword arguments are placed in the environment for a command, not just those that precede the command name.

-n Read commands but do not execute them.

-t Exit after reading and executing one command.

-u Treat unset variables as an error when substituting.

-v Print shell input lines as they are read.

-x Print commands and their arguments as they are executed.

- Turn off the **-x** and **-v** options.

These flags can also be used upon invocation of the shell. The current set of flags may be found in **\$-**.

Remaining arguments are positional parameters and are assigned, in order, to **\$1**, **\$2**, etc. If no arguments are given, the values of all names are printed.

shift The positional parameters from **\$2...** are renamed **\$1...**

times Print the accumulated user and system times for processes run from the shell.

trap [*arg*] [*n*] ...

Arg is a command to be read and executed when the shell receives signal(s) *n*. (Note that *arg* is scanned once when the trap is set and once when the trap is taken.) Trap commands are executed in order of signal number. If *arg* is absent, all trap(s) *n* are reset to their original values. If *arg* is the null string, this signal is ignored by the shell and by invoked commands. If *n* is 0, the command *arg* is executed on exit from the shell, otherwise upon receipt of signal *n* as numbered in *sigvec*(2). *Trap* with no arguments prints a list of commands associated with each signal number.

umask [*nnn*]

The user file creation mask is set to the octal value *nnn* (see *umask(2)*). If *nnn* is omitted, the current value of the mask is printed.

wait [*n*]

Wait for the specified process and report its termination status. If *n* is not given, all currently active child processes are waited for. The return code from this command is that of the process waited for.

Invocation.

If the first character of argument zero is **-**, commands are read from **\$HOME/.profile**, if such a file exists. Commands are then read as described below. The following flags are interpreted by the shell when it is invoked.

- c** *string* If the **-c** flag is present, commands are read from *string*.
- s** If the **-s** flag is present or if no arguments remain then commands are read from the standard input. Shell output is written to file descriptor 2.
- i** If the **-i** flag is present or if the shell input and output are attached to a terminal (as told by *tty*) then this shell is *interactive*. In this case the terminate signal SIGTERM (see *sigvec(2)*) is ignored (so that 'kill 0' does not kill an interactive shell) and the interrupt signal SIGINT is caught and ignored (so that **wait** is interruptible). In all cases SIGQUIT is ignored by the shell.

The remaining flags and arguments are described under the **set** command.

FILES

\$HOME/.profile
/tmp/sh*
/dev/null

SEE ALSO

csh(1), **test(1)**, **execve(2)**, **environ(7)**

DIAGNOSTICS

Errors detected by the shell, such as syntax errors cause the shell to return a non zero exit status. If the shell is being used non interactively then execution of the shell file is abandoned. Otherwise, the shell returns the exit status of the last command executed (see also **exit**).

BUGS

If **<<** is used to provide standard input to an asynchronous process invoked by **&**, the shell gets mixed up about naming the input document. A garbage file **/tmp/sh*** is created, and the shell complains about not being able to find the file by another name.

NAME

size — size of an object file

SYNOPSIS

size [object ...]

DESCRIPTION

Size prints the (decimal) number of bytes required by the text, data, and bss portions, and their sum in hex and decimal, of each object-file argument. If no file is specified, **a.out** is used.

SEE ALSO

a.out(5)

NAME

sleep — suspend execution for an interval

SYNOPSIS

sleep time

DESCRIPTION

Sleep suspends execution for *time* seconds. It is used to execute a command after a certain amount of time as in:

```
(sleep 105; command)&
```

or to execute a command every so often, as in:

```
while true
do
    command
    sleep 37
done
```

SEE ALSO

setitimer(2), alarm(3C), sleep(3)

BUGS

Time must be less than 2,147,483,647 seconds.

NAME

sort — sort or merge files

SYNOPSIS

```
sort [ -mubdfinrtx ] [ +pos1 [ -pos2 ] ] ... [ -o name ] [ -T directory ] [ name ] ...
```

DESCRIPTION

Sort sorts lines of all the named files together and writes the result on the standard output. The name ‘—’ means the standard input. If no input files are named, the standard input is sorted.

The default sort key is an entire line. Default ordering is lexicographic by bytes in machine collating sequence. The ordering is affected globally by the following options, one or more of which may appear.

- b** Ignore leading blanks (spaces and tabs) in field comparisons.
- d** ‘Dictionary’ order: only letters, digits and blanks are significant in comparisons.
- f** Fold upper case letters onto lower case.
- i** Ignore characters outside the ASCII range 040-0176 in nonnumeric comparisons.
- n** An initial numeric string, consisting of optional blanks, optional minus sign, and zero or more digits with optional decimal point, is sorted by arithmetic value. Option **n** implies option **b**.
- r** Reverse the sense of comparisons.
- tx** ‘Tab character’ separating fields is *x*.

The notation *+pos1 -pos2* restricts a sort key to a field beginning at *pos1* and ending just before *pos2*. *Pos1* and *pos2* each have the form *m.n*, optionally followed by one or more of the flags **bdfinr**, where *m* tells a number of fields to skip from the beginning of the line and *n* tells a number of characters to skip further. If any flags are present they override all the global ordering options for this key. If the **b** option is in effect *n* is counted from the first nonblank in the field; **b** is attached independently to *pos2*. A missing *.n* means *.0*; a missing *-pos2* means the end of the line. Under the **-tx** option, fields are strings separated by *x*; otherwise fields are nonempty nonblank strings separated by blanks.

When there are multiple sort keys, later keys are compared only after all earlier keys compare equal. Lines that otherwise compare equal are ordered with all bytes significant.

These option arguments are also understood:

- c** Check that the input file is sorted according to the ordering rules; give no output unless the file is out of sort.
- m** Merge only, the input files are already sorted.
- o** The next argument is the name of an output file to use instead of the standard output. This file may be the same as one of the inputs.
- T** The next argument is the name of a directory in which temporary files should be made.
- u** Suppress all but one in each set of equal lines. Ignored bytes and bytes outside keys do not participate in this comparison.

EXAMPLES

Print in alphabetical order all the unique spellings in a list of words. Capitalized words differ from uncapitalized.

```
sort -u +0f +0 list
```

Print the password file (*passwd(5)*) sorted by user id number (the 3rd colon-separated field).

```
sort -t: +2n /etc/passwd
```

Print the first instance of each month in an already sorted file of (month day) entries. The options `-um` with just one input file make the choice of a unique representative from a set of equal lines predictable.

```
sort -um +0 -1 dates
```

FILES

/usr/tmp/stm*, /tmp/* first and second tries for temporary files

SEE ALSO

uniq(1), comm(1), rev(1), join(1)

DIAGNOSTICS

Comments and exits with nonzero status for various trouble conditions and for disorder discovered under option `-c`.

BUGS

Very long lines are silently truncated.

NAME

spell, spellin, spellout — find spelling errors

SYNOPSIS

```
spell [ -v ] [ -b ] [ -x ] [ -d hlist ] [ -s hstop ] [ -h spellhist ] [ file ] ...
spellin [ list ]
spellout [ -d ] list
```

DESCRIPTION

Spell collects words from the named documents, and looks them up in a spelling list. Words that neither occur among nor are derivable (by applying certain inflections, prefixes or suffixes) from words in the spelling list are printed on the standard output. If no files are named, words are collected from the standard input.

Spell ignores most *troff*, *tbl* and *eqn*(1) constructions.

Under the *-v* option, all words not literally in the spelling list are printed, and plausible derivations from spelling list words are indicated.

Under the *-b* option, British spelling is checked. Besides preferring *centre*, *colour*, *speciality*, *travelled*, etc., this option insists upon *-ise* in words like *standardise*, Fowler and the OED to the contrary notwithstanding.

Under the *-x* option, every plausible stem is printed with '=' for each word.

The spelling list is based on many sources. While it is more haphazard than an ordinary dictionary, it is also more effective with proper names and popular technical words. Coverage of the specialized vocabularies of biology, medicine and chemistry is light.

The auxiliary files used for the spelling list, stop list, and history file may be specified by arguments following the *-d*, *-s*, and *-h* options. The default files are indicated below. Copies of all output may be accumulated in the history file. The stop list filters out misspellings (e.g. *thier=thy-y+ier*) that would otherwise pass.

Two routines help maintain the hash lists used by *spell*. Both expect a set of words, one per line, from the standard input. *Spellin* combines the words from the standard input and the preexisting *list* file and places a new list on the standard output. If no *list* file is specified, the new list is created from scratch. *Spellout* looks up each word from the standard input and prints on the standard output those that are missing from (or present on, with option *-d*) the hashed *list* file. For example, to verify that *hookey* is not on the default spelling list, add it to your own private list, and then use it with *spell*,

```
echo hookey | spellout /usr/dict/hlista
echo hookey | spellin /usr/dict/hlista > myhlist
spell -d myhlist huckfinn
```

FILES

/usr/dict/hlist[ab]	hashed spelling lists, American & British, default for <i>-d</i>
/usr/dict/hstop	hashed stop list, default for <i>-s</i>
/dev/null	history file, default for <i>-h</i>
/tmp/spell.\$\$*	temporary files
/usr/lib/spell	

SEE ALSO

deroff(1), sort(1), tee(1), sed(1)

BUGS

The spelling list's coverage is uneven; new installations will probably wish to monitor the output for several months to gather local additions.
British spelling was done by an American.

NAME

split — split a file into pieces

SYNOPSIS

split [-n] [file [name]]

DESCRIPTION

Split reads *file* and writes it in *n*-line pieces (default 1000), as many as necessary, onto a set of output files. The name of the first output file is *name* with **aa** appended, and so on lexicographically. If no output name is given, **x** is default.

If no input file is given, or if **-** is given in its stead, then the standard input file is used.

NAME

strings — find the printable strings in a object, or other binary, file

SYNOPSIS

strings [-] [-o] [-number] file ...

DESCRIPTION

Strings looks for ascii strings in a binary file. A string is any sequence of 4 or more printing characters ending with a newline or a null. Unless the **-** flag is given, *strings* only looks in the initialized data space of object files. If the **-o** flag is given, then each string is preceded by its offset in the file (in octal). If the **-number** flag is given then number is used as the minimum string length rather than 4.

Strings is useful for identifying random object files and many other things.

SEE ALSO

od(1)

BUGS

The algorithm for identifying strings is extremely primitive

NAME

strip — remove symbols and relocation bits

SYNOPSIS

strip name ...

DESCRIPTION

Strip removes the symbol table and relocation bits ordinarily attached to the output of the assembler and loader. This is useful to save space after a program has been debugged.

The effect of *strip* is the same as use of the *-s* option of *ld*.

FILES

/tmp/stm? temporary file

SEE ALSO

ld(1)

NAME

stty — set terminal options

SYNOPSIS

stty [option ...]

DESCRIPTION

Stty sets certain I/O options on the current output terminal, placing its output on the diagnostic output. With no argument, it reports the speed of the terminal and the settings of the options which are different from their defaults. With the argument “all”, all normally used option settings are reported. With the argument “everything”, everything *stty* knows about is printed. The option strings are selected from the following set:

even allow even parity input
-even disallow even parity input
odd allow odd parity input
-odd disallow odd parity input
raw raw mode input (**no** input processing (erase, kill, interrupt, ...); parity bit passed back)
-raw negate raw mode
cooked same as ‘-raw’
cbreak make each character available to *read(2)* as received; no erase and kill processing, but all other processing (interrupt, suspend, ...) is performed
-cbreak make characters available to *read* only when newline is received
-nl allow carriage return for new-line, and output CR-LF for carriage return or new-line
nl accept only new-line to end lines
echo echo back every character typed
-echo do not echo characters
lcase map upper case to lower case
-lcase do not map case
tandem enable flow control, so that the system sends out the stop character when its internal queue is in danger of overflowing on input, and sends the start character when it is ready to accept further input
-tandem disable flow control
-tabs replace tabs by spaces when printing
tabs preserve tabs
ek set erase and kill characters to # and @

For the following commands which take a character argument *c*, you may also specify *c* as the “u” or “undef”, to set the value to be undefined. A value of “^x”, a 2 character sequence, is also interpreted as a control character, with “^?” representing delete.

erase *c* set erase character to *c* (default ‘#’, but often reset to ^H.)
kill *c* set kill character to *c* (default ‘@’, but often reset to ^U.)
intr *c* set interrupt character to *c* (default DEL or ^? (delete), but often reset to ^C.)
quit *c* set quit character to *c* (default control \.)
start *c* set start character to *c* (default control Q.)
stop *c* set stop character to *c* (default control S.)
eof *c* set end of file character to *c* (default control D.)
brk *c* set break character to *c* (default undefined.) This character is an extra wakeup causing character.
cr0 cr1 cr2 cr3
select style of delay for carriage return (see *ioctl(2)*)
nl0 nl1 nl2 nl3
select style of delay for linefeed
tab0 tab1 tab2 tab3

	select style of delay for tab
ff0 ff1	select style of delay for form feed
bs0 bs1	select style of delay for backspace
tty33	set all modes suitable for the Teletype Corporation Model 33 terminal.
tty37	set all modes suitable for the Teletype Corporation Model 37 terminal.
vt05	set all modes suitable for Digital Equipment Corp. VT05 terminal
dec	set all modes suitable for Digital Equipment Corp. operating systems users; (erase, kill, and interrupt characters to ^? , ^U , and ^C , decctlq and "newcrt".)
tn300	set all modes suitable for a General Electric TermiNet 300
ti700	set all modes suitable for Texas Instruments 700 series terminal
tek	set all modes suitable for Tektronix 4014 terminal
0	hang up phone line immediately
50 75 110 134 150 200 300 600 1200 1800 2400 4800 9600 exta extb	Set terminal baud rate to the number given, if possible. (These are the speeds supported by the DH-11 interface).

A teletype driver which supports the job control processing of *cs*(1) and more functionality than the basic driver is fully described in *ty*(4). The following options apply only to it.

new	Use new driver (switching flushes typeahead).
crt	Set options for a CRT (crtbs , ctlecho and, if ≥ 1200 baud, crterase and crtkill .)
crtbs	Echo backspaces on erase characters.
prterase	For printing terminal echo erased characters backwards within " \ " and " / ".
crterase	Wipe out erased characters with "backspace-space-backspace."
-crterase	Leave erased characters visible; just backspace.
crtkill	Wipe out input on like kill ala crterase .
-crtkill	Just echo line kill character and a newline on line kill.
ctlecho	Echo control characters as " ^x " (and delete as " ^? ".) Print two backspaces following the EOT character (control D).
-ctlecho	Control characters echo as themselves; in cooked mode EOT (control-D) is not echoed.
decctlq	After output is suspended (normally by ^S), only a start character (normally ^Q) will restart it. This is compatible with DEC's vendor supplied systems.
-decctlq	After output is suspended, any character typed will restart it; the start character will restart output without providing any input. (This is the default.)
tostop	Background jobs stop if they attempt terminal output.
-tostop	Output from background jobs to the terminal is allowed.
tilde	Convert " ~ " to " ^~ " on output (for Hazeltine terminals).
-tilde	Leave poor " ~ " alone.
flusho	Output is being discarded usually because user hit control O (internal state bit).
-flusho	Output is not being discarded.
pendin	Input is pending after a switch from cbreak to cooked and will be re-input when a read becomes pending or more input arrives (internal state bit).
-pendin	Input is not pending.
intrup	Send a signal (SIGTINT) to the terminal control process group whenever an input record (line in cooked mode, character in cbreak or raw mode) is available for reading.
-intrup	Don't send input available interrupts.
mdmbuf	Start/stop output on carrier transitions (not implemented).
-mdmbuf	
	Return error if write attempted after carrier drops.
litout	Send output characters without any processing.

- litout** Do normal output processing, inserting delays, etc.
- nohang** Don't send hangup signal if carrier drops.
- nohang** Send hangup signal to control process group when carrier drops.
- etxack** Diablo style etx/ack handshaking (not implemented).

The following special characters are applicable only to the new teletype driver and are not normally changed.

- susp** *c* set suspend process character to *c* (default control Z).
- dsusp** *c* set delayed suspend process character to *c* (default control Y).
- rprnt** *c* set reprint line character to *c* (default control R).
- flush** *c* set flush output character to *c* (default control O).
- werase** *c* set word erase character to *c* (default control W).
- lnext** *c* set literal next character to *c* (default control V).

SEE ALSO

ioctl(2), tabs(1), tset(1), tty(4)

NAME

su — substitute user id temporarily

SYNOPSIS

su [userid]

DESCRIPTION

Su demands the password of the specified *userid*, and if it is given, changes to that *userid* and invokes the Shell *sh*(1) without changing the current directory. The user environment is unchanged except for HOME and SHELL, which are taken from the password file for the user being substituted (see *environ*(7)). The new user ID stays in force until the Shell exits.

If no *userid* is specified, 'root' is assumed. To remind the super-user of his responsibilities, the Shell substitutes '#' for its usual prompt.

SEE ALSO

sh(1)

BUGS

Local administrative rules cause restrictions to be placed on who can *su* to 'root', even with the root password. These rules vary from site to site.

NAME

sum — sum and count blocks in a file

SYNOPSIS

sum file

DESCRIPTION

Sum calculates and prints a 16-bit checksum for the named file, and also prints the number of blocks in the file. It is typically used to look for bad spots, or to validate a file communicated over some transmission line.

SEE ALSO

wc(1)

DIAGNOSTICS

'Read error' is indistinguishable from end of file on most devices; check the block count.

NAME

tail — deliver the last part of a file

SYNOPSIS

tail [\pm number[lbc][fr]] [file]

DESCRIPTION

Tail copies the named file to the standard output beginning at a designated place. If no file is named, the standard input is used.

Copying begins at distance *+number* from the beginning, or *-number* from the end of the input. *Number* is counted in units of lines, blocks or characters, according to the appended option **l**, **b** or **c**. When no units are specified, counting is by lines.

Specifying **r** causes *tail* to print lines from the end of the file in reverse order. The default for **r** is to print the entire file this way. Specifying **f** causes *tail* to not quit at end of file, but rather wait and try to read repeatedly in hopes that the file will grow.

SEE ALSO

dd(1)

BUGS

Tails relative to the end of the file are treasured up in a buffer, and thus are limited in length.

Various kinds of anomalous behavior may happen with character special files.

NAME

`tar` — tape archiver

SYNOPSIS

`tar` [*key*] [*name ...*]

DESCRIPTION

Tar saves and restores multiple files on a single file (usually a magnetic tape, but it can be any file). *Tar*'s actions are controlled by the *key* argument. The *key* is a string of characters containing at most one function letter and possibly one or more function modifiers. Other arguments to *tar* are file or directory names specifying which files to dump or restore. In all cases, appearance of a directory name refers to the files and (recursively) subdirectories of that directory.

The function portion of the key is specified by one of the following letters:

- r** The named files are written on the end of the tape. The **c** function implies this.
- x** The named files are extracted from the tape. If the named file matches a directory whose contents had been written onto the tape, this directory is (recursively) extracted. The owner, modification time, and mode are restored (if possible). If no file argument is given, the entire content of the tape is extracted. Note that if multiple entries specifying the same file are on the tape, the last one overwrites all earlier.
- t** The names of the specified files are listed each time they occur on the tape. If no file argument is given, all of the names on the tape are listed.
- u** The named files are added to the tape if either they are not already there or have been modified since last put on the tape.
- c** Create a new tape; writing begins on the beginning of the tape instead of after the last file. This command implies **r**.
- o** On output, *tar* normally places information specifying owner and modes of directories in the archive. Former versions of *tar*, when encountering this information will give error message of the form
 "<name>/: cannot create".
 This option will suppress the directory information.
- p** This option says to restore files to their original modes, ignoring the present *umask*(2). Setuid and sticky information will also be restored to the super-user.

The following characters may be used in addition to the letter which selects the function desired.

- 0, ..., 9** This modifier selects an alternate drive on which the tape is mounted. The default is drive 0 at 1600 bpi, which is normally `/dev/rmt8`.
- v** Normally *tar* does its work silently. The **v** (verbose) option make *tar* type the name of each file it treats preceded by the function letter. With the **t** function, the verbose option gives more information about the tape entries than just their names.
- w** *Tar* prints the action to be taken followed by file name, then wait for user confirmation. If a word beginning with 'y' is given, the action is done. Any other input means don't do it.
- f** *Tar* uses the next argument as the name of the archive instead of `/dev/rmt?`. If the name of the file is '-', *tar* writes to standard output or reads from standard input, whichever is appropriate. Thus, *tar* can be used as the head or tail of a filter chain. *Tar* can also be used to move hierarchies with the command
 `cd fromdir; tar cf - . | (cd todir; tar xf -)`

- b** *Tar* uses the next argument as the blocking factor for tape records. The default is 20 (the maximum). This option should only be used with raw magnetic tape archives (See **f** above). The block size is determined automatically when reading tapes (key letters 'x' and 't').
- l** tells *tar* to complain if it cannot resolve all of the links to the files dumped. If this is not specified, no error messages are printed.
- m** tells *tar* not to restore the modification times. The modification time will be the time of extraction.
- h** Force *tar* to follow symbolic links as if they were normal files or directories. Normally, *tar* does not follow symbolic links.
- B** Forces input and output blocking to 20 blocks per record. This option was added so that *tar* can work across a communications channel where the blocking may not be maintained.

If a file name is preceded by **-C**, then *tar* will perform a *chdir(2)* to that file name. This allows multiple directories not related by a close common parent to be archived using short relative path names. For example, to archive files from `/usr/include` and from `/etc`, one might use

```
tar c -C /usr include -C / etc
```

Previous restrictions dealing with *tar*'s inability to properly handle blocked archives have been lifted.

FILES

`/dev/rmt?`
`/tmp/tar*`

DIAGNOSTICS

Complaints about bad key characters and tape read/write errors.
 Complaints if enough memory is not available to hold the link tables.

BUGS

There is no way to ask for the *n*-th occurrence of a file.
 Tape errors are handled ungracefully.
 The **u** option can be slow.
 The current limit on file name length is 10⁷ characters.
 There is no way to selectively follow symbolic links.

NAME

tee — pipe fitting

SYNOPSIS

tee [**-i**] [**-a**] [file] ...

DESCRIPTION

Tee transcribes the standard input to the standard output and makes copies in the *files*. Option **-i** ignores interrupts; option **-a** causes the output to be appended to the *files* rather than overwriting them.

NAME

`test` — condition command

SYNOPSIS

`test expr`

DESCRIPTION

`test` evaluates the expression *expr*, and if its value is true then returns zero exit status; otherwise, a non zero exit status is returned. `test` returns a non zero exit if there are no arguments.

The following primitives are used to construct *expr*.

- `-r file` true if the file exists and is readable.
- `-w file` true if the file exists and is writable.
- `-f file` true if the file exists and is not a directory.
- `-d file` true if the file exists exists and is a directory.
- `-s file` true if the file exists and has a size greater than zero.
- `-t [fildes]`
true if the open file whose file descriptor number is *fildes* (1 by default) is associated with a terminal device.
- `-z s1` true if the length of string *s1* is zero.
- `-n s1` true if the length of the string *s1* is nonzero.
- `s1 = s2` true if the strings *s1* and *s2* are equal.
- `s1 != s2` true if the strings *s1* and *s2* are not equal.
- `s1` true if *s1* is not the null string.
- `n1 -eq n2`
true if the integers *n1* and *n2* are algebraically equal. Any of the comparisons `-ne`, `-gt`, `-ge`, `-lt`, or `-le` may be used in place of `-eq`.

These primaries may be combined with the following operators:

- `!` unary negation operator
 - `-a` binary *and* operator
 - `-o` binary *or* operator
 - `(expr)`
parentheses for grouping.
- `-a` has higher precedence than `-o`. Notice that all the operators and flags are separate arguments to `test`. Notice also that parentheses are meaningful to the Shell and must be escaped.

SEE ALSO

`sh(1)`, `find(1)`

NAME

time — time a command

SYNOPSIS

time command

DESCRIPTION

The given command is executed; after it is complete, *time* prints the elapsed time during the command, the time spent in the system, and the time spent in execution of the command. Times are reported in seconds.

On a PDP-11, the execution time can depend on what kind of memory the program happens to land in; the user time in MOS is often half what it is in core.

The times are printed on the diagnostic output stream.

Time is built in to *cs**h*(1), using a different output format.

BUGS

Elapsed time is accurate to the second, while the CPU times are measured to the 100th second. Thus the sum of the CPU times can be up to a second larger than the elapsed time.

Time is a built-in command to *cs**h*(1), with a much different syntax. This command is available as “/bin/time” to *cs**h* users.

NAME

touch — update date last modified of a file

SYNOPSIS

touch [-c] [-f] file ...

DESCRIPTION

Touch attempts to set the modified date of each *file*. If a *file* exists, this is done by reading a character from the file and writing it back. If a *file* does not exist, an attempt will be made to create it unless the **-c** option is specified. The **-f** option will attempt to force the touch in spite of read and write permissions on a *file*.

SEE ALSO

utimes(2)

NAME

`tp` — manipulate tape archive

SYNOPSIS

`tp [key] [name ...]`

DESCRIPTION

Tp saves and restores files on DECTape or magtape. Its actions are controlled by the *key* argument. The key is a string of characters containing at most one function letter and possibly one or more function modifiers. Other arguments to the command are file or directory names specifying which files are to be dumped, restored, or listed. In all cases, appearance of a directory name refers to the files and (recursively) subdirectories of that directory.

The function portion of the key is specified by one of the following letters:

- r** The named files are written on the tape. If files with the same names already exist, they are replaced. 'Same' is determined by string comparison, so './abc' can never be the same as '/usr/dmr/abc' even if '/usr/dmr' is the current directory. If no file argument is given, '.' is the default.
- u** updates the tape. **u** is like **r**, but a file is replaced only if its modification date is later than the date stored on the tape; that is to say, if it has changed since it was dumped. **u** is the default command if none is given.
- d** deletes the named files from the tape. At least one name argument must be given. This function is not permitted on magtapes.
- x** extracts the named files from the tape to the file system. The owner and mode are restored. If no file argument is given, the entire contents of the tape are extracted.
- t** lists the names of the specified files. If no file argument is given, the entire contents of the tape is listed.

The following characters may be used in addition to the letter which selects the function desired.

- m** Specifies magtape as opposed to DECTape.
- 0,...,7** This modifier selects the drive on which the tape is mounted. For DECTape, **x** is default; for magtape '0' is the default.
- v** Normally *tp* does its work silently. The **v** (verbose) option causes it to type the name of each file it treats preceded by the function letter. With the **t** function, **v** gives more information about the tape entries than just the name.
- c** means a fresh dump is being created; the tape directory is cleared before beginning. Usable only with **r** and **u**. This option is assumed with magtape since it is impossible to selectively overwrite magtape.
- i** Errors reading and writing the tape are noted, but no action is taken. Normally, errors cause a return to the command level.
- f** Use the first named file, rather than a tape, as the archive. This option currently acts like **m**; *i.e.* **r** implies **c**, and neither **d** nor **u** are permitted.
- w** causes *tp* to pause before treating each file, type the indicative letter and the file name (as with **v**) and await the user's response. Response **y** means 'yes', so the file is treated. Null response means 'no', and the file does not take part in whatever is being done. Response **x** means 'exit'; the *tp* command terminates immediately. In the **x** function, files previously asked about have been extracted already. With **r**, **u**, and **d** no change has been made to the tape.

FILES

/dev/tap?
/dev/rmt?

SEE ALSO

ar(1), tar(1)

DIAGNOSTICS

Several; the non-obvious one is 'Phase error', which means the file changed after it was selected for dumping but before it was dumped.

BUGS

A single file with several links to it is treated like several files.

Binary-coded control information makes magnetic tapes written by *tp* difficult to carry to other machines; *tar*(1) avoids the problem.

NAME

tr — translate characters

SYNOPSIS

```
tr [ -cds ] [ string1 [ string2 ] ]
```

DESCRIPTION

Tr copies the standard input to the standard output with substitution or deletion of selected characters. Input characters found in *string1* are mapped into the corresponding characters of *string2*. When *string2* is short it is padded to the length of *string1* by duplicating its last character. Any combination of the options *-cds* may be used: *-c* complements the set of characters in *string1* with respect to the universe of characters whose ASCII codes are 01 through 0377 octal; *-d* deletes all input characters in *string1*; *-s* squeezes all strings of repeated output characters that are in *string2* to single characters.

In either string the notation *a-b* means a range of characters from *a* to *b* in increasing ASCII order. The character ** followed by 1, 2 or 3 octal digits stands for the character whose ASCII code is given by those digits. A ** followed by any other character stands for that character.

The following example creates a list of all the words in 'file1' one per line in 'file2', where a word is taken to be a maximal string of alphabets. The second string is quoted to protect ** from the Shell. 012 is the ASCII code for newline.

```
tr -cs A-Za-z '\012' <file1 >file2
```

SEE ALSO

ed(1), ascii(7), expand(1)

BUGS

Won't handle ASCII NUL in *string1* or *string2*; always deletes NUL from input.

NAME

troff, **nroff** — text formatting and typesetting

SYNOPSIS

troff [option] ... [file] ...

nroff [option] ... [file] ...

DESCRIPTION

Troff formats text in the named *files* for printing on a Graphic Systems C/A/T phototypesetter; *nroff* is used for typewriter-like devices. Their capabilities are described in the *Nroff/Troff user's manual*.

If no *file* argument is present, the standard input is read. An argument consisting of a single minus (-) is taken to be a file name corresponding to the standard input. The options, which may appear in any order so long as they appear before the files, are:

- olist* Print only pages whose page numbers appear in the comma-separated *list* of numbers and ranges. A range *N-M* means pages *N* through *M*; an initial *-N* means from the beginning to page *N*; and a final *N-* means from *N* to the end.
- nN* Number first generated page *N*.
- sN* Stop every *N* pages. *Nroff* will halt prior to every *N* pages (default *N=1*) to allow paper loading or changing, and will resume upon receipt of a newline. *Troff* will stop the phototypesetter every *N* pages, produce a trailer to allow changing cassettes, and resume when the typesetter's start button is pressed.
- mname* Prepend the macro file */usr/lib/tmac/tmac.name* to the input *files*.
- raN* Set register *a* (one-character) to *N*.
- i* Read standard input after the input files are exhausted.
- q* Invoke the simultaneous input-output mode of the *rd* request.

Troff only

- t* Direct output to the standard output instead of the phototypesetter.
- f* Refrain from feeding out paper and stopping phototypesetter at the end of the run.
- w* Wait until phototypesetter is available, if currently busy.
- b* Report whether the phototypesetter is busy or available. No text processing is done.
- a* Send a printable ASCII approximation of the results to the standard output.
- pN* Print all characters in point size *N* while retaining all prescribed spacings and motions, to reduce phototypesetter elapsed time.

If the file */usr/adm/tracct* is writable, *troff* keeps phototypesetter accounting records there. The integrity of that file may be secured by making *troff* a 'set user-id' program.

FILES

<i>/tmp/ta*</i>	temporary file
<i>/usr/lib/tmac/tmac.*</i>	standard macro files
<i>/usr/lib/term/*</i>	terminal driving tables for <i>nroff</i>
<i>/usr/lib/font/*</i>	font width tables for <i>troff</i>
<i>/dev/cat</i>	phototypesetter
<i>/usr/adm/tracct</i>	accounting statistics for <i>/dev/cat</i>

SEE ALSO

J. F. Ossanna, *Nroff/Troff user's manual*
 B. W. Kernighan, *A TROFF Tutorial*
 eqn(1), tbl(1), ms(7), me(7), man(7), col(1)

NAME

true, false — provide truth values

SYNOPSIS

true

false

DESCRIPTION

True and *false* are usually used in a Bourne shell script. They test for the appropriate status "true" or "false" before running (or failing to run) a list of commands.

EXAMPLE

```
while true
do
    command list
done
```

SEE ALSO

csh(1), sh(1), false(1)

DIAGNOSTICS

True has exit status zero.

NAME

`tsort` — topological sort

SYNOPSIS

`tsort [file]`

DESCRIPTION

Tsort produces on the standard output a totally ordered list of items consistent with a partial ordering of items mentioned in the input *file*. If no *file* is specified, the standard input is understood.

The input consists of pairs of items (nonempty strings) separated by blanks. Pairs of different items indicate ordering. Pairs of identical items indicate presence, but not ordering.

SEE ALSO

`lorder(1)`

DIAGNOSTICS

Odd data: there is an odd number of fields in the input file.

BUGS

Uses a quadratic algorithm; not worth fixing for the typical use of ordering a library archive file.

NAME

tty — get terminal name

SYNOPSIS

tty [-s]

DESCRIPTION

Tty prints the pathname of the user's terminal unless the *-s* (silent) is given. In either case, the exit value is zero if the standard input is a terminal and one if it is not.

DIAGNOSTICS

'not a tty' if the standard input file is not a terminal.

NAME

`unget` — undo a previous `get` of an SCCS file

SYNOPSIS

`unget` [`-rSID`] [`-s`] [`-n`] files

DESCRIPTION

`Unget` undoes the effect of a `get -e` done prior to creating the intended new delta. If a directory is named, `unget` behaves as though each file in the directory were specified as a named file, except that non-SCCS files and unreadable files are silently ignored. If a name of `-` is given, the standard input is read with each line being taken as the name of an SCCS file to be processed.

Keyletter arguments apply independently to each named file.

- `-rSID` Uniquely identifies which delta is no longer intended. (This would have been specified by `get` as the "new delta"). The use of this keyletter is necessary only if two or more outstanding `gets` for editing on the same SCCS file were done by the same person (login name). A diagnostic results if the specified `SID` is ambiguous, or if it is necessary and omitted on the command line.
- `-s` Suppresses the printout, on the standard output, of the intended delta's `SID`.
- `-n` Causes the retention of the gotten file which would normally be removed from the current directory.

SEE ALSO

`delta(1)`, `get(1)`, `sact(1)`.

DIAGNOSTICS

Use `help(1)` for explanations.

NAME

uniq — report repeated lines in a file

SYNOPSIS

uniq [-udc [+n] [-n]] [input [output]]

DESCRIPTION

Uniq reads the input file comparing adjacent lines. In the normal case, the second and succeeding copies of repeated lines are removed; the remainder is written on the output file. Note that repeated lines must be adjacent in order to be found; see *sort(1)*. If the **-u** flag is used, just the lines that are not repeated in the original file are output. The **-d** option specifies that one copy of just the repeated lines is to be written. The normal mode output is the union of the **-u** and **-d** mode outputs.

The **-c** option supersedes **-u** and **-d** and generates an output report in default style but with each line preceded by a count of the number of times it occurred.

The *n* arguments specify skipping an initial portion of each line in the comparison:

- n** The first *n* fields together with any blanks before each are ignored. A field is defined as a string of non-space, non-tab characters separated by tabs and spaces from its neighbors.
- +n** The first *n* characters are ignored. Fields are skipped before characters.

SEE ALSO

sort(1), comm(1)

NAME

units — conversion program

SYNOPSIS

units

DESCRIPTION

Units converts quantities expressed in various standard scales to their equivalents in other scales. It works interactively in this fashion:

```

You have: inch
You want: cm
    * 2.54000e+00
    / 3.93701e-01
  
```

A quantity is specified as a multiplicative combination of units optionally preceded by a numeric multiplier. Powers are indicated by suffixed positive integers, division by the usual sign:

```

You have: 15 pounds force/in2
You want: atm
    * 1.02069e+00
    / 9.79730e-01
  
```

Units only does multiplicative scale changes. Thus it can convert Kelvin to Rankine, but not Centigrade to Fahrenheit. Most familiar units, abbreviations, and metric prefixes are recognized, together with a generous leavening of exotica and a few constants of nature including:

```

pi      ratio of circumference to diameter
c       speed of light
e       charge on an electron
g       acceleration of gravity
force   same as g
mole    Avogadro's number
water   pressure head per unit height of water
au      astronomical unit
  
```

'Pound' is a unit of mass. Compound names are run together, e.g. 'lightyear'. British units that differ from their US counterparts are prefixed thus: 'brgallon'. Currency is denoted 'belgiumfranc', 'britainpound', ...

For a complete list of units, 'cat /usr/lib/units'.

FILES

/usr/lib/units

BUGS

Don't base your financial plans on the currency conversions.

NAME

vfontinfo — inspect and print out information about UNIX fonts

SYNOPSIS

vfontinfo [**-v**] fontname [characters]

DESCRIPTION

Vfontinfo allows you to examine a font in the UNIX format. It prints out all the information in the font header and information about every non-null (width > 0) glyph. This can be used to make sure the font is consistent with the format.

The *fontname* argument is the name of the font you wish to inspect. It writes to standard output. If it can't find the file in your working directory, it looks in */usr/lib/vfont* (the place most of the fonts are kept).

The *characters*, if given, specify certain characters to show. If omitted, the entire font is shown.

If the **-v** (verbose) flag is used, the bits of the glyph itself are shown as an array of X's and spaces, in addition to the header information.

SEE ALSO

vpr(1), vfont(5)
The Berkeley Font Catalog

AUTHORS

Mark Horton
Andy Hertzfeld

NAME

vi — screen oriented (visual) display editor based on **ex**

SYNOPSIS

vi [**-t tag**] [**-r**] [**+command**] [**-l**] [**-wn**] name ...

DESCRIPTION

Vi (visual) is a display oriented text editor based on *ex*(1). *Ex* and *vi* run the same code; it is possible to get to the command mode of *ex* from within *vi* and vice-versa.

The *Vi Quick Reference* card and the *Introduction to Display Editing with Vi* provide full details on using *vi*.

FILES

See *ex*(1).

SEE ALSO

ex (1), *edit* (1), “*Vi Quick Reference*” card, “*An Introduction to Display Editing with Vi*”.

AUTHOR

William Joy

Mark Horton added macros to *visual* mode and is maintaining version 3.

BUGS

Software tabs using **^T** work only immediately after the *autoindent*.

Left and right shifts on intelligent terminals don't make use of insert and delete character operations in the terminal.

The *wrapmargin* option can be fooled since it looks at output columns when blanks are typed. If a long word passes through the margin and onto the next line without a break, then the line won't be broken.

Insert/delete within a line can be slow if tabs are present on intelligent terminals, since the terminals need help in doing this correctly.

Saving text on deletes in the named buffers is somewhat inefficient.

The *source* command does not work when executed as **:source**; there is no way to use the **:append**, **:change**, and **:insert** commands, since it is not possible to give more than one line of input to a **:** escape. To use these on a **:global** you must **Q** to *ex* command mode, execute them, and then reenter the screen editor with *vi* or *open*.

NAME

`vpl` - send plot file to plotter

SYNOPSIS

`vpl` name ...

DESCRIPTION

vpl is a shell script which calls and passes it's command line arguments to *Vpr*. *vpl* is used to plot files of the type created by *ged*, the graphics editor program.

EXAMPLE

`vpl /u0/chris/vw.spool`

SEE ALSO

`vpr(1)`

NAME

vpr, *vprm*, *vpq*, *Vpr*, *Vprm*, *Vpq* – Versatec spooler

SYNOPSIS

vpr [**-m**] [**-r**] [**name ...**]

Vpr [**-m**] [**-r**] [**name ...**]

DESCRIPTION

vpr causes the named text files to be queued for printing on the Versatec plotter.

Vpr spools the named plot files of the type created by *ged*, the graphics editor, for plotting on the Verstec.

The **-m** option causes notification via *mail*(1) to be sent when the job completes. The **-r** option causes the file to be removed when printing is complete.

[vV]prm removes an entry from the plotter queue. The *id*, filename or owner should be that reported by *[vV]pq*. All appropriate files will be removed. The *id* of each file removed from the queue will be printed.

[vV]pq shows the files in the Versatec queue. The *id* printed is useful for removing specific files from the queue via *[vV]prm*.

FILES

<i>/usr/lib/[vV]pd</i>	plotter daemon
<i>/usr/lib/[vV]pf</i>	plotter filter
<i>/u0/spool/[vV]pd/*</i>	spool area
<i>/u0/spool/[vV]pd]/cf*</i>	daemon control files
<i>/u0/spool/[vV]pd]/df*</i>	data files specified in "cf" files
<i>/u0/spool/[vV]pd]/tf*</i>	temporary copies of "cf" files

SEE ALSO

lpr(1)

NAME

vtroff, *rvtroff* - troff to a raster plotter

SYNOPSIS

vtroff [troff arguments] name ...

rvtroff [troff arguments] name ...

DESCRIPTION

vtroff runs *troff*(1) sending its output through various programs to produce typeset output on a raster plotter such as a Benson-Varian or or a Versatec.

rvtroff is identical to *vtroff* except that the output is rotated by ninety degrees.

FILES

<i>/usr/lib/tmac/tmac.vcat</i>	default font mounts and bug fixes
<i>/usr/lib/fontinfo/*</i>	fixes for other fonts
<i>/usr/lib/vfont</i>	directory containing fonts

SEE ALSO

troff(1), *vfont*(5), *vpr*(1)

NAME

wall — write to all users

SYNOPSIS

wall

DESCRIPTION

Wall reads its standard input until an end-of-file. It then sends this message, preceded by 'Broadcast Message ...', to all logged in users.

The sender should be super-user to override any protections the users may have invoked.

FILES

/dev/tty?
/etc/utmp

SEE ALSO

mesg(1), write(1)

DIAGNOSTICS

'Cannot send to ...' when the open on a user's tty file fails.

NAME

`wc` — word count

SYNOPSIS

`wc [-lwc] [name ...]`

DESCRIPTION

Wc counts lines, words and characters in the named files, or in the standard input if no name appears. A word is a maximal string of characters delimited by spaces, tabs or newlines.

If an argument beginning with one of “lwc” is present, the specified counts (lines, words, or characters) are selected by the letters l, w, or c. The default is `-lwc`.

BUGS

NAME

what — show what versions of object modules were used to construct a file

SYNOPSIS

what name ...

DESCRIPTION

What reads each file and searches for sequences of the form “@(#)” as inserted by the source code control system. It then prints the remainder of the string after this marker, up to a null character, newline, double quote, or “>” character.

BUGS

As SCCS is not licensed with UNIX/32V, this is a rewrite of the *what* command which is part of SCCS, and may not behave exactly the same as that command does.

NAME

who — who is on the system

SYNOPSIS

who [who-file] [am I]

DESCRIPTION

Who, without an argument, lists the login name, terminal name, and login time for each current UNIX user.

Without an argument, *who* examines the */etc/utmp* file to obtain its information. If a file is given, that file is examined. Typically the given file will be */usr/adm/wtmp*, which contains a record of all the logins since it was created. Then *who* lists logins, logouts, and crashes since the creation of the *wtmp* file. Each login is listed with user name, terminal name (with *'/dev/'* suppressed), and date and time. When an argument is given, logouts produce a similar line without a user name. Reboots produce a line with *'x'* in the place of the device name, and a fossil time indicative of when the system went down.

With two arguments, as in *'who am I'* (and also *'who are you'*), *who* tells who you are logged in as.

FILES

/etc/utmp

SEE ALSO

getuid(2), *utmp(5)*

NAME

whoami — print effective current user id

SYNOPSIS

whoami

DESCRIPTION

Whoami prints who you are. It works even if you are su'd, while 'who am i' does not since it uses /etc/utmp.

FILES

/etc/passwd Name data base

SEE ALSO

who (1)

NAME

write — write to another user

SYNOPSIS

write user [ttyname]

DESCRIPTION

Write copies lines from your terminal to that of another user. When first called, it sends the message

Message from yoursystem!yourname yourttyname...

The recipient of the message should write back at this point. Communication continues until an end of file is read from the terminal or an interrupt is sent. At that point *write* writes 'EOT' on the other terminal and exits.

If you want to write to a user who is logged in more than once, the *ttyname* argument may be used to indicate the appropriate terminal name.

Permission to write may be denied or granted by use of the *mesg* command. At the outset writing is allowed. Certain commands, in particular *nroff* and *pr(1)* disallow messages in order to prevent messy output.

If the character '!' is found at the beginning of a line, *write* calls the shell to execute the rest of the line as a command.

The following protocol is suggested for using *write*: when you first write to another user, wait for him to write back before starting to send. Each party should end each message with a distinctive signal—(o) for 'over' is conventional—that the other may reply. (oo) for 'over and out' is suggested when conversation is about to be terminated.

FILES

/etc/utmp	to find user
/bin/sh	to execute '!'

SEE ALSO

mesg(1), *who(1)*, *mail(1)*

NAME

xstr — extract strings from C programs to implement shared strings

SYNOPSIS

xstr [-c] [-] [file]

DESCRIPTION

Xstr maintains a file *strings* into which strings in component parts of a large program are hashed. These strings are replaced with references to this common area. This serves to implement shared constant strings, most useful if they are also read-only.

The command

```
xstr -c name
```

will extract the strings from the C source in *name*, replacing string references by expressions of the form (&*xstr*[number]) for some number. An appropriate declaration of *xstr* is prepended to the file. The resulting C text is placed in the file *x.c*, to then be compiled. The strings from this file are placed in the *strings* data base if they are not there already. Repeated strings and strings which are suffixes of existing strings do not cause changes to the data base.

After all components of a large program have been compiled a file *xs.c* declaring the common *xstr* space can be created by a command of the form

```
xstr
```

This *xs.c* file should then be compiled and loaded with the rest of the program. If possible, the array can be made read-only (shared) saving space and swap overhead.

Xstr can also be used on a single file. A command

```
xstr name
```

creates files *x.c* and *xs.c* as before, without using or affecting any *strings* file in the same directory.

It may be useful to run *xstr* after the C preprocessor if any macro definitions yield strings or if there is conditional code which contains strings which may not, in fact, be needed. *Xstr* reads from its standard input when the argument '-' is given. An appropriate command sequence for running *xstr* after the C preprocessor is:

```
cc -E name.c | xstr -c -  
cc -c x.c  
mv x.o name.o
```

Xstr does not touch the file *strings* unless new items are added, thus *make* can avoid remaking *xs.o* unless truly necessary.

FILES

<i>strings</i>	Data base of strings
<i>x.c</i>	Massaged C source
<i>xs.c</i>	C source for definition of array 'xstr'
/tmp/xs*	Temp file when 'xstr name' doesn't touch <i>strings</i>

SEE ALSO

mkstr(1)

AUTHOR

William Joy

BUGS

If a string is a suffix of another string in the data base, but the shorter string is seen first by *xstr* both strings will be placed in the data base, when just placing the longer one there will do.

NAME

yacc — yet another compiler-compiler

SYNOPSIS

yacc [-vd] grammar

DESCRIPTION

Yacc converts a context-free grammar into a set of tables for a simple automaton which executes an LR(1) parsing algorithm. The grammar may be ambiguous; specified precedence rules are used to break ambiguities.

The output file, *y.tab.c*, must be compiled by the C compiler to produce a program *yyparse*. This program must be loaded with the lexical analyzer program, *yylex*, as well as *main* and *yyerror*, an error handling routine. These routines must be supplied by the user; *Lex(1)* is useful for creating lexical analyzers usable by *yacc*.

If the *-v* flag is given, the file *y.output* is prepared, which contains a description of the parsing tables and a report on conflicts generated by ambiguities in the grammar.

If the *-d* flag is used, the file *y.tab.h* is generated with the *define* statements that associate the *yacc*-assigned 'token codes' with the user-declared 'token names'. This allows source files other than *y.tab.c* to access the token codes.

FILES

y.output
y.tab.c
y.tab.h defines for token names
yacc.tmp, *yacc.acts* temporary files
/usr/lib/yaccpar parser prototype for C programs

SEE ALSO

lex(1)
LR Parsing by A. V. Aho and S. C. Johnson, Computing Surveys, June, 1974.
YACC — Yet Another Compiler Compiler by S. C. Johnson.

DIAGNOSTICS

The number of reduce-reduce and shift-reduce conflicts is reported on the standard output; a more detailed report is found in the *y.output* file. Similarly, if some rules are not reachable from the start symbol, this is also reported.

BUGS

Because file names are fixed, at most one *yacc* process can be active in a given directory at a time.

NAME

yes — be repetitively affirmative

SYNOPSIS

yes [*expletive*]

DESCRIPTION

Yes repeatedly outputs “y”, or if *expletive* is given, that is output repeatedly. Termination is by rubout.

NAME

intro — introduction to system calls and error numbers

SYNOPSIS

```
#include <errno.h>
```

DESCRIPTION

This section describes all of the system calls. Most of these calls have one or more error returns. An error condition is indicated by an otherwise impossible return value. This is almost always `-1`; the individual descriptions specify the details.

As with normal arguments, all return codes and values from functions are of type integer unless otherwise noted. An error number is also made available in the external variable `errno`, which is not cleared on successful calls. Thus `errno` should be tested only after an error has occurred.

The following is a complete list of the errors and their names as given in `<errno.h>`.

- 0 Error 0
Unused.
- 1 EPERM Not owner
Typically this error indicates an attempt to modify a file in some way forbidden except to its owner or super-user. It is also returned for attempts by ordinary users to do things allowed only to the super-user.
- 2 ENOENT No such file or directory
This error occurs when a file name is specified and the file should exist but doesn't, or when one of the directories in a path name does not exist.
- 3 ESRCH No such process
The process whose number was given to `kill` and `ptrace` does not exist, or is already dead.
- 4 EINTR Interrupted system call
An asynchronous signal (such as interrupt or quit), which the user has elected to catch, occurred during a system call. If execution is resumed after processing the signal, it will appear as if the interrupted system call returned this error condition.
- 5 EIO I/O error
Some physical I/O error occurred during a `read` or `write`. This error may in some cases occur on a call following the one to which it actually applies.
- 6 ENXIO No such device or address
I/O on a special file refers to a subdevice which does not exist, or beyond the limits of the device. It may also occur when, for example, an illegal tape drive unit number is selected or a disk pack is not loaded on a drive.
- 7 E2BIG Arg list too long
An argument list longer than 10240 bytes is presented to `execve`.
- 8 ENOEXEC Exec format error
A request is made to execute a file which, although it has the appropriate permissions, does not start with a valid magic number, see `a.out(5)`.
- 9 EBADF Bad file number
Either a file descriptor refers to no open file, or a read (resp. write) request is made to a file which is open only for writing (resp. reading).
- 10 ECHILD No children
`Wait` and the process has no living or unwaited-for children.

- 11 EAGAIN No more processes
In a *fork*, the system's process table is full or the user is not allowed to create any more processes.
- 12 ENOMEM Not enough core
During an *execve* or *break*, a program asks for more core or swap space than the system is able to supply. A lack of swap space is normally a temporary condition, however a lack of core is not a temporary condition; the maximum size of the text, data, and stack segments is a system parameter.
- 13 EACCES Permission denied
An attempt was made to access a file in a way forbidden by the protection system.
- 14 EFAULT Bad address
The system encountered a hardware fault in attempting to access the arguments of a system call.
- 15 ENOTBLK Block device required
A plain file was mentioned where a block device was required, e.g. in *mount*.
- 16 EBUSY Mount device busy
An attempt to mount a device that was already mounted or an attempt was made to dismount a device on which there is an active file directory. (open file, current directory, mounted-on file, active text segment).
- 17 EEXIST File exists
An existing file was mentioned in an inappropriate context, e.g. *link*.
- 18 EXDEV Cross-device link
A hard link to a file on another device was attempted.
- 19 ENODEV No such device
An attempt was made to apply an inappropriate system call to a device; e.g. read a write-only device.
- 20 ENOTDIR Not a directory
A non-directory was specified where a directory is required, for example in a path name or as an argument to *chdir*.
- 21 EISDIR Is a directory
An attempt to write on a directory.
- 22 EINVAL Invalid argument
Some invalid argument: dismounting a non-mounted device, mentioning an unknown signal in *signal*, reading or writing a file for which *seek* has generated a negative pointer. Also set by math functions, see *intro(3)*.
- 23 ENFILE File table overflow
The system's table of open files is full, and temporarily no more *opens* can be accepted.
- 24 EMFILE Too many open files
Customary configuration limit is 20 per process.
- 25 ENOTTY Not a typewriter
The file mentioned in an *ioctl* is not a terminal or one of the other devices to which these calls apply.
- 26 ETXTBSY Text file busy
An attempt to execute a pure-procedure program which is currently open for writing (or reading!). Also an attempt to open for writing a pure-procedure program that is being executed.

- 27 EFBIG File too large
The size of a file exceeded the maximum (about 10^9 bytes).
- 28 ENOSPC No space left on device
During a *write* to an ordinary file, there is no free space left on the device.
- 29 ESPIPE Illegal seek
An *lseek* was issued to a pipe. This error may also be issued for other non-seekable devices.
- 30 EROFS Read-only file system
An attempt to modify a file or directory was made on a device mounted read-only.
- 31 EMLINK Too many links
An attempt to make more than 32767 hard links to a file.
- 32 EPIPE Broken pipe
A write on a pipe or socket for which there is no process to read the data. This condition normally generates a signal; the error is returned if the signal is ignored.
- 33 EDOM Math argument
The argument of a function in the math package (3M) is out of the domain of the function.
- 34 ERANGE Result too large
The value of a function in the math package (3M) is unrepresentable within machine precision.
- 35 EWOULDBLOCK Operation would block
An operation which would cause a process to block was attempted on a object in non-blocking mode (see *ioctl* (2)).
- 36 EINPROGRESS Operation now in progress
An operation which takes a long time to complete (such as a *connect* (2)) was attempted on a non-blocking object (see *ioctl* (2)).
- 37 EALREADY Operation already in progress
An operation was attempted on a non-blocking object which already had an operation in progress.
- 38 ENOTSOCK Socket operation on non-socket
Self-explanatory.
- 39 EDESTADDRREQ Destination address required
A required address was omitted from an operation on a socket.
- 40 EMSGSIZE Message too long
A message sent on a socket was larger than the internal message buffer.
- 41 EPROTOTYPE Protocol wrong type for socket
A protocol was specified which does not support the semantics of the socket type requested. For example you cannot use the ARPA Internet UDP protocol with type `SOCK_STREAM`.
- 42 ENOPROTOPT Bad protocol option
A bad option was specified in a *getsockopt*(2) or *setsockopt*(2) call.
- 43 EPROTONOSUPPORT Protocol not supported
The protocol has not been configured into the system or no implementation for it exists.

- 44 ESOCKTNOSUPPORT Socket type not supported
The support for the socket type has not been configured into the system or no implementation for it exists.
- 45 EOPNOTSUPP Operation not supported on socket
For example, trying to *accept* a connection on a datagram socket.
- 46 EPFNOSUPPORT Protocol family not supported
The protocol family has not been configured into the system or no implementation for it exists.
- 47 EAFNOSUPPORT Address family not supported by protocol family
An address incompatible with the requested protocol was used. For example, you shouldn't necessarily expect to be able to use PUP Internet addresses with ARPA Internet protocols.
- 48 EADDRINUSE Address already in use
Only one usage of each address is normally permitted.
- 49 EADDRNOTAVAIL Can't assign requested address
Normally results from an attempt to create a socket with an address not on this machine.
- 50 ENETDOWN Network is down
A socket operation encountered a dead network.
- 51 ENETUNREACH Network is unreachable
A socket operation was attempted to an unreachable network.
- 52 ENETRESET Network dropped connection on reset
The host you were connected to crashed and rebooted.
- 53 ECONNABORTED Software caused connection abort
A connection abort was caused internal to your host machine.
- 54 ECONNRESET Connection reset by peer
A connection was forcibly closed by a peer. This normally results from the peer executing a *shutdown* (2) call.
- 55 ENOBUFS No buffer space available
An operation on a socket or pipe was not performed because the system lacked sufficient buffer space.
- 56 EISCONN Socket is already connected
A *connect* request was made on an already connected socket; or, a *sendto* or *sendmsg* request on a connected socket specified a destination other than the connected party.
- 57 ENOTCONN Socket is not connected
An request to send or receive data was disallowed because the socket is not connected.
- 58 ESHUTDOWN Can't send after socket shutdown
A request to send data was disallowed because the socket had already been shut down with a previous *shutdown*(2) call.
- 59 *unused*
- 60 ETIMEDOUT Connection timed out
A *connect* request failed because the connected party did not properly respond after a period of time. (The timeout period is dependent on the communication protocol.)
- 61 ECONNREFUSED Connection refused
No connection could be made because the target machine actively refused it. This usually results from trying to connect to a service which is inactive on the foreign host.

- 62 ELOOP Too many levels of symbolic links
A path name lookup involved more than 8 symbolic links.
- 63 ENAMETOOLONG File name too long
A component of a path name exceeded 255 characters, or an entire path name exceeded 1023 characters.
- 64 ENOTEMPTY Directory not empty
A directory with entries other than “.” and “..” was supplied to a remove directory or rename call.

DEFINITIONS

Process ID

Each active process in the system is uniquely identified by a positive integer called a process ID. The range of this ID is from 0 to {PROC_MAX}.

Parent process ID

A new process is created by a currently active process; see *fork(2)*. The parent process ID of a process is the process ID of its creator.

Process Group ID

Each active process is a member of a process group that is identified by a positive integer called the process group ID. This is the process ID of the group leader. This grouping permits the signalling of related processes (see *killpg(2)*) and the job control mechanisms of *cs(1)*.

Tty Group ID

Each active process can be a member of a terminal group that is identified by a positive integer called the tty group ID. This grouping is used to arbitrate between multiple jobs contending for the same terminal; see *cs(1)*, and *ty(4)*.

Real User ID and Real Group ID

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

Each user is also a member of one or more groups. One of these groups is distinguished from others and used in implementing accounting facilities. The positive integer corresponding to this distinguished group is termed the real group ID.

All processes have a real user ID and real group ID. These are initialized from the equivalent attributes of the process which created it.

Effective User Id, Effective Group Id, and Access Groups

Access to system resources is governed by three values: the effective user ID, the effective group ID, and the group access list.

The effective user ID and effective group ID are initially the process's real user ID and real group ID respectively. Either may be modified through execution of a set-user-ID or set-group-ID file (possibly by one its ancestors); see *execve(2)*.

The group access list is an additional set of group ID's used only in determining resource accessibility. Access checks are performed as described below in “File Access Permissions”.

Super-user

A process is recognized as a *super-user* process and is granted special privileges if its effective user ID is 0.

Special Processes

The processes with a process ID's of 0, 1, and 2 are special. Process 0 is the scheduler. Process 1 is the initialization process *init*, and is the ancestor of every other process in the system. It is used to control the process structure. Process 2 is the paging daemon.

Descriptor

An integer assigned by the system when a file is referenced by *open(2)*, *dup(2)*, or *pipe(2)* or a socket is referenced by *socket(2)* or *socketpair(2)* which uniquely identifies an access path to that file or socket from a given process or any of its children.

File Name

Names consisting of up to {FILENAME_MAX} characters may be used to name an ordinary file, special file, or directory.

These characters may be selected from the set of all ASCII character excluding 0 (null) and the ASCII code for / (slash). (The parity bit, bit 8, must be 0.)

Note that it is generally unwise to use *, ?, [or] as part of file names because of the special meaning attached to these characters by the shell.

Path Name

A path name is a null-terminated character string starting with an optional slash (/), followed by zero or more directory names separated by slashes, optionally followed by a file name. The total length of a path name must be less than {PATHNAME_MAX} characters.

If a path name begins with a slash, the path search begins at the *root* directory. Otherwise, the search begins from the current working directory. A slash by itself names the root directory. A null pathname refers to the current directory.

Directory

A directory is a special type of file which contains entries which are references to other files. Directory entries are called links. By convention, a directory contains at least two links, . and .., referred to as *dot* and *dot-dot* respectively. Dot refers to the directory itself and dot-dot refers to its parent directory.

Root Directory and Current Working Directory

Each process has associated with it a concept of a root directory and a current working directory for the purpose of resolving path name searches. A process's root directory need not be the root directory of the root file system.

File Access Permissions

Every file in the file system has a set of access permissions. These permissions are used in determining whether a process may perform a requested operation on the file (such as opening a file for writing). Access permissions are established at the time a file is created. They may be changed at some later time through the *chmod(2)* call.

File access is broken down according to whether a file may be: read, written, or executed. Directory files use the execute permission to control if the directory may be searched.

File access permissions are interpreted by the system as they apply to three different classes of users: the owner of the file, those users in the file's group, anyone else. Every file has an independent set of access permissions for each of these classes. When an access check is made, the system decides if permission should be granted by checking the access information applicable to the caller.

Read, write, and execute/search permissions on a file are granted to a process if:

The process's effective user ID is that of the super-user.

The process's effective user ID matches the user ID of the owner of the file and the owner permissions allow the access.

The process's effective user ID does not match the user ID of the owner of the file, and either the process's effective group ID matches the group ID of the file, or the group ID of the file is in the process's group access list, and the group permissions allow the access.

Neither the effective user ID nor effective group ID and group access list of the process match the corresponding user ID and group ID of the file, but the permissions for "other users" allow access.

Otherwise, permission is denied.

Sockets and Address Families

A socket is an endpoint for communication between processes. Each socket has queues for sending and receiving data.

Sockets are typed according to their communications properties. These properties include whether messages sent and received at a socket require the name of the partner, whether communication is reliable, the format used in naming message recipients, etc.

Each instance of the system supports some collection of socket types; consult *socket(2)* for more information about the types available and their properties.

Each instance of the system supports some number of sets of communications protocols. Each protocol set supports addresses of a certain format. An Address Family is the set of addresses for a specific group of protocols. Each socket has an address chosen from the address family in which the socket was created.

SEE ALSO

intro(3), perror(3)

NAME

`accept` — accept a connection on a socket

SYNOPSIS

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

ns = accept(s, addr, addrlen)
int ns, s;
struct sockaddr *addr;
int *addrlen;
```

DESCRIPTION

The argument *s* is a socket which has been created with *socket(2)*, bound to an address with *bind(2)*, and is listening for connections after a *listen(2)*. *Accept* extracts the first connection on the queue of pending connections, creates a new socket with the same properties of *s* and allocates a new file descriptor, *ns*, for the socket. If no pending connections are present on the queue, and the socket is not marked as non-blocking, *accept* blocks the caller until a connection is present. If the socket is marked non-blocking and no pending connections are present on the queue, *accept* returns an error as described below. The accepted socket, *ns*, may not be used to accept more connections. The original socket *s* remains open.

The argument *addr* is a result parameter which is filled in with the address of the connecting entity, as known to the communications layer. The exact format of the *addr* parameter is determined by the domain in which the communication is occurring. The *addrlen* is a value-result parameter; it should initially contain the amount of space pointed to by *addr*, on return it will contain the actual length (in bytes) of the address returned. This call is used with connection-based socket types, currently with `SOCK_STREAM`.

It is possible to *select(2)* a socket for the purposes of doing an *accept* by selecting it for read.

RETURN VALUE

The call returns `-1` on error. If it succeeds it returns a non-negative integer which is a descriptor for the accepted socket.

ERRORS

The *accept* will fail if:

- | | |
|---------------|--|
| [EBADF] | The descriptor is invalid. |
| [ENOTSOCK] | The descriptor references a file, not a socket. |
| [EOPNOTSUPP] | The referenced socket is not of type <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. |

SEE ALSO

bind(2), *connect(2)*, *listen(2)*, *select(2)*, *socket(2)*

NAME

`access` — determine accessibility of file

SYNOPSIS

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

#define R_OK    4    /* test for read permission */
#define W_OK    2    /* test for write permission */
#define X_OK    1    /* test for execute (search) permission */
#define F_OK    0    /* test for presence of file */

accessible = access(path, mode)
int accessible;
char *path;
int mode;
```

DESCRIPTION

`Access` checks the given file *path* for accessibility according to *mode*, which is an inclusive or of the bits `R_OK`, `W_OK` and `X_OK`. Specifying *mode* as `F_OK` (i.e. 0) tests whether the directories leading to the file can be searched and the file exists.

The real user ID and the group access list (including the real group ID) are used in verifying permission, so this call is useful to set-UID programs.

Notice that only access bits are checked. A directory may be indicated as writable by `access`, but an attempt to open it for writing will fail (although files may be created there); a file may look executable, but `execve` will fail unless it is in proper format.

RETURN VALUE

If *path* cannot be found or if any of the desired access modes would not be granted, then a `-1` value is returned; otherwise a `0` value is returned.

ERRORS

Access to the file is denied if one or more of the following are true:

- [ENOTDIR] A component of the path prefix is not a directory.
- [ENOENT] The argument path name was too long.
- [ENOENT] Read, write, or execute (search) permission is requested for a null path name or the named file does not exist.
- [EPERM] The argument contains a byte with the high-order bit set.
- [ELOOP] Too many symbolic links were encountered in translating the pathname.
- [EROFS] Write access is requested for a file on a read-only file system.
- [ETXTBSY] Write access is requested for a pure procedure (shared text) file that is being executed.
- [EACCES] Permission bits of the file mode do not permit the requested access; or search permission is denied on a component of the path prefix. The owner of a file has permission checked with respect to the "owner" read, write, and execute mode bits, members of the file's group other than the owner have permission checked with respect to the "group" mode bits, and all others have permissions checked with respect to the "other" mode bits.
- [EFAULT] *Path* points outside the process's allocated address space.

SEE ALSO

`chmod(2)`, `stat(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;
```

DESCRIPTION

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

NOTES

Binding a name in the UNIX domain creates a socket in the file system which must be deleted by the caller when it is no longer needed (using *unlink(2)*). The file created is a side-effect of the current implementation, and will not be created in future versions of the UNIX ipc domain.

The rules used in name binding vary between communication domains. Consult the manual entries in section 4 for detailed information.

RETURN VALUE

If the *bind* is successful, a 0 value is returned. A return value of `-1` indicates an error, which is further specified in the global *errno*.

ERRORS

The *bind* call will fail if:

[EBADF]	<i>S</i> is not a valid descriptor.
[ENOTSOCK]	<i>S</i> is not a socket.
[EADDRNOTAVAIL]	The specified address is not available from the local machine.
[EADDRINUSE]	The specified address is already in use.
[EINVAL]	The socket is already bound to an address.
[EACCESS]	The requested address is protected, and the current user has inadequate permission to access it.
[EFAULT]	The <i>name</i> parameter is not in a valid part of the user address space.

SEE ALSO

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

NAME

brk, *sbrk* — change data segment size

SYNOPSIS

```
caddr_t brk(addr)  
caddr_t addr;  
caddr_t sbrk(incr)  
int incr;
```

DESCRIPTION

Brk sets the system's idea of the lowest data segment location not used by the program (called the break) to *addr* (rounded up to the next multiple of the system's page size). Locations greater than *addr* and below the stack pointer are not in the address space and will thus cause a memory violation if accessed.

In the alternate function *sbrk*, *incr* more bytes are added to the program's data space and a pointer to the start of the new area is returned.

When a program begins execution via *execve* the break is set at the highest location defined by the program and data storage areas. Ordinarily, therefore, only programs with growing data areas need to use *sbrk*.

The *getrlimit*(2) system call may be used to determine the maximum permissible size of the *data* segment; it will not be possible to set the break beyond the *rlim_max* value returned from a call to *getrlimit*, e.g. "*etext* + *rlp* → *rlim_max*." (See *end*(3) for the definition of *etext*.)

RETURN VALUE

Zero is returned if the *brk* could be set; -1 if the program requests more memory than the system limit. *Sbrk* returns -1 if the break could not be set.

ERRORS

Sbrk will fail and no additional memory will be allocated if one of the following are true:

- [ENOMEM] The limit, as set by *setrlimit*(2), was exceeded.
- [ENOMEM] The maximum possible size of a data segment (compiled into the system) was exceeded.
- [ENOMEM] Insufficient space existed in the swap area to support the expansion.

SEE ALSO

execve(2), *getrlimit*(2), *malloc*(3), *end*(3)

BUGS

Setting the break may fail due to a temporary lack of swap space. It is not possible to distinguish this from a failure caused by exceeding the maximum size of the data segment without consulting *getrlimit*.

NAME

chdir — change current working directory

SYNOPSIS

chdir(path)

char *path;

DESCRIPTION

Path is the pathname of a directory. *Chdir* causes this directory to become the current working directory, the starting point for path names not beginning with “/”.

In order for a directory to become the current directory, a process must have execute (search) access to the directory.

RETURN VALUE

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

ERRORS

Chdir will fail and the current working directory will be unchanged if one or more of the following are true:

- | | |
|-----------|---|
| [ENOTDIR] | A component of the pathname is not a directory. |
| [ENOENT] | The named directory does not exist. |
| [ENOENT] | The argument path name was too long. |
| [EPERM] | The argument contains a byte with the high-order bit set. |
| [EACCES] | Search permission is denied for any component of the path name. |
| [EFAULT] | <i>Path</i> points outside the process's allocated address space. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |

SEE ALSO

chroot(2)

[EROFS] The file resides on a read-only file system.

SEE ALSO

open(2), chown(2)

NAME

`chmod` — change mode of file

SYNOPSIS

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

fchmod(fd, mode)
int fd, mode;
```

DESCRIPTION

The file whose name is given by *path* or referenced by the descriptor *fd* has its mode changed to *mode*. Modes are constructed by *or*'ing together some combination of the following:

```
04000 set user ID on execution
02000 set group ID on execution
01000 save text image after execution
00400 read by owner
00200 write by owner
00100 execute (search on directory) by owner
00070 read, write, execute (search) by group
00007 read, write, execute (search) by others
```

If an executable file is set up for sharing (this is the default) then mode 1000 prevents the system from abandoning the swap-space image of the program-text portion of the file when its last user terminates. Ability to set this bit is restricted to the super-user.

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

Writing or changing the owner of a file turns off the set-user-id and set-group-id bits. This makes the system somewhat more secure by protecting set-user-id (set-group-id) files from remaining set-user-id (set-group-id) if they are modified, at the expense of a degree of compatibility.

RETURN VALUE

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

ERRORS

Chmod will fail and the file mode will be unchanged if:

[EPERM]	The argument contains a byte with the high-order bit set.
[ENOTDIR]	A component of the path prefix is not a directory.
[ENOENT]	The pathname was too long.
[ENOENT]	The named file does not exist.
[EACCES]	Search permission is denied on a component of the path prefix.
[EPERM]	The effective user ID does not match the owner of the file and the effective user ID is not the super-user.
[EROFS]	The named file resides on a read-only file system.
[EFAULT]	<i>Path</i> points outside the process's allocated address space.
[ELOOP]	Too many symbolic links were encountered in translating the pathname.

Fchmod will fail if:

[EBADF]	The descriptor is not valid.
[EINVAL]	<i>Fd</i> refers to a socket, not to a file.

NAME

`chown` — change owner and group of a file

SYNOPSIS

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

fchown(fd, owner, group)
int fd, owner, group;
```

DESCRIPTION

The file which is named by *path* or referenced by *fd* has its *owner* and *group* changed as specified. Only the super-user may execute this call, because if users were able to give files away, they could defeat the file-space accounting procedures.

On some systems, *chown* clears the set-user-id and set-group-id bits on the file to prevent accidental creation of set-user-id and set-group-id programs owned by the super-user.

Fchown is particularly useful when used in conjunction with the file locking primitives (see *flock(2)*).

Only one of the owner and group id's may be set by specifying the other as `-1`.

RETURN VALUE

Zero is returned if the operation was successful; `-1` is returned if an error occurs, with a more specific error code being placed in the global variable *errno*.

ERRORS

Chown will fail and the file will be unchanged if:

- [EINVAL] The argument *path* does not refer to a file.
- [ENOTDIR] A component of the path prefix is not a directory.
- [ENOENT] The argument pathname is too long.
- [EPERM] The argument contains a byte with the high-order bit set.
- [ENOENT] The named file does not exist.
- [EACCES] Search permission is denied on a component of the path prefix.
- [EPERM] The effective user ID does not match the owner of the file and the effective user ID is not the super-user.
- [EROFS] The named file resides on a read-only file system.
- [EFAULT] *Path* points outside the process's allocated address space.
- [ELOOP] Too many symbolic links were encountered in translating the pathname.

Fchown will fail if:

- [EBADF] *Fd* does not refer to a valid descriptor.
- [EINVAL] *Fd* refers to a socket, not a file.

SEE ALSO

chmod(2), *flock(2)*

NAME

chroot — change root directory

SYNOPSIS

```
chroot(dirname)
char *dirname;
```

DESCRIPTION

Dirname is the address of the pathname of a directory, terminated by a null byte. *Chroot* causes this directory to become the root directory, the starting point for path names beginning with “/”.

In order for a directory to become the root directory a process must have execute (search) access to the directory.

This call is restricted to the super-user.

RETURN VALUE

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

ERRORS

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

- [ENOTDIR] A component of the path name is not a directory.
- [ENOENT] The pathname was too long.
- [EPERM] The argument contains a byte with the high-order bit set.
- [ENOENT] The named directory does not exist.
- [EACCES] Search permission is denied for any component of the path name.
- [EFAULT] *Path* points outside the process's allocated address space.
- [ELOOP] Too many symbolic links were encountered in translating the pathname.

SEE ALSO

chdir(2)

NAME

`close` — delete a descriptor

SYNOPSIS

```
close(d)  
int d;
```

DESCRIPTION

The `close` call deletes a descriptor from the per-process object reference table. If this is the last reference to the underlying object, then it will be deactivated. For example, on the last close of a file the current `seek` pointer associated with the file is lost; on the last close of a `socket(2)` associated naming information and queued data are discarded; on the last close of a file holding an advisory lock the lock is released; see further `flock(2)`.

A close of all of a process's descriptors is automatic on `exit`, but since there is a limit on the number of active descriptors per process, `close` is necessary for programs which deal with many descriptors.

When a process forks (see `fork(2)`), all descriptors for the new child process reference the same objects as they did in the parent before the fork. If a new process is then to be run using `execve(2)`, the process would normally inherit these descriptors. Most of the descriptors can be rearranged with `dup2(2)` or deleted with `close` before the `execve` is attempted, but if some of these descriptors will still be needed if the `execve` fails, it is necessary to arrange for them to be closed if the `execve` succeeds. For this reason, the call "`fcntl(d, F_SETFD, 1)`" is provided which arranges that a descriptor will be closed after a successful `execve`; the call "`fcntl(d, F_SETFD, 0)`" restores the default, which is to not close the descriptor.

RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of `-1` is returned and the global integer variable `errno` is set to indicate the error.

ERRORS

`Close` will fail if:

[EBADF] `D` is not an active descriptor.

SEE ALSO

`accept(2)`, `flock(2)`, `open(2)`, `pipe(2)`, `socket(2)`, `socketpair(2)`, `execve(2)`, `fcntl(2)`

NAME

connect — initiate a connection on a socket

SYNOPSIS

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

connect(s, name, namelen)
int s;
struct sockaddr *name;
int namelen;
```

DESCRIPTION

The parameter *s* is a socket. If it is of type `SOCK_DGRAM`, then this call permanently specifies the peer to which datagrams are to be sent; if it is of type `SOCK_STREAM`, then this call attempts to make a connection to another socket. The other socket is specified by *name* which is an address in the communications space of the socket. Each communications space interprets the *name* parameter in its own way.

RETURN VALUE

If the connection or binding succeeds, then 0 is returned. Otherwise a `-1` is returned, and a more specific error code is stored in *errno*.

ERRORS

The call fails if:

- | | |
|-----------------|--|
| [EBADF] | <i>S</i> is not a valid descriptor. |
| [ENOTSOCK] | <i>S</i> is a descriptor for a file, not a socket. |
| [EADDRNOTAVAIL] | The specified address is not available on this machine. |
| [EAFNOSUPPORT] | Addresses in the specified address family cannot be used with this socket. |
| [EISCONN] | The socket is already connected. |
| [ETIMEDOUT] | Connection establishment timed out without establishing a connection. |
| [ECONNREFUSED] | The attempt to connect was forcefully rejected. |
| [ENETUNREACH] | The network isn't reachable from this host. |
| [EADDRINUSE] | The address is already in use. |
| [EFAULT] | The <i>name</i> parameter specifies an area outside the process address space. |
| [EWOULDBLOCK] | The socket is non-blocking and the connection cannot be completed immediately. It is possible to <i>select(2)</i> the socket while it is connecting by selecting it for writing. |

SEE ALSO

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

NAME

`creat` — create a new file

SYNOPSIS

```
creat(name, mode)
char *name;
```

DESCRIPTION

This interface is obsoleted by `open(2)`.

`Creat` creates a new file or prepares to rewrite an existing file called *name*, given as the address of a null-terminated string. If the file did not exist, it is given mode *mode*, as modified by the process's mode mask (see `umask(2)`). Also see `chmod(2)` for the construction of the *mode* argument.

If the file did exist, its mode and owner remain unchanged but it is truncated to 0 length.

The file is also opened for writing, and its file descriptor is returned.

NOTES

The *mode* given is arbitrary; it need not allow writing. This feature has been used in the past by programs to construct a simple exclusive locking mechanism. It is replaced by the `O_EXCL` open mode, or `flock(2)` facility.

RETURN VALUE

The value `-1` is returned if an error occurs. Otherwise, the call returns a non-negative descriptor which only permits writing.

ERRORS

`Creat` will fail and the file will not be created or truncated if one of the following occur:

- | | |
|--------------|--|
| [EPERM] | The argument contains a byte with the high-order bit set. |
| [ENOTDIR] | A component of the path prefix is not a directory. |
| [EACCES] | A needed directory does not have search permission. |
| [EACCES] | The file does not exist and the directory in which it is to be created is not writable. |
| [EACCES] | The file exists, but it is unwritable. |
| [EISDIR] | The file is a directory. |
| [EMFILE] | There are already too many files open. |
| [EROFS] | The named file resides on a read-only file system. |
| [ENXIO] | The file is a character special or block special file, and the associated device does not exist. |
| [ETXTBSY] | The file is a pure procedure (shared text) file that is being executed. |
| [EFAULT] | <i>Name</i> points outside the process's allocated address space. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |
| [EOPNOTSUPP] | The file was a socket (not currently implemented). |

SEE ALSO

`open(2)`, `write(2)`, `close(2)`, `chmod(2)`, `umask(2)`

NAME

`dup`, `dup2` — duplicate a descriptor

SYNOPSIS

```
newd = dup(oldd)
```

```
int newd, oldd;
```

```
dup2(oldd, newd)
```

```
int oldd, newd;
```

DESCRIPTION

Dup duplicates an existing object descriptor. The argument *oldd* is a small non-negative integer index in the per-process descriptor table. The value must be less than the size of the table, which is returned by *getdtablesize(2)*. The new descriptor *newd* returned by the call is the lowest numbered descriptor which is not currently in use by the process.

The object referenced by the descriptor does not distinguish between references using *oldd* and *newd* in any way. Thus if *newd* and *oldd* are duplicate references to an open file, *read(2)*, *write(2)* and *lseek(2)* calls all move a single pointer into the file. If a separate pointer into the file is desired, a different object reference to the file must be obtained by issuing an additional *open(2)* call.

In the second form of the call, the value of *newd* desired is specified. If this descriptor is already in use, the descriptor is first deallocated as if a *close(2)* call had been done first.

RETURN VALUE

The value `-1` is returned if an error occurs in either call. The external variable *errno* indicates the cause of the error.

ERRORS

Dup and *dup2* fail if:

[EBADF] *Oldd* or *newd* is not a valid active descriptor

[EMFILE] Too many descriptors are active.

SEE ALSO

accept(2), *open(2)*, *close(2)*, *pipe(2)*, *socket(2)*, *socketpair(2)*, *getdtablesize(2)*

NAME

`execve` — execute a file

SYNOPSIS

```
execve(name, argv, envp)
char *name, *argv[], *envp[];
```

DESCRIPTION

Execve transforms the calling process into a new process. The new process is constructed from an ordinary file called the *new process file*. This file is either an executable object file, or a file of data for an interpreter. An executable object file consists of an identifying header, followed by pages of data representing the initial program (text) and initialized data pages. Additional pages may be specified by the header to be initialize with zero data. See *a.out(5)*.

An interpreter file begins with a line of the form “#! *interpreter*”; When an interpreter file is *execve*'d, the system *execve*'s the specified *interpreter*, giving it the name of the originally *exec*'d file as an argument, shifting over the rest of the original arguments.

There can be no return from a successful *execve* because the calling core image is lost. This is the mechanism whereby different process images become active.

The argument *argv* is an array of character pointers to null-terminated character strings. These strings constitute the argument list to be made available to the new process. By convention, at least one argument must be present in this array, and the first element of this array should be the name of the executed program (i.e. the last component of *name*).

The argument *envp* is also an array of character pointers to null-terminated strings. These strings pass information to the new process which are not directly arguments to the command, see *environ(7)*.

Descriptors open in the calling process remain open in the new process, except for those for which the close-on-exec flag is set; see *close(2)*. Descriptors which remain open are unaffected by *execve*.

Ignored signals remain ignored across an *execve*, but signals that are caught are reset to their default values. The signal stack is reset to be undefined; see *sigvec(2)* for more information.

Each process has *real* user and group IDs and a *effective* user and group IDs. The *real* ID identifies the person using the system; the *effective* ID determines his access privileges. *Execve* changes the effective user and group ID to the owner of the executed file if the file has the “set-user-ID” or “set-group-ID” modes. The *real* user ID is not affected.

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

process ID	see <i>getpid(2)</i>
parent process ID	see <i>getppid(2)</i>
process group ID	see <i>getpgrp(2)</i>
access groups	see <i>getgroups(2)</i>
working directory	see <i>chdir(2)</i>
root directory	see <i>chroot(2)</i>
control terminal	see <i>tty(4)</i>
resource usages	see <i>getrusage(2)</i>
interval timers	see <i>getitimer(2)</i>
resource limits	see <i>getrlimit(2)</i>
file mode mask	see <i>umask(2)</i>
signal mask	see <i>sigvec(2)</i>

When the executed program begins, it is called as follows:

```

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

```

where *argc* is the number of elements in *argv* (the “arg count”) and *argv* is the array of character pointers to the arguments themselves.

Envp is a pointer to an array of strings that constitute the *environment* of the process. A pointer to this array is also stored in the global variable “*environ*”. Each string consists of a name, an “=”, and a null-terminated value. The array of pointers is terminated by a null pointer. The shell *sh*(1) passes an environment entry for each global shell variable defined when the program is called. See *environ*(7) for some conventionally used names.

RETURN VALUE

If *execve* returns to the calling process an error has occurred; the return value will be -1 and the global variable *errno* will contain an error code.

ERRORS

Execve will fail and return to the calling process if one or more of the following are true:

- [ENOENT] One or more components of the new process file's path name do not exist.
- [ENOTDIR] A component of the new process file is not a directory.
- [EACCES] Search permission is denied for a directory listed in the new process file's path prefix.
- [EACCES] The new process file is not an ordinary file.
- [EACCES] The new process file mode denies execute permission.
- [ENOEXEC] The new process file has the appropriate access permission, but has an invalid magic number in its header.
- [ETXTBSY] The new process file is a pure procedure (shared text) file that is currently open for writing or reading by some process.
- [ENOMEM] The new process requires more virtual memory than is allowed by the imposed maximum (*getrlimit*(2)).
- [E2BIG] The number of bytes in the new process's argument list is larger than the system-imposed limit of {*ARG_MAX*} bytes.
- [EFAULT] The new process file is not as long as indicated by the size values in its header.
- [EFAULT] *Path*, *argv*, or *envp* point to an illegal address.

CAVEATS

If a program is *setuid* to a non-super-user, but is executed when the real *uid* is “root”, then the program has the powers of a super-user as well.

SEE ALSO

exit(2), *fork*(2), *execl*(3), *environ*(7)

NAME

`_exit` — terminate a process

SYNOPSIS

```
_exit(status)
int status;
```

DESCRIPTION

`_exit` terminates a process with the following consequences:

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

If the parent process of the calling process is executing a `wait` or is interested in the SIGCHLD signal, then it is notified of the calling process's termination and the low-order eight bits of `status` are made available to it; see `wait(2)`.

The parent process ID of all of the calling process's existing child processes are also set to 1. This means that the initialization process (see `intro(2)`) inherits each of these processes as well.

Most C programs call the library routine `exit(3)` which performs cleanup actions in the standard i/o library before calling `_exit`.

RETURN VALUE

This call never returns.

SEE ALSO

`fork(2)`, `wait(2)`, `exit(3)`

NAME

`flock` — apply or remove an advisory lock on an open file

SYNOPSIS

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

#define LOCK_SH  1    /* shared lock */
#define LOCK_EX  2    /* exclusive lock */
#define LOCK_NB  4    /* don't block when locking */
#define LOCK_UN  8    /* unlock */

flock(fd, operation)
int fd, operation;
```

DESCRIPTION

Flock applies or removes an *advisory* lock on the file associated with the file descriptor *fd*. A lock is applied by specifying an *operation* parameter which is the inclusive or of `LOCK_SH` or `LOCK_EX` and, possibly, `LOCK_NB`. To unlock an existing lock *operation* should be `LOCK_UN`.

Advisory locks allow cooperating processes to perform consistent operations on files, but do not guarantee consistency (i.e. processes may still access files without using advisory locks possibly resulting in inconsistencies).

The locking mechanism allows two types of locks: *shared* locks and *exclusive* locks. At any time multiple shared locks may be applied to a file, but at no time are multiple exclusive, or both shared and exclusive, locks allowed simultaneously on a file.

A shared lock may be *upgraded* to an exclusive lock, and vice versa, simply by specifying the appropriate lock type; this results in the previous lock being released and the new lock applied (possibly after other processes have gained and released the lock).

Requesting a lock on an object which is already locked normally causes the caller to be blocked until the lock may be acquired. If `LOCK_NB` is included in *operation*, then this will not happen; instead the call will fail and the error `EWouldBlock` will be returned.

NOTES

Locks are on files, not file descriptors. That is, file descriptors duplicated through `dup(2)` or `fork(2)` do not result in multiple instances of a lock, but rather multiple references to a single lock. If a process holding a lock on a file forks and the child explicitly unlocks the file, the parent will lose its lock.

Processes blocked awaiting a lock may be awakened by signals.

RETURN VALUE

Zero is returned if the operation was successful; on an error a `-1` is returned and an error code is left in the global location `errno`.

ERRORS

The *flock* call fails if:

- [`EWouldBlock`] The file is locked and the `LOCK_NB` option was specified.
- [`EBADF`] The argument *fd* is an invalid descriptor.
- [`EINVAL`] The argument *fd* refers to an object other than a file.

SEE ALSO

`open(2)`, `close(2)`, `dup(2)`, `execve(2)`, `fork(2)`

NAME

`fork` — create a new process

SYNOPSIS

```
pid = fork()
int pid;
```

DESCRIPTION

Fork causes creation of a new process. The new process (child process) is an exact copy of the calling process except for the following:

The child process has a unique process ID.

The child process has a different parent process ID (i.e., the process ID of the parent process).

The child process has its own copy of the parent's descriptors. These descriptors reference the same underlying objects, so that, for instance, file pointers in file objects are shared between the child and the parent, so that a *lseek(2)* on a descriptor in the child process can affect a subsequent *read* or *write* by the parent. This descriptor copying is also used by the shell to establish standard input and output for newly created processes as well as to set up pipes.

The child processes resource utilizations are set to 0; see *setrlimit(2)*.

RETURN VALUE

Upon successful completion, *fork* returns a value of 0 to the child process and returns the process ID of the child process to the parent process. Otherwise, a value of -1 is returned to the parent process, no child process is created, and the global variable *errno* is set to indicate the error.

ERRORS

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

- [EAGAIN] The system-imposed limit {PROC_MAX} on the total number of processes under execution would be exceeded.
- [EAGAIN] The system-imposed limit {KID_MAX} on the total number of processes under execution by a single user would be exceeded.

SEE ALSO

execve(2), *wait(2)*

NAME

fsync — synchronize a file's in-core state with that on disk

SYNOPSIS

```
fsync(fd)  
int fd;
```

DESCRIPTION

Fsync causes all modified data and attributes of *fd* to be moved to a permanent storage device. This normally results in all in-core modified copies of buffers for the associated file to be written to a disk.

Fsync should be used by programs which require a file to be in a known state; for example in building a simple transaction facility.

RETURN VALUE

A 0 value is returned on success. A -1 value indicates an error.

ERRORS

The *fsync* fails if:

[EBADF] *Fd* is not a valid descriptor.

[EINVAL] *Fd* refers to a socket, not to a file.

SEE ALSO

sync(2), *sync*(8), *update*(8)

BUGS

The current implementation of this call is expensive for large files.

NAME

getdtablesize — get descriptor table size

SYNOPSIS

```
nds = getdtablesize()
int nds;
```

DESCRIPTION

Each process has a fixed size descriptor table which is guaranteed to have at least 20 slots. The entries in the descriptor table are numbered with small integers starting at 0. The call *getdtablesize* returns the size of this table.

SEE ALSO

close(2), dup(2), open(2)

NAME

`getgid`, `getegid` — get group identity

SYNOPSIS

```
gid = getgid()
```

```
int gid;
```

```
egid = getegid()
```

```
int egid;
```

DESCRIPTION

Getgid returns the real group ID of the current process, *getegid* the effective group ID.

The real group ID is specified at login time.

The effective group ID is more transient, and determines additional access permission during execution of a “set-group-ID” process, and it is for such processes that *getgid* is most useful.

SEE ALSO

`getuid(2)`, `setregid(2)`, `setgid(3)`

NAME

getgroups — get group access list

SYNOPSIS

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

getgroups(ngroups, gidset)
int *ngroups, *gidset;
```

DESCRIPTION

Getgroups gets the current group access list of the user process and stores it in the array *gidset*. The parameter *ngroups* indicates the number of entries which may be placed in *gidset* and is modified on return to indicate the actual number of groups returned. No more than NGRPS, as defined in *<sys/param.h>*, will ever be returned.

RETURN VALUE

A value of 0 indicates that the call succeeded, and that the number of elements of *gidset* and the set itself were returned. A value of -1 indicates that an error occurred, and the error code is stored in the global variable *errno*.

ERRORS

The possible errors for *getgroup* are:

[EFAULT] The arguments *ngroups* or *gidset* specify invalid addresses.

SEE ALSO

setgroups(2), initgroups(3)

NAME

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

SYNOPSIS

```
hostid = gethostid()
int hostid;

sethostid(hostid)
int hostid;
```

DESCRIPTION

Sethostid establishes a 32-bit identifier for the current processor which is intended to be unique among all UNIX systems in existence. This is normally a DARPA Internet address for the local machine. This call is allowed only to the super-user and is normally performed at boot time.

Gethostid returns the 32-bit identifier for the current processor.

SEE ALSO

hostid(1), gethostname(2)

BUGS

32 bits for the identifier is too small.

NAME

gethostname, sethostname — get/set name of current host

SYNOPSIS

```
gethostname(name, namelen)
char *name;
int namelen;

sethostname(name, namelen)
char *name;
int namelen;
```

DESCRIPTION

Gethostname returns the standard host name for the current processor, as previously set by *sethostname*. The parameter *namelen* specifies the size of the *name* array. The returned name is null-terminated unless insufficient space is provided.

Sethostname sets the name of the host machine to be *name*, which has length *namelen*. This call is restricted to the super-user and is normally used only when the system is bootstrapped.

RETURN VALUE

If the call succeeds a value of 0 is returned. If the call fails, then a value of -1 is returned and an error code is placed into the global location *errno*.

ERRORS

The following errors may be returned by these calls:

[EFAULT] The *name* or *namelen* parameter gave an invalid address.
[EPERM] The caller was not the super-user.

SEE ALSO

gethostid(2)

BUGS

Host names are limited to 255 characters.

NAME

getitimer, setitimer — get/set value of interval timer

SYNOPSIS

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

#define ITIMER_REAL    0    /* real time intervals */
#define ITIMER_VIRTUAL 1    /* virtual time intervals */
#define ITIMER_PROF    2    /* user and system virtual time */

getitimer(which, value)
int which;
struct itimerval *value;

setitimer(which, value, ovalue)
int which;
struct itimerval *value, *ovalue;
```

DESCRIPTION

The system provides each process with three interval timers, defined in `<sys/time.h>`. The `getitimer` call returns the current value for the timer specified in `which`, while the `setitimer` call sets the value of a timer (optionally returning the previous value of the timer).

A timer value is defined by the `itimerval` structure:

```
struct itimerval {
    struct timeval it_interval;    /* timer interval */
    struct timeval it_value;      /* current value */
};
```

If `it_value` is non-zero, it indicates the time to the next timer expiration. If `it_interval` is non-zero, it specifies a value to be used in reloading `it_value` when the timer expires. Setting `it_value` to 0 disables a timer. Setting `it_interval` to 0 causes a timer to be disabled after its next expiration (assuming `it_value` is non-zero).

Time values smaller than the resolution of the system clock are rounded up to this resolution (on the VAX, 10 microseconds).

The `ITIMER_REAL` timer decrements in real time. A `SIGALRM` signal is delivered when this timer expires.

The `ITIMER_VIRTUAL` timer decrements in process virtual time. It runs only when the process is executing. A `SIGVTALRM` signal is delivered when it expires.

The `ITIMER_PROF` timer decrements both in process virtual time and when the system is running on behalf of the process. It is designed to be used by interpreters in statistically profiling the execution of interpreted programs. Each time the `ITIMER_PROF` timer expires, the `SIGPROF` signal is delivered. Because this signal may interrupt in-progress system calls, programs using this timer must be prepared to restart interrupted system calls.

NOTES

Three macros for manipulating time values are defined in `<sys/time.h>`. `Timerclear` sets a time value to zero, `timerisset` tests if a time value is non-zero, and `timercmp` compares two time values (beware that `>=` and `<=` do not work with this macro).

RETURN VALUE

If the calls succeed, a value of 0 is returned. If an error occurs, the value `-1` is returned, and a more precise error code is placed in the global variable `errno`.

ERRORS

The possible errors are:

[EFAULT] The *value* structure specified a bad address.

[EINVAL] A *value* structure specified a time was too large to be handled.

SEE ALSO

sigvec(2), gettimeofday(2)

NAME

getpagesize — get system page size

SYNOPSIS

```
pagesize = getpagesize()
int pagesize;
```

DESCRIPTION

Getpagesize returns the number of bytes in a page. Page granularity is the granularity of many of the memory management calls.

The page size is a *system* page size and may not be the same as the underlying hardware page size.

SEE ALSO

sbrk(2), pagesize(1)

NAME

getpgrp — get process group

SYNOPSIS

```
pgrp = getpgrp(pid)
int pgrp;
int pid;
```

DESCRIPTION

The process group of the specified process is returned by *getpgrp*. If *pid* is zero, then the call applies to the current process.

Process groups are used for distribution of signals, and by terminals to arbitrate requests for their input: processes which have the same process group as the terminal are foreground and may read, while others will block with a signal if they attempt to read.

This call is thus used by programs such as *cs**h*(1) to create process groups in implementing job control. The *TIOCGPGRP* and *TIOCSPGRP* calls described in *tty*(4) are used to get/set the process group of the control terminal.

SEE ALSO

setpgrp(2), getuid(2), tty(4)

NAME

`getpid`, `getppid` — get process identification

SYNOPSIS

```
pid = getpid()  
long pid;
```

```
ppid = getppid()  
long ppid;
```

DESCRIPTION

Getpid returns the process ID of the current process. Most often it is used with the host identifier *gethostid(2)* to generate uniquely-named temporary files.

Getppid returns the process ID of the parent of the current process.

SEE ALSO

gethostid(2)

NAME

getpriority, setpriority — get/set program scheduling priority

SYNOPSIS

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

#define PRIO_PROCESS    0    /* process */
#define PRIO_PGRP      1    /* process group */
#define PRIO_USER      2    /* user id */

prio = getpriority(which, who)
int prio, which, who;

setpriority(which, who, prio)
int which, who, prio;
```

DESCRIPTION

The scheduling priority of the process, process group, or user, as indicated by *which* and *who* is obtained with the *getpriority* call and set with the *setpriority* call. *Which* is one of PRIO_PROCESS, PRIO_PGRP, or PRIO_USER, and *who* is interpreted relative to *which* (a process identifier for PRIO_PROCESS, process group identifier for PRIO_PGRP, and a user ID for PRIO_USER). *Prio* is a value in the range -20 to 20 . The default priority is 0 ; lower priorities cause more favorable scheduling.

The *getpriority* call returns the highest priority (lowest numerical value) enjoyed by any of the specified processes. The *setpriority* call sets the priorities of all of the specified processes to the specified value. Only the super-user may lower priorities.

RETURN VALUE

Since *getpriority* can legitimately return the value -1 , it is necessary to clear the external variable *errno* prior to the call, then check it afterward to determine if a -1 is an error or a legitimate value. The *setpriority* call returns 0 if there is no error, or -1 if there is.

ERRORS

Getpriority and *setpriority* may return one of the following errors:

- [ESRCH] No process(es) were located using the *which* and *who* values specified.
- [EINVAL] *Which* was not one of PRIO_PROCESS, PRIO_PGRP, or PRIO_USER.

In addition to the errors indicated above, *setpriority* may fail with one of the following errors returned:

- [EACCES] A process was located, but neither its effective nor real user ID matched the effective user ID of the caller.
- [EACCES] A non super-user attempted to change a process priority to a negative value.

SEE ALSO

nice(1), fork(2), renice(8)

NAME

getrlimit, setrlimit — control maximum system resource consumption

SYNOPSIS

```
#include <sys/time.h>
#include <sys/resource.h>

getrlimit(resource, rlp)
int resource;
struct rlimit *rlp;

setrlimit(resource, rlp)
int resource;
struct rlimit *rlp;
```

DESCRIPTION

Limits on the consumption of system resources by the current process and each process it creates may be obtained with the *getrlimit* call, and set with the *setrlimit* call.

The *resource* parameter is one of the following:

- RLIMIT_CPU the maximum amount of cpu time (in milliseconds) to be used by each process.
- RLIMIT_FSIZE the largest size, in bytes, of any single file which may be created.
- RLIMIT_DATA the maximum size, in bytes, of the data segment for a process; this defines how far a program may extend its break with the *sbrk(2)* system call.
- RLIMIT_STACK the maximum size, in bytes, of the stack segment for a process; this defines how far a program's stack segment may be extended, either automatically by the system, or explicitly by a user with the *sbrk(2)* system call.
- RLIMIT_CORE the largest size, in bytes, of a *core* file which may be created.
- RLIMIT_RSS the maximum size, in bytes, a process's resident set size may grow to. This imposes a limit on the amount of physical memory to be given to a process; if memory is tight, the system will prefer to take memory from processes which are exceeding their declared resident set size.

A resource limit is specified as a soft limit and a hard limit. When a soft limit is exceeded a process may receive a signal (for example, if the cpu time is exceeded), but it will be allowed to continue execution until it reaches the hard limit (or modifies its resource limit). The *rlimit* structure is used to specify the hard and soft limits on a resource,

```
struct rlimit {
    int    rlim_cur;    /* current (soft) limit */
    int    rlim_max;    /* hard limit */
};
```

Only the super-user may raise the maximum limits. Other users may only alter *rlim_cur* within the range from 0 to *rlim_max* or (irreversibly) lower *rlim_max*.

An "infinite" value for a limit is defined as RLIMIT_INFINITY (0x7fffffff).

Because this information is stored in the per-process information, this system call must be executed directly by the shell if it is to affect all future processes created by the shell; *limit* is thus a built-in command to *cs(1)*.

The system refuses to extend the data or stack space when the limits would be exceeded in the normal way: a *break* call fails if the data space limit is reached, or the process is killed when the stack limit is reached (since the stack cannot be extended, there is no way to send a signal!).

A file i/o operation which would create a file which is too large will cause a signal SIGXFSZ to be generated, this normally terminates the process, but may be caught. When the soft cpu time limit is exceeded, a signal SIGXCPU is sent to the offending process.

RETURN VALUE

A 0 return value indicates that the call succeeded, changing or returning the resource limit. A return value of -1 indicates that an error occurred, and an error code is stored in the global location *errno*.

ERRORS

The possible errors are:

- [EFAULT] The address specified for *rlp* is invalid.
- [EPERM] The limit specified to *setrlimit* would have raised the maximum limit value, and the caller is not the super-user.

SEE ALSO

csh(1), *quota*(2)

BUGS

There should be *limit* and *unlimit* commands in *sh*(1) as well as in *csh*.

NAME

getrusage — get information about resource utilization

SYNOPSIS

```
#include <sys/time.h>
#include <sys/resource.h>

#define RUSAGE_SELF      0      /* calling process */
#define RUSAGE_CHILDREN -1     /* terminated child processes */

getrusage(who, rusage)
int who;
struct rusage *rusage;
```

DESCRIPTION

Getrusage returns information describing the resources utilized by the current process, or all its terminated child processes. The *who* parameter is one of `RUSAGE_SELF` and `RUSAGE_CHILDREN`. If *rusage* is non-zero, the buffer it points to will be filled in with the following structure:

```
struct rusage {
    struct timeval ru_utime;    /* user time used */
    struct timeval ru_stime;    /* system time used */
    int ru_maxrss;
    int ru_ixrss;              /* integral shared memory size */
    int ru_idrss;              /* integral unshared data size */
    int ru_isrss;              /* integral unshared stack size */
    int ru_minflt;            /* page reclaims */
    int ru_majflt;            /* page faults */
    int ru_nswap;              /* swaps */
    int ru_inblock;            /* block input operations */
    int ru_oublock;            /* block output operations */
    int ru_msgsnd;             /* messages sent */
    int ru_msrvcv;             /* messages received */
    int ru_nsignals;           /* signals received */
    int ru_nvcsw;              /* voluntary context switches */
    int ru_nivcsw;             /* involuntary context switches */
};
```

The fields are interpreted as follows:

<code>ru_utime</code>	the total amount of time spent executing in user mode.
<code>ru_stime</code>	the total amount of time spent in the system executing on behalf of the process(es).
<code>ru_maxrss</code>	the maximum resident set size utilized (in kilobytes).
<code>ru_ixrss</code>	an “integral” value indicating the amount of memory used which was also shared among other processes. This value is expressed in units of kilobytes * seconds-of-execution and is calculated by summing the number of shared memory pages in use each time the internal system clock ticks and then averaging over 1 second intervals.
<code>ru_idrss</code>	an integral value of the amount of unshared memory residing in the data segment of a process (expressed in units of kilobytes * seconds-of-execution).
<code>ru_isrss</code>	an integral value of the amount of unshared memory residing in the stack segment of a process (expressed in units of kilobytes * seconds-of-execution).
<code>ru_minflt</code>	the number of page faults serviced without any i/o activity; here i/o activity is

avoided by "reclaiming" a page frame from the list of pages awaiting reallocation.

<code>ru_majflt</code>	the number of page faults serviced which required i/o activity.
<code>ru_nswap</code>	the number of times a process was "swapped" out of main memory.
<code>ru_inblock</code>	the number of times the file system had to perform input.
<code>ru_outblock</code>	the number of times the file system had to perform output.
<code>ru_msgsnd</code>	the number of ipc messages sent.
<code>ru_msgrcv</code>	the number of ipc messages received.
<code>ru_nsignals</code>	the number of signals delivered.
<code>ru_nvcsw</code>	the number of times a context switch resulted due to a process voluntarily giving up the processor before its time slice was completed (usually to await availability of a resource).
<code>ru_nivcsw</code>	the number of times a context switch resulted due to a higher priority process becoming runnable or because the current process exceeded its time slice.

NOTES

The numbers `ru_inblock` and `ru_outblock` account only for real i/o; data supplied by the caching mechanism is charged only to the first process to read or write the data.

SEE ALSO

`gettimeofday(2)`, `wait(2)`

BUGS

There is no way to obtain information about a child process which has not yet terminated.

NAME

getsockopt, setsockopt — get and set options on sockets

SYNOPSIS

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

```
getsockopt(s, level, optname, optval, optlen)
int s, level, optname;
char *optval;
int *optlen;
```

```
setsockopt(s, level, optname, optval, optlen)
int s, level, optname;
char *optval;
int optlen;
```

DESCRIPTION

Getsockopt and *setsockopt* manipulate *options* associated with a socket. Options may exist at multiple protocol levels; they are always present at the uppermost “socket” level.

When manipulating socket options the level at which the option resides and the name of the option must be specified. To manipulate options at the “socket” level, *level* is specified as SOL_SOCKET. To manipulate options at any other level the protocol number of the appropriate protocol controlling the option is supplied. For example, to indicate an option is to be interpreted by the TCP protocol, *level* should be set to the protocol number of TCP; see *getprotoent(3N)*.

The parameters *optval* and *optlen* are used to access option values for *setsockopt*. For *getsockopt* they identify a buffer in which the value for the requested option(s) are to be returned. For *getsockopt*, *optlen* is a value-result parameter, initially containing the size of the buffer pointed to by *optval*, and modified on return to indicate the actual size of the value returned. If no option value is to be supplied or returned, *optval* may be supplied as 0.

Optname and any specified options are passed uninterpreted to the appropriate protocol module for interpretation. The include file *<sys/socket.h>* contains definitions for “socket” level options; see *socket(2)*. Options at other protocol levels vary in format and name, consult the appropriate entries in (4P).

RETURN VALUE

A 0 is returned if the call succeeds, -1 if it fails.

ERRORS

The call succeeds unless:

[EBADF]	The argument <i>s</i> is not a valid descriptor.
[ENOTSOCK]	The argument <i>s</i> is a file, not a socket.
[ENOPROTOOPT]	The option is unknown.
[EFAULT]	The options are not in a valid part of the process address space.

SEE ALSO

socket(2), *getprotoent(3N)*

NAME

gettimeofday, settimeofday — get/set date and time

SYNOPSIS

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

```
gettimeofday(tp, tzp)
struct timeval *tp;
struct timezone *tzp;
```

```
settimeofday(tp, tzp)
struct timeval *tp;
struct timezone *tzp;
```

DESCRIPTION

Gettimeofday returns the system's notion of the current Greenwich time and the current time zone. Time returned is expressed relative in seconds and microseconds since midnight January 1, 1970.

The structures pointed to by *tp* and *tzp* are defined in *<sys/time.h>* as:

```
struct timeval {
    u_long tv_sec;          /* seconds since Jan. 1, 1970 */
    long tv_usec;         /* and microseconds */
};

struct timezone {
    int tz_minuteswest; /* of Greenwich */
    int tz_dsttime;     /* type of dst correction to apply */
};
```

The *timezone* structure indicates the local time zone (measured in minutes of time westward from Greenwich), and a flag that, if nonzero, indicates that Daylight Saving time applies locally during the appropriate part of the year.

Only the super-user may set the time of day.

RETURN

A 0 return value indicates that the call succeeded. A -1 return value indicates an error occurred, and in this case an error code is stored into the global variable *errno*.

ERRORS

The following error codes may be set in *errno*:

```
[EFAULT]    An argument address referenced invalid memory.
[EPERM]     A user other than the super-user attempted to set the time.
```

SEE ALSO

date(1), ctime(3)

BUGS

Time is never correct enough to believe the microsecond values. There should a mechanism by which, at least, local clusters of systems might synchronize their clocks to millisecond granularity.

NAME

`getuid`, `geteuid` — get user identity

SYNOPSIS

```
uid = getuid()
int uid;

euid = geteuid()
int euid;
```

DESCRIPTION

Getuid returns the real user ID of the current process, *geteuid* the effective user ID.

The real user ID identifies the person who is logged in. The effective user ID gives the process additional permissions during execution of “set-user-ID” mode processes, which use *getuid* to determine the real-user-id of the process which invoked them.

SEE ALSO

`getgid(2)`, `setreuid(2)`

NAME

`ioctl` — control device

SYNOPSIS

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

```
ioctl(d, request, argp)
```

```
int d, request;
```

```
char *argp;
```

DESCRIPTION

ioctl performs a variety of functions on open descriptors. In particular, many operating characteristics of character special files (e.g. terminals) may be controlled with *ioctl* requests. The writeups of various devices in section 4 discuss how *ioctl* applies to them.

An *ioctl request* has encoded in it whether the argument is an “in” parameter or “out” parameter, and the size of the argument *argp* in bytes. Macros and defines used in specifying an *ioctl request* are located in the file `<sys/ioctl.h>`.

RETURN VALUE

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

ERRORS

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

[EBADF] *D* is not a valid descriptor.

[ENOTTY] *D* is not associated with a character special device.

[ENOTTY] The specified request does not apply to the kind of object which the descriptor *d* references.

[EINVAL] *Request* or *argp* is not valid.

SEE ALSO

`execve(2)`, `fcntl(2)`, `mt(4)`, `tty(4)`, `intro(4N)`

NAME

kill — send signal to a process

SYNOPSIS

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

DESCRIPTION

Kill sends the signal *sig* to a process, specified by the process number *pid*. *Sig* may be one of the signals specified in *sigvec(2)*, or it may be 0, in which case error checking is performed but no signal is actually sent. This can be used to check the validity of *pid*.

The sending and receiving processes must have the same effective user ID, otherwise this call is restricted to the super-user. A single exception is the signal SIGCONT which may always be sent to any child or grandchild of the current process.

If the process number is 0, the signal is sent to all other processes in the sender's process group; this is a variant of *killpg(2)*.

If the process number is -1, and the user is the super-user, the signal is broadcast universally except to system processes and the process sending the signal.

Processes may send signals to themselves.

RETURN VALUE

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

ERRORS

Kill will fail and no signal will be sent if any of the following occur:

- [EINVAL] *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 its effective user id does not match the effective user-id of the receiving process.

SEE ALSO

getpid(2), getpgrp(2), killpg(2), sigvec(2)

NAME

`killpg` — send signal to a process group

SYNOPSIS

```
killpg(pgrp, sig)
int pgrp, sig;
```

DESCRIPTION

Killpg sends the signal *sig* to the process group *pgrp*. See *sigvec(2)* for a list of signals.

The sending process and members of the process group must have the same effective user ID, otherwise this call is restricted to the super-user. As a single special case the continue signal SIGCONT may be sent to any process which is a descendant of the current process.

RETURN VALUE

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

ERRORS

Killpg will fail and no signal will be sent if any of the following occur:

- [EINVAL] *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.

SEE ALSO

`kill(2)`, `getpgrp(2)`, `sigvec(2)`

NAME

link — make a hard link to a file

SYNOPSIS

```
link(name1, name2)
char *name1, *name2;
```

DESCRIPTION

A hard link to *name1* is created; the link has the name *name2*. *Name1* must exist.

With hard links, both *name1* and *name2* must be in the same file system. Unless the caller is the super-user, *name1* must not be a directory. Both the old and the new *link* share equal access and rights to the underlying object.

RETURN VALUE

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

ERRORS

Link will fail and no link will be created if one or more of the following are true:

- | | |
|-----------|--|
| [EPERM] | Either pathname contains a byte with the high-order bit set. |
| [ENOENT] | Either pathname was too long. |
| [ENOTDIR] | A component of either path prefix is not a directory. |
| [ENOENT] | A component of either path prefix does not exist. |
| [EACCES] | A component of either path prefix denies search permission. |
| [ENOENT] | The file named by <i>name1</i> does not exist. |
| [EEXIST] | The link named by <i>name2</i> does exist. |
| [EPERM] | The file named by <i>name1</i> is a directory and the effective user ID is not super-user. |
| [EXDEV] | The link named by <i>name2</i> and the file named by <i>name1</i> are on different file systems. |
| [EACCES] | The requested link requires writing in a directory with a mode that denies write permission. |
| [EROFS] | The requested link requires writing in a directory on a read-only file system. |
| [EFAULT] | One of the pathnames specified is outside the process's allocated address space. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |

SEE ALSO

symlink(2), unlink(2)

NAME

`listen` — listen for connections on a socket.

SYNOPSIS

```
listen(s, backlog)  
int s, backlog;
```

DESCRIPTION

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

SEE ALSO

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

BUGS

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

NAME

`lseek` — move read/write pointer

SYNOPSIS

```
#define L_SET    0    /* set the seek pointer */
#define L_INCR  1    /* increment the seek pointer */
#define L_XTND  2    /* extend the file size */

pos = lseek(d, offset, whence)
int pos;
int d, offset, whence;
```

DESCRIPTION

The descriptor *d* refers to a file or device open for reading and/or writing. *Lseek* sets the file pointer of *d* as follows:

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

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

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

Upon successful completion, the resulting pointer location as measured in bytes from beginning of the file is returned. Some devices are incapable of seeking. The value of the pointer associated with such a device is undefined.

NOTES

Seeking far beyond the end of a file, then writing, creates a gap or “hole”, which occupies no physical space and reads as zeros.

RETURN VALUE

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

ERRORS

Lseek will fail and the file pointer will remain unchanged if:

- [EBADF] *Fildes* is not an open file descriptor.
- [ESPIPE] *Fildes* is associated with a pipe or a socket.
- [EINVAL] *Whence* is not a proper value.
- [EINVAL] The resulting file pointer would be negative.

SEE ALSO

`dup(2)`, `open(2)`

BUGS

This document's use of *whence* is incorrect English, but maintained for historical reasons.

NAME

`mkdir` — make a directory file

SYNOPSIS

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

DESCRIPTION

Mkdir creates a new directory file with name *path*. The mode of the new file is initialized from *mode*. (The protection part of the mode is modified by the process's mode mask; see *umask(2)*).

The directory's owner ID is set to the process's effective user ID. The directory's group ID is set to that of the parent directory in which it is created.

The low-order 9 bits of mode are modified by the process's file mode creation mask: all bits set in the process's file mode creation mask are cleared. See *umask(2)*.

RETURN VALUE

A 0 return value indicates success. A -1 return value indicates an error, and an error code is stored in *errno*.

ERRORS

Mkdir will fail and no directory will be created if:

- | | |
|-----------|---|
| [EPERM] | The process's effective user ID is not super-user. |
| [EPERM] | The <i>path</i> argument contains a byte with the high-order bit set. |
| [ENOTDIR] | A component of the path prefix is not a directory. |
| [ENOENT] | A component of the path prefix does not exist. |
| [EROFS] | The named file resides on a read-only file system. |
| [EEXIST] | The named file exists. |
| [EFAULT] | <i>Path</i> points outside the process's allocated address space. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |
| [EIO] | An I/O error occurred while writing to the file system. |

SEE ALSO

chmod(2), *stat(2)*, *umask(2)*

NAME

`mknod` — make a special file

SYNOPSIS

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

DESCRIPTION

Mknod creates a new file whose name is *path*. The mode of the new file (including special file bits) is initialized from *mode*. (The protection part of the mode is modified by the process's mode mask; see *umask*(2)). The first block pointer of the i-node is initialized from *dev* and is used to specify which device the special file refers to.

If *mode* indicates a block or character special file, *dev* is a configuration dependent specification of a character or block I/O device. If *mode* does not indicate a block special or character special device, *dev* is ignored.

Mknod may be invoked only by the super-user.

RETURN VALUE

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

ERRORS

Mknod will fail and the file mode will be unchanged if:

- | | |
|-----------|---|
| [EPERM] | The process's effective user ID is not super-user. |
| [EPERM] | The pathname contains a character with the high-order bit set. |
| [ENOTDIR] | A component of the path prefix is not a directory. |
| [ENOENT] | A component of the path prefix does not exist. |
| [EROFS] | The named file resides on a read-only file system. |
| [EEXIST] | The named file exists. |
| [EFAULT] | <i>Path</i> points outside the process's allocated address space. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |

SEE ALSO

chmod(2), *stat*(2), *umask*(2)

NAME

mount, umount — mount or remove file system

SYNOPSIS

```
mount(special, name, rwflag)
char *special, *name;
int rwflag;

umount(special)
char *special;
```

DESCRIPTION

Mount announces to the system that a removable file system has been mounted on the block-structured special file *special*; from now on, references to file *name* will refer to the root file on the newly mounted file system. *Special* and *name* are pointers to null-terminated strings containing the appropriate path names.

Name must exist already. *Name* must be a directory. Its old contents are inaccessible while the file system is mounted.

The *rwflag* argument determines whether the file system can be written on; if it is 0 writing is allowed, if non-zero no writing is done. Physically write-protected and magnetic tape file systems must be mounted read-only or errors will occur when access times are updated, whether or not any explicit write is attempted.

Umount announces to the system that the *special* file is no longer to contain a removable file system. The associated file reverts to its ordinary interpretation.

RETURN VALUE

Mount returns 0 if the action occurred, -1 if *special* is inaccessible or not an appropriate file, if *name* does not exist, if *special* is already mounted, if *name* is in use, or if there are already too many file systems mounted.

Umount returns 0 if the action occurred; -1 if the special file is inaccessible or does not have a mounted file system, or if there are active files in the mounted file system.

ERRORS

Mount will fail when one of the following occurs:

- | | |
|-----------|---|
| [NODEV] | The caller is not the super-user. |
| [NODEV] | <i>Special</i> does not exist. |
| [ENOTBLK] | <i>Special</i> is not a block device. |
| [ENXIO] | The major device number of <i>special</i> is out of range (this indicates no device driver exists for the associated hardware). |
| [EPERM] | The pathname contains a character with the high-order bit set. |
| [ENOTDIR] | A component of the path prefix in <i>name</i> is not a directory. |
| [EROFS] | <i>Name</i> resides on a read-only file system. |
| [EBUSY] | <i>Name</i> is not a directory, or another process currently holds a reference to it. |
| [EBUSY] | No space remains in the mount table. |
| [EBUSY] | The super block for the file system had a bad magic number or an out of range block size. |
| [EBUSY] | Not enough memory was available to read the cylinder group information for the file system. |
| [EBUSY] | An i/o error occurred while reading the super block or cylinder group information. |

Umount may fail with one of the following errors:

- [NODEV] The caller is not the super-user.
- [NODEV] *Special* does not exist.
- [ENOTBLK] *Special* is not a block device.
- [ENXIO] The major device number of *special* is out of range (this indicates no device driver exists for the associated hardware).
- [EINVAL] The requested device is not in the mount table.
- [EBUSY] A process is holding a reference to a file located on the file system.

SEE ALSO

mount(8), umount(8)

BUGS

The error codes are in a state of disarray; too many errors appear to the caller as one value.

NAME

`open` — open a file for reading or writing, or create a new file

SYNOPSIS

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

```
open(path, flags, mode)
```

```
char *path;
```

```
int flags, mode;
```

DESCRIPTION

`Open` opens the file *path* for reading and/or writing, as specified by the *flags* argument and returns a descriptor for that file. The *flags* argument may indicate the file is to be created if it does not already exist (by specifying the `O_CREAT` flag), in which case the file is created with mode *mode* as described in `chmod(2)` and modified by the process' umask value (see `umask(2)`).

Path is the address of a string of ASCII characters representing a path name, terminated by a null character. The flags specified are formed by *or*'ing the following values

<code>O_RDONLY</code>	open for reading only
<code>O_WRONLY</code>	open for writing only
<code>O_RDWR</code>	open for reading and writing
<code>O_NDELAY</code>	do not block on open
<code>O_APPEND</code>	append on each write
<code>O_CREAT</code>	create file if it does not exist
<code>O_TRUNC</code>	truncate size to 0
<code>O_EXCL</code>	error if create and file exists

Opening a file with `O_APPEND` set causes each write on the file to be appended to the end. If `O_TRUNC` is specified and the file exists, the file is truncated to zero length. If `O_EXCL` is set with `O_CREAT`, then if the file already exists, the open returns an error. This can be used to implement a simple exclusive access locking mechanism. If the `O_NDELAY` flag is specified and the open call would result in the process being blocked for some reason (e.g. waiting for carrier on a dialup line), the open returns immediately. The first time the process attempts to perform i/o on the open file it will block (not currently implemented).

Upon successful completion a non-negative integer termed a file descriptor is returned. The file pointer used to mark the current position within the file is set to the beginning of the file.

The new descriptor is set to remain open across `execve` system calls; see `close(2)`.

No process may have more than `{OPEN_MAX}` file descriptors open simultaneously.

ERRORS

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

[EPERM]	The pathname contains a character with the high-order bit set.
[ENOTDIR]	A component of the path prefix is not a directory.
[ENOENT]	<code>O_CREAT</code> is not set and the named file does not exist.
[EACCES]	A component of the path prefix denies search permission.
[EACCES]	The required permissions (for reading and/or writing) are denied for the named flag.
[EISDIR]	The named file is a directory, and the arguments specify it is to be opened for writing.
[EROFS]	The named file resides on a read-only file system, and the file is to be modified.

- [EMFILE] {OPEN_MAX} file descriptors are currently open.
- [ENXIO] The named file is a character special or block special file, and the device associated with this special file does not exist.
- [ETXTBSY] The file is a pure procedure (shared text) file that is being executed and the *open* call requests write access.
- [EFAULT] *Path* points outside the process's allocated address space.
- [ELOOP] Too many symbolic links were encountered in translating the pathname.
- [EEXIST] O_EXCL was specified and the file exists.
- [ENXIO] The O_NDELAY flag is given, and the file is a communications device on which there is no carrier present.
- [EOPNOTSUPP] An attempt was made to open a socket (not currently implemented).

SEE ALSO

chmod(2), close(2), dup(2), lseek(2), read(2), write(2), umask(2)

NAME

pipe — create an interprocess communication channel

SYNOPSIS

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

DESCRIPTION

The *pipe* system call creates an I/O mechanism called a pipe. The file descriptors returned can be used in read and write operations. When the pipe is written using the descriptor *fildes*[1] up to 4096 bytes of data are buffered before the writing process is suspended. A read using the descriptor *fildes*[0] will pick up the data.

It is assumed that after the pipe has been set up, two (or more) cooperating processes (created by subsequent *fork* calls) will pass data through the pipe with *read* and *write* calls.

The shell has a syntax to set up a linear array of processes connected by pipes.

Read calls on an empty pipe (no buffered data) with only one end (all write file descriptors closed) returns an end-of-file.

Pipes are really a special case of the *socketpair*(2) call and, in fact, are implemented as such in the system.

A signal is generated if a write on a pipe with only one end is attempted.

RETURN VALUE

The function value zero is returned if the pipe was created; -1 if an error occurred.

ERRORS

The *pipe* call will fail if:

[EMFILE] Too many descriptors are active.

[EFAULT] The *fildes* buffer is in an invalid area of the process's address space.

SEE ALSO

sh(1), read(2), write(2), fork(2), socketpair(2)

BUGS

Should more than 4096 bytes be necessary in any pipe among a loop of processes, deadlock will occur.

NAME

`profil` — execution time profile

SYNOPSIS

```
profil(buff, bufsiz, offset, scale)  
char *buff;  
int bufsiz, offset, scale;
```

DESCRIPTION

Buff points to an area of core whose length (in bytes) is given by *bufsiz*. After this call, the user's program counter (*pc*) is examined each clock tick (10 milliseconds); *offset* is subtracted from it, and the result multiplied by *scale*. If the resulting number corresponds to a word inside *buff*, that word is incremented.

The scale is interpreted as an unsigned, fixed-point fraction with binary point at the left: 0x10000 gives a 1-1 mapping of *pc*'s to words in *buff*; 0x8000 maps each pair of instruction words together. 0x2 maps all instructions onto the beginning of *buff* (producing a non-interrupting core clock).

Profiling is turned off by giving a *scale* of 0 or 1. It is rendered ineffective by giving a *bufsiz* of 0. Profiling is turned off when an *execve* is executed, but remains on in child and parent both after a *fork*. Profiling is turned off if an update in *buff* would cause a memory fault.

RETURN VALUE

A 0, indicating success, is always returned.

SEE ALSO

`gprof`(1), `setitimer`(2), `monitor`(3)

NAME

`ptrace` — process trace

SYNOPSIS

```
#include <signal.h>
```

```
ptrace(request, pid, addr, data)
```

```
int request, pid, *addr, data;
```

DESCRIPTION

Ptrace provides a means by which a parent process may control the execution of a child process, and examine and change its core image. Its primary use is for the implementation of breakpoint debugging. There are four arguments whose interpretation depends on a *request* argument. Generally, *pid* is the process ID of the traced process, which must be a child (no more distant descendant) of the tracing process. A process being traced behaves normally until it encounters some signal whether internally generated like “illegal instruction” or externally generated like “interrupt”. See *sigvec(2)* for the list. Then the traced process enters a stopped state and its parent is notified via *wait(2)*. When the child is in the stopped state, its core image can be examined and modified using *ptrace*. If desired, another *ptrace* request can then cause the child either to terminate or to continue, possibly ignoring the signal.

The value of the *request* argument determines the precise action of the call:

- 0 This request is the only one used by the child process; it declares that the process is to be traced by its parent. All the other arguments are ignored. Peculiar results will ensue if the parent does not expect to trace the child.
- 1,2 The word in the child process's address space at *addr* is returned. If I and D space are separated (e.g. historically on a pdp-11), request 1 indicates I space, 2 D space. *Addr* must be even. The child must be stopped. The input *data* is ignored.
- 3 The word of the system's per-process data area corresponding to *addr* is returned. *Addr* must be even and less than 512. This space contains the registers and other information about the process; its layout corresponds to the *user* structure in the system.
- 4,5 The given *data* is written at the word in the process's address space corresponding to *addr*, which must be even. No useful value is returned. If I and D space are separated, request 4 indicates I space, 5 D space. Attempts to write in pure procedure fail if another process is executing the same file.
- 6 The process's system data is written, as it is read with request 3. Only a few locations can be written in this way: the general registers, the floating point status and registers, and certain bits of the processor status word.
- 7 The *data* argument is taken as a signal number and the child's execution continues at location *addr* as if it had incurred that signal. Normally the signal number will be either 0 to indicate that the signal that caused the stop should be ignored, or that value fetched out of the process's image indicating which signal caused the stop. If *addr* is (int *)1 then execution continues from where it stopped.
- 8 The traced process terminates.
- 9 Execution continues as in request 7; however, as soon as possible after execution of at least one instruction, execution stops again. The signal number from the stop is SIGTRAP. (On the VAX-11 the T-bit is used and just one instruction is executed.) This is part of the mechanism for implementing breakpoints.

As indicated, these calls (except for request 0) can be used only when the subject process has stopped. The *wait* call is used to determine when a process stops; in such a case the “termination” status returned by *wait* has the value 0177 to indicate stoppage rather than genuine termination.

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

On a VAX-11, "word" also means a 32-bit integer, but the "even" restriction does not apply.

RETURN VALUE

A 0 value is returned if the call succeeds. If the call fails then a -1 is returned and the global variable *errno* is set to indicate the error.

ERRORS

[EINVAL]	The request code is invalid.
[EINVAL]	The specified process does not exist.
[EINVAL]	The given signal number is invalid.
[EFAULT]	The specified address is out of bounds.
[EPERM]	The specified process cannot be traced.

SEE ALSO

wait(2), *sigvec(2)*, *adb(1)*

BUGS

Ptrace is unique and arcane; it should be replaced with a special file which can be opened and read and written. The control functions could then be implemented with *ioctl(2)* calls on this file. This would be simpler to understand and have much higher performance.

The request 0 call should be able to specify signals which are to be treated normally and not cause a stop. In this way, for example, programs with simulated floating point (which use "illegal instruction" signals at a very high rate) could be efficiently debugged.

The error indication, -1 , is a legitimate function value; *errno*, see *intro(2)*, can be used to disambiguate.

It should be possible to stop a process on occurrence of a system call; in this way a completely controlled environment could be provided.

NAME

read, readv — read input

SYNOPSIS

```
cc = read(d, buf, nbytes)
int cc, d;
char *buf;
int nbytes;

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

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

DESCRIPTION

Read attempts to read *nbytes* of data from the object referenced by the descriptor *d* into the buffer pointed to by *buf*. *Readv* performs the same action, but scatters the input data into the *iovcnt* buffers specified by the members of the *iovec* array: *iov*[0], *iov*[1], ..., *iov*[*iovcnt* - 1].

For *readv*, the *iovec* structure is defined as

```
struct iovec {
    caddr_t iov_base;
    int     iov_len;
};
```

Each *iovec* entry specifies the base address and length of an area in memory where data should be placed. *Readv* will always fill an area completely before proceeding to the next.

On objects capable of seeking, the *read* starts at a position given by the pointer associated with *d*, see *lseek*(2). Upon return from *read*, the pointer is incremented by the number of bytes actually read.

Objects that are not capable of seeking always read from the current position. The value of the pointer associated with such a object is undefined.

Upon successful completion, *read* and *readv* return the number of bytes actually read and placed in the buffer. The system guarantees to read the number of bytes requested if the descriptor references a file which has that many bytes left before the end-of-file, but in no other cases.

If the returned value is 0, then end-of-file has been reached.

RETURN VALUE

If successful, the number of bytes actually read is returned. Otherwise, a -1 is returned and the global variable *errno* is set to indicate the error.

ERRORS

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

- [EBADF] *Fildes* is not a valid file descriptor open for reading.
- [EFAULT] *Buf* points outside the allocated address space.
- [EINTR] A read from a slow device was interrupted before any data arrived by the delivery of a signal.

In addition, *readv* may return one of the following errors:

- [EINVAL] *iovcnt* was less than or equal to 0, or greater than 16.
- [EINVAL] One of the *iov_len* values in the *iov* array was negative.

[EINVAL] The sum of the *iov_len* values in the *iov* array overflowed a 32-bit integer.

SEE ALSO

dup(2), open(2), pipe(2), socket(2), socketpair(2)

NAME

readlink — read value of a symbolic link

SYNOPSIS

```
cc = readlink(path, buf, bufsiz)
int cc;
char *path, *buf;
int bufsiz;
```

DESCRIPTION

Readlink places the contents of the symbolic link *name* in the buffer *buf* which has size *bufsiz*. The contents of the link are not null terminated when returned.

RETURN VALUE

The call returns the count of characters placed in the buffer if it succeeds, or a `-1` if an error occurs, placing the error code in the global variable *errno*.

ERRORS

Readlink will fail and the file mode will be unchanged if:

- [EPERM] The *path* argument contained a byte with the high-order bit set.
- [ENOENT] The pathname was too long.
- [ENOTDIR] A component of the path prefix is not a directory.
- [ENOENT] The named file does not exist.
- [ENXIO] The named file is not a symbolic link.
- [EACCES] Search permission is denied on a component of the path prefix.
- [EPERM] The effective user ID does not match the owner of the file and the effective user ID is not the super-user.
- [EINVAL] The named file is not a symbolic link.
- [EFAULT] *Buf* extends outside the process's allocated address space.
- [ELOOP] Too many symbolic links were encountered in translating the pathname.

SEE ALSO

stat(2), lstat(2), symlink(2)

NAME

reboot — reboot system or halt processor

SYNOPSIS

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

reboot(howto)
int howto;
```

DESCRIPTION

Reboot reboots the system, and is invoked automatically in the event of unrecoverable system failures. *Howto* is a mask of options passed to the bootstrap program. The system call interface permits only RB_HALT or RB_AUTOBOOT to be passed to the reboot program; the other flags are used in scripts stored on the console storage media, or used in manual bootstrap procedures. When none of these options (e.g. RB_AUTOBOOT) is given, the system is rebooted from file "vmunix" in the root file system of unit 0 of a disk chosen in a processor specific way. An automatic consistency check of the disks is then normally performed.

The bits of *howto* are:

RB_HALT

the processor is simply halted; no reboot takes place. RB_HALT should be used with caution.

RB_ASKNAME

Interpreted by the bootstrap program itself, causing it to inquire as to what file should be booted. Normally, the system is booted from the file "xx(0,0)vmunix" without asking.

RB_SINGLE

Normally, the reboot procedure involves an automatic disk consistency check and then multi-user operations. RB_SINGLE prevents the consistency check, rather simply booting the system with a single-user shell on the console. RB_SINGLE is interpreted by the *init(8)* program in the newly booted system. This switch is not available from the system call interface.

Only the super-user may *reboot* a machine.

RETURN VALUES

If successful, this call never returns. Otherwise, a -1 is returned and an error is returned in the global variable *errno*.

ERRORS

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

SEE ALSO

crash(8), halt(8), init(8), reboot(8)

BUGS

The notion of "console medium", among other things, is specific to the VAX.

NAME

`recv`, `recvfrom`, `recvmsg` — receive a message from a socket

SYNOPSIS

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

cc = recv(s, buf, len, flags)
int cc, s;
char *buf;
int len, flags;

cc = recvfrom(s, buf, len, flags, from, fromlen)
int cc, s;
char *buf;
int len, flags;
struct sockaddr *from;
int *fromlen;

cc = recvmsg(s, msg, flags)
int cc, s;
struct msghdr msg[];
int flags;
```

DESCRIPTION

`Recv`, `recvfrom`, and `recvmsg` are used to receive messages from a socket.

The `recv` call may be used only on a *connected* socket (see `connect(2)`), while `recvfrom` and `recvmsg` may be used to receive data on a socket whether it is in a connected state or not.

If `from` is non-zero, the source address of the message is filled in. `Fromlen` is a value-result parameter, initialized to the size of the buffer associated with `from`, and modified on return to indicate the actual size of the address stored there. The length of the message is returned in `cc`. If a message is too long to fit in the supplied buffer, excess bytes may be discarded depending on the type of socket the message is received from; see `socket(2)`.

If no messages are available at the socket, the receive call waits for a message to arrive, unless the socket is nonblocking (see `ioctl(2)`) in which case a `cc` of `-1` is returned with the external variable `errno` set to `EWOULDBLOCK`.

The `select(2)` call may be used to determine when more data arrives.

The `flags` argument to a send call is formed by *or*'ing one or more of the values,

```
#define MSG_PEEK    0x1    /* peek at incoming message */
#define MSG_OOB    0x2    /* process out-of-band data */
```

The `recvmsg` call uses a `msghdr` structure to minimize the number of directly supplied parameters. This structure has the following form, as defined in `<sys/socket.h>`:

```
struct msghdr {
    caddr_t msg_name;           /* optional address */
    int msg_namelen;          /* size of address */
    struct iovec *msg_iov;     /* scatter/gather array */
    int msg_iovlen;           /* # elements in msg_iov */
    caddr_t msg_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, as described in *read(2)*. Access rights to be sent along with the message are specified in *msg_accrights*, which has length *msg_accrightslen*.

RETURN VALUE

These calls return the number of bytes received, or -1 if an error occurred.

ERRORS

The calls fail if:

- | | |
|---------------|---|
| [EBADF] | The argument <i>s</i> is an invalid descriptor. |
| [ENOTSOCK] | The argument <i>s</i> is not a socket. |
| [EWOULDBLOCK] | The socket is marked non-blocking and the receive operation would block. |
| [EINTR] | The receive was interrupted by delivery of a signal before any data was available for the receive. |
| [EFAULT] | The data was specified to be received into a non-existent or protected part of the process address space. |

SEE ALSO

read(2), *send(2)*, *socket(2)*

NAME

rename — change the name of a file

SYNOPSIS

```
rename(from, to)
char *from, *to;
```

DESCRIPTION

Rename causes the link named *from* to be renamed as *to*. If *to* exists, then it is first removed. Both *from* and *to* must be of the same type (that is, both directories or both non-directories), and must reside on the same file system.

Rename guarantees that an instance of *to* will always exist, even if the system should crash in the middle of the operation.

CAVEAT

The system can deadlock if a loop in the file system graph is present. This loop takes the form of an entry in directory “a”, say “a/foo”, being a hard link to directory “b”, and an entry in directory “b”, say “b/bar”, being a hard link to directory “a”. When such a loop exists and two separate processes attempt to perform “rename a/foo b/bar” and “rename b/bar a/foo”, respectively, the system may deadlock attempting to lock both directories for modification. Hard links to directories should be replaced by symbolic links by the system administrator.

RETURN VALUE

A 0 value is returned if the operation succeeds, otherwise *rename* returns -1 and the global variable *errno* indicates the reason for the failure.

ERRORS

Rename will fail and neither of the argument files will be affected if any of the following are true:

- [ENOTDIR] A component of either path prefix is not a directory.
- [ENOENT] A component of either path prefix does not exist.
- [EACCES] A component of either path prefix denies search permission.
- [ENOENT] The file named by *from* does not exist.
- [EPERM] The file named by *from* is a directory and the effective user ID is not super-user.
- [EXDEV] The link named by *to* and the file named by *from* are on different logical devices (file systems). Note that this error code will not be returned if the implementation permits cross-device links.
- [EACCES] The requested link requires writing in a directory with a mode that denies write permission.
- [EROFS] The requested link requires writing in a directory on a read-only file system.
- [EFAULT] *Path* points outside the process's allocated address space.
- [EINVAL] *From* is a parent directory of *to*.

SEE ALSO

open(2)

NAME

`rmdir` — remove a directory file

SYNOPSIS

```
rmdir(path)
char *path;
```

DESCRIPTION

Rmdir removes a directory file whose name is given by *path*. The directory must not have any entries other than “.” and “..”.

RETURN VALUE

A 0 is returned if the remove succeeds; otherwise a `-1` is returned and an error code is stored in the global location *errno*.

ERRORS

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

- [ENOTEMPTY] The named directory contains files other than “.” and “..” in it.
- [EPERM] The pathname contains a character with the high-order bit set.
- [ENOENT] The pathname was too long.
- [ENOTDIR] A component of the path prefix is not a directory.
- [ENOENT] The named file does not exist.
- [EACCES] A component of the path prefix denies search permission.
- [EACCES] Write permission is denied on the directory containing the link to be removed.
- [EBUSY] The directory to be removed is the mount point for a mounted file system.
- [EROFS] The directory entry to be removed resides on a read-only file system.
- [EFAULT] *Path* points outside the process's allocated address space.
- [ELOOP] Too many symbolic links were encountered in translating the pathname.

SEE ALSO

`mkdir(2)`, `unlink(2)`

NAME

`send`, `sendto`, `sendmsg` — send a message from a socket

SYNOPSIS

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

cc = send(s, msg, len, flags)
int cc, s;
char *msg;
int len, flags;

cc = sendto(s, msg, len, flags, to, tolen)
int cc, s;
char *msg;
int len, flags;
struct sockaddr *to;
int tolen;

cc = sendmsg(s, msg, flags)
int cc, s;
struct msghdr msg[];
int flags;
```

DESCRIPTION

Send, *sendto*, and *sendmsg* are used to transmit a message to another socket. *Send* may be used only when the socket is in a *connected* state, while *sendto* and *sendmsg* may be used at any time.

The address of the target is given by *to* with *tolen* specifying its size. The length of the message is given by *len*. If the message is too long to pass atomically through the underlying protocol, then the error EMSGSIZE is returned, and the message is not transmitted.

No indication of failure to deliver is implicit in a *send*. Return values of -1 indicate some locally detected errors.

If no messages space is available at the socket to hold the message to be transmitted, then *send* normally blocks, unless the socket has been placed in non-blocking i/o mode. The *select(2)* call may be used to determine when it is possible to send more data.

The *flags* parameter may be set to SOF_OOB to send “out-of-band” data on sockets which support this notion (e.g. SOCK_STREAM).

See *recv(2)* for a description of the *msghdr* structure.

RETURN VALUE

The call returns the number of characters sent, or -1 if an error occurred.

ERRORS

[EBADF]	An invalid descriptor was specified.
[ENOTSOCK]	The argument <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.

SEE ALSO

recv(2), *socket(2)*

NAME

setgroups — set group access list

SYNOPSIS

```
#include <sys/param.h>
setgroups(ngroups, gidset)
int ngroups, *gidset;
```

DESCRIPTION

Setgroups sets the group access list of the current user process according to the array *gidset*. The parameter *ngroups* indicates the number of entries in the array and must be no more than *NGRPS*, as defined in *<sys/param.h>*.

Only the super-user may set new groups.

RETURN VALUE

A 0 value is returned on success, -1 on error, with a error code stored in *errno*.

ERRORS

The *setgroups* call will fail if:

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

[EFAULT] The address specified for *gidset* is outside the process address space.

SEE ALSO

getgroups(2), initgroups(3X)

NAME

setpgrp — set process group

SYNOPSIS

```
setpgrp(pid, pgrp)
int pid, pgrp;
```

DESCRIPTION

Setpgrp sets the process group of the specified process *pid* to the specified *pgrp*. If *pid* is zero, then the call applies to the current process.

If the invoker is not the super-user, then the affected process must have the same effective user-id as the invoker or be a descendant of the invoking process.

RETURN VALUE

Setpgrp returns when the operation was successful. If the request failed, -1 is returned and the global variable *errno* indicates the reason.

ERRORS

Setpgrp will fail and the process group will not be altered if one of the following occur:

- | | |
|---------|---|
| [ESRCH] | The requested process does not exist. |
| [EPERM] | The effective user ID of the requested process is different from that of the caller and the process is not a descendent of the calling process. |

SEE ALSO

getpgrp(2)

NAME

setregid — set real and effective group ID

SYNOPSIS

```
setregid(rgid, egid)
int rgid, egid;
```

DESCRIPTION

The real and effective group ID's of the current process are set to the arguments. Only the super-user may change the real group ID of a process. Unprivileged users may change the effective group ID to the real group ID, but to no other.

Supplying a value of -1 for either the real or effective group ID forces the system to substitute the current ID in place of the -1 parameter.

RETURN VALUE

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

ERRORS

[EPERM] The current process is not the super-user and a change other than changing the effective group-id to the real group-id was specified.

SEE ALSO

getgid(2), setreuid(2), setgid(3)

NAME

setreuid — set real and effective user ID's

SYNOPSIS

```
setreuid(ruid, euid)
int ruid, euid;
```

DESCRIPTION

The real and effective user ID's of the current process are set according to the arguments. If *ruid* or *euid* is -1 , the current uid is filled in by the system. Only the super-user may modify the real uid of a process. Users other than the super-user may change the effective uid of a process only to the real uid.

RETURN VALUE

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

ERRORS

[EPERM] The current process is not the super-user and a change other than changing the effective user-id to the real user-id was specified.

SEE ALSO

getuid(2), setregid(2), setuid(3)

NAME

`shutdown` — shut down part of a full-duplex connection

SYNOPSIS

```
shutdown(s, how)
int s, how;
```

DESCRIPTION

The *shutdown* call causes all or part of a full-duplex connection on the socket associated with *s* to be shut down. If *how* is 0, then further receives will be disallowed. If *how* is 1, then further sends will be disallowed. If *how* is 2, then further sends and receives will be disallowed.

DIAGNOSTICS

A 0 is returned if the call succeeds, -1 if it fails.

ERRORS

The call succeeds unless:

- [EBADF] *S* is not a valid descriptor.
- [ENOTSOCK] *S* is a file, not a socket.
- [ENOTCONN] The specified socket is not connected.

SEE ALSO

`connect(2)`, `socket(2)`

NAME

socket — create an endpoint for communication

SYNOPSIS

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

s = socket(af, type, protocol)
int s, af, type, protocol;
```

DESCRIPTION

Socket creates an endpoint for communication and returns a descriptor.

The *af* parameter specifies an address format with which addresses specified in later operations using the socket should be interpreted. These formats are defined in the include file *<sys/socket.h>*. The currently understood formats are

AF_UNIX	(UNIX path names),
AF_INET	(ARPA Internet addresses),
AF_PUP	(Xerox PUP-I Internet addresses), 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(3N)* and *protocols(3N)*.

Sockets of type SOCK_STREAM are full-duplex byte streams, similar to pipes. A stream socket must be in a *connected* state before any data may be sent or received on it. A connection to another socket is created with a *connect(2)* call. Once connected, data may be transferred using *read(2)* and *write(2)* calls or some variant of the *send(2)* and *recv(2)* calls. When a session has been completed a *close(2)* may be performed. Out-of-band data may also be transmitted as described in *send(2)* and received as described in *recv(2)*.

The communications protocols used to implement a SOCK_STREAM insure that data is not lost or duplicated. If a piece of data for which the peer protocol has buffer space cannot be successfully transmitted within a reasonable length of time, then the connection is considered broken and calls will indicate an error with -1 returns and with ETIMEDOUT as the specific code in the global variable *errno*. The protocols optionally keep sockets "warm" by forcing transmissions roughly every minute in the absence of other activity. An error is then indicated if no response can be elicited on an otherwise idle connection for 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(2)` calls. It is also possible to receive datagrams at such a socket with `recv(2)`.

An `fcntl(2)` call can be used to specify a process group to receive a `SIGURG` signal when the out-of-band data arrives.

The operation of sockets is controlled by socket level *options*. These options are defined in the file `<sys/socket.h>` and explained below. `setsockopt` and `getsockopt(2)` are used to set and get options, respectively.

<code>SO_DEBUG</code>	turn on recording of debugging information
<code>SO_REUSEADDR</code>	allow local address reuse
<code>SO_KEEPALIVE</code>	keep connections alive
<code>SO_DONTROUTE</code>	do not 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(2)` call should allow reuse of local addresses. `SO_KEEPALIVE` enables the periodic transmission of messages on a connected socket. Should the connected party fail to respond to these messages, the connection is considered broken and processes using the socket are notified via a `SIGPIPE` signal. `SO_DONTROUTE` indicates that outgoing messages should bypass the standard routing facilities. Instead, messages are directed to the appropriate network interface according to the network portion of the destination address. `SO_LINGER` and `SO_DONTLINGER` control the actions taken when unsent messages are queued on socket and a `close(2)` is performed. If the socket promises reliable delivery of data and `SO_LINGER` is set, the system will block the process on the `close` attempt until it is able to transmit the data or until it decides it is unable to deliver the information (a timeout period, termed the linger interval, is specified in the `setsockopt` call when `SO_LINGER` is requested). If `SO_DONTLINGER` is specified and a `close` is issued, the system will process the close in a manner which allows the process to continue as quickly as possible.

RETURN VALUE

A `-1` is returned if an error occurs, otherwise the return value is a descriptor referencing the socket.

ERRORS

The `socket` call fails if:

- [`EAFNOSUPPORT`] The specified address family is not supported in this version of the system.
- [`ESOCKTNOSUPPORT`] The specified socket type is not supported in this address family.
- [`EPROTONOSUPPORT`] The specified protocol is not supported.
- [`EMFILE`] The per-process descriptor table is full.
- [`ENOBUFS`] No buffer space is available. The socket cannot be created.

SEE ALSO

`accept(2)`, `bind(2)`, `connect(2)`, `getsockname(2)`, `getsockopt(2)`, `ioctl(2)`, `listen(2)`, `recv(2)`, `select(2)`, `send(2)`, `shutdown(2)`, `socketpair(2)`
 "A 4.2BSD Interprocess Communication Primer".

BUGS

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

NAME

stat, lstat, fstat — get file status

SYNOPSIS

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

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

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

```
lstat(path, buf)
char *path;
struct stat *buf;
```

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

DESCRIPTION

Stat obtains information about the file *path*. Read, write or execute permission of the named file is not required, but all directories listed in the path name leading to the file must be reachable.

Lstat is like *stat* except in the case where the named file is a symbolic link, in which case *lstat* returns information about the link, while *stat* returns information about the file the link references.

Fstat obtains the same information about an open file referenced by the argument descriptor, such as would be obtained by an *open* call.

Buf is a pointer to a *stat* structure into which information is placed concerning the file. The contents of the structure pointed to by *buf*

```
struct stat {
    dev_t    st_dev;    /* device inode resides on */
    ino_t    st_ino;   /* this inode's number */
    u_short  st_mode;   /* protection */
    short    st_nlink; /* number or hard links to the file */
    short    st_uid;   /* user-id of owner */
    short    st_gid;   /* group-id of owner */
    dev_t    st_rdev;  /* the device type, for inode that is device */
    off_t    st_size;  /* total size of file */
    time_t   st_atime; /* file last access time */
    int      st_spare1;
    time_t   st_mtime; /* file last modify time */
    int      st_spare2;
    time_t   st_ctime; /* file last status change time */
    int      st_spare3;
    long     st_blksize; /* optimal blocksize for file system i/o ops */
    long     st_blocks; /* actual number of blocks allocated */
    long     st_spare4[2];
};
```

st_atime Time when file data was last read or modified. Changed by the following system calls: *mknod(2)*, *utimes(2)*, *read(2)*, and *write(2)*. For reasons of efficiency, *st_atime* is not set when a directory is searched, although this would be more logical.

`st_mtime` Time when data was last modified. It is not set by changes of owner, group, link count, or mode. Changed by the following system calls: `mknod(2)`, `utimes(2)`, `write(2)`.

`st_ctime` Time when file status was last changed. It is set both both by writing and changing the i-node. Changed by the following system calls: `chmod(2)`, `chown(2)`, `link(2)`, `mknod(2)`, `unlink(2)`, `utimes(2)`, `write(2)`.

The status information word `st_mode` has bits:

```
#define S_IFMT      0170000 /* type of file */
#define S_IFDIR    0040000 /* directory */
#define S_IFCHR    0020000 /* character special */
#define S_IFBLK    0060000 /* block special */
#define S_IFREG    0100000 /* regular */
#define S_IFLNK    0120000 /* symbolic link */
#define S_IFSOCK   0140000 /* socket */
#define S_ISUID    0004000 /* set user id on execution */
#define S_ISGID    0002000 /* set group id on execution */
#define S_ISVTX    0001000 /* save swapped text even after use */
#define S_IRREAD   0000400 /* read permission, owner */
#define S_IWRITE   0000200 /* write permission, owner */
#define S_IXEXEC   0000100 /* execute/search permission, owner */
```

The mode bits 0000070 and 0000007 encode group and others permissions (see `chmod(2)`).

When `fd` is associated with a pipe, `fstat` reports an ordinary file with an i-node number, restricted permissions, and a not necessarily meaningful length.

RETURN VALUE

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

ERRORS

`Stat` and `lstat` will fail if one or more of the following are true:

[ENOTDIR] A component of the path prefix is not a directory.

[EPERM] The pathname contains a character with the high-order bit set.

[ENOENT] The pathname was too long.

[ENOENT] The named file does not exist.

[EACCES] Search permission is denied for a component of the path prefix.

[EFAULT] `Buf` or `name` points to an invalid address.

`Fstat` will fail if one or both of the following are true:

[EBADF] `Fildes` is not a valid open file descriptor.

[EFAULT] `Buf` points to an invalid address.

[ELOOP] Too many symbolic links were encountered in translating the pathname.

CAVEAT

The fields in the `stat` structure currently marked `st_spare1`, `st_spare2`, and `st_spare3` are present in preparation for inode time stamps expanding to 64 bits. This, however, can break certain programs which depend on the time stamps being contiguous (in calls to `utimes(2)`).

SEE ALSO

`chmod(2)`, `chown(2)`, `utimes(2)`

BUGS

Applying *fstat* to a socket returns a zero'd buffer.

The list of calls which modify the various fields should be carefully checked with reality.

NAME

symlink — make symbolic link to a file

SYNOPSIS

```
symlink(name1, name2)
char *name1, *name2;
```

DESCRIPTION

A symbolic link *name2* is created to *name1* (*name2* is the name of the file created, *name1* is the string used in creating the symbolic link). Either name may be an arbitrary path name; the files need not be on the same file system.

RETURN VALUE

Upon successful completion, a zero value is returned. If an error occurs, the error code is stored in *errno* and a -1 value is returned.

ERRORS

The symbolic link is made unless on or more of the following are true:

- [EPERM] Either *name1* or *name2* contains a character with the high-order bit set.
- [ENOENT] One of the pathnames specified was too long.
- [ENOTDIR] A component of the *name2* prefix is not a directory.
- [EEXIST] *Name2* already exists.
- [EACCES] A component of the *name2* path prefix denies search permission.
- [EROFS] The file *name2* would reside on a read-only file system.
- [EFAULT] *Name1* or *name2* points outside the process's allocated address space.
- [ELOOP] Too many symbolic links were encountered in translating the pathname.

SEE ALSO

link(2), ln(1), unlink(2)

NAME

sync — update super-block

SYNOPSIS

sync()

DESCRIPTION

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

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

SEE ALSO

fsync(2), sync(8), update(8)

BUGS

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

NAME

`syscall` — indirect system call

SYNOPSIS

`syscall(number, arg, ...)` (VAX-11)

DESCRIPTION

Syscall performs the system call whose assembly language interface has the specified *number*, register arguments *r0* and *r1* and further arguments *arg*.

The *r0* value of the system call is returned.

DIAGNOSTICS

When the C-bit is set, *syscall* returns `-1` and sets the external variable *errno* (see *intro(2)*).

BUGS

There is no way to simulate system calls such as *pipe(2)*, which return values in register *r1*.

NAME

`truncate` — truncate a file to a specified length

SYNOPSIS

`truncate(path, length)`

`char *path;`

`int length;`

`ftruncate(fd, length)`

`int fd, length;`

DESCRIPTION

Truncate causes the file named by *path* or referenced by *fd* to be truncated to at most *length* bytes in size. If the file previously was larger than this size, the extra data is lost. With *ftruncate*, the file must be open for writing.

RETURN VALUES

A value of 0 is returned if the call succeeds. If the call fails a `-1` is returned, and the global variable *errno* specifies the error.

ERRORS

Truncate succeeds unless:

[EPERM] The pathname contains a character with the high-order bit set.

[ENOENT] The pathname was too long.

[ENOTDIR] A component of the path prefix of *path* is not a directory.

[ENOENT] The named file does not exist.

[EACCES] A component of the *path* prefix denies search permission.

[EISDIR] The named file is a directory.

[EROFS] The named file resides on a read-only file system.

[ETXTBSY] The file is a pure procedure (shared text) file that is being executed.

[EFAULT] *Name* points outside the process's allocated address space.

Ftruncate succeeds unless:

[EBADF] The *fd* is not a valid descriptor.

[EINVAL] The *fd* references a socket, not a file.

SEE ALSO

`open(2)`

BUGS

Partial blocks discarded as the result of truncation are not zero filled; this can result in holes in files which do not read as zero.

These calls should be generalized to allow ranges of bytes in a file to be discarded.

NAME

umask — set file creation mode mask

SYNOPSIS

```
oumask = umask(numask)  
int oumask, numask;
```

DESCRIPTION

Umask sets the process's file mode creation mask to *numask* and returns the previous value of the mask. The low-order 9 bits of *numask* are used whenever a file is created, clearing corresponding bits in the file mode (see *chmod(2)*). This clearing allows each user to restrict the default access to his files.

The value is initially 022 (write access for owner only). The mask is inherited by child processes.

RETURN VALUE

The previous value of the file mode mask is returned by the call.

SEE ALSO

chmod(2), *mknod(2)*, *open(2)*

NAME

`unlink` — remove directory entry

SYNOPSIS

`unlink(path)`

`char *path;`

DESCRIPTION

Unlink removes the entry for the file *path* from its directory. If this entry was the last link to the file, and no process has the file open, then all resources associated with the file are reclaimed. If, however, the file was open in any process, the actual resource reclamation is delayed until it is closed, even though the directory entry has disappeared.

RETURN VALUE

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

ERRORS

The *unlink* succeeds unless:

- | | |
|-----------|---|
| [EPERM] | The path contains a character with the high-order bit set. |
| [ENOENT] | The path name is too long. |
| [ENOTDIR] | A component of the path prefix is not a directory. |
| [ENOENT] | The named file does not exist. |
| [EACCES] | Search permission is denied for a component of the path prefix. |
| [EACCES] | Write permission is denied on the directory containing the link to be removed. |
| [EPERM] | The named file is a directory and the effective user ID of the process is not the super-user. |
| [EBUSY] | The entry to be unlinked is the mount point for a mounted file system. |
| [EROFS] | The named file resides on a read-only file system. |
| [EFAULT] | <i>Path</i> points outside the process's allocated address space. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |

SEE ALSO

`close(2)`, `link(2)`, `rmdir(2)`

NAME

`utimes` — set file times

SYNOPSIS

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

utimes(file, tvp)
char *file;
struct timeval *tvp[2];
```

DESCRIPTION

The *utimes* call uses the “accessed” and “updated” times in that order from the *tvp* vector to set the corresponding recorded times for *file*.

The caller must be the owner of the file or the super-user. The “inode-changed” time of the file is set to the current time.

RETURN VALUE

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

ERRORS

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

- | | |
|-----------|--|
| [EPERM] | The pathname contained a character with the high-order bit set. |
| [ENOENT] | The pathname was too long. |
| [ENOENT] | The named file does not exist. |
| [ENOTDIR] | A component of the path prefix is not a directory. |
| [EACCES] | A component of the path prefix denies search permission. |
| [EPERM] | The process is not super-user and not the owner of the file. |
| [EACCES] | The effective user ID is not super-user and not the owner of the file and <i>times</i> is NULL and write access is denied. |
| [EROFS] | The file system containing the file is mounted read-only. |
| [EFAULT] | <i>Tvp</i> points outside the process's allocated address space. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |

SEE ALSO

`stat(2)`

NAME

vfork — spawn new process in a virtual memory efficient way

SYNOPSIS

```
pid = vfork()
int pid;
```

DESCRIPTION

Vfork can be used to create new processes without fully copying the address space of the old process, which is horrendously inefficient in a paged environment. It is useful when the purpose of *fork(2)* would have been to create a new system context for an *execve*. *Vfork* differs from *fork* in that the child borrows the parent's memory and thread of control until a call to *execve(2)* or an exit (either by a call to *exit(2)* or abnormally.) The parent process is suspended while the child is using its resources.

Vfork returns 0 in the child's context and (later) the pid of the child in the parent's context.

Vfork can normally be used just like *fork*. It does not work, however, to return while running in the child's context from the procedure which called *vfork* since the eventual return from *vfork* would then return to a no longer existent stack frame. Be careful, also, to call *_exit* rather than *exit* if you can't *execve*, since *exit* will flush and close standard I/O channels, and thereby mess up the parent process's standard I/O data structures. (Even with *fork* it is wrong to call *exit* since buffered data would then be flushed twice.)

SEE ALSO

fork(2), *execve(2)*, *sigvec(2)*, *wait(2)*,

DIAGNOSTICS

Same as for *fork*.

BUGS

This system call will be eliminated when proper system sharing mechanisms are implemented. Users should not depend on the memory sharing semantics of *vfork* as it will, in that case, be made synonymous to *fork*.

To avoid a possible deadlock situation, processes which are children in the middle of a *vfork* are never sent SIGTTOU or SIGTTIN signals; rather, output or *ioctl*s are allowed and input attempts result in an end-of-file indication.

NAME

vhangup — virtually “hangup” the current control terminal

SYNOPSIS

vhangup()

DESCRIPTION

Vhangup is used by the initialization process *init*(8) (among others) to arrange that users are given “clean” terminals at login, by revoking access of the previous users' processes to the terminal. To effect this, *vhangup* searches the system tables for references to the control terminal of the invoking process, revoking access permissions on each instance of the terminal which it finds. Further attempts to access the terminal by the affected processes will yield i/o errors (EBADF). Finally, a hangup signal (SIGHUP) is sent to the process group of the control terminal.

SEE ALSO

init (8)

BUGS

Access to the control terminal via */dev/tty* is still possible.

This call should be replaced by an automatic mechanism which takes place on process exit.

NAME

`wait`, `wait3` — wait for process to terminate

SYNOPSIS

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

pid = wait(status)
int pid;
union wait *status;

pid = wait(0)
int pid;

#include <sys/time.h>
#include <sys/resource.h>

pid = wait3(status, options, rusage)
int pid;
union wait *status;
int options;
struct rusage *rusage;
```

DESCRIPTION

Wait causes its caller to delay until a signal is received or one of its child processes terminates. If any child has died since the last *wait*, return is immediate, returning the process id and exit status of one of the terminated children. If there are no children, return is immediate with the value `-1` returned.

On return from a successful *wait* call, *status* is nonzero, and the high byte of *status* contains the low byte of the argument to *exit* supplied by the child process; the low byte of *status* contains the termination status of the process. A more precise definition of the *status* word is given in `<sys/wait.h>`.

Wait3 provides an alternate interface for programs which must not block when collecting the status of child processes. The *status* parameter is defined as above. The *options* parameter is used to indicate the call should not block if there are no processes which wish to report status (WNOHANG), and/or that only children of the current process which are stopped due to a SIGTTIN, SIGTTOU, SIGTSTP, or SIGSTOP signal should have their status reported (WUNTRACED). If *rusage* is non-zero, a summary of the resources used by the terminated process and all its children is returned (this information is currently not available for stopped processes).

When the WNOHANG option is specified and no processes wish to report status, *wait3* returns a *pid* of 0. The WNOHANG and WUNTRACED options may be combined by *or*'ing the two values.

NOTES

See *sigvec*(2) for a list of termination statuses (signals); 0 status indicates normal termination. A special status (0177) is returned for a stopped process which has not terminated and can be restarted; see *ptrace*(2). If the 0200 bit of the termination status is set, a core image of the process was produced by the system.

If the parent process terminates without waiting on its children, the initialization process (process ID = 1) inherits the children.

Wait and *wait3* are automatically restarted when a process receives a signal while awaiting termination of a child process.

RETURN VALUE

If *wait* returns due to a stopped or terminated child process, the process ID of the child is returned to the calling process. Otherwise, a value of `-1` is returned and *errno* is set to

indicate the error.

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.

ERRORS

Wait will fail and return immediately if one or more of the following are true:

- [ECHILD] The calling process has no existing unwaited-for child processes.
- [EFAULT] The *status* or *rusage* arguments point to an illegal address.

SEE ALSO

exit(2)

NAME

write, writev — write on a file

SYNOPSIS

```
write(d, buf, nbytes)
int d;
char *buf;
int nbytes;

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

writev(d, iov, ioveclen)
int d;
struct iovec *iov;
int ioveclen;
```

DESCRIPTION

Write attempts to write *nbytes* of data to the object referenced by the descriptor *d* from the buffer pointed to by *buf*. *Writev* performs the same action, but gathers the output data from the *iovcn* buffers specified by the members of the *iovec* array: *iov*[0], *iov*[1], etc.

On objects capable of seeking, the *write* starts at a position given by the pointer associated with *d*, see *lseek*(2). Upon return from *write*, the pointer is incremented by the number of bytes actually written.

Objects that are not capable of seeking always write from the current position. The value of the pointer associated with such an object is undefined.

If the real user is not the super-user, then *write* clears the set-user-id bit on a file. This prevents penetration of system security by a user who “captures” a writable set-user-id file owned by the super-user.

RETURN VALUE

Upon successful completion the number of bytes actually written is returned. Otherwise a *-1* is returned and *errno* is set to indicate the error.

ERRORS

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

- | | |
|----------|--|
| [EBADF] | <i>D</i> is not a valid descriptor open for writing. |
| [EPIPE] | An attempt is made to write to a pipe that is not open for reading by any process. |
| [EPIPE] | An attempt is made to write to a socket of type SOCK_STREAM which is not connected to a peer socket. |
| [EFBIG] | An attempt was made to write a file that exceeds the process's file size limit or the maximum file size. |
| [EFAULT] | Part of <i>iov</i> or data to be written to the file points outside the process's allocated address space. |

SEE ALSO

lseek(2), *open*(2), *pipe*(2)

NAME

intro — introduction to library functions

DESCRIPTION

This section describes functions that may be found in various libraries. The library functions are those other than the functions which directly invoke UNIX system primitives, described in section 2. This section has the libraries physically grouped together. This is a departure from older versions of the UNIX Programmer's Reference Manual, which did not group functions by library. The functions described in this section are grouped into various libraries:

(3) and (3S)

The straight "3" functions are the standard C library functions. The C library also includes all the functions described in section 2. The 3S functions comprise the standard I/O library. Together with the (3N), (3X), and (3C) routines, these functions constitute library *libc*, which is automatically loaded by the C compiler *cc(1)*, the Pascal compiler *pc(1)*, and the Fortran compiler *f77(1)*. The link editor *ld(1)* searches this library under the '-lc' option. Declarations for some of these functions may be obtained from include files indicated on the appropriate pages.

(3F) The 3F functions are all functions callable from FORTRAN. These functions perform the same jobs as do the straight "3" functions.

(3M) These functions constitute the math library, *libm*. They are automatically loaded as needed by the Pascal compiler *pc(1)* and the Fortran compiler *f77(1)*. The link editor searches this library under the '-lm' option. Declarations for these functions may be obtained from the include file *<math.h>*.

(3N) These functions constitute the internet network library,

(3S) These functions constitute the 'standard I/O package', see *intro(3S)*. These functions are in the library *libc* already mentioned. Declarations for these functions may be obtained from the include file *<stdio.h>*.

(3X) Various specialized libraries have not been given distinctive captions. Files in which such libraries are found are named on appropriate pages.

(3C) Routines included for compatibility with other systems. In particular, a number of system call interfaces provided in previous releases of 4BSD have been included for source code compatibility. The manual page entry for each compatibility routine indicates the proper interface to use.

FILES

/lib/libc.a
 /usr/lib/libm.a
 /usr/lib/libc_p.a
 /usr/lib/libm_p.a

SEE ALSO

intro(3C), intro(3S), intro(3F), intro(3M), intro(3N), nm(1), ld(1), cc(1), f77(1), intro(2)

DIAGNOSTICS

Functions in the math library (3M) may return conventional values 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 (domain error) or ERANGE (range error). The values of EDOM and ERANGE are defined in the include file *<math.h>*.

LIST OF FUNCTIONS

<i>Name</i>	<i>Appears on Page</i>	<i>Description</i>
abort	abort.3	generate a fault
abort	abort.3f	terminate abruptly with memory image

abs	abs.3	integer absolute value
access	access.3f	determine accessibility of a file
acos	sin.3m	trigonometric functions
alarm	alarm.3c	schedule signal after specified time
alarm	alarm.3f	execute a subroutine after a specified time
alloca	malloc.3	memory allocator
arc	plot.3x	graphics interface
asctime	ctime.3	convert date and time to ASCII
asin	sin.3m	trigonometric functions
assert	assert.3x	program verification
atan	sin.3m	trigonometric functions
atan2	sin.3m	trigonometric functions
atof	atof.3	convert ASCII to numbers
atoi	atof.3	convert ASCII to numbers
atol	atof.3	convert ASCII to numbers
bcmp	bstring.3	bit and byte string operations
bcopy	bstring.3	bit and byte string operations
bessel	bessel.3f	of two kinds for integer orders
bit	bit.3f	and, or, xor, not, rshift, lshift bitwise functions
bzero	bstring.3	bit and byte string operations
cabs	hypot.3m	Euclidean distance
calloc	malloc.3	memory allocator
ceil	floor.3m	absolute value, floor, ceiling functions
chdir	chdir.3f	change default directory
chmod	chmod.3f	change mode of a file
circle	plot.3x	graphics interface
clearerr	ferror.3s	stream status inquiries
closedir	directory.3	directory operations
closelog	syslog.3	control system log
closepl	plot.3x	graphics interface
cont	plot.3x	graphics interface
cos	sin.3m	trigonometric functions
cosh	sinh.3m	hyperbolic functions
crypt	crypt.3	DES encryption
ctime	ctime.3	convert date and time to ASCII
ctime	time.3f	return system time
curses	curses.3x	screen functions with "optimal" cursor motion
dbminit	dbm.3x	data base subroutines
delete	dbm.3x	data base subroutines
dfrac	flmin.3f	return extreme values
dflmax	flmin.3f	return extreme values
dflmax	range.3f	return extreme values
dflmin	flmin.3f	return extreme values
dflmin	range.3f	return extreme values
drand	rand.3f	return random values
dtime	etime.3f	return elapsed execution time
ecvt	ecvt.3	output conversion
edata	end.3	last locations in program
encrypt	crypt.3	DES encryption
end	end.3	last locations in program
endsent	getfsent.3x	get file system descriptor file entry
endgrent	getgrent.3	get group file entry

endhostent	gethostent.3n	get network host entry
endnetent	getnetent.3n	get network entry
endprotoent	getprotoent.3n	get protocol entry
endpwent	getpwent.3	get password file entry
endservent	getservent.3n	get service entry
environ	execl.3	execute a file
erase	plot.3x	graphics interface
etext	end.3	last locations in program
etime	etime.3f	return elapsed execution time
exec	execl.3	execute a file
exece	execl.3	execute a file
execl	execl.3	execute a file
execle	execl.3	execute a file
execlp	execl.3	execute a file
exect	execl.3	execute a file
execv	execl.3	execute a file
execvp	execl.3	execute a file
exit	exit.3	terminate a process after flushing any pending output
exit	exit.3f	terminate process with status
exp	exp.3m	exponential, logarithm, power, square root
fabs	floor.3m	absolute value, floor, ceiling functions
fclose	fclose.3s	close or flush a stream
fcvt	ecvt.3	output conversion
fdate	fdate.3f	return date and time in an ASCII string
feof	ferror.3s	stream status inquiries
ferror	ferror.3s	stream status inquiries
fetch	dbm.3x	data base subroutines
flush	fclose.3s	close or flush a stream
frac	flmin.3f	return extreme values
ffs	bstring.3	bit and byte string operations
fgetc	getc.3f	get a character from a logical unit
fgetc	getc.3s	get character or word from stream
fgets	gets.3s	get a string from a stream
fileno	ferror.3s	stream status inquiries
firstkey	dbm.3x	data base subroutines
flmax	flmin.3f	return extreme values
flmax	range.3f	return extreme values
flmin	flmin.3f	return extreme values
flmin	range.3f	return extreme values
floor	floor.3m	absolute value, floor, ceiling functions
flush	flush.3f	flush output to a logical unit
fork	fork.3f	create a copy of this process
fpecnt	trpfpe.3f	trap and repair floating point faults
fprintf	printf.3s	formatted output conversion
fputc	putc.3f	write a character to a fortran logical unit
fputc	putc.3s	put character or word on a stream
fputs	puts.3s	put a string on a stream
fread	fread.3s	buffered binary input/output
free	malloc.3	memory allocator
frexp	frexp.3	split into mantissa and exponent
fscanf	scanf.3s	formatted input conversion
fseek	fseek.3f	reposition a file on a logical unit

fseek	fseek.3s	reposition a stream
fstat	stat.3f	get file status
ftell	fseek.3f	reposition a file on a logical unit
ftell	fseek.3s	reposition a stream
ftime	time.3c	get date and time
fwrite	fread.3s	buffered binary input/output
gamma	gamma.3m	log gamma function
gcv	ecvt.3	output conversion
gerror	perror.3f	get system error messages
getarg	getarg.3f	return command line arguments
getc	getc.3f	get a character from a logical unit
getc	getc.3s	get character or word from stream
getchar	getc.3s	get character or word from stream
getcwd	getcwd.3f	get pathname of current working directory
getdiskbyname	getdisk.3x	get disk description by its name
getenv	getenv.3	value for environment name
getenv	getenv.3f	get value of environment variables
getfsent	getfsent.3x	get file system descriptor file entry
getfsfile	getfsent.3x	get file system descriptor file entry
getfsspec	getfsent.3x	get file system descriptor file entry
getfstype	getfsent.3x	get file system descriptor file entry
getgid	getuid.3f	get user or group ID of the caller
getgrent	getgrent.3	get group file entry
getgrgid	getgrent.3	get group file entry
getgrnam	getgrent.3	get group file entry
gethostbyaddr	gethostent.3n	get network host entry
gethostbyname	gethostent.3n	get network host entry
gethostent	gethostent.3n	get network host entry
getlog	getlog.3f	get user's login name
getlogin	getlogin.3	get login name
getnetbyaddr	getnetent.3n	get network entry
getnetbyname	getnetent.3n	get network entry
getnetent	getnetent.3n	get network entry
getpass	getpass.3	read a password
getpid	getpid.3f	get process id
getprotobyname	getprotoent.3n	get protocol entry
getprotobyname	getprotoent.3n	get protocol entry
getprotoent	getprotoent.3n	get protocol entry
getpw	getpw.3	get name from uid
getpwent	getpwent.3	get password file entry
getpwnam	getpwent.3	get password file entry
getpwuid	getpwent.3	get password file entry
gets	gets.3s	get a string from a stream
getservbyname	getservent.3n	get service entry
getservbyport	getservent.3n	get service entry
getservent	getservent.3n	get service entry
getuid	getuid.3f	get user or group ID of the caller
getw	getc.3s	get character or word from stream
getwd	getwd.3	get current working directory pathname
gmtime	ctime.3	convert date and time to ASCII
gmtime	time.3f	return system time
gtty	stty.3c	set and get terminal state (defunct)

endhostent	gethostent.3n	get network host entry
endnetent	getnetent.3n	get network entry
endprotoent	getprotoent.3n	get protocol entry
endpwent	getpwent.3	get password file entry
endservent	getservent.3n	get service entry
environ	execl.3	execute a file
erase	plot.3x	graphics interface
etext	end.3	last locations in program
etime	etime.3f	return elapsed execution time
exec	execl.3	execute a file
exece	execl.3	execute a file
execl	execl.3	execute a file
execle	execl.3	execute a file
execlp	execl.3	execute a file
exect	execl.3	execute a file
execv	execl.3	execute a file
execvp	execl.3	execute a file
exit	exit.3	terminate a process after flushing any pending output
exit	exit.3f	terminate process with status
exp	exp.3m	exponential, logarithm, power, square root
fabs	floor.3m	absolute value, floor, ceiling functions
fclose	fclose.3s	close or flush a stream
fcvt	ecvt.3	output conversion
fdate	fdate.3f	return date and time in an ASCII string
feof	ferror.3s	stream status inquiries
ferror	ferror.3s	stream status inquiries
fetch	dbm.3x	data base subroutines
fflush	fclose.3s	close or flush a stream
ffrac	flmin.3f	return extreme values
ffs	bstring.3	bit and byte string operations
fgetc	getc.3f	get a character from a logical unit
fgetc	getc.3s	get character or word from stream
fgets	gets.3s	get a string from a stream
fileno	ferror.3s	stream status inquiries
firstkey	dbm.3x	data base subroutines
flmax	flmin.3f	return extreme values
flmax	range.3f	return extreme values
flmin	flmin.3f	return extreme values
flmin	range.3f	return extreme values
floor	floor.3m	absolute value, floor, ceiling functions
flush	flush.3f	flush output to a logical unit
fork	fork.3f	create a copy of this process
fpecnt	trpfpe.3f	trap and repair floating point faults
fprintf	printf.3s	formatted output conversion
fputc	putc.3f	write a character to a fortran logical unit
fputc	putc.3s	put character or word on a stream
fputs	puts.3s	put a string on a stream
fread	fread.3s	buffered binary input/output
free	malloc.3	memory allocator
frexp	frexp.3	split into mantissa and exponent
fscanf	scanf.3s	formatted input conversion
fseek	fseek.3f	reposition a file on a logical unit

fseek	fseek.3s	reposition a stream
fstat	stat.3f	get file status
ftell	fseek.3f	reposition a file on a logical unit
ftell	fseek.3s	reposition a stream
ftime	time.3c	get date and time
fwrite	fread.3s	buffered binary input/output
gamma	gamma.3m	log gamma function
gcv	ecvt.3	output conversion
gerror	perror.3f	get system error messages
getarg	getarg.3f	return command line arguments
getc	getc.3f	get a character from a logical unit
getc	getc.3s	get character or word from stream
getchar	getc.3s	get character or word from stream
getcwd	getcwd.3f	get pathname of current working directory
getdiskbyname	getdisk.3x	get disk description by its name
getenv	getenv.3	value for environment name
getenv	getenv.3f	get value of environment variables
getfsent	getfsent.3x	get file system descriptor file entry
getfsfile	getfsent.3x	get file system descriptor file entry
getfsspec	getfsent.3x	get file system descriptor file entry
getfstype	getfsent.3x	get file system descriptor file entry
getgid	getuid.3f	get user or group ID of the caller
getgrent	getgrent.3	get group file entry
getgrgid	getgrent.3	get group file entry
getgrnam	getgrent.3	get group file entry
gethostbyaddr	gethostent.3n	get network host entry
gethostbyname	gethostent.3n	get network host entry
gethostent	gethostent.3n	get network host entry
getlog	getlog.3f	get user's login name
getlogin	getlogin.3	get login name
getnetbyaddr	getnetent.3n	get network entry
getnetbyname	getnetent.3n	get network entry
getnetent	getnetent.3n	get network entry
getpass	getpass.3	read a password
getpid	getpid.3f	get process id
getprotobyname	getprotoent.3n	get protocol entry
getprotobyname	getprotoent.3n	get protocol entry
getprotoent	getprotoent.3n	get protocol entry
getpw	getpw.3	get name from uid
getpwent	getpwent.3	get password file entry
getpwnam	getpwent.3	get password file entry
getpwuid	getpwent.3	get password file entry
gets	gets.3s	get a string from a stream
getservbyname	getservent.3n	get service entry
getservbyport	getservent.3n	get service entry
getservent	getservent.3n	get service entry
getuid	getuid.3f	get user or group ID of the caller
getw	getc.3s	get character or word from stream
getwd	getwd.3	get current working directory pathname
gmtime	ctime.3	convert date and time to ASCII
gmtime	time.3f	return system time
gtty	stty.3c	set and get terminal state (defunct)

hostnm	hostnm.3f	get name of current host
htonl	byteorder.3n	convert values between host and network byte order
htons	byteorder.3n	convert values between host and network byte order
hypot	hypot.3m	Euclidean distance
iargc	getarg.3f	return command line arguments
idate	idate.3f	return date or time in numerical form
ierrno	perror.3f	get system error messages
index	index.3f	tell about character objects
index	string.3	string operations
inet_addr	inet.3n	Internet address manipulation routines
inet_lnaof	inet.3n	Internet address manipulation routines
inet_makeaddr	inet.3n	Internet address manipulation routines
inet_netof	inet.3n	Internet address manipulation routines
inet_network	inet.3n	Internet address manipulation routines
initgroups	initgroups.3x	initialize group access list
initstate	random.3	better random number generator
inmax	flmin.3f	return extreme values
inmax	range.3f	return extreme values
insque	insque.3	insert/remove element from a queue
ioinit	ioinit.3f	change f77 I/O initialization
irand	rand.3f	return random values
isalnum	ctype.3	character classification macros
isalpha	ctype.3	character classification macros
isascii	ctype.3	character classification macros
isatty	ttynam.3f	find name of a terminal port
isatty	ttyname.3	find name of a terminal
isctrl	ctype.3	character classification macros
isdigit	ctype.3	character classification macros
islower	ctype.3	character classification macros
isprint	ctype.3	character classification macros
ispunct	ctype.3	character classification macros
isspace	ctype.3	character classification macros
isupper	ctype.3	character classification macros
itime	idate.3f	return date or time in numerical form
j0	j0.3m	bessel functions
j1	j0.3m	bessel functions
jn	j0.3m	bessel functions
kill	kill.3f	send a signal to a process
label	plot.3x	graphics interface
ldexp	frexp.3	split into mantissa and exponent
len	index.3f	tell about character objects
lib2648	lib2648.3x	subroutines for the HP 2648 graphics terminal
line	plot.3x	graphics interface
linemod	plot.3x	graphics interface
link	link.3f	make a link to an existing file
lnblk	index.3f	tell about character objects
loc	loc.3f	return the address of an object
localtime	ctime.3	convert date and time to ASCII
log	exp.3m	exponential, logarithm, power, square root
log10	exp.3m	exponential, logarithm, power, square root
long	long.3f	integer object conversion
longjmp	setjmp.3	non-local goto

lstat	stat.3f	get file status
ltime	time.3f	return system time
malloc	malloc.3	memory allocator
mktemp	mktemp.3	make a unique file name
modf	frexp.3	split into mantissa and exponent
moncontrol	monitor.3	prepare execution profile
monitor	monitor.3	prepare execution profile
monstartup	monitor.3	prepare execution profile
move	plot.3x	graphics interface
nextkey	dbm.3x	data base subroutines
nice	nice.3c	set program priority
nlist	nlist.3	get entries from name list
ntohl	byteorder.3n	convert values between host and network byte order
ntohs	byteorder.3n	convert values between host and network byte order
opendir	directory.3	directory operations
openlog	syslog.3	control system log
pause	pause.3c	stop until signal
pclose	popen.3	initiate I/O to/from a process
perror	perror.3	system error messages
perror	perror.3f	get system error messages
plot: openpl	plot.3x	graphics interface
point	plot.3x	graphics interface
popen	popen.3	initiate I/O to/from a process
pow	exp.3m	exponential, logarithm, power, square root
printf	printf.3s	formatted output conversion
psignal	psignal.3	system signal messages
putc	putc.3f	write a character to a fortran logical unit
putc	putc.3s	put character or word on a stream
putchar	putc.3s	put character or word on a stream
puts	puts.3s	put a string on a stream
putw	putc.3s	put character or word on a stream
qsort	qsort.3	quicker sort
qsort	qsort.3f	quick sort
rand	rand.3c	random number generator
rand	rand.3f	return random values
random	random.3	better random number generator
rcmd	rcmd.3x	routines for returning a stream to a remote command
re_comp	regex.3	regular expression handler
re_exec	regex.3	regular expression handler
readdir	directory.3	directory operations
realloc	malloc.3	memory allocator
remque	insque.3	insert/remove element from a queue
rename	rename.3f	rename a file
rewind	fseek.3s	reposition a stream
rewinddir	directory.3	directory operations
rexec	rexec.3x	return stream to a remote command
rindex	index.3f	tell about character objects
rindex	string.3	string operations
rresvport	rcmd.3x	routines for returning a stream to a remote command
ruserok	rcmd.3x	routines for returning a stream to a remote command
scandir	scandir.3	scan a directory
scanf	scanf.3s	formatted input conversion

seekdir	directory.3	directory operations
setbuf	setbuf.3s	assign buffering to a stream
setbuffer	setbuf.3s	assign buffering to a stream
setegid	setuid.3	set user and group ID
seteuid	setuid.3	set user and group ID
setfsent	getfsent.3x	get file system descriptor file entry
setgid	setuid.3	set user and group ID
setgrent	getgrent.3	get group file entry
sethostent	gethostent.3n	get network host entry
setjmp	setjmp.3	non-local goto
setkey	crypt.3	DES encryption
setlinebuf	setbuf.3s	assign buffering to a stream
setnetent	getnetent.3n	get network entry
setprotoent	getprotoent.3n	get protocol entry
setpwent	getpwent.3	get password file entry
setrgid	setuid.3	set user and group ID
setruid	setuid.3	set user and group ID
setservent	getservent.3n	get service entry
setstate	random.3	better random number generator
setuid	setuid.3	set user and group ID
short	long.3f	integer object conversion
signal	signal.3	simplified software signal facilities
signal	signal.3f	change the action for a signal
sin	sin.3m	trigonometric functions
sinh	sinh.3m	hyperbolic functions
sleep	sleep.3	suspend execution for interval
sleep	sleep.3f	suspend execution for an interval
space	plot.3x	graphics interface
sprintf	printf.3s	formatted output conversion
sqrt	exp.3m	exponential, logarithm, power, square root
srand	rand.3c	random number generator
random	random.3	better random number generator
sscanf	scanf.3s	formatted input conversion
stat	stat.3f	get file status
stdio	intro.3s	standard buffered input/output package
store	dbm.3x	data base subroutines
strcat	string.3	string operations
strcmp	string.3	string operations
strcpy	string.3	string operations
strlen	string.3	string operations
strncat	string.3	string operations
strncmp	string.3	string operations
strncpy	string.3	string operations
stty	stty.3c	set and get terminal state (defunct)
swab	swab.3	swap bytes
sys_errlist	perror.3	system error messages
sys_nerr	perror.3	system error messages
sys_siglist	psignal.3	system signal messages
syslog	syslog.3	control system log
system	system.3	issue a shell command
system	system.3f	execute a UNIX command
tan	sin.3m	trigonometric functions

tanh	sinh.3m	hyperbolic functions
tclose	topen.3f	f77 tape I/O
tellidir	directory.3	directory operations
tgetent	termcap.3x	terminal independent operation routines
tgetflag	termcap.3x	terminal independent operation routines
tgetnum	termcap.3x	terminal independent operation routines
tgetstr	termcap.3x	terminal independent operation routines
tgoto	termcap.3x	terminal independent operation routines
time	time.3c	get date and time
time	time.3f	return system time
times	times.3c	get process times
timezone	ctime.3	convert date and time to ASCII
topen	topen.3f	f77 tape I/O
tputs	termcap.3x	terminal independent operation routines
traper	traper.3f	trap arithmetic errors
trapov	trapov.3f	trap and repair floating point overflow
tread	topen.3f	f77 tape I/O
trewin	topen.3f	f77 tape I/O
trpfpe	trpfpe.3f	trap and repair floating point faults
tskipf	topen.3f	f77 tape I/O
tstate	topen.3f	f77 tape I/O
ttynam	ttynam.3f	find name of a terminal port
ttynam	ttynam.3	find name of a terminal
ttyslot	ttynam.3	find name of a terminal
twrite	topen.3f	f77 tape I/O
ungetc	ungetc.3s	push character back into input stream
unlink	unlink.3f	remove a directory entry
utime	utime.3c	set file times
valloc	valloc.3	aligned memory allocator
varargs	varargs.3	variable argument list
vlimit	vlimit.3c	control maximum system resource consumption
vtimes	vtimes.3c	get information about resource utilization
wait	wait.3f	wait for a process to terminate
y0	j0.3m	bessel functions
y1	j0.3m	bessel functions
yn	j0.3m	bessel functions

NAME

abort — generate a fault

DESCRIPTION

Abort executes an instruction which is illegal in user mode. This causes a signal that normally terminates the process with a core dump, which may be used for debugging.

SEE ALSO

adb(1), sigvec(2), exit(2)

DIAGNOSTICS

Usually 'IOT trap — core dumped' from the shell.

BUGS

The abort() function does not flush standard I/O buffers. Use *fflush* (3S).

NAME

`abs` — integer absolute value

SYNOPSIS

```
abs(i)  
int i;
```

DESCRIPTION

Abs returns the absolute value of its integer operand.

SEE ALSO

`floor(3M)` for *fabs*

BUGS

Applying the *abs* function to the most negative integer generates a result which is the most negative integer. That is,

```
abs(0x80000000)
```

returns `0x80000000` as a result.

NAME

atof, *atoi*, *atol* — convert ASCII to numbers

SYNOPSIS

double *atof*(*nptr*)

char **nptr*;

int *atoi*(*nptr*)

char **nptr*;

long *atol*(*nptr*)

char **nptr*;

DESCRIPTION

These functions convert a string pointed to by *nptr* to floating, integer, and long integer representation respectively. The first unrecognized character ends the string.

Atof recognizes an optional string of spaces, 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.

Atoi and *atol* recognize an optional string of spaces, then an optional sign, then a string of digits.

SEE ALSO

scanf(3S)

BUGS

There are no provisions for overflow.

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*;

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(3)* 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.

BUGS

The *bcmp* and *bcopy* routines take parameters backwards from *strcmp* and *strcpy*.

NAME

crypt, *setkey*, *encrypt* — DES encryption

SYNOPSIS

```
char *crypt(key, salt)
char *key, *salt;

setkey(key)
char *key;

encrypt(block, edflag)
char *block;
```

DESCRIPTION

Crypt is the password encryption routine. It is based on the NBS Data Encryption Standard, with variations intended (among other things) to frustrate use of hardware implementations of the DES for key search.

The first argument to *crypt* is normally a user's typed password. The second is a 2-character string chosen from the set [a-zA-Z0-9./]. The *salt* string is used to perturb the DES algorithm in one of 4096 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, in the same alphabet as the salt. The first two characters are the salt itself.

The other 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, leading to a 56-bit key which is set into the machine.

The argument to the *encrypt* entry is likewise a character array of length 64 containing 0's and 1's. 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 0, the argument is encrypted; if non-zero, it is decrypted.

SEE ALSO

passwd(1), *passwd*(5), *login*(1), *getpass*(3)

BUGS

The return value points to static data whose content is overwritten by each call.

NAME

`ctime`, `localtime`, `gmtime`, `asctime`, `timezone` — convert date and time to ASCII

SYNOPSIS

```
char *ctime(clock)
long *clock;
#include <sys/time.h>
struct tm *localtime(clock)
long *clock;
struct tm *gmtime(clock)
long *clock;
char *asctime(tm)
struct tm *tm;
char *timezone(zone, dst)
```

DESCRIPTION

Ctime converts a time pointed to by *clock* such as returned by *time(2)* into ASCII 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 structures containing the broken-down time. *Localtime* corrects for the time zone and possible daylight savings time; *gmtime* converts directly to GMT, which is the time UNIX uses. *asctime* converts a broken-down time to ASCII and returns a pointer to a 26-character string.

The structure declaration from the include file is:

```
struct tm {
    int    tm_sec;
    int    tm_min;
    int    tm_hour;
    int    tm_mday;
    int    tm_mon;
    int    tm_year;
    int    tm_wday;
    int    tm_yday;
    int    tm_isdst;
};
```

These quantities give the time on a 24-hour clock, day of month (1-31), month of year (0-11), day of week (Sunday = 0), year — 1900, day of year (0-365), and a flag that is nonzero if daylight saving time is in effect.

When local time is called for, the program consults the system to determine the time zone and whether the U.S.A., Australian, Eastern European, Middle European, or Western European daylight saving time adjustment is appropriate. The program knows about various peculiarities in time conversion over the past 10-20 years; if necessary, this understanding can be extended.

Timezone returns the name of the time zone associated with its first argument, which is measured in minutes westward from Greenwich. If the second argument is 0, the standard name is used, otherwise the Daylight Saving version. If the required name does not appear in a table built into the routine, the difference from GMT is produced; e.g. in Afghanistan *timezone(-60*4+30, 0)* is appropriate because it is 4:30 ahead of GMT and the string *GMT+4:30* is produced.

SEE ALSO

gettimeofday(2), time(3)

BUGS

The return values point to static data whose content is overwritten by each call.

NAME

isalpha, *isupper*, *islower*, *isdigit*, *isalnum*, *isspace*, *ispunct*, *isprint*, *isctrl*, *isascii* — character classification macros

SYNOPSIS

```
#include <ctype.h>
```

```
isalpha(c)
```

```
...
```

DESCRIPTION

These macros classify ASCII-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 (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
<i>isalnum</i>	<i>c</i> is an alphanumeric character
<i>isspace</i>	<i>c</i> is a space, tab, carriage return, newline, or formfeed
<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(8) (space) through 0176 (tilde)
<i>isctrl</i>	<i>c</i> is a delete character (0177) or ordinary control character (less than 040).
<i>isascii</i>	<i>c</i> is an ASCII character, code less than 0200

SEE ALSO

ascii(7)

NAME

opendir, readdir, telldir, seekdir, rewinddir, closedir — 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;
```

DESCRIPTION

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 if it cannot *malloc(3)* enough memory to hold the whole thing.

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 closes the named *directory stream* and frees the structure associated with the DIR pointer.

Sample code which searches a 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;
```

SEE ALSO

open(2), close(2), read(2), lseek(2), dir(5)

NAME

ecvt, *fcvt*, *gcvt* — output conversion

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;
char *buf;
```

DESCRIPTION

Ecvt converts the *value* to a null-terminated string of *ndigit* ASCII digits and returns a pointer thereto. 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). If the sign of the result is negative, the word pointed to by *sign* is non-zero, otherwise it is zero. The low-order digit is rounded.

Fcvt is identical to *ecvt*, except that the correct digit has been rounded for Fortran F-format output of the number of digits specified by *ndigits*.

Gcvt converts the *value* to a null-terminated ASCII string in *buf* and returns a pointer to *buf*. It attempts to produce *ndigit* significant digits in Fortran F format if possible, otherwise E format, ready for printing. Trailing zeros may be suppressed.

SEE ALSO

printf(3)

BUGS

The return values point to static data whose content is overwritten by each call.

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 coincides with *end*, but it is reset by the routines *brk(2)*, *malloc(3)*, standard input/output (*stdio(3)*), the profile (**-p**) option of *cc(1)*, etc. The current value of the program break is reliably returned by 'sbrk(0)', see *brk(2)*.

SEE ALSO

brk(2), *malloc(3)*

NAME

`execl`, `execv`, `execle`, `execlp`, `execvp`, `exec`, `exece`, `exec`, `exec`, `environ` — execute a file

SYNOPSIS

```

execl(name, arg0, arg1, ..., argn, 0)
char *name, *arg0, *arg1, ..., *argn;

execv(name, argv)
char *name, *argv[];

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

exec(name, argv, envp)
char *name, *argv[], *envp[];

extern char **environ;

```

DESCRIPTION

These routines provide various interfaces to the *execve* system call. Refer to *execve*(2) for a description of their properties; only brief descriptions are provided here.

Exec in all its forms overlays the calling process with the named file, then transfers to the entry point of the core image of the file. There can be no return from a successful *exec*; the calling core image is lost.

The *name* argument is a pointer to the name of the file to be executed. The pointers *arg*[0], *arg*[1] ... address null-terminated strings. Conventionally *arg*[0] is the name of the file.

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.

The *execv* version is useful when the number of arguments is unknown in advance; the arguments to *execv* are the name of the file to be executed and a vector of strings containing the arguments. The last argument string must be followed by a 0 pointer.

The *exec* version is used when the executed file is to be manipulated with *ptrace*(2). The program is forced to single step a single instruction giving the parent an opportunity to manipulate its state. On the VAX-11 this is done by setting the trace bit in the process status longword.

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.

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

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*(7) 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.

Execlp 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.

FILES

/bin/sh shell, invoked if command file found by *execlp* or *execvp*

SEE ALSO

execve(2), *fork*(2), *environ*(7), *csh*(1)

DIAGNOSTICS

If the file cannot be found, if it is not executable, if it does not start with a valid magic number (see *a.out*(5)), if maximum memory is exceeded, or if the arguments require too much space, a return constitutes the diagnostic; the return value is -1 . Even for the super-user, at least one of the execute-permission bits must be set for a file to be executed.

BUGS

If *execvp* is called to execute a file that turns out to be a shell command file, and if it is impossible to execute the shell, the values of *argv*[0] and *argv*[-1] will be modified before return.

NAME

exit — terminate a process after flushing any pending output

SYNOPSIS

```
exit(status)  
int status;
```

DESCRIPTION

Exit terminates a process after calling the Standard I/O library function *_cleanup* to flush any buffered output. *Exit* never returns.

SEE ALSO

exit(2), *intro(3S)*

NAME

frexp, *ldexp*, *modf* — split into mantissa and exponent

SYNOPSIS

double *frexp*(value, eptr)

double value;

int *eptr;

double *ldexp*(value, exp)

double value;

double *modf*(value, iptr)

double value, *iptr;

DESCRIPTION

Frexp returns the mantissa of a double *value* as a double quantity, *x*, of magnitude less than 1 and stores an integer *n* such that $value = x \cdot 2^n$ indirectly through *eptr*.

Ldexp returns the quantity $value \cdot 2^{exp}$.

Modf returns the positive fractional part of *value* and stores the integer part indirectly through *iptr*.

NAME

getenv — value for environment name

SYNOPSIS

```
char *getenv(name)
char *name;
```

DESCRIPTION

Getenv searches the environment list (see *environ*(7)) for a string of the form *name=value* and returns a pointer to the string *value* if such a string is present, otherwise *getenv* returns the value 0 (NULL).

SEE ALSO

environ(7), *execve*(2)

NAME

getgrent, *getgrgid*, *getgrnam*, *setgrent*, *endgrent* — get group file entry

SYNOPSIS

```
#include <grp.h>

struct group *getgrent()
struct group *getgrgid(gid)
int gid;
struct group *getgrnam(name)
char *name;
setgrent()
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 group file.

```
struct group { /* see getgrent(3) */
    char *gr_name;
    char *gr_passwd;
    int gr_gid;
    char **gr_mem;
};
```

```
struct group *getgrent(), *getgrgid(), *getgrnam();
```

The members of this structure are:

gr_name The name of the group.
gr_passwd The encrypted password of the group.
gr_gid The numerical group-ID.
gr_mem Null-terminated vector of pointers to the individual member names.

Getgrent simply reads the next line while *getgrgid* and *getgrnam* search until a matching *gid* or *name* is found (or until EOF is encountered). Each routine picks up where the others leave off so successive calls may be used to search the entire file.

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.

FILES

/etc/group

SEE ALSO

getlogin(3), *getpwent*(3), *group*(5)

DIAGNOSTICS

A null pointer (0) is 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.

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 userid is shared by several login names.

If *getlogin* is called within a process that is not attached to a typewriter, it returns NULL. The correct procedure for determining the login name is to first call *getlogin* and if it fails, to call *getpw(getuid())*.

FILES

/etc/utmp

SEE ALSO

getpwent(3), getgrent(3), utmp(5), getpw(3)

DIAGNOSTICS

Returns NULL (0) if name not found.

BUGS

The return values point to static data whose content is overwritten by each call.

NAME

getpass — read a password

SYNOPSIS

```
char *getpass(prompt)
char *prompt;
```

DESCRIPTION

Getpass reads a password from the file */dev/tty*, or if that cannot be opened, from the standard input, after prompting with the null-terminated string *prompt* and disabling echoing. A pointer is returned to a null-terminated string of at most 8 characters.

FILES

/dev/tty

SEE ALSO

crypt(3)

BUGS

The return value points to static data whose content is overwritten by each call.

NAME

getpwent, *getpwuid*, *getpwnam*, *setpwent*, *endpwent* — get password file entry

SYNOPSIS

```
#include <pwd.h>

struct passwd *getpwent()

struct passwd *getpwuid(uid)
int uid;

struct passwd *getpwnam(name)
char *name;

int setpwent()

int endpwent()
```

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 password file.

```
struct passwd { /* see getpwent(3) */
    char *pw_name;
    char *pw_passwd;
    int pw_uid;
    int pw_gid;
    int pw_quota;
    char *pw_comment;
    char *pw_gecos;
    char *pw_dir;
    char *pw_shell;
};

struct passwd *getpwent(), *getpwuid(), *getpwnam();
```

The fields *pw_quota* and *pw_comment* are unused; the others have meanings described in *passwd(5)*.

Getpwent reads the next line (opening the file if necessary); *setpwent* rewinds the file; *endpwent* closes it.

Getpwuid and *getpwnam* search from the beginning until a matching *uid* or *name* is found (or until EOF is encountered).

FILES

/etc/passwd

SEE ALSO

getlogin(3), *getgrent(3)*, *passwd(5)*

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.

NAME

`getwd` — get current working directory pathname

SYNOPSIS

```
char *getwd(pathname)
char *pathname;
```

DESCRIPTION

Getwd copies the absolute pathname of the current working directory to *pathname* and returns a pointer to the result.

LIMITATIONS

Maximum pathname length is MAXPATHLEN characters (1024).

DIAGNOSTICS

Getwd returns zero and places a message in *pathname* if an error occurs.

BUGS

Getwd may fail to return to the current directory if an error occurs.

NAME

malloc, *free*, *realloc*, *calloc*, *alloca* — memory allocator

SYNOPSIS

```
char *malloc(size)
unsigned size;

free(ptr)
char *ptr;

char *realloc(ptr, size)
char *ptr;
unsigned size;

char *calloc(nelem, elsize)
unsigned nelem, elsize;

char *alloca(size)
int size;
```

DESCRIPTION

Malloc and *free* provide a simple general-purpose memory allocation package. *Malloc* returns a pointer to a block of at least *size* bytes beginning on a word boundary.

The argument to *free* is a pointer to a block previously allocated by *malloc*; this space is made available for further allocation, but its contents are left undisturbed.

Needless to say, grave disorder will result if the space assigned by *malloc* is overrun or if some random number is handed to *free*.

Malloc maintains multiple lists of free blocks according to size, allocating space from the appropriate list. 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 will be unchanged up to the lesser of the new and old sizes.

In order to be compatible with older versions, *realloc* also works if *ptr* points to a block freed since the last call of *malloc*, *realloc* or *calloc*; sequences of *free*, *malloc* and *realloc* were previously used to attempt storage compaction. This procedure is no longer recommended.

Calloc allocates space for an array of *nelem* elements of size *elsize*. The space is initialized to zeros.

Alloca allocates *size* bytes of space in the stack frame of the caller. This temporary space is automatically freed on return.

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 (0) if there is no available memory or if the arena has been detectably corrupted by storing outside the bounds of a block. *Malloc* may be recompiled to check the arena very stringently on every transaction; those sites with a source code license may check the source code to see how this can be done.

BUGS

When *realloc* returns 0, the block pointed to by *ptr* may be destroyed.

Alloca is machine dependent; it's use is discouraged.

NAME

mktemp — make a unique file name

SYNOPSIS

```
char *mktemp(template)
char *template;
```

DESCRIPTION

Mktemp replaces *template* by a unique file name, and returns the address of the template. The template should look like a file name with six trailing X's, which will be replaced with the current process id and a unique letter.

SEE ALSO

getpid(2)

NAME

monitor, monstartup, moncontrol — prepare execution profile

SYNOPSIS

```
monitor(lowpc, highpc, buffer, bufsize, nfunc)
int (*lowpc)(), (*highpc)();
short buffer[];
```

```
monstartup(lowpc, highpc)
int (*lowpc)(), (*highpc)();
```

```
moncontrol(mode)
```

DESCRIPTION

There are two different forms of monitoring available: An executable program created by:

```
cc -p . . .
```

automatically includes calls for the *prof*(1) monitor and includes an initial call to its start-up routine *monstartup* with default parameters; *monitor* need not be called explicitly except to gain fine control over profil buffer allocation. An executable program created by:

```
cc -pg . . .
```

automatically includes calls for the *gprof*(1) monitor.

Monstartup is a high level interface to *profil*(2). *Lowpc* and *highpc* specify the address range that is to be sampled; the lowest address sampled is that of *lowpc* and the highest is just below *highpc*. *Monstartup* allocates space using *sbrk*(2) and passes it to *monitor* (see below) to record a histogram of periodically sampled values of the program counter, and of counts of calls of certain functions, in the buffer. Only calls of functions compiled with the profiling option *-p* of *cc*(1) are recorded.

To profile the entire program, it is sufficient to use

```
extern etext();
. . .
monstartup((int) 2, etext);
```

Etect lies just above all the program text, see *end*(3).

To stop execution monitoring and write the results on the file *mon.out*, use

```
monitor(0);
```

then *prof*(1) can be used to examine the results.

Moncontrol is used to selectively control profiling within a program. This works with either *prof*(1) or *gprof*(1) type profiling. When the program starts, profiling begins. To stop the collection of histogram ticks and call counts use *moncontrol*(0); to resume the collection of histogram ticks and call counts use *moncontrol*(1). This allows the cost of particular operations to be measured. Note that an output file will be produced upon program exit irregardless of the state of *moncontrol*.

Monitor is a low level interface to *profil*(2). *Lowpc* and *highpc* are the addresses of two functions; *buffer* is the address of a (user supplied) array of *bufsize* short integers. At most *nfunc* call counts can be kept. 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. *Monitor* divides the buffer into space to record the histogram of program counter samples over the range *lowpc* to *highpc*, and space to record call counts of functions compiled with the *-p* option to *cc*(1).

To profile the entire program, it is sufficient to use

```
extern etext();  
...  
monitor((int) 2, etext, buf, bufsize, nfunc);
```

FILES

mon.out

SEE ALSO

cc(1), prof(1), gprof(1), profil(2), sbrk(2)

NAME

nlist - get entries from name list

SYNOPSIS

```
#include <nlist.h>

nlist(filename, nl)
char *filename;
struct nlist nll[];
```

DESCRIPTION

Nlist examines the name list in the given executable output file and selectively extracts a list of values. The name list consists of an array of structures containing names, types and values. The list is terminated with a null name. Each 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. If the name is not found, both entries are set to 0. See *a.out(5)* for the structure declaration.

This subroutine is useful for examining the system name list kept in the file */vmunix*. In this way programs can obtain system addresses that are up to date.

SEE ALSO

a.out(5)

DIAGNOSTICS

All type entries are set to 0 if the file cannot be found or if it is not a valid namelist.

NAME

perror, sys_errlist, sys_nerr — system error messages

SYNOPSIS

```
perror(s)  
char *s;  
  
int sys_nerr;  
char *sys_errlist[];
```

DESCRIPTION

Perror produces a short error message on the standard error file describing the last error encountered during a call to the system from a C program. First the argument string *s* is printed, then a colon, then the message and a new-line. Most usefully, the argument string is the name of the program which incurred the error. The error number is taken from the external variable *errno* (see *intro(2)*), which is set when errors occur but not cleared when non-erroneous calls are made.

To simplify variant formatting of messages, the vector of message strings *sys_errlist* is provided; *errno* can be used as an index in this table to get the message string without the newline. *Sys_nerr* is the number of messages 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), *psignal(3)*

NAME

popen, *pclose* — initiate I/O 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 containing, respectively, a shell command line and an I/O mode, either *r* for reading or *w* for writing. *Popen* creates a pipe between the calling process and the command to be executed. The value returned is a stream pointer that can be used (as appropriate) to write to the standard input of the command or read from its standard output.

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 “*popen ed*” command.

BUGS

Only one stream opened by *popen* can be in use at once.

Buffered reading before opening an input filter may leave the standard input of that filter mispositioned. Similar problems with an output filter may be forestalled by careful buffer flushing, e.g. with *fflush*; see *fclose(3S)*.

NAME

qsort — quicker sort

SYNOPSIS

```
qsort(base, nel, width, compar)
char *base;
int (*compar)();
```

DESCRIPTION

Qsort is an implementation of the quicker-sort algorithm. The first argument is a pointer to the base of the data; the second is the number of elements; the third is the width of an element in bytes; the last is the name of the comparison routine to be called with two arguments which are pointers to the elements being compared. The routine must return an integer less than, equal to, or greater than 0 according as the first argument is to be considered less than, equal to, or greater than the second.

SEE ALSO

sort(1)

NAME

random, srandom, initstate, setstate — better random number generator; routines for changing generators

SYNOPSIS

```
long random()
srandom(seed)
int seed;
char *initstate(seed, state, n)
unsigned seed;
char *state;
int n;
char *setstate(state)
char *state;
```

DESCRIPTION

Random uses a non-linear additive feedback random number generator employing a default table of size 31 long integers to return successive pseudo-random numbers in the range from 0 to $2^{31}-1$. The period of this random number generator is very large, approximately $16 \cdot (2^{31}-1)$.

Random/srandom have (almost) the same calling sequence and initialization properties as *rand/srand*. The difference is that *rand(3)* produces a much less random sequence -- in fact, the low dozen bits generated by *rand* go through a cyclic pattern. All the bits generated by *random* are usable. For example, "random()&01" will produce a random binary value.

Unlike *srand*, *srandom* does not return the old seed; the reason for this is that the amount of state information used is much more than a single word. (Two other routines are provided to deal with restarting/changing random number generators). Like *rand(3)*, however, *random* will by default produce a sequence of numbers that can be duplicated by calling *srandom* with 1 as the seed.

The *initstate* routine allows a state array, passed in as an argument, to be initialized for future use. The size of the state array (in bytes) is used by *initstate* to decide how sophisticated a random number generator it should use -- the more state, the better the random numbers will be. (Current "optimal" values for the amount of state information are 8, 32, 64, 128, and 256 bytes; other amounts will be rounded down to the nearest known amount. Using less than 8 bytes will cause an error). The seed for the initialization (which specifies a starting point for the random number sequence, and provides for restarting at the same point) is also an argument. *Initstate* returns a pointer to the previous state information array.

Once a state has been initialized, the *setstate* routine provides for rapid switching between states. *Setstate* returns a pointer to the argument state array is used for further random number generation until the next call to *initstate* or *setstate*.

Once a state array has been initialized, it may be restarted at a different point either by calling *initstate* (with the desired seed, the state array, and its size) or by calling both *setstate* (with the state array) and *srandom* (with the desired seed). The advantage of calling both *setstate* and *srandom* is that the size of the state array does not have to be remembered after it is initialized.

With 256 bytes of state information, the period of the random number generator is greater than 2^{69} , which should be sufficient for most purposes.

AUTHOR

Earl T. Cohen

DIAGNOSTICS

If *initstate* is called with less than 8 bytes of state information, or if *setstate* detects that the state information has been garbled, error messages are printed on the standard error output.

SEE ALSO

rand(3)

BUGS

About 2/3 the speed of *rand*(3C).

NAME

`re_comp`, `re_exec` — regular expression handler

SYNOPSIS

`char *re_comp(s)`

`char *s;`

`re_exec(s)`

`char *s;`

DESCRIPTION

Re_comp compiles a string into an internal form suitable for pattern matching. *Re_exec* checks the argument string against the last string passed to *re_comp*.

Re_comp returns 0 if the string *s* was compiled successfully; otherwise a string containing an error message is returned. If *re_comp* is passed 0 or a null string, it returns without changing the currently compiled regular expression.

Re_exec returns 1 if the string *s* matches the last compiled regular expression, 0 if the string *s* failed to match the last compiled regular expression, and -1 if the compiled regular expression was invalid (indicating an internal error).

The strings passed to both *re_comp* and *re_exec* may have trailing or embedded newline characters; they are terminated by nulls. The regular expressions recognized are described in the manual entry for *ed(1)*, given the above difference.

SEE ALSO

ed(1), *ex(1)*, *egrep(1)*, *fgrep(1)*, *grep(1)*

DIAGNOSTICS

Re_exec returns -1 for an internal error.

Re_comp returns one of the following strings if an error occurs:

No previous regular expression,

Regular expression too long,

unmatched \[,

missing],

too many \(\) pairs,

unmatched \).

NAME

`scandir` — scan a directory

SYNOPSIS

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

scandir(dirname, namelist, select, compar)
char *dirname;
struct direct *(*namelist[]);
int (*select)();
int (*compar)();

alphasort(d1, d2)
struct direct **d1, **d2;
```

DESCRIPTION

Scandir reads the directory *dirname* and builds an array of pointers to directory entries using *malloc(3)*. It returns the number of entries in the array and a pointer to the array through *namelist*.

The *select* parameter is a pointer to a user supplied subroutine which is called by *scandir* to select which entries are to be included in the array. The select routine is passed a pointer to a directory entry and should return a non-zero value if the directory entry is to be included in the array. If *select* is null, then all the directory entries will be included.

The *compar* parameter is a pointer to a user supplied subroutine which is passed to *qsort(3)* to sort the completed array. If this pointer is null, the array is not sorted. *Alphasort* is a routine which can be used for the *compar* parameter to sort the array alphabetically.

The memory allocated for the array can be deallocated with *free* (see *malloc(3)*) by freeing each pointer in the array and the array itself.

SEE ALSO

directory(3), *malloc(3)*, *qsort(3)*, *dir(5)*

DIAGNOSTICS

Returns `-1` if the directory cannot be opened for reading or if *malloc(3)* cannot allocate enough memory to hold all the data structures.

NAME

setjmp, *longjmp* — non-local goto

SYNOPSIS

```
#include <setjmp.h>
```

```
setjmp(env)
```

```
jmp_buf env;
```

```
longjmp(env, val)
```

```
jmp_buf env;
```

```
_setjmp(env)
```

```
jmp_buf env;
```

```
_longjmp(env, val)
```

```
jmp_buf env;
```

DESCRIPTION

These routines 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*. It returns value 0.

Longjmp restores the environment saved by the last call of *setjmp*. It then returns in such a way that execution continues as if the call of *setjmp* had just returned the value *val* to the function that invoked *setjmp*, which must not itself have returned in the interim. All accessible data have values as of the time *longjmp* was called.

Setjmp and *longjmp* save and restore the signal mask *sigmask(2)*, while *_setjmp* and *_longjmp* manipulate only the C stack and registers.

SEE ALSO

sigvec(2), *sigstack(2)*, *signal(3)*

BUGS

Setjmp does not save current notion of whether the process is executing on the signal stack. The result is that a *longjmp* to some place on the signal stack leaves the signal stack state incorrect.

NAME

setuid, seteuid, setruid, setgid, setegid, setrgid — set user and group ID

SYNOPSIS

setuid(uid)
seteuid(euid)
setruid(ruid)
setgid(gid)
setegid(egid)
setrgid(rgid)

DESCRIPTION

Setuid (setgid) sets both the real and effective user ID (group ID) of the current process to as specified.

Seteuid (setegid) sets the effective user ID (group ID) of the current process.

Setruid (setruid) sets the real user ID (group ID) of the current process.

These calls are only permitted to the super-user or if the argument is the real or effective ID.

SEE ALSO

setreuid(2), setregid(2), getuid(2), getgid(2)

DIAGNOSTICS

Zero is returned if the user (group) ID is set; -1 is returned otherwise.

NAME

sleep — suspend execution for interval

SYNOPSIS

sleep(seconds)
unsigned seconds;

DESCRIPTION

The current process is suspended from execution for the number of seconds specified by the argument. The actual suspension time may be up to 1 second less than that requested, because scheduled wakeups occur at fixed 1-second intervals, and an arbitrary amount longer because of other activity in the system.

The routine is implemented by setting an interval timer and pausing until it occurs. The previous state of this timer is saved and restored. If the sleep time exceeds the time to the expiration of the previous timer, the process sleeps only until the signal would have occurred, and the signal is sent 1 second later.

SEE ALSO

setitimer(2), sigpause(2)

BUGS

An interface with finer resolution is needed.

NAME

strcat, *strncat*, *strcmp*, *strncmp*, *strcpy*, *strncpy*, *strlen*, *index*, *rindex* — string operations

SYNOPSIS

```
#include <strings.h>
char *strcat(s1, s2)
char *s1, *s2;
char *strncat(s1, s2, n)
char *s1, *s2;
strcmp(s1, s2)
char *s1, *s2;
strncmp(s1, s2, n)
char *s1, *s2;
char *strcpy(s1, s2)
char *s1, *s2;
char *strncpy(s1, s2, n)
char *s1, *s2;
strlen(s)
char *s;
char *index(s, c)
char *s, c;
char *rindex(s, c)
char *s, c;
```

DESCRIPTION

These functions operate on null-terminated strings. They do not check for overflow of any receiving string.

Strcat appends a copy of string *s2* to the end of string *s1*. *Strncat* copies at most *n* characters. Both return a pointer to the null-terminated result.

Strcmp compares its arguments and returns an integer greater than, equal to, or less than 0, according as *s1* is lexicographically greater than, equal to, or less than *s2*. *Strncmp* makes the same comparison but looks at at most *n* characters.

Strcpy copies string *s2* to *s1*, stopping after the null character has been moved. *Strncpy* copies exactly *n* characters, truncating or null-padding *s2*; the target may not be null-terminated if the length of *s2* is *n* or more. Both return *s1*.

Strlen returns the number of non-null characters in *s*.

Index (*rindex*) returns a pointer to the first (last) occurrence of character *c* in string *s*, or zero if *c* does not occur in the string.

NAME

swab — swap bytes

SYNOPSIS

swab(*from*, *to*, *nbytes*)
char **from*, **to*;

DESCRIPTION

Swab copies *nbytes* bytes pointed to by *from* to the position pointed to by *to*, exchanging adjacent even and odd bytes. It is useful for carrying binary data between PDP11's and other machines. *Nbytes* should be even.

NAME

system — issue a shell command

SYNOPSIS

```
system(string)  
char *string;
```

DESCRIPTION

System causes the *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.

SEE ALSO

popen(3S), *execve*(2), *wait*(2)

DIAGNOSTICS

Exit status 127 indicates the shell couldn't be executed.

NAME

ttyname, *isatty*, *ttyslot* — find name of a terminal

SYNOPSIS

char **ttyname*(*filedes*)

***isatty*(*filedes*)**

***ttyslot*()**

DESCRIPTION

Ttyname returns a pointer to the null-terminated path name of the terminal device associated with file descriptor *filedes* (this is a system file descriptor and has nothing to do with the standard I/O FILE typedef).

Isatty returns 1 if *filedes* is associated with a terminal device, 0 otherwise.

Ttyslot returns the number of the entry in the *ttys(5)* file for the control terminal of the current process.

FILES

*/dev/**

/etc/ttys

SEE ALSO

ioctl(2), *ttys(5)*

DIAGNOSTICS

Ttyname returns a null pointer (0) if *filedes* does not describe a terminal device in directory *'/dev'*.

Ttyslot returns 0 if *'/etc/ttys'* is inaccessible or if it cannot determine the control terminal.

BUGS

The return value points to static data whose content is overwritten by each call.

NAME

varargs — variable argument list

SYNOPSIS

```
#include <varargs.h>

function(va_alist)
va_dcl
va_list pvar;
va_start(pvar);
f = va_arg(pvar, type);
va_end(pvar);
```

DESCRIPTION

This set of macros provides a means of writing portable procedures that accept variable argument lists. Routines having variable argument lists (such as *printf(3)*) that do not use varargs are inherently nonportable, since different machines use different argument passing conventions.

va_alist is used in a function header to declare a variable argument list.

va_dcl is a declaration for *va_alist*. Note that there is no semicolon after *va_dcl*.

va_list is a type which can be used for the variable *pvar*, which is used to traverse the list. One such variable must always be declared.

va_start(pvar) is called to initialize *pvar* to the beginning of the list.

va_arg(pvar, type) 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, since it cannot be determined at runtime.

va_end(pvar) is used to finish up.

Multiple traversals, each bracketed by *va_start ... va_end*, are possible.

EXAMPLE

```
#include <varargs.h>
exec1(va_alist)
va_dcl
{
    va_list ap;
    char *file;
    char *args[100];
    int argno = 0;

    va_start(ap);
    file = va_arg(ap, char *);
    while (args[argno++] = va_arg(ap, char *))
        ;
    va_end(ap);
    return execv(file, args);
}
```

BUGS

It is up to the calling routine to determine how many arguments there are, since it is not possible to determine this from the stack frame. For example, *exec1* passes a 0 to signal the end of the list. *Printf* can tell how many arguments are supposed to be there by the format.

NAME

intro — introduction to mathematical library functions

DESCRIPTION

These functions constitute the math library, *libm*. They are automatically loaded as needed by the Fortran compiler *f77(1)*. The link editor searches this library under the “-lm” option. Declarations for these functions may be obtained from the include file *<math.h>*.

LIST OF FUNCTIONS

<i>Name</i>	<i>Appears on Page</i>	<i>Description</i>
acos	sin.3m	trigonometric functions
asin	sin.3m	trigonometric functions
atan	sin.3m	trigonometric functions
atan2	sin.3m	trigonometric functions
cabs	hypot.3m	Euclidean distance
ceil	floor.3m	absolute value, floor, ceiling functions
cos	sin.3m	trigonometric functions
cosh	sinh.3m	hyperbolic functions
exp	exp.3m	exponential, logarithm, power, square root
fabs	floor.3m	absolute value, floor, ceiling functions
floor	floor.3m	absolute value, floor, ceiling functions
gamma	gamma.3m	log gamma function
hypot	hypot.3m	Euclidean distance
j0	j0.3m	bessel functions
j1	j0.3m	bessel functions
jn	j0.3m	bessel functions
log	exp.3m	exponential, logarithm, power, square root
log10	exp.3m	exponential, logarithm, power, square root
pow	exp.3m	exponential, logarithm, power, square root
sin	sin.3m	trigonometric functions
sinh	sinh.3m	hyperbolic functions
sqrt	exp.3m	exponential, logarithm, power, square root
tan	sin.3m	trigonometric functions
tanh	sinh.3m	hyperbolic functions
y0	j0.3m	bessel functions
y1	j0.3m	bessel functions
yn	j0.3m	bessel functions

NAME

`exp`, `log`, `log10`, `pow`, `sqrt` — exponential, logarithm, power, square root

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 the exponential function of x .

Log returns the natural logarithm of x ; *log10* returns the base 10 logarithm.

Pow returns x^y .

Sqrt returns the square root of x .

SEE ALSO

`hypot(3M)`, `sinh(3M)`, `intro(3M)`

DIAGNOSTICS

Exp and *pow* return a huge value when the correct value would overflow; *errno* is set to ERANGE. *Pow* returns 0 and sets *errno* to EDOM when the second argument is negative and non-integral and when both arguments are 0.

Log returns 0 when x is zero or negative; *errno* is set to EDOM.

Sqrt returns 0 when x is negative; *errno* is set to EDOM.

NAME

fabs, *floor*, *ceil* — absolute value, floor, ceiling functions

SYNOPSIS

```
#include <math.h>
```

```
double floor(x)
```

```
double x;
```

```
double ceil(x)
```

```
double x;
```

```
double fabs(x)
```

```
double x;
```

DESCRIPTION

Fabs returns the absolute value $|x|$.

Floor returns the largest integer not greater than x .

Ceil returns the smallest integer not less than x .

SEE ALSO

abs(3)

NAME

gamma — log gamma function

SYNOPSIS

```
#include <math.h>
double gamma(x)
double x;
```

DESCRIPTION

Gamma returns $\ln |\Gamma(|x|)|$. The sign of $\Gamma(|x|)$ is returned in the external integer *signgam*. The following C program might be used to calculate Γ :

```
y = gamma(x);
if (y > 88.0)
    error();
y = exp(y);
if (signgam)
    y = -y;
```

DIAGNOSTICS

A huge value is returned for negative integer arguments.

BUGS

There should be a positive indication of error.

NAME

hypot, cabs — Euclidean distance

SYNOPSIS

```
#include <math.h>
double hypot(x, y)
double x, y;
double cabs(z)
struct { double x, y;} z;
```

DESCRIPTION

Hypot and *cabs* return
 $\text{sqrt}(x*x + y*y)$,
taking precautions against unwarranted overflows.

SEE ALSO

exp(3M) for *sqrt*

NAME

`j0`, `j1`, `jn`, `y0`, `y1`, `yn` — Bessel functions

SYNOPSIS

```
#include <math.h>

double j0(x)
double x;

double j1(x)
double x;

double jn(n, x)
double x;

double y0(x)
double x;

double y1(x)
double x;

double yn(n, x)
double x;
```

DESCRIPTION

These functions calculate Bessel functions of the first and second kinds for real arguments and integer orders.

DIAGNOSTICS

Negative arguments cause `y0`, `y1`, and `yn` to return a huge negative value and set `errno` to `EDOM`.

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 asin(x)
double x;
double acos(x)
double x;
double atan(x)
double x;
double atan2(x, y)
double x, y;
```

DESCRIPTION

Sin, *cos* and *tan* return trigonometric functions of radian arguments. The magnitude of the argument should be checked by the caller to make sure the result is meaningful.

Asin returns the arc sin in the range $-\pi/2$ to $\pi/2$.

Acos returns the arc cosine in the range 0 to π .

Atan returns the arc tangent of x in the range $-\pi/2$ to $\pi/2$.

Atan2 returns the arc tangent of x/y in the range $-\pi$ to π .

DIAGNOSTICS

Arguments of magnitude greater than 1 cause *asin* and *acos* to return value 0; *errno* is set to EDOM. The value of *tan* at its singular points is a huge number, and *errno* is set to ERANGE.

BUGS

The value of *tan* for arguments greater than about $2^{*}31$ is garbage.

NAME

`sinh`, `cosh`, `tanh` — hyperbolic functions

SYNOPSIS

```
#include <math.h>
```

```
double sinh(x)
```

```
double cosh(x)
```

```
double x;
```

```
double tanh(x)
```

```
double x;
```

DESCRIPTION

These functions compute the designated hyperbolic functions for real arguments.

DIAGNOSTICS

Sinh and *cosh* return a huge value of appropriate sign when the correct value would overflow.

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;
```

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)*.

SEE ALSO

gethostent(3N), *getservent(3N)*

BUGS

The VAX handles bytes backwards from most everyone else in the world. This is not expected to be fixed in the near future.

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()
```

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

SEE ALSO

hosts(5)

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

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()
```

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

SEE ALSO

networks(5)

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

getprotoent, getprotobynumber, getprotobyname, setprotoent, endprotoent — get protocol entry

SYNOPSIS

```
#include <netdb.h>

struct protoent *getprotoent()
struct protoent *getprotobyname(name)
char *name;
struct protoent *getprotobynumber(proto)
int proto;
setprotoent(stayopen)
int stayopen
endprotoent()
```

DESCRIPTION

Getprotoent, *getprotobyname*, 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.

Getprotobyname 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

SEE ALSO

protocols(5)

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

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()
```

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

SEE ALSO

getprotoent(3N), services(5)

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

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;
```

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. Note that when an Internet address is viewed as a 32-bit integer quantity on the VAX the bytes referred to above appear as “d.c.b.a”. That is, VAX bytes are ordered from right to left.

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).

SEE ALSO

gethostent(3N), getnetent(3N), hosts(5), networks(5),

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

stdio — standard buffered input/output package

SYNOPSIS

```
#include <stdio.h>
```

```
FILE *stdin;
```

```
FILE *stdout;
```

```
FILE *stderr;
```

DESCRIPTION

The functions described in section 3S constitute a user-level buffering scheme. The in-line macros *getc* and *putc*(3S) handle characters quickly. The higher level routines *gets*, *fgets*, *scanf*, *fscanf*, *fread*, *puts*, *fputs*, *printf*, *sprintf*, *fwrite* 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. There are three normally open streams with constant pointers declared in the include file and associated with the standard open files:

```
stdin    standard input file
stdout   standard output file
stderr   standard error file
```

A constant 'pointer' **NULL** (0) designates no stream at all.

An integer constant **EOF** (-1) is returned upon end of file or error by integer functions that deal with streams.

Any routine that uses the standard input/output package must include the header file *<stdio.h>* of pertinent macro definitions. The functions and constants mentioned in sections labeled 3S are declared in the include file and need no further declaration. The constants, and the following 'functions' are implemented as macros; redeclaration of these names is perilous: *getc*, *getchar*, *putc*, *putchar*, *feof*, *ferror*, *fileno*.

SEE ALSO

open(2), *close*(2), *read*(2), *write*(2), *fread*(3S), *fseek*(3S), *f**(3S)

DIAGNOSTICS

The value **EOF** is returned uniformly to indicate that a **FILE** pointer has not been initialized with *fopen*, input (output) has been attempted on an output (input) stream, or a **FILE** pointer designates corrupt or otherwise unintelligible **FILE** data.

For purposes of efficiency, this implementation of the standard library has been changed to line buffer output to a terminal by default and attempts to do this transparently by flushing the output whenever a *read*(2) from the standard input is necessary. This is almost always transparent, but may cause confusion or malfunctioning of programs which use standard i/o routines but use *read*(2) themselves to read from the standard input.

In cases where a large amount of computation is done after printing part of a line on an output terminal, it is necessary to *fflush*(3S) the standard output before going off and computing so that the output will appear.

BUGS

The standard buffered functions do not interact well with certain other library and system functions, especially *vfork* and *abort*.

LIST OF FUNCTIONS

<i>Name</i>	<i>Appears on Page</i>	<i>Description</i>
<i>clearerr</i>	<i>ferror.3s</i>	stream status inquiries
<i>fclose</i>	<i>fclose.3s</i>	close or flush a stream

feof	ferror.3s	stream status inquiries
ferror	ferror.3s	stream status inquiries
fflush	fclose.3s	close or flush a stream
fgetc	getc.3s	get character or word from stream
fgets	gets.3s	get a string from a stream
fileno	ferror.3s	stream status inquiries
fprintf	printf.3s	formatted output conversion
fputc	putc.3s	put character or word on a stream
fputs	puts.3s	put a string on a stream
fread	fread.3s	buffered binary input/output
fscanf	scanf.3s	formatted input conversion
fseek	fseek.3s	reposition a stream
ftell	fseek.3s	reposition a stream
fwrite	fread.3s	buffered binary input/output
getc	getc.3s	get character or word from stream
getchar	getc.3s	get character or word from stream
gets	gets.3s	get a string from a stream
getw	getc.3s	get character or word from stream
printf	printf.3s	formatted output conversion
putc	putc.3s	put character or word on a stream
putchar	putc.3s	put character or word on a stream
puts	puts.3s	put a string on a stream
putw	putc.3s	put character or word on a stream
rewind	fseek.3s	reposition a stream
scanf	scanf.3s	formatted input conversion
setbuf	setbuf.3s	assign buffering to a stream
setbuffer	setbuf.3s	assign buffering to a stream
setlinebuf	setbuf.3s	assign buffering to a stream
sprintf	printf.3s	formatted output conversion
sscanf	scanf.3s	formatted input conversion
ungetc	ungetc.3s	push character back into input stream

NAME

fclose, *fflush* — close or flush a stream

SYNOPSIS

```
#include <stdio.h>
```

```
fclose(stream)
```

```
FILE *stream;
```

```
fflush(stream)
```

```
FILE *stream;
```

DESCRIPTION

Fclose causes any buffers for the named *stream* to be emptied, and the file to be closed. Buffers allocated by the standard input/output system are freed.

Fclose is performed automatically upon calling *exit*(3).

Fflush causes any buffered data for the named output *stream* to be written to that file. The stream remains open.

SEE ALSO

close(2), *fopen*(3S), *setbuf*(3S)

DIAGNOSTICS

These routines return EOF if *stream* is not associated with an output file, or if buffered data cannot be transferred to that file.

NAME

ferror, *feof*, *clearerr*, *fileno* — stream status inquiries

SYNOPSIS

```
#include <stdio.h>
```

```
feof(stream)
```

```
FILE *stream;
```

```
ferror(stream)
```

```
FILE *stream
```

```
clearerr(stream)
```

```
FILE *stream
```

```
fileno(stream)
```

```
FILE *stream;
```

DESCRIPTION

Feof returns non-zero when end of file is read on the named input *stream*, otherwise zero.

Ferror returns non-zero when an error has occurred reading or writing the named *stream*, otherwise zero. Unless cleared by *clearerr*, the error indication lasts until the stream is closed.

Clrerr resets the error indication on the named *stream*.

Fileno returns the integer file descriptor associated with the *stream*, see *open*(2).

These functions are implemented as macros; they cannot be redeclared.

SEE ALSO

fopen(3S), *open*(2)

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)
```

```
char *type;
```

DESCRIPTION

Fopen opens the file named by *filename* and associates a stream with it. *Fopen* returns a pointer to be used to identify the stream in subsequent operations.

Type is a character string having one of the following values:

"r" open for reading

"w" create for writing

"a" append: open for writing at end of file, or create for writing

In addition, each *type* may be followed by a '+' to have the file opened for reading and writing. "r+" positions the stream at the beginning of the file, "w+" creates or truncates it, and "a+" positions it at the end. Both reads and writes may be used on read/write streams, with the limitation that an *fseek*, *rewind*, or reading an end-of-file must be used between a read and a write or vice-versa.

Freopen substitutes the named file in place of the open *stream*. It returns the original value of *stream*. The original stream is closed.

Freopen is typically used to attach the preopened constant names, *stdin*, *stdout*, *stderr*, to specified files.

Fdopen associates a stream with a file descriptor obtained from *open*, *dup*, *creat*, or *pipe(2)*. The *type* of the stream must agree with the mode of the open file.

SEE ALSO

open(2), *fclose(3)*

DIAGNOSTICS

Fopen and *freopen* return the pointer NULL if *filename* cannot be accessed.

BUGS

Fdopen is not portable to systems other than UNIX.

The read/write *types* do not exist on all systems. Those systems without read/write modes will probably treat the *type* as if the '+' was not present. These are unreliable in any event.

NAME

fread, *fwrite* — buffered binary input/output

SYNOPSIS

```
#include <stdio.h>
```

```
fread(ptr, sizeof(*ptr), nitems, stream)
```

```
FILE *stream;
```

```
fwrite(ptr, sizeof(*ptr), nitems, stream)
```

```
FILE *stream;
```

DESCRIPTION

Fread reads, into a block beginning at *ptr*, *nitems* of data of the type of **ptr* from the named input *stream*. It returns the number of items actually read.

If *stream* is *stdin* and the standard output is line buffered, then any partial output line will be flushed before any call to *read(2)* to satisfy the *fread*.

Fwrite appends at most *nitems* of data of the type of **ptr* beginning at *ptr* to the named output *stream*. It returns the number of items actually written.

SEE ALSO

read(2), *write(2)*, *fopen(3S)*, *getc(3S)*, *putc(3S)*, *gets(3S)*, *puts(3S)*, *printf(3S)*, *scanf(3S)*

DIAGNOSTICS

Fread and *fwrite* return 0 upon end of file or error.

NAME

fseek, *ftell*, *rewind* — reposition a stream

SYNOPSIS

```
#include <stdio.h>

fseek(stream, offset, ptrname)
FILE *stream;
long offset;

long ftell(stream)
FILE *stream;

rewind(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, according as *ptrname* has the value 0, 1, or 2.

Fseek undoes any effects of *ungetc*(3S).

Ftell returns the current value of the offset relative to the beginning of the file associated with the named *stream*. It is measured in bytes on UNIX; on some other systems it is a magic cookie, and the only foolproof way to obtain an *offset* for *fseek*.

Rewind(*stream*) is equivalent to *fseek*(*stream*, 0L, 0).

SEE ALSO

lseek(2), *fopen*(3S)

DIAGNOSTICS

Fseek returns -1 for improper seeks.

NAME

getc, *getchar*, *fgetc*, *getw* — get character or word from 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 from the named input *stream*.

Getchar() is identical to *getc(stdin)*.

Fgetc behaves like *getc*, but is a genuine function, not a macro; it may be used to save object text.

Getw returns the next word (in a 32-bit integer on a VAX-11) from the named input *stream*. It returns the constant EOF upon end of file or error, but since that is a good integer value, *feof* and *ferror(3S)* should be used to check the success of *getw*. *Getw* assumes no special alignment in the file.

SEE ALSO

fopen(3S), *putc(3S)*, *gets(3S)*, *scanf(3S)*, *fread(3S)*, *ungetc(3S)*

DIAGNOSTICS

These functions return the integer constant EOF at end of file or upon read error.

A stop with message, 'Reading bad file', means an attempt has been made to read from a stream that has not been opened for reading by *fopen*.

BUGS

The end-of-file return from *getchar* is incompatible with that in UNIX editions 1-6.

Because it is implemented as a macro, *getc* treats a *stream* argument with side effects incorrectly. In particular, '*getc(*f++);*' doesn't work sensibly.

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;
```

```
FILE *stream;
```

DESCRIPTION

Gets reads a string into *s* from the standard input stream **stdin**. The string is terminated by a newline character, which is replaced in *s* by a null character. *Gets* returns its argument.

Fgets reads *n*−1 characters, or up to a newline character, whichever comes first, from the *stream* into the string *s*. The last character read into *s* is followed by a null character. *Fgets* returns its first argument.

SEE ALSO

puts(3S), getc(3S), scanf(3S), fread(3S), ferror(3S)

DIAGNOSTICS

Gets and *fgets* return the constant pointer **NULL** upon end of file or error.

BUGS

Gets deletes a newline, *fgets* keeps it, all in the name of backward compatibility.

NAME

printf, fprintf, sprintf — formatted output conversion

SYNOPSIS

```
#include <stdio.h>
printf(format [, arg ] ... )
char *format;

fprintf(stream, format [, arg ] ... )
FILE *stream;
char *format;

sprintf(s, format [, arg ] ... )
char *s, format;

#include <varargs.h>
_doprnt(format, args, stream)
char *format;
va_list *args;
FILE *stream;
```

DESCRIPTION

Printf places output on the standard output stream `stdout`. *Fprintf* places output on the named output *stream*. *Sprintf* places ‘output’ in the string *s*, followed by the character ‘\0’. All of these routines work by calling the internal routine `_doprnt`, using the variable-length argument facilities of *varargs*(3).

Each of these functions converts, formats, and prints its arguments after the first under control of the first argument. The first argument is a character string which contains two types of objects: plain characters, which are simply copied to the output stream, and conversion specifications, each of which causes conversion and printing of the next successive *arg printf*.

Each conversion specification is introduced by the character `%`. Following the `%`, there may be

- an optional minus sign ‘-’ which specifies *left adjustment* of the converted value in the indicated field;
- an optional digit string specifying a *field width*; if the converted value has fewer characters than the field width it will be blank-padded on the left (or right, if the left-adjustment indicator has been given) to make up the field width; if the field width begins with a zero, zero-padding will be done instead of blank-padding;
- an optional period ‘.’ which serves to separate the field width from the next digit string;
- an optional digit string specifying a *precision* which specifies the number of digits to appear after the decimal point, for e- and f-conversion, or the maximum number of characters to be printed from a string;
- an optional ‘#’ character specifying that the value should be converted to an “alternate form”. For c, d, s, and u, conversions, this option has no effect. For o conversions, the precision of the number is increased to force the first character of the output string to a zero. For x(X) conversion, a non-zero result has the string `0x(0X)` prepended 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 only appears in the results of those conversions if a digit follows the decimal point). For g and G conversions, trailing zeros are not removed from the result as they would otherwise be.
- the character l specifying that a following d, o, x, or u corresponds to a long integer *arg*.
- a character which indicates the type of conversion to be applied.

A field width or precision may be '*' instead of a digit string. In this case an integer *arg* supplies the field width or precision.

The conversion characters and their meanings are

- d** **ox** The integer *arg* is converted to decimal, octal, or hexadecimal notation respectively.
- f** The float or double *arg* is converted to decimal notation in the style '[−]ddd.ddd' where the number of d's after the decimal point is equal to the precision specification for the argument. If the precision is missing, 6 digits are given; if the precision is explicitly 0, no digits and no decimal point are printed.
- 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 after is equal to the precision specification for the argument; when the precision is missing, 6 digits are produced.
- g** The float or double *arg* is printed in style **d**, in style **f**, or in style **e**, whichever gives full precision in minimum space.
- c** The character *arg* is printed.
- s** *Arg* is taken to be a string (character pointer) and characters from the string are printed until a null character or until the number of characters indicated by the precision specification is reached; however if the precision is 0 or missing all characters up to a null are printed.
- u** The unsigned integer *arg* is converted to decimal and printed (the result will be in the range 0 through MAXUINT, where MAXUINT equals 4294967295 on a VAX-11 and 65535 on a PDP-11).
- %** Print a '%'; no argument is converted.

In no case does a non-existent or small field width cause truncation of a field; padding takes place only if the specified field width exceeds the actual width. Characters generated by *printf* are printed by *putc*(3S).

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, %02d:%02d", weekday, month, day, hour, min);
```

To print π to 5 decimals:

```
printf("pi = %.5f", 4*atan(1.0));
```

SEE ALSO

putc(3S), *scanf*(3S), *ecvt*(3)

BUGS

Very wide fields (>128 characters) fail.

NAME

`putc`, `putchar`, `fputc`, `putw` — put character or word on a stream

SYNOPSIS

```
#include <stdio.h>
```

```
int putc(c, stream)
```

```
char c;
```

```
FILE *stream;
```

```
putchar(c)
```

```
fputc(c, stream)
```

```
FILE *stream;
```

```
putw(w, stream)
```

```
FILE *stream;
```

DESCRIPTION

Putc appends the character *c* to the named output *stream*. It returns the character written.

Putchar(c) is defined as *putc(c, stdout)*.

Fputc behaves like *putc*, but is a genuine function rather than a macro.

Putw appends word (that is, `int`) *w* to the output *stream*. It returns the word written. *Putw* neither assumes nor causes special alignment in the file.

SEE ALSO

`fopen(3S)`, `fclose(3S)`, `getc(3S)`, `puts(3S)`, `printf(3S)`, `fread(3S)`

DIAGNOSTICS

These functions return the constant `EOF` upon error. Since this is a good integer, `ferror(3S)` should be used to detect *putw* errors.

BUGS

Because it is implemented as a macro, *putc* treats a *stream* argument with side effects improperly. In particular

```
putc(c, *f++);
```

doesn't work sensibly.

Errors can occur long after the call to *putc*.

NAME

puts, *fputs* — put a string on a stream

SYNOPSIS

```
#include <stdio.h>
```

```
puts(s)
```

```
char *s;
```

```
fputs(s, stream)
```

```
char *s;
```

```
FILE *stream;
```

DESCRIPTION

Puts copies the null-terminated string *s* to the standard output stream *stdout* and appends a newline character.

Fputs copies the null-terminated string *s* to the named output *stream*.

Neither routine copies the terminal null character.

SEE ALSO

fopen(3S), *gets*(3S), *putc*(3S), *printf*(3S), *ferror*(3S)

fread(3S) for *fwrite*

BUGS

Puts appends a newline, *fputs* does not, all in the name of backward compatibility.

NAME

scanf, fscanf, sscanf — formatted input conversion

SYNOPSIS

```
#include <stdio.h>

scanf(format [ , pointer ] . . . )
char *format;

fscanf(stream, format [ , pointer ] . . . )
FILE *stream;
char *format;

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 a format, and stores the results in its arguments. Each expects as 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. Blanks, tabs or newlines, which match optional white space in the input.
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 suppressing character *, an optional numerical maximum field width, and a conversion character.

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 was indicated by *. An input field is defined as a string of non-space characters; it extends to the next inappropriate character or until the field width, if specified, is exhausted.

The conversion character indicates the interpretation of the input field; the corresponding pointer argument must usually be of a restricted type. The following conversion characters 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.
- 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.
- s** a character string is expected; the corresponding argument should be a character pointer pointing to an array of characters large enough to accept the string and a terminating '\0', which will be added. The input field is terminated by a space character or a newline.
- c** a character is expected; the corresponding argument should be a character pointer. The normal skip over space characters is suppressed in this case; to read the next non-space character, try '%ls'. If a field width is given, the corresponding argument should refer to a character array, and the indicated number of characters is read.
- e** 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 e followed by

an optionally signed integer.

- [indicates a string not to be delimited by space characters. The left bracket is followed by a set of characters and a right bracket; the characters between the brackets define a set of characters making up the string. If the first character is not circumflex (^), the input field is all characters until the first character not in the set between the brackets; if the first character after the left bracket is ^, the input field is all characters until the first character which is in the remaining set of characters between the brackets. The corresponding argument must point to a character array.

The conversion characters **d**, **o** and **x** may be capitalized or preceded by **l** to indicate that a pointer to **long** rather than to **int** is in the argument list. Similarly, the conversion characters **e** or **f** may be capitalized or preceded by **l** to indicate a pointer to **double** rather than to **float**. The conversion characters **d**, **o** and **x** may be preceded by **h** to indicate a pointer to **short** rather than to **int**.

The *scanf* functions return the number of successfully matched and assigned input items. This can be used to decide how many input items were found. The constant **EOF** is returned upon end of input; note that this is different from 0, which means that no conversion was done; if conversion was intended, it was frustrated by an inappropriate character in the input.

For example, the call

```
int i; float x; char name[50];
scanf("%d%f%s", &i, &x, name);
```

with the input line

```
25 54.32E-1 thompson
```

will assign to *i* the value 25, *x* the value 5.432, and *name* will contain 'thompson\0'. Or,

```
int i; float x; char name[50];
scanf("%2d%f%d%[1234567890]", &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* will return 'a'.

SEE ALSO

atof(3), *getc*(3S), *printf*(3S)

DIAGNOSTICS

The *scanf* 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`, `setbuffer`, `setlinebuf` — assign buffering to a stream

SYNOPSIS

```
#include <stdio.h>

setbuf(stream, buf)
FILE *stream;
char *buf;

setbuffer(stream, buf, size)
FILE *stream;
char *buf;
int size;

setlinebuf(stream)
FILE *stream;
```

DESCRIPTION

The three types of buffering available are unbuffered, block buffered, and line buffered. When an output stream is unbuffered, information appears on the destination file or terminal as soon as written; when it is block buffered many characters are saved up and written as a block; when it is line buffered characters are saved up until a newline is encountered or input is read from `stdin`. *Flush* (see `fclose(3S)`) may be used to force the block out early. Normally all files are block buffered. A buffer is obtained from `malloc(3)` upon the first `getc` or `putc(3S)` on the file. If the standard stream `stdout` refers to a terminal it is line buffered. The standard stream `stderr` is always unbuffered.

Setbuf is used after a stream has been opened but before it is read or written. The character array *buf* is used instead of an automatically allocated buffer. If *buf* is the constant pointer `NULL`, input/output will be completely unbuffered. A manifest constant `BUFSIZ` tells how big an array is needed:

```
char buf[BUFSIZ];
```

Setbuffer, an alternate form of *setbuf*, is used after a stream has been opened but before it is read or written. The character array *buf* whose size is determined by the *size* argument is used instead of an automatically allocated buffer. If *buf* is the constant pointer `NULL`, input/output will be completely unbuffered.

Setlinebuf is used to change `stdout` or `stderr` from block buffered or unbuffered to line buffered. Unlike *setbuf* and *setbuffer* it can be used at any time that the file descriptor is active.

A file can be changed from unbuffered or line buffered to block buffered by using *freopen* (see `fopen(3S)`). A file can be changed from block buffered or line buffered to unbuffered by using *freopen* followed by *setbuf* with a buffer argument of `NULL`.

SEE ALSO

`fopen(3S)`, `getc(3S)`, `putc(3S)`, `malloc(3)`, `fclose(3S)`, `puts(3S)`, `printf(3S)`, `fread(3S)`

BUGS

The standard error stream should be line buffered by default.

The *setbuffer* and *setlinebuf* functions are not portable to non 4.2 BSD versions of UNIX.

NAME

`ungetc` — push character back into input stream

SYNOPSIS

```
#include <stdio.h>
```

```
ungetc(c, stream)
```

```
FILE *stream;
```

DESCRIPTION

Ungetc pushes the character *c* back on an input stream. That character will be returned by the next *getc* call on that stream. *Ungetc* returns *c*.

One character of pushback is guaranteed provided something has been read from the stream and the stream is actually buffered. Attempts to push EOF are rejected.

Fseek(3S) erases all memory of pushed back characters.

SEE ALSO

getc(3S), *setbuf*(3S), *fseek*(3S)

DIAGNOSTICS

Ungetc returns EOF if it can't push a character back.

NAME

intro — introduction to miscellaneous library functions

DESCRIPTION

These functions constitute minor libraries and other miscellaneous run-time facilities. Most are available only when programming in C. The list below includes libraries which provide device independent plotting functions, terminal independent screen management routines for two dimensional non-bitmap display terminals, functions for managing data bases with inverted indexes, and sundry routines used in executing commands on remote machines. The routines *getdiskbyname*, *rcmd*, *rresvport*, *ruserok*, and *rexec* reside in the standard C run-time library “-lc”. All other functions are located in separate libraries indicated in each manual entry.

FILES

/lib/libc.a
 /usr/lib/libdbm.a
 /usr/lib/libtermcap.a
 /usr/lib/libcurses.a
 /usr/lib/lib2648.a
 /usr/lib/libplot.a

LIST OF FUNCTIONS

<i>Name</i>	<i>Appears on Page</i>	<i>Description</i>
arc	plot.3x	graphics interface
assert	assert.3x	program verification
circle	plot.3x	graphics interface
closepi	plot.3x	graphics interface
cont	plot.3x	graphics interface
curses	curses.3x	screen functions with “optimal” cursor motion
dbminit	dbm.3x	data base subroutines
delete	dbm.3x	data base subroutines
endsent	getfsent.3x	get file system descriptor file entry
erase	plot.3x	graphics interface
fetch	dbm.3x	data base subroutines
firstkey	dbm.3x	data base subroutines
getdiskbyname	getdisk.3x	get disk description by its name
getfsent	getfsent.3x	get file system descriptor file entry
getfsfile	getfsent.3x	get file system descriptor file entry
getfsspec	getfsent.3x	get file system descriptor file entry
getfstype	getfsent.3x	get file system descriptor file entry
initgroups	initgroups.3x	initialize group access list
label	plot.3x	graphics interface
lib2648	lib2648.3x	subroutines for the HP 2648 graphics terminal
line	plot.3x	graphics interface
linemod	plot.3x	graphics interface
move	plot.3x	graphics interface
nextkey	dbm.3x	data base subroutines
plot: openpl	plot.3x	graphics interface
point	plot.3x	graphics interface
rcmd	rcmd.3x	routines for returning a stream to a remote command
rexec	rexec.3x	return stream to a remote command
rresvport	rcmd.3x	routines for returning a stream to a remote command
ruserok	rcmd.3x	routines for returning a stream to a remote command
setfsent	getfsent.3x	get file system descriptor file entry
space	plot.3x	graphics interface

store	dbm.3x	data base subroutines
tgetent	termcap.3x	terminal independent operation routines
tgetflag	termcap.3x	terminal independent operation routines
tgetnum	termcap.3x	terminal independent operation routines
tgetstr	termcap.3x	terminal independent operation routines
tgoto	termcap.3x	terminal independent operation routines
tputs	termcap.3x	terminal independent operation routines

NAME

curse_s — screen functions with “optimal” cursor motion

SYNOPSIS

cc [flags] files -lcurse_s -ltermcap [libraries]

DESCRIPTION

These routines give the user a method of updating screens with reasonable optimization. They keep an image of the current screen, and the user sets up an image of a new one. Then the *refresh()* tells the routines to make the current screen look like the new one. 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.

SEE ALSO

Screen Updating and Cursor Movement Optimization: A Library Package, Ken Arnold, *ioctl(2)*, *getenv(3)*, *tty(4)*, *termcap(5)*

AUTHOR

Ken Arnold

FUNCTIONS

<i>addch(ch)</i>	add a character to <i>stdscr</i>
<i>addstr(str)</i>	add a string to <i>stdscr</i>
<i>box(win,vert,hor)</i>	draw a box around a window
<i>crmode()</i>	set cbreak mode
<i>clear()</i>	clear <i>stdscr</i>
<i>clearok(scr,boolf)</i>	set clear flag for <i>scr</i>
<i>clrtobot()</i>	clear to bottom on <i>stdscr</i>
<i>clrtoeol()</i>	clear to end of line on <i>stdscr</i>
<i>delch()</i>	delete a character
<i>deleteln()</i>	delete a line
<i>delwin(win)</i>	delete <i>win</i>
<i>echo()</i>	set echo mode
<i>endwin()</i>	end window modes
<i>erase()</i>	erase <i>stdscr</i>
<i>getch()</i>	get a char through <i>stdscr</i>
<i>getcap(name)</i>	get terminal capability <i>name</i>
<i>getstr(str)</i>	get a string through <i>stdscr</i>
<i>gettmode()</i>	get tty modes
<i>getyx(win,y,x)</i>	get (y,x) co-ordinates
<i>inch()</i>	get char at current (y,x) co-ordinates
<i>initscr()</i>	initialize screens
<i>insch(c)</i>	insert a char
<i>insertln()</i>	insert a line
<i>leaveok(win,boolf)</i>	set leave flag for <i>win</i>
<i>longname(termbuf,name)</i>	get long name from <i>termbuf</i>
<i>move(y,x)</i>	move to (y,x) on <i>stdscr</i>
<i>mvcur(lasty,lastx,newy,newx)</i>	actually move cursor
<i>newwin(lines,cols,begin_y,begin_x)</i>	create a new window
<i>nl()</i>	set newline mapping
<i>nocrmode()</i>	unset cbreak mode
<i>noecho()</i>	unset echo mode
<i>nonl()</i>	unset newline mapping
<i>noraw()</i>	unset raw mode
<i>overlay(win1,win2)</i>	overlay win1 on win2
<i>overwrite(win1,win2)</i>	overwrite win1 on top of win2

printw(fmt, arg1, arg2, ...)
 raw()
 refresh()
 resetty()
 savetty()
 scanw(fmt, arg1, arg2, ...)
 scroll(win)
 scrollok(win, boolf)
 setterm(name)
 standend()
 standout()
 subwin(win, lines, cols, begin_y, begin_x)
 touchwin(win)
 unctrl(ch)
 waddch(win, ch)
 waddstr(win, str)
 wclear(win)
 wclrtoebot(win)
 wclrtoeol(win)
 wdelch(win, c)
 wdeleteln(win)
 werase(win)
 wgetch(win)
 wgetstr(win, str)
 winch(win)
 winsch(win, c)
 winsertln(win)
 wmove(win, y, x)
 wprintw(win, fmt, arg1, arg2, ...)
 wrefresh(win)
 wscanw(win, fmt, arg1, arg2, ...)
 wstandend(win)
 wstandout(win)

printf on *stdscr*
 set raw mode
 make current screen look like *stdscr*
 reset tty flags to stored value
 stored current tty flags
 scanf through *stdscr*
 scroll *win* one line
 set scroll flag
 set term variables for name
 end standout mode
 start standout mode
 create a subwindow
 "change" all of *win*
 printable version of *ch*
 add char to *win*
 add string to *win*
 clear *win*
 clear to bottom of *win*
 clear to end of line on *win*
 delete char from *win*
 delete line from *win*
 erase *win*
 get a char through *win*
 get a string through *win*
 get char at current (y,x) in *win*
 insert char into *win*
 insert line into *win*
 set current (y,x) co-ordinates on *win*
 printf on *win*
 make screen look like *win*
 scanf through *win*
 end standout mode on *win*
 start standout mode on *win*

BUGS

NAME

getfsent, *getfsspec*, *getfsfile*, *getfstype*, *setfsent*, *endfsent* — get file system descriptor file entry

SYNOPSIS

```
#include <fstab.h>

struct fstab *getfsent()
struct fstab *getfsspec(spec)
char *spec;

struct fstab *getfsfile(file)
char *file;

struct fstab *getfstype(type)
char *type;

int setfsent()
int endfsent()
```

DESCRIPTION

Getfsent, *getfsspec*, *getfstype*, and *getfsfile* each return a pointer to an object with the following structure containing the broken-out fields of a line in the file system description file, <fstab.h>.

```
struct fstab{
    char    *fs_spec;
    char    *fs_file;
    char    *fs_type;
    int     fs_freq;
    int     fs_passno;
};
```

The fields have meanings described in *fstab(5)*.

Getfsent reads the next line of the file, opening the file if necessary.

Setfsent opens and rewinds the file.

Endfsent closes the file.

Getfsspec and *getfsfile* sequentially search from the beginning of the file until a matching special file name or file system file name is found, or until EOF is encountered. *Getfstype* does likewise, matching on the file system type field.

FILES

/etc/fstab

SEE ALSO

fstab(5)

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.

NAME

`initgroups` — initialize group access list

SYNOPSIS

```
initgroups(name, basegid)
char *name;
int basegid;
```

DESCRIPTION

Initgroups reads through the group file and sets up, using the *setgroups*(2) call, the group access list for the user specified in *name*. The *basegid* is automatically included in the groups list. Typically this value is given as the group number from the password file.

FILES

`/etc/group`

SEE ALSO

`setgroups`(2)

DIAGNOSTICS

Initgroups returns `-1` if it was not invoked by the super-user.

BUGS

Initgroups uses the routines based on *getgrent*(3). If the invoking program uses any of these routines, the group structure will be overwritten in the call to *initgroups*.

Noone seems to keep `/etc/group` up to date.

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)*. These are low level routines; see *curses(3X)* for a higher level package.

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 all goes well. It will look in the environment for a TERMCAP variable. If found, and the 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 it does begin with a slash, the string is used as a path name rather than *letc/termcap*. This can speed up entry into programs that call *tgetent*, as well as to help debug new terminal descriptions or to make one for your terminal if you can't write the file *letc/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 $\backslash n$, $\backslash D$ or $\backslash @$ in the returned string. (Programs which 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, *ouic* is a routine which is called with each character in turn. The external variable *ospeed* should contain the output speed of the terminal as encoded by *stty*(3). 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), curses(3X), termcap(5)

AUTHOR

William Joy

NAME

intro — introduction to compatibility library functions

DESCRIPTION

These functions constitute the compatibility library portion of *libc*. They are automatically loaded as needed by the C compiler *cc*(1). The link editor searches this library under the “-lc” option. Use of these routines should, for the most part, be avoided. Manual entries for the functions in this library describe the proper routine to use.

LIST OF FUNCTIONS

<i>Name</i>	<i>Appears on Page</i>	<i>Description</i>
alarm	alarm.3c	schedule signal after specified time
ftime	time.3c	get date and time
getpw	getpw.3c	get name from uid
gtty	stty.3c	set and get terminal state (defunct)
nice	nice.3c	set program priority
pause	pause.3c	stop until signal
rand	rand.3c	random number generator
signal	signal.3c	simplified software signal facilities
srand	rand.3c	random number generator
stty	stty.3c	set and get terminal state (defunct)
time	time.3c	get date and time
times	times.3c	get process times
utime	utime.3c	set file times
vlimit	vlimit.3c	control maximum system resource consumption
vtimes	vtimes.3c	get information about resource utilization

NAME

alarm — schedule signal after specified time

SYNOPSIS

alarm(seconds)
unsigned seconds;

DESCRIPTION

This interface is obsoleted by setitimer(2).

Alarm causes signal SIGALRM, see *signal(3C)*, to be sent to the invoking process in a number of seconds given by the argument. Unless caught or ignored, the signal terminates the process.

Alarm requests are not stacked; successive calls reset the alarm clock. If the argument is 0, any alarm request is canceled. 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 2147483647 seconds.

The return value is the amount of time previously remaining in the alarm clock.

SEE ALSO

sigpause(2), sigvec(2), signal(3C), sleep(3)

NAME

`getpw` — get name from uid

SYNOPSIS

```
getpw(uid, buf)
char *buf;
```

DESCRIPTION

`Getpw` is obsoleted by `getpwuid(3)`.

`Getpw` searches the password file for the (numerical) *uid*, and fills in *buf* with the corresponding line; it returns non-zero if *uid* could not be found. The line is null-terminated.

FILES

`/etc/passwd`

SEE ALSO

`getpwent(3)`, `passwd(5)`

DIAGNOSTICS

Non-zero return on error.

NAME

nice — set program priority

SYNOPSIS

nice(*incr*)

DESCRIPTION

This interface is obsoleted by *setpriority*(2).

The scheduling priority of the process is augmented by *incr*. Positive priorities get less service than normal. Priority 10 is recommended to users who wish to execute long-running programs without flak from the administration.

Negative increments are ignored except on behalf of the super-user. The priority is limited to the range -20 (most urgent) to 20 (least).

The priority of a process is passed to a child process by *fork*(2). For a privileged process to return to normal priority from an unknown state, *nice* should be called successively with arguments -40 (goes to priority -20 because of truncation), 20 (to get to 0), then 0 (to maintain compatibility with previous versions of this call).

SEE ALSO

nice(1), *setpriority*(2), *fork*(2), *renice*(8)

NAME

pause — stop until signal

SYNOPSIS

pause()

DESCRIPTION

Pause never returns normally. It is used to give up control while waiting for a signal from *kill(2)* or an interval timer, see *setitimer(2)*. Upon termination of a signal handler started during a *pause*, the *pause* call will return.

RETURN VALUE

Always returns -1.

ERRORS

Pause always returns:

[EINTR] The call was interrupted.

SEE ALSO

kill(2), select(2), sigpause(2)

NAME

rand, *srand* — random number generator

SYNOPSIS

```
srand(seed)
int seed;
rand()
```

DESCRIPTION

The newer *random(3)* should be used in new applications; *rand* remains for compatibility.

Rand uses a multiplicative congruential random number generator with period 2^{32} to return successive pseudo-random numbers in the range from 0 to $2^{31}-1$.

The generator is reinitialized by calling *srand* with 1 as argument. It can be set to a random starting point by calling *srand* with whatever you like as argument.

SEE ALSO

random(3)

NAME

signal — simplified software signal facilities

SYNOPSIS

```
#include <signal.h>

(*signal(sig, func)) ()
void (*func) ();
```

DESCRIPTION

Signal is a simplified interface to the more general *sigvec*(2) facility.

A signal is generated by some abnormal event, initiated by a user at a terminal (quit, interrupt, stop), by a program error (bus error, etc.), by request of another program (kill), or when a process is stopped because it wishes to access its control terminal while in the background (see *tty*(4)). Signals are optionally generated when a process resumes after being stopped, when the status of child processes changes, or when input is ready at the control terminal. Most signals cause termination of the receiving process if no action is taken; some signals instead cause the process receiving them to be stopped, or are simply discarded if the process has not requested otherwise. Except for the SIGKILL and SIGSTOP signals, the *signal* call allows signals either to be ignored or to cause an interrupt to a specified location. The following is a list of all signals with names as in the include file *<signal.h>*:

SIGHUP	1	hangup
SIGINT	2	interrupt
SIGQUIT	3*	quit
SIGILL	4*	illegal instruction
SIGTRAP	5*	trace trap
SIGIOT	6*	IOT instruction
SIGEMT	7*	EMT instruction
SIGFPE	8*	floating point exception
SIGKILL	9	kill (cannot be caught or ignored)
SIGBUS	10*	bus error
SIGSEGV	11*	segmentation violation
SIGSYS	12*	bad argument to system call
SIGPIPE	13	write on a pipe with no one to read it
SIGALRM	14	alarm clock
SIGTERM	15	software termination signal
SIGURG	16●	urgent condition present on socket
SIGSTOP	17†	stop (cannot be caught or ignored)
SIGTSTP	18†	stop signal generated from keyboard
SIGCONT	19●	continue after stop
SIGCHLD	20●	child status has changed
SIGTTIN	21†	background read attempted from control terminal
SIGTTOU	22†	background write attempted to control terminal
SIGIO	23●	i/o is possible on a descriptor (see <i>fcntl</i> (2))
SIGXCPU	24	cpu time limit exceeded (see <i>setrlimit</i> (2))
SIGXFSZ	25	file size limit exceeded (see <i>setrlimit</i> (2))
SIGVTALRM	26	virtual time alarm (see <i>setitimer</i> (2))
SIGPROF	27	profiling timer alarm (see <i>setitimer</i> (2))

The starred signals in the list above cause a core image if not caught or ignored.

If *func* is SIG_DFL, the default action for signal *sig* is reinstated; this default is termination (with a core image for starred signals) except for signals marked with ● or †. Signals marked with ● are discarded if the action is SIG_DFL; signals marked with † cause the process to stop. If *func* is SIG_IGN the signal is subsequently ignored and pending instances of the signal are

discarded. Otherwise, when the signal occurs further occurrences of the signal are automatically blocked and *func* is called.

A return from the function unblocks the handled signal and continues the process at the point it was interrupted. Unlike previous signal facilities, the handler *func* remains installed after a signal has been delivered.

If a caught signal occurs during certain system calls, causing the call to terminate prematurely, the call is automatically restarted. In particular this can occur during a *read* or *write(2)* on a slow device (such as a terminal; but not a file) and during a *wait(2)*.

The value of *signal* is the previous (or initial) value of *func* for the particular signal.

After a *fork(2)* or *vfork(2)* the child inherits all signals. *Execve(2)* resets all caught signals to the default action; ignored signals remain ignored.

RETURN VALUE

The previous action is returned on a successful call. Otherwise, -1 is returned and *errno* is set to indicate the error.

ERRORS

Signal will fail and no action will take place if one of the following occur:

- [EINVAL] *Sig* is not a valid signal number.
- [EINVAL] An attempt is made to ignore or supply a handler for SIGKILL or SIGSTOP.
- [EINVAL] An attempt is made to ignore SIGCONT (by default SIGCONT is ignored).

SEE ALSO

kill(1), *ptrace(2)*, *kill(2)*, *sigvec(2)*, *sigblock(2)*, *sigsetmask(2)*, *sigpause(2)*, *sigstack(2)*, *setjmp(3)*, *tty(4)*

NOTES (VAX-11)

The handler routine can be declared:

```
handler(sig, code, scp)
```

Here *sig* is the signal number, into which the hardware faults and traps are mapped as defined below. Code is a parameter which is either a constant as given below or, for compatibility mode faults, the code provided by the hardware. *Scp* is a pointer to the *struct sigcontext* used by the system to restore the process context from before the signal. Compatibility mode faults are distinguished from the other SIGILL traps by having PSL_CM set in the *psl*.

The following defines the mapping of hardware traps to signals and codes. All of these symbols are defined in *<signal.h>*:

Hardware condition	Signal	Code
Arithmetic traps:		
Integer overflow	SIGFPE	FPE_INTOVF_TRAP
Integer division by zero	SIGFPE	FPE_INTDIV_TRAP
Floating overflow trap	SIGFPE	FPE_FLTOVF_TRAP
Floating/decimal division by zero	SIGFPE	FPE_FLTDIV_TRAP
Floating underflow trap	SIGFPE	FPE_FLTUND_TRAP
Decimal overflow trap	SIGFPE	FPE_DECOVF_TRAP
Subscript-range	SIGFPE	FPE_SUBRNG_TRAP
Floating overflow fault	SIGFPE	FPE_FLTOVF_FAULT
Floating divide by zero fault	SIGFPE	FPE_FLTDIV_FAULT
Floating underflow fault	SIGFPE	FPE_FLTUND_FAULT
Length access control	SIGSEGV	
Protection violation	SIGBUS	

Reserved instruction	SIGILL	ILL_RESAD_FAULT
Customer-reserved instr.	SIGEMT	
Reserved operand	SIGILL	ILL_PRIVIN_FAULT
Reserved addressing	SIGILL	ILL_RESOP_FAULT
Trace pending	SIGTRAP	
Bpt instruction	SIGTRAP	
Compatibility-mode	SIGILL	hardware supplied code
Chme	SIGSEGV	
Chms	SIGSEGV	
Chmu	SIGSEGV	

NAME

stty, *gtty* — set and get terminal state (defunct)

SYNOPSIS

```
#include <sgtty.h>
```

```
stty(fd, buf)  
int fd;  
struct sgttyb *buf;
```

```
gtty(fd, buf)  
int fd;  
struct sgttyb *buf;
```

DESCRIPTION

This interface is obsoleted by *ioctl*(2).

Stty sets the state of the terminal associated with *fd*. *Gtty* retrieves the state of the terminal associated with *fd*. To set the state of a terminal the call must have write permission.

The *stty* call is actually “*ioctl*(*fd*, TIOCSETP, *buf*)”, while the *gtty* call is “*ioctl*(*fd*, TIOCGETP, *buf*)”. See *ioctl*(2) and *tty*(4) for an explanation.

DIAGNOSTICS

If the call is successful 0 is returned, otherwise -1 is returned and the global variable *errno* contains the reason for the failure.

SEE ALSO

ioctl(2), *tty*(4)

NAME

time, ftime — get date and time

SYNOPSIS

```
long time(0)
long time(tloc)
long *tloc;
#include <sys/types.h>
#include <sys/timeb.h>
ftime(tp)
struct timeb *tp;
```

DESCRIPTION

These interfaces are obsoleted by `gettimeofday(2)`.

Time returns the time since 00:00:00 GMT, Jan. 1, 1970, measured in seconds.

If *tloc* is nonnull, the return value is also stored in the place to which *tloc* points.

The *ftime* entry fills in a structure pointed to by its argument, as defined by `<sys/timeb.h>`:

```
/* timeb.h 6.183/07/29*/

/*
 * Structure returned by ftime system call
 */
struct timeb
{
    time_t    time;
    unsigned short millitm;
    short     timezone;
    short     dstflag;
};
```

The structure contains the time since the epoch in seconds, up to 1000 milliseconds of more-precise interval, the local time zone (measured in minutes of time westward from Greenwich), and a flag that, if nonzero, indicates that Daylight Saving time applies locally during the appropriate part of the year.

SEE ALSO

`date(1)`, `gettimeofday(2)`, `settimeofday(2)`, `ctime(3)`

NAME

times — get process times

SYNOPSIS

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

times(buffer)
struct tms *buffer;
```

DESCRIPTION

This interface is obsoleted by `getrusage(2)`.

Times returns time-accounting information for the current process and for the terminated child processes of the current process. All times are in 1/HZ seconds, where HZ is 60.

This is the structure returned by *times*:

```
/*    times.h 6.1    83/07/29    */

/*
 * Structure returned by times()
 */
struct tms {
    time_t tms_utime;        /* user time */
    time_t tms_stime;        /* system time */
    time_t tms_cutime;       /* user time, children */
    time_t tms_cstime;       /* system time, children */
};
```

The children times are the sum of the children's process times and their children's times.

SEE ALSO

`time(1)`, `getrusage(2)`, `wait3(2)`, `time(3)`

NAME

`utime` — set file times

SYNOPSIS

```
#include <sys/types.h>
utime(file, timep)
char *file;
time_t timep[2];
```

DESCRIPTION

This interface is **obsoleted by `utimes(2)`**.

The `utime` call uses the 'accessed' and 'updated' times in that order from the `timep` vector to set the corresponding recorded times for `file`.

The caller must be the owner of the file or the super-user. The 'inode-changed' time of the file is set to the current time.

SEE ALSO

`utimes(2)`, `stat(2)`

NAME

vlimit — control maximum system resource consumption

SYNOPSIS

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

```
vlimit(resource, value)
```

DESCRIPTION

This facility is superseded by `getrlimit(2)`.

Limits the consumption by the current process and each process it creates to not individually exceed *value* on the specified *resource*. If *value* is specified as `-1`, then the current limit is returned and the limit is unchanged. The resources which are currently controllable are:

LIM_NORAISE A pseudo-limit; if set non-zero then the limits may not be raised. Only the super-user may remove the *noraise* restriction.

LIM_CPU the maximum number of cpu-seconds to be used by each process

LIM_FSIZE the largest single file which can be created

LIM_DATA the maximum growth of the data+stack region via `sbrk(2)` beyond the end of the program text

LIM_STACK the maximum size of the automatically-extended stack region

LIM_CORE the size of the largest core dump that will be created.

LIM_MAXRSS a soft limit for the amount of physical memory (in bytes) to be given to the program. If memory is tight, the system will prefer to take memory from processes which are exceeding their declared **LIM_MAXRSS**.

Because this information is stored in the per-process information this system call must be executed directly by the shell if it is to affect all future processes created by the shell; *limit* is thus a built-in command to `csh(1)`.

The system refuses to extend the data or stack space when the limits would be exceeded in the normal way; a *break* call fails if the data space limit is reached, or the process is killed when the stack limit is reached (since the stack cannot be extended, there is no way to send a signal!).

A file i/o operation which would create a file which is too large will cause a signal **SIGXFSZ** to be generated, this normally terminates the process, but may be caught. When the cpu time limit is exceeded, a signal **SIGXCPU** is sent to the offending process; to allow it time to process the signal it is given 5 seconds grace by raising the cpu time limit.

SEE ALSO

`csh(1)`

BUGS

If **LIM_NORAISE** is set, then no grace should be given when the cpu time limit is exceeded.

There should be *limit* and *unlimit* commands in `sh(1)` as well as in `csh`.

This call is peculiar to this version of UNIX. The options and specifications of this system call and even the call itself are subject to change. It may be extended or replaced by other facilities in future versions of the system.

NAME

`vtimes` — get information about resource utilization

SYNOPSIS

```
vtimes(par_vm, ch_vm)
struct vtimes *par_vm, *ch_vm;
```

DESCRIPTION

This facility is superseded by `getrusage(2)`.

`Vtimes` returns accounting information for the current process and for the terminated child processes of the current process. Either `par_vm` or `ch_vm` or both may be 0, in which case only the information for the pointers which are non-zero is returned.

After the call, each buffer contains information as defined by the contents of the include file `usr/include/sys/vtimes.h`:

```
struct vtimes {
    int    vm_utime;           /* user time (*HZ) */
    int    vm_stime;         /* system time (*HZ) */
    /* divide next two by utime+stime to get averages */
    unsigned vm_idrss;       /* integral of d+s rss */
    unsigned vm_ixrss;       /* integral of text rss */
    int    vm_maxrss;        /* maximum rss */
    int    vm_majflt;        /* major page faults */
    int    vm_minflt;        /* minor page faults */
    int    vm_nswap;         /* number of swaps */
    int    vm_inblk;         /* block reads */
    int    vm_oublk;         /* block writes */
};
```

The `vm_utime` and `vm_stime` fields give the user and system time respectively in 60ths of a second (or 50ths if that is the frequency of wall current in your locality.) The `vm_idrss` and `vm_ixrss` measure memory usage. They are computed by integrating the number of memory pages in use each over cpu time. They are reported as though computed discretely, adding the current memory usage (in 512 byte pages) each time the clock ticks. If a process used 5 core pages over 1 cpu-second for its data and stack, then `vm_idrss` would have the value $5 \cdot 60$, where $vm_utime + vm_stime$ would be the 60. `vm_idrss` integrates data and stack segment usage, while `vm_ixrss` integrates text segment usage. `vm_maxrss` reports the maximum instantaneous sum of the text+data+stack core-resident page count.

The `vm_majflt` field gives the number of page faults which resulted in disk activity; the `vm_minflt` field gives the number of page faults incurred in simulation of reference bits; `vm_nswap` is the number of swaps which occurred. The number of file system input/output events are reported in `vm_inblk` and `vm_oublk`. These numbers account only for real i/o; data supplied by the caching mechanism is charged only to the first process to read or write the data.

SEE ALSO

`time(2)`, `wait3(2)`

BUGS

This call is peculiar to this version of UNIX. The options and specifications of this system call are subject to change. It may be extended to include additional information in future versions of the system.

NAME

intro — introduction to special files and hardware support

DESCRIPTION

This section describes the special files, related driver functions, and networking support available in the system. In this part of the manual, the SYNOPSIS section of each configurable device gives a sample specification for use in constructing a system description for the *config(8)* program. The DIAGNOSTICS section lists messages which may appear on the console and in the system error log */usr/adm/messages* due to errors in device operation.

This section contains both devices which may be configured into the system, "4" entries, and network related information, "4N", "4P", and "4F" entries; The networking support is introduced in *intro(4N)*.

VAX DEVICE SUPPORT

This section describes the hardware supported on the DEC VAX-11. Software support for these devices comes in two forms. A hardware device may be supported with a character or block *device driver*, or it may be used within the networking subsystem and have a *network interface driver*. Block and character devices are accessed through files in the file system of a special type; c.f. *mknod(8)*. Network interfaces are indirectly accessed through the interprocess communication facilities provided by the system; see *socket(2)*.

A hardware device is identified to the system at configuration time and the appropriate device or network interface driver is then compiled into the system. When the resultant system is booted, the autoconfiguration facilities in the system probe for the device on either the UNIBUS or MASSBUS and, if found, enable the software support for it. If a UNIBUS device does not respond at autoconfiguration time it is not accessible at any time afterwards. To enable a UNIBUS device which did not autoconfigure, the system will have to be rebooted. If a MASSBUS device comes "on-line" after the autoconfiguration sequence it will be dynamically autoconfigured into the running system.

The autoconfiguration system is described in *autoconf(4)*. VAX specific device support is described in "4V" entries. A list of the supported devices is given below.

SEE ALSO

intro(4), intro(4N), autoconf(4), config(8)

LIST OF DEVICES

The devices listed below are supported in this incarnation of the system. Devices are indicated by their functional interface. If second vendor products provide functionally identical interfaces they should be usable with the supplied software. (Beware however that we promise the software works **ONLY** with the hardware indicated on the appropriate manual page.)

acc	ACC LH/DH IMP communications interface
ad	Data translation A/D interface
css	DEC IMP-11A communications interface
ct	C/A/T phototypesetter
dh	DH-11 emulators, terminal multiplexor
dmc	DEC DMC-11/DMR-11 point-to-point communications device
dmf	DEC DMF-32 terminal multiplexor
dn	DEC DN-11 autodialer interface
dz	DZ-11 terminal multiplexor
ec	3Com 10Mb/s Ethernet controller
en	Xerox 3Mb/s Ethernet controller (obsolete)
kg	KL-11/DL-11W line clock
fl	VAX-11/780 console floppy interface
hk	RK6-11/RK06 and RK07 moving head disk

hp	MASSBUS disk interface (with RP06, RM03, RM05, etc.)
ht	TM03 MASSBUS tape drive interface (with TE-16, TU-45, TU-77)
hy	DR-11B or GI-13 interface to an NSC Hyperchannel
ik	Ikonas frame buffer graphics device interface
il	Interlan 10Mb/s Ethernet controller
lp	LP-11 parallel line printer interface
mt	TM78 MASSBUS tape drive interface
pcl	DEC PCL-11 communications interface
ps	Evans and Sutherland Picture System 2 graphics interface
rx	DEC RX02 floppy interface
tm	TM-11/TE-10 tape drive interface
ts	TS-11 tape drive interface
tu	VAX-11/730 TU58 console cassette interface
uda	DEC UDA-50 disk controller
un	DR-11W interface to Ungermann-Bass
up	Emulex SC-21V UNIBUS disk controller
ut	UNIBUS TU-45 tape drive interface
uu	TU58 dual cassette drive interface (DL11)
va	Benson-Varian printer/plotter interface
vp	Versatec printer/plotter interface
vv	Proteon proNET 10Mb/s ring network interface

NAME

networking — introduction to networking facilities

SYNOPSIS

```
#include <sys/socket.h>
#include <net/route.h>
#include <net/if.h>
```

DESCRIPTION

This section briefly describes the networking facilities available in the system. Documentation in this part of section 4 is broken up into three areas: *protocol-families*, *protocols*, and *network interfaces*. Entries describing a protocol-family are marked “4F”, while entries describing protocol use are marked “4P”. Hardware support for network interfaces are found among the standard “4” 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(2)* 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(2)*. 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. The SYNOPSIS section of each network interface entry gives a sample specification of the related drivers for use in providing a system description to the *config(8)* program. The DIAGNOSTICS section lists messages which may appear on the console and in the system error log */usr/adm/messages* due to errors in device operation.

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-1 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 `<net/route.h>`;

```
struct rtenry {
    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 (net 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(4)*, 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 *ifrequest* structure as its parameter. This structure has the form

```
struct ifreq {
    char    ifr_name[16];          /* name of interface (e.g. "ec0") */
    union {
        struct sockaddr ifru_addr;
        struct sockaddr ifru_dstaddr;
        short   ifru_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(2), ioctl(2), intro(4), config(8), routed(8C)

NAME

drum — paging device

DESCRIPTION

This file refers to the paging device in use by the system. This may actually be a subdevice of one of the disk drivers, but in a system with paging interleaved across multiple disk drives it provides an indirect driver for the multiple drives.

FILES

/dev/drum

BUGS

Reads from the drum are not allowed across the interleaving boundaries. Since these only occur every .5Mbytes or so, and since the system never allocates blocks across the boundary, this is usually not a problem.

NAME

ec — 3Com 10 Mb/s Ethernet interface

SYNOPSIS

device *ec0* at uba0 csr 161000 vector *ecrint eccollide ecxint*

DESCRIPTION

The *ec* interface provides access to a 10 Mb/s Ethernet network through a 3com controller.

The hardware has 32 kilobytes of dual-ported memory on the UNIBUS. This memory is used for internal buffering by the board, and the interface code reads the buffer contents directly through the UNIBUS.

The host's Internet address is specified at boot time with an SIOCSIFADDR ioctl. The *ec* interface employs the address resolution protocol described in *arp(4P)* to dynamically map between Internet and Ethernet addresses on the local network.

The interface software implements an exponential backoff algorithm when notified of a collision on the cable. This algorithm utilizes a 16-bit mask and the VAX-11's interval timer in calculating a series of random backoff values. The algorithm is as follows:

1. Initialize the mask to be all 1's.
2. If the mask is zero, 16 retries have been made and we give up.
3. Shift the mask left one bit and formulate a backoff by masking the interval timer with the mask (this is actually the two's complement of the value).
4. Use the value calculated in step 3 to delay before retransmitting the packet. The delay is done in a software busy loop.

The interface normally tries to use a "trailer" encapsulation to minimize copying data on input and output. This may be disabled, on a per-interface basis, by setting the IFF_NOTRAILERS flag with an SIOCSIFFLAGS ioctl.

DIAGNOSTICS

ec%d: send error. After 16 retransmissions using the exponential backoff algorithm described above, the packet was dropped.

ec%d: input error (offset=%d). The hardware indicated an error in reading a packet off the cable or an illegally sized packet. The buffer offset value is printed for debugging purposes.

ec%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(4N), *inet(4F)*, *arp(4P)*

BUGS

The PUP protocol family should be added.

The hardware is not capable of talking to itself. The software implements local sending and broadcast by sending such packets to the loop interface. This is a kludge.

Backoff delays are done in a software busy loop. This can degrade the system if the network experiences frequent collisions.

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(2)* call may also be used to fix the destination for future packets (in which case the *read(2)* or *recv(2)* and *write(2)* or *send(2)* 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(2), *recv(2)*, *intro(4N)*, *inet(4F)*

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 SIOCSI-FADDR 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(4N), inet(4F)

BUGS

It should handle all address and protocol families. An approved network address should be reserved for this interface.

NAME

lp — line printer

SYNOPSIS

device lp0 at uba0 csr 0177514 vector lpintr

DESCRIPTION

Lp provides the interface to any of the standard DEC line printers on an LP-11 parallel interface. When it is opened or closed, a suitable number of page ejects is generated. Bytes written are printed.

The unit number of the printer is specified by the minor device after removing the low 3 bits, which act as per-device parameters. Currently only the lowest of the low three bits is interpreted: if it is set, the device is treated as having a 64-character set, rather than a full 96-character set. In the resulting half-ASCII mode, lower case letters are turned into upper case and certain characters are escaped according to the following table:

{	{
}	}
`	~
	+
-	^

The driver correctly interprets carriage returns, backspaces, tabs, and form feeds. Lines longer than the maximum page width are truncated. The default page width is 132 columns. This may be overridden by specifying, for example, "flags 256".

FILES

/dev/lp

SEE ALSO

lpr(1)

DIAGNOSTICS

None.

NAME

mem, *kmem* — main memory

DESCRIPTION

Mem is a special file that is an image of the main memory of the computer. It may be used, for example, to examine (and even to patch) the system.

Byte addresses in *mem* are interpreted as physical memory addresses. References to non-existent locations cause errors to be returned.

Examining and patching device registers is likely to lead to unexpected results when read-only or write-only bits are present.

The file *kmem* is the same as *mem* except that kernel virtual memory rather than physical memory is accessed.

On PDP11's, the I/O page begins at location 0160000 of *kmem* and per-process data for the current process begins at 0140000. On VAX 11/780 the I/O space begins at physical address 20000000(16); on an 11/750 I/O space addresses are of the form *fxxxxx*(16); on all VAX'en per-process data for the current process is at virtual *7ffff000*(16).

FILES

/dev/mem
/dev/kmem

BUGS

On PDP11's and VAX's, memory files are accessed one byte at a time, an inappropriate method for some device registers.

NAME

mtio — UNIX magtape interface

DESCRIPTION

The files *mt0*, ..., *mt15* refer to the UNIX magtape drives, which may be on the MASSBUS using the TM03 formatter *ht(4)*, or TM78 formatter, *mt(4)*, or on the UNIBUS using either the TM11 or TS11 formatters *tm(4)*, TU45 compatible formatters, *ut(4)*, or *ts(4)*. The following description applies to any of the transport/controller pairs. The files *mt0*, ..., *mt7* are 800bpi, *mt8*, ..., *mt15* are 1600bpi, and *mt16*, ..., *mt23* are 6250bpi. (But note that only 1600 bpi is available with the TS11.) The files *mt0*, ..., *mt3*, *mt8*, ..., *mt11*, and *mt16*, ..., *mt19* are rewound when closed; the others are not. When a file open for writing is closed, two end-of-files are written. If the tape is not to be rewound it is positioned with the head between the two tapemarks.

A standard tape consists of a series of 1024 byte records terminated by an end-of-file. To the extent possible, the system makes it possible, if inefficient, to treat the tape like any other file. Seeks have their usual meaning and it is possible to read or write a byte at a time. Writing in very small units is inadvisable, however, because it tends to create monstrous record gaps.

The *mt* files discussed above are useful when it is desired to access the tape in a way compatible with ordinary files. When foreign tapes are to be dealt with, and especially when long records are to be read or written, the 'raw' interface is appropriate. The associated files are named *rmt0*, ..., *rmt23*, but the same minor-device considerations as for the regular files still apply. A number of other ioctl operations are available on raw magnetic tape. The following definitions are from `<sys/mtio.h>`:

```
/*
 * Structures and definitions for mag tape io control commands
 */

/* structure for MTIOCTOP - mag tape op command */
struct mtop {
    short  mt_op;          /* operations defined below */
    daddr_t mt_count;     /* how many of them */
};

/* operations */
#define MTWEOF  0  /* write an end-of-file record */
#define MTFSF  1  /* forward space file */
#define MTBSF  2  /* backward space file */
#define MTFSR  3  /* forward space record */
#define MTBSR  4  /* backward space record */
#define MTREW  5  /* rewind */
#define MTOFFL 6  /* rewind and put the drive offline */
#define MTNOP  7  /* no operation, sets status only */

/* structure for MTIOCGET - mag tape get status command */

struct mtget {
    short  mt_type;       /* type of magtape device */
    /* the following two registers are grossly device dependent */
    short  mt_dsreg;     /* "drive status" register */
    short  mt_erreg;     /* "error" register */
    /* end device-dependent registers */
    short  mt_resid;     /* residual count */
};
```

```

/* the following two are not yet implemented */
    daddr_t mt_fileno;    /* file number of current position */
    daddr_t mt_blkno;    /* block number of current position */
/* end not yet implemented */
};

/*
 * Constants for mt_type byte
 */
#define MT_ISTS          0x01
#define MT_ISHT          0x02
#define MT_ISTM          0x03
#define MT_ISMT          0x04
#define MT_ISUT          0x05
#define MT_ISCPC         0x06
#define MT_ISAR          0x07

/* mag tape io control commands */
#define MTIOCTOP         _IOW(m, 1, struct mtop)    /* do a mag tape op */
#define MTIOCGET         _IOR(m, 2, struct mtget)   /* get tape status */

#ifdef KERNEL
#define DEFTAPE          "/dev/rmt12"
#endif

```

Each *read* or *write* call reads or writes the next record on the tape. In the write case the record has the same length as the buffer given. During a read, the record size is passed back as the number of bytes read, provided it is no greater than the buffer size; if the record is long, an error is indicated. In raw tape I/O seeks are ignored. A zero byte count is returned when a tape mark is read, but another read will fetch the first record of the new tape file.

FILES

```

/dev/mt?
/dev/rmt?

```

SEE ALSO

mt(1), tar(1), tp(1), ht(4), tm(4), ts(4), mt(4), ut(4)

BUGS

The status should be returned in a device independent format.

NAME

null — data sink

DESCRIPTION

Data written on a null special file is discarded.

Reads from a null special file always return 0 bytes.

FILES

/dev/null

NAME

pty — pseudo terminal driver

SYNOPSIS

pseudo-device pty

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 *tty(4)*. However, whereas all other devices which provide the interface described in *tty(4)* 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.

In configuring, if no optional "count" is given in the specification, 16 pseudo terminal pairs are configured.

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_stopc* 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(1C)* and *rlogind(8C)* 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

tty — general terminal interface

SYNOPSIS

```
#include <sgtty.h>
```

DESCRIPTION

This section describes both a particular special file `/dev/tty` and the terminal drivers used for conversational computing.

Line disciplines.

The system provides different *line disciplines* for controlling communications lines. In this version of the system there are three disciplines available:

- old The old (standard) terminal driver. This is used when using the standard shell `sh(1)` and for compatibility with other standard version 7 UNIX systems.
- new A newer terminal driver, with features for job control; this must be used when using `cs(1)`.
- net A line discipline used for networking and loading data into the system over communications lines. It allows high speed input at very low overhead, and is described in `bk(4)`.

Line discipline switching is accomplished with the TIOCSETD *ioctl*:

```
int ldisc = LDISC; ioctl(filedes, TIOCSETD, &ldisc);
```

where LDISC is OTTYDISC for the standard tty driver, NTTYDISC for the new driver and NETLDISC for the networking discipline. The standard (currently old) tty driver is discipline 0 by convention. The current line discipline can be obtained with the TIOCGETD *ioctl*. Pending input is discarded when the line discipline is changed.

All of the low-speed asynchronous communications ports can use any of the available line disciplines, no matter what hardware is involved. The remainder of this section discusses the "old" and "new" disciplines.

The control terminal.

When a terminal file is opened, it causes the process to wait until a connection is established. In practice, user programs seldom open these files; they are opened by `init(8)` and become a user's standard input and output file.

If a process which has no control terminal opens a terminal file, then that terminal file becomes the control terminal for that process. The control terminal is thereafter inherited by a child process during a `fork(2)`, even if the control terminal is closed.

The file `/dev/tty` is, in each process, a synonym for a *control terminal* associated with that process. It is useful for programs that wish to be sure of writing messages on the terminal no matter how output has been redirected. It can also be used for programs that demand a file name for output, when typed output is desired and it is tiresome to find out which terminal is currently in use.

Process groups.

Command processors such as `cs(1)` can arbitrate the terminal between different *jobs* by placing related jobs in a single process group and associating this process group with the terminal. A terminals associated process group may be set using the TIOCSPGRP *ioctl(2)*:

```
ioctl(filedes, TIOCSPGRP, &pgrp)
```


or examined using TIOCGPGRP rather than TIOCSPGRP, returning the current process group in *pgrp*. The new terminal driver aids in this arbitration by restricting access to the terminal by processes which are not in the current process group; see **Job access control** below.

Modes.

The terminal drivers have three major modes, characterized by the amount of processing on the input and output characters:

- cooked** The normal mode. In this mode lines of input are collected and input editing is done. The edited line is made available when it is completed by a newline or when an EOT (control-D, hereafter **^D**) is entered. A carriage return is usually made synonymous with newline in this mode, and replaced with a newline whenever it is typed. All driver functions (input editing, interrupt generation, output processing such as delay generation and tab expansion, etc.) are available in this mode.
- CBREAK** This mode eliminates the character, word, and line editing input facilities, making the input character available to the user program as it is typed. Flow control, literal-next and interrupt processing are still done in this mode. Output processing is done.
- RAW** This mode eliminates all input processing and makes all input characters available as they are typed; no output processing is done either.

The style of input processing can also be very different when the terminal is put in non-blocking i/o mode; see *fcntl(2)*. In this case a *read(2)* from the control terminal will never block, but rather return an error indication (EWOULDBLOCK) if there is no input available.

A process may also request a SIGIO signal be sent it whenever input is present. To enable this mode the FASYNC flag should be set using *fcntl(2)*.

Input editing.

A UNIX terminal ordinarily operates in full-duplex mode. Characters may be typed at any time, even while output is occurring, and are only lost when the system's character input buffers become completely choked, which is rare, or when the user has accumulated the maximum allowed number of input characters that have not yet been read by some program. Currently this limit is 256 characters. In the old terminal driver all the saved characters are thrown away when the limit is reached, without notice; the new driver simply refuses to accept any further input, and rings the terminal bell.

Input characters are normally accepted in either even or odd parity with the parity bit being stripped off before the character is given to the program. By clearing either the EVEN or ODD bit in the flags word it is possible to have input characters with that parity discarded (see the **Summary** below.)

In all of the line disciplines, it is possible to simulate terminal input using the TIOCSTI ioctl, which takes, as its third argument, the address of a character. The system pretends that this character was typed on the argument terminal, which must be the control terminal except for the super-user (this call is not in standard version 7 UNIX).

Input characters are normally echoed by putting them in an output queue as they arrive. This may be disabled by clearing the ECHO bit in the flags word using the *stty(3)* call or the TIOCSETN or TIOCSETP ioctls (see the **Summary** below).

In cooked mode, terminal input is processed in units of lines. A program attempting to read will normally be suspended until an entire line has been received (but see the description of SIGTTIN in **Modes** above and FIONREAD in **Summary** below.) No matter how many characters are requested in the read call, at most one line will be returned. It is not, however, necessary to read a whole line at once; any number of characters may be requested in a read, even one, without losing information.

During input, line editing is normally done, with the character '#' logically erasing the last character typed and the character '@' logically erasing the entire current input line. These are often reset on crt's, with ^H replacing #, and ^U replacing @. These characters never erase beyond the beginning of the current input line or an ^D. These characters may be entered literally by preceding them with '\'; in the old teletype driver both the '\' and the character entered literally will appear on the screen; in the new driver the '\' will normally disappear.

The drivers normally treat either a carriage return or a newline character as terminating an input line, replacing the return with a newline and echoing a return and a line feed. If the CRMOD bit is cleared in the local mode word then the processing for carriage return is disabled, and it is simply echoed as a return, and does not terminate cooked mode input.

In the new driver there is a literal-next character ^V which can be typed in both cooked and CBREAK mode preceding any character to prevent its special meaning. This is to be preferred to the use of '\' escaping erase and kill characters, but '\' is (at least temporarily) retained with its old function in the new driver for historical reasons.

The new terminal driver also provides two other editing characters in normal mode. The word-erase character, normally ^W, erases the preceding word, but not any spaces before it. For the purposes of ^W, a word is defined as a sequence of non-blank characters, with tabs counted as blanks. Finally, the reprint character, normally ^R, retypes the pending input beginning on a new line. Retyping occurs automatically in cooked mode if characters which would normally be erased from the screen are fouled by program output.

Input echoing and redisplay

In the old terminal driver, nothing special occurs when an erase character is typed; the erase character is simply echoed. When a kill character is typed it is echoed followed by a new-line (even if the character is not killing the line, because it was preceded by a '\'.)

The new terminal driver has several modes for handling the echoing of terminal input, controlled by bits in a local mode word.

Hardcopy terminals. When a hardcopy terminal is in use, the LPRTERA bit is normally set in the local mode word. Characters which are logically erased are then printed out backwards preceded by '\' and followed by '/' in this mode.

Crt terminals. When a crt terminal is in use, the LCRTBS bit is normally set in the local mode word. The terminal driver then echoes the proper number of erase characters when input is erased; in the normal case where the erase character is a ^H this causes the cursor of the terminal to back up to where it was before the logically erased character was typed. If the input has become fouled due to interspersed asynchronous output, the input is automatically retyped.

Erasing characters from a crt. When a crt terminal is in use, the LCRTERA bit may be set to cause input to be erased from the screen with a "backspace-space-backspace" sequence when character or word deleting sequences are used. A LCRTKIL bit may be set as well, causing the input to be erased in this manner on line kill sequences as well.

Echoing of control characters. If the LCTLECH bit is set in the local state word, then non-printing (control) characters are normally echoed as ^X (for some X) rather than being echoed unmodified; delete is echoed as ^?.

The normal modes for using the new terminal driver on crt terminals are speed dependent. At speeds less than 1200 baud, the LCRTERA and LCRTKILL processing is painfully slow, so *stty(1)* normally just sets LCRTBS and LCTLECH; at speeds of 1200 baud or greater all of these bits are normally set. *Stty(1)* summarizes these option settings and the use of the new terminal driver as "newcrt."

Output processing.

When one or more characters are written, they are actually transmitted to the terminal as soon as previously-written characters have finished typing. (As noted above, input characters are normally echoed by putting them in the output queue as they arrive.) When a process produces characters more rapidly than they can be typed, it will be suspended when its output queue exceeds some limit. When the queue has drained down to some threshold the program is resumed. Even parity is normally generated on output. The EOT character is not transmitted in cooked mode to prevent terminals that respond to it from hanging up; programs using raw or cbreak mode should be careful.

The terminal drivers provide necessary processing for cooked and CBREAK mode output including delay generation for certain special characters and parity generation. Delays are available after backspaces `^H`, form feeds `^L`, carriage returns `^M`, tabs `^I` and newlines `^J`. The driver will also optionally expand tabs into spaces, where the tab stops are assumed to be set every eight columns. These functions are controlled by bits in the tty flags word; see **Summary** below.

The terminal drivers provide for mapping between upper and lower case on terminals lacking lower case, and for other special processing on deficient terminals.

Finally, in the new terminal driver, there is a output flush character, normally `^O`, which sets the LFLUSHO bit in the local mode word, causing subsequent output to be flushed until it is cleared by a program or more input is typed. This character has effect in both cooked and CBREAK modes and causes pending input to be retyped if there is any pending input. An `ioctl` to flush the characters in the input and output queues `TIOCFLUSH`, is also available.

Upper case terminals and Hazeltines

If the LCASE bit is set in the tty flags, then all upper-case letters are mapped into the corresponding lower-case letter. The upper-case letter may be generated by preceding it by `\'`. If the new terminal driver is being used, then upper case letters are preceded by a `\'` when output. In addition, the following escape sequences can be generated on output and accepted on input:

```
for   `      |      -      {      }
use   \`     \|     \^     \(\     \)
```

To deal with Hazeltine terminals, which do not understand that `~` has been made into an ASCII character, the LTI LDE bit may be set in the local mode word when using the new terminal driver; in this case the character `~` will be replaced with the character ``` on output.

Flow control.

There are two characters (the stop character, normally `^S`, and the start character, normally `^Q`) which cause output to be suspended and resumed respectively. Extra stop characters typed when output is already stopped have no effect, unless the start and stop characters are made the same, in which case output resumes.

A bit in the flags word may be set to put the terminal into TANDEM mode. In this mode the system produces a stop character (default `^S`) when the input queue is in danger of overflowing, and a start character (default `^Q`) when the input has drained sufficiently. This mode is useful when the terminal is actually another machine that obeys the conventions.

Line control and breaks.

There are several `ioctl` calls available to control the state of the terminal line. The `TIOCSBRK` `ioctl` will set the break bit in the hardware interface causing a break condition to exist; this can be cleared (usually after a delay with `sleep(3)`) by `TIOCCBRK`. Break conditions in the input are reflected as a null character in RAW mode or as the interrupt character in cooked or CBREAK mode. The `TIOCCDTR` `ioctl` will clear the data terminal ready condition; it can be

set again by TIOCSDTR.

When the carrier signal from the dataset drops (usually because the user has hung up his terminal) a SIGHUP hangup signal is sent to the processes in the distinguished process group of the terminal; this usually causes them to terminate (the SIGHUP can be suppressed by setting the LNOHANG bit in the local state word of the driver.) Access to the terminal by other processes is then normally revoked, so any further reads will fail, and programs that read a terminal and test for end-of-file on their input will terminate appropriately.

When using an ACU it is possible to ask that the phone line be hung up on the last close with the TIOCHPCL ioctl; this is normally done on the outgoing line.

Interrupt characters.

There are several characters that generate interrupts in cooked and CBREAK mode; all are sent the processes in the control group of the terminal, as if a TIOCGPGRP ioctl were done to get the process group and then a *killpg(2)* system call were done, except that these characters also flush pending input and output when typed at a terminal (*à la* TIOCFLUSH). The characters shown here are the defaults; the field names in the structures (given below) are also shown. The characters may be changed, although this is not often done.

- ^? **t_intrc** (Delete) generates a SIGINT signal. This is the normal way to stop a process which is no longer interesting, or to regain control in an interactive program.
- ^\
 t_quite (FS) generates a SIGQUIT signal. This is used to cause a program to terminate and produce a core image, if possible, in the file **core** in the current directory.
- ^Z **t_suspc** (EM) generates a SIGTSTP signal, which is used to suspend the current process group.
- ^Y **t_dsuspc** (SUB) generates a SIGTSTP signal as ^Z does, but the signal is sent when a program attempts to read the ^Y, rather than when it is typed.

Job access control.

When using the new terminal driver, if a process which is not in the distinguished process group of its control terminal attempts to read from that terminal its process group is sent a SIGTTIN signal. This signal normally causes the members of that process group to stop. If, however, the process is ignoring SIGTTIN, has SIGTTIN blocked, is an *orphan process*, or is in the middle of process creation using *vfork(2)*, it is instead returned an end-of-file. (An *orphan process* is a process whose parent has exited and has been inherited by the *init(8)* process.) Under older UNIX systems these processes would typically have had their input files reset to */dev/null*, so this is a compatible change.

When using the new terminal driver with the LTOSTOP bit set in the local modes, a process is prohibited from writing on its control terminal if it is not in the distinguished process group for that terminal. Processes which are holding or ignoring SIGTTOU signals, which are orphans, or which are in the middle of a *vfork(2)* are excepted and allowed to produce output.

Summary of modes.

Unfortunately, due to the evolution of the terminal driver, there are 4 different structures which contain various portions of the driver data. The first of these (**sgttyb**) contains that part of the information largely common between version 6 and version 7 UNIX systems. The second contains additional control characters added in version 7. The third is a word of local state peculiar to the new terminal driver, and the fourth is another structure of special characters added for the new driver. In the future a single structure may be made available to programs which need to access all this information; most programs need not concern themselves with all this state.

Basic modes: sgtty.

The basic *ioctls* use the structure defined in `<sgtty.h>`:

```
struct sgtyb {
    char    sg_ispeed;
    char    sg_ospeed;
    char    sg_erase;
    char    sg_kill;
    short   sg_flags;
};
```

The *sg_ispeed* and *sg_ospeed* fields describe the input and output speeds of the device according to the following table, which corresponds to the DEC DH-11 interface. If other hardware is used, impossible speed changes are ignored. Symbolic values in the table are as defined in `<sgtty.h>`.

B0	0	(hang up dataphone)
B50	1	50 baud
B75	2	75 baud
B110	3	110 baud
B134	4	134.5 baud
B150	5	150 baud
B200	6	200 baud
B300	7	300 baud
B600	8	600 baud
B1200	9	1200 baud
B1800	10	1800 baud
B2400	11	2400 baud
B4800	12	4800 baud
B9600	13	9600 baud
EXTA	14	External A
EXTB	15	External B

In the current configuration, only 110, 150, 300 and 1200 baud are really supported on dial-up lines. Code conversion and line control required for IBM 2741's (134.5 baud) must be implemented by the user's program. The half-duplex line discipline required for the 202 dataset (1200 baud) is not supplied; full-duplex 212 datasets work fine.

The *sg_erase* and *sg_kill* fields of the argument structure specify the erase and kill characters respectively. (Defaults are # and @.)

The *sg_flags* field of the argument structure contains several bits that determine the system's treatment of the terminal:

```
ALLDELAY 0177400 Delay algorithm selection
BSDELAY  0100000 Select backspace delays (not implemented):
BS0      0
BS1      0100000
VTDELAY  0040000 Select form-feed and vertical-tab delays:
FF0      0
FF1      0100000
CRDELAY  0030000 Select carriage-return delays:
CR0      0
CR1      0010000
CR2      0020000
CR3      0030000
```

TBDELAY	0006000	Select tab delays:
TAB0	0	
TAB1	0001000	
TAB2	0004000	
XTABS	0006000	
NLDELAY	0001400	Select new-line delays:
NL0	0	
NL1	0000400	
NL2	0001000	
NL3	0001400	
EVENP	0000200	Even parity allowed on input (most terminals)
ODDP	0000100	Odd parity allowed on input
RAW	0000040	Raw mode: wake up on all characters, 8-bit interface
CRMOD	0000020	Map CR into LF; echo LF or CR as CR-LF
ECHO	0000010	Echo (full duplex)
LCASE	0000004	Map upper case to lower on input
CBREAK	0000002	Return each character as soon as typed
TANDEM	0000001	Automatic flow control

The delay bits specify how long transmission stops to allow for mechanical or other movement when certain characters are sent to the terminal. In all cases a value of 0 indicates no delay.

Backspace delays are currently ignored but might be used for Terminet 300's.

If a form-feed/vertical tab delay is specified, it lasts for about 2 seconds.

Carriage-return delay type 1 lasts about .08 seconds and is suitable for the Terminet 300. Delay type 2 lasts about .16 seconds and is suitable for the VT05 and the TI 700. Delay type 3 is suitable for the concept-100 and pads lines to be at least 9 characters at 9600 baud.

New-line delay type 1 is dependent on the current column and is tuned for Teletype model 37's. Type 2 is useful for the VT05 and is about .10 seconds. Type 3 is unimplemented and is 0.

Tab delay type 1 is dependent on the amount of movement and is tuned to the Teletype model 37. Type 3, called XTABS, is not a delay at all but causes tabs to be replaced by the appropriate number of spaces on output.

Input characters with the wrong parity, as determined by bits 200 and 100, are ignored in cooked and CBREAK mode.

RAW disables all processing save output flushing with LFLUSHO; full 8 bits of input are given as soon as it is available; all 8 bits are passed on output. A break condition in the input is reported as a null character. If the input queue overflows in raw mode it is discarded; this applies to both new and old drivers.

CRMOD causes input carriage returns to be turned into new-lines; input of either CR or LF causes LF-CR both to be echoed (for terminals with a new-line function).

CBREAK is a sort of half-cooked (rare?) mode. Programs can read each character as soon as typed, instead of waiting for a full line; all processing is done except the input editing: character and word erase and line kill, input reprint, and the special treatment of \ or EOT are disabled.

TANDEM mode causes the system to produce a stop character (default ^S) whenever the input queue is in danger of overflowing, and a start character (default ^Q) when the input queue has drained sufficiently. It is useful for flow control when the 'terminal' is really another computer which understands the conventions.

Basic ioctls

In addition to the TIOCSETD and TIOCGTD disciplines discussed in **Line disciplines** above, a large number of other *ioctl(2)* calls apply to terminals, and have the general form:

```
#include <sgtty.h>
```

```
ioctl(fildes, code, arg)
```

```
struct sgttyb *arg;
```

The applicable codes are:

- | | |
|----------|--|
| TIOCGTTP | Fetch the basic parameters associated with the terminal, and store in the pointed-to <i>sgttyb</i> structure. |
| TIOCSETP | Set the parameters according to the pointed-to <i>sgttyb</i> structure. The interface delays until output is quiescent, then throws away any unread characters, before changing the modes. |
| TIOCSETN | Set the parameters like TIOCSETP but do not delay or flush input. Input is not preserved, however, when changing to or from RAW. |

With the following codes the *arg* is ignored.

- | | |
|-----------|---|
| TIOCEXCL | Set "exclusive-use" mode: no further opens are permitted until the file has been closed. |
| TIOCNXCL | Turn off "exclusive-use" mode. |
| TIOCHPCL | When the file is closed for the last time, hang up the terminal. This is useful when the line is associated with an ACU used to place outgoing calls. |
| TIOCFLUSH | All characters waiting in input or output queues are flushed. |

The remaining calls are not available in vanilla version 7 UNIX. In cases where arguments are required, they are described; *arg* should otherwise be given as 0.

- | | |
|-----------|---|
| TIOCSTI | the argument is the address of a character which the system pretends was typed on the terminal. |
| TIOCSBRK | the break bit is set in the terminal. |
| TIOCCBRK | the break bit is cleared. |
| TIOCSDTR | data terminal ready is set. |
| TIOCCDTR | data terminal ready is cleared. |
| TIOCGPGRP | <i>arg</i> is the address of a word into which is placed the process group number of the control terminal. |
| TIOCSPGRP | <i>arg</i> is a word (typically a process id) which becomes the process group for the control terminal. |
| FIONREAD | returns in the long integer whose address is <i>arg</i> the number of immediately readable characters from the argument unit. This works for files, pipes, and terminals, but not (yet) for multiplexed channels. |

Tchars

The second structure associated with each terminal specifies characters that are special in both the old and new terminal interfaces: The following structure is defined in *<sys/ioctl.h>*, which is automatically included in *<sgtty.h>*:

```
struct tchars {
    char  t_intrc;      /* interrupt */
    char  t_quite;     /* quit */
};
```

```

char   t_startc;    /* start output */
char   t_stopc;     /* stop output */
char   t_eofc;      /* end-of-file */
char   t_brkc;      /* input delimiter (like nl) */
};

```

The default values for these characters are `^?`, `^\\`, `^Q`, `^S`, `^D`, and `-1`. A character value of `-1` eliminates the effect of that character. The `t_brkc` character, by default `-1`, acts like a new-line in that it terminates a 'line,' is echoed, and is passed to the program. The 'stop' and 'start' characters may be the same, to produce a toggle effect. It is probably counterproductive to make other special characters (including erase and kill) identical. The applicable `ioctl` calls are:

TIOCGETC Get the special characters and put them in the specified structure.

TIOCSETC Set the special characters to those given in the structure.

Local mode

The third structure associated with each terminal is a local mode word; except for the `LNOHANG` bit, this word is interpreted only when the new driver is in use. The bits of the local mode word are:

LCRTBS	000001	Backspace on erase rather than echoing erase
LPRTERA	000002	Printing terminal erase mode
LCRTERA	000004	Erase character echoes as backspace-space-backspace
LTILDE	000010	Convert <code>~</code> to <code>`</code> on output (for Hazeltine terminals)
LMDMBUF	000020	Stop/start output when carrier drops
LLITOUT	000040	Suppress output translations
LTOSTOP	000100	Send <code>SIGTTOU</code> for background output
LFLUSHO	000200	Output is being flushed
LNOHANG	000400	Don't send hangup when carrier drops
LETXACK	001000	Diablo style buffer hacking (unimplemented)
LCRTKIL	002000	BS-space-BS erase entire line on line kill
LINTRUP	004000	Generate interrupt <code>SIGTINT</code> when input ready to read
LCTLECH	010000	Echo input control chars as <code>^X</code> , delete as <code>^?</code>
LPENDIN	020000	Retype pending input at next read or input character
LDECCTQ	040000	Only <code>^Q</code> restarts output after <code>^S</code> , like DEC systems

The applicable `ioctl` functions are:

TIOCLBIS arg is the address of a mask which is the bits to be set in the local mode word.

TIOCLBIC arg is the address of a mask of bits to be cleared in the local mode word.

TIOCLSET arg is the address of a mask to be placed in the local mode word.

TIOCLGET arg is the address of a word into which the current mask is placed.

Local special chars

The final structure associated with each terminal is the `lchars` structure which defines interrupt characters for the new terminal driver. Its structure is:

```

struct lchars {
char   t_suspc;     /* stop process signal */
char   t_dsuspc;    /* delayed stop process signal */
char   t_rprntc;    /* reprint line */
char   t_flushc;    /* flush output (toggles) */
};

```



```
    char  t_werasc;    /* word erase */
    char  t_lnextc;    /* literal next character */
};
```

The default values for these characters are `^Z`, `^Y`, `^R`, `^O`, `^W`, and `^V`. A value of `-1` disables the character.

The applicable *ioctl* functions are:

TIOCSLTC *args* is the address of a *lchars* structure which defines the new local special characters.

TIOCGLTC *args* is the address of a *lchars* structure into which is placed the current set of local special characters.

FILES

`/dev/tty`
`/dev/tty*`
`/dev/console`

SEE ALSO

`csh(1)`, `stty(1)`, `ioctl(2)`, `sigvec(2)`, `stty(3C)`, `getty(8)`, `init(8)`

BUGS

Half-duplex terminals are not supported.

NAME

va — Benson-Varian interface

SYNOPSIS

controller va0 at uba0 csr 0164000 vector vaintr
disk vz0 at va0 drive 0

DESCRIPTION

(NOTE: the configuration description, while counter-intuitive, is actually as shown above.)

The Benson-Varian printer/plotter is normally used with the programs *vpr(1)*, *vprint(1)* or *vtroff(1)*. This description is designed for those who wish to drive the Benson-Varian directly.

In print mode, the Benson-Varian uses a modified ASCII character set. Most control characters print various non-ASCII graphics such as daggers, sigmas, copyright symbols, etc. Only LF and FF are used as format effectors. LF acts as a newline, advancing to the beginning of the next line, and FF advances to the top of the next page.

In plot mode, the Benson-Varian prints one raster line at a time. An entire raster line of bits (2112 bits = 264 bytes) is sent, and then the Benson-Varian advances to the next raster line.

Note: The Benson-Varian must be sent an even number of bytes. If an odd number is sent, the last byte will be lost. Nulls can be used in print mode to pad to an even number of bytes.

To use the Benson-Varian yourself, you must realize that you cannot open the device, */dev/va0* if there is a daemon active. You can see if there is an active daemon by doing a *lpq(1)* and seeing if there are any files being printed.

To set the Benson-Varian into plot mode include the file *<sys/vcmd.h>* and use the following *ioctl(2)* call

```
ioctl(fileno(va), VSETSTATE, plotmd);
```

where *plotmd* is defined to be

```
int plotmd[] = { VPLOT, 0, 0 };
```

and *va* is the result of a call to *fopen* on *stdio*. When you finish using the Benson-Varian in plot mode you should advance to a new page by sending it a FF after putting it back into print mode, i.e. by

```
int prtmd[] = { VPRINT, 0, 0 };
...
flush(va);
ioctl(fileno(va), VSETSTATE, prtmd);
write(fileno(va), "\f\0", 2);
```

N.B.: If you use the standard I/O library with the Benson-Varian you **must** do

```
setbuf(vp, vbuf);
```

where *vbuf* is declared

```
char vbuf[BUFSIZ];
```

otherwise the standard I/O library, thinking that the Benson-Varian is a terminal (since it is a character special file) will not adequately buffer the data you are sending to the Benson-Varian. This will cause it to run **extremely** slowly and tend to grind the system to a halt.

FILES

/dev/va0

SEE ALSO

vfont(5), *lpr(1)*, *lpd(8)*, *vtroff(1)*, *vp(4)*

DIAGNOSTICS

The following error numbers are significant at the time the device is opened.

[ENXIO] The device is already in use.

[EIO] The device is offline.

The following message may be printed on the console.

va%d: npr timeout. The device was not able to get data from the UNIBUS within the timeout period, most likely because some other device was hogging the bus. (But see BUGS below).

BUGS

The l's (one's) and l's (lower-case el's) in the Benson-Varian's standard character set look very similar; caution is advised.

The interface hardware is rumored to have problems which can play havoc with the UNIBUS. We have intermittent minor problems on the UNIBUS where our va lives, but haven't ever been able to pin them down completely.

NAME

`vp` — Versatec interface

SYNOPSIS

`device vp0 at uba0 csr 0177510 vector vpintr vpintr`

DESCRIPTION

The Versatec printer/plotter is normally used with the programs `vpr(1)`, `vprint(1)` or `vtroff(1)`. This description is designed for those who wish to drive the Versatec directly.

To use the Versatec yourself, you must realize that you cannot open the device, `/dev/vp0` if there is a daemon active. You can see if there is a daemon active by doing a `lpq(1)`, and seeing if there are any files being sent.

To set the Versatec into plot mode you should include `<sys/vcmd.h>` and use the `ioctl(2)` call

```
ioctl(fileno(vp), VSETSTATE, plotmd);
```

where `plotmd` is defined to be

```
int plotmd[] = { VPLOT, 0, 0 };
```

and `vp` is the result of a call to `fopen` on `stdio`. When you finish using the Versatec in plot mode you should eject paper by sending it a EOT after putting it back into print mode, i.e. by

```
int prtmd[] = { VPRINT, 0, 0 };
```

```
...
```

```
flush(vp);
```

```
ioctl(fileno(vp), VSETSTATE, prtmd);
```

```
write(fileno(vp), "\04", 1);
```

N.B.: If you use the standard I/O library with the Versatec you **must** do

```
setbuf(vp, vpbuf);
```

where `vpbuf` is declared

```
char vpbuf[BUFSIZ];
```

otherwise the standard I/O library, thinking that the Versatec is a terminal (since it is a character special file) will not adequately buffer the data you are sending to the Versatec. This will cause it to run **extremely** slowly and tends to grind the system to a halt.

FILES

`/dev/vp0`

SEE ALSO

`vfont(5)`, `lpr(1)`, `lpd(8)`, `vtroff(1)`, `va(4)`

DIAGNOSTICS

The following error numbers are significant at the time the device is opened.

[ENXIO] The device is already in use.

[EIO] The device is offline.

BUGS

The configuration part of the driver assumes that the device is set up to vector print mode through 0174 and plot mode through 0200. As the configuration program can't be sure which vector interrupted at boot time, we specify that it has two interrupt vectors, and if an interrupt comes through 0200 it is reset to 0174. This is safe for devices with one or two vectors at these two addresses. Other configurations with 2 vectors may require changes in the driver.

NAME

a.out — assembler and link editor output

SYNOPSIS

```
#include <a.out.h>
```

DESCRIPTION

A.out is the output file of the assembler *as(1)* and the link editor *ld(1)*. Both programs make *a.out* executable if there were no errors and no unresolved external references. Layout information as given in the include file for the VAX-11 is:

```
/*
 * Header prepended to each a.out file.
 */
struct exec {
    long      a_magic; /* magic number */
    unsigned  a_text; /* size of text segment */
    unsigned  a_data; /* size of initialized data */
    unsigned  a_bss; /* size of uninitialized data */
    unsigned  a_syms; /* size of symbol table */
    unsigned  a_entry; /* entry point */
    unsigned  a_trsize; /* size of text relocation */
    unsigned  a_drsize; /* size of data relocation */
};

#define OMAGIC 0407 /* old impure format */
#define NMAGIC 0410 /* read-only text */
#define ZMAGIC 0413 /* demand load format */

/*
 * Macros which take exec structures as arguments and tell whether
 * the file has a reasonable magic number or offsets to text|symbols|strings.
 */
#define N_BADMAG(x) \
    (((x).a_magic)!=OMAGIC && ((x).a_magic)!=NMAGIC && ((x).a_magic)!=ZMAGIC)

#define N_TXTOFF(x) \
    ((x).a_magic==ZMAGIC ? 1024 : sizeof (struct exec))
#define N_SYMOFF(x) \
    (N_TXTOFF(x) + (x).a_text+(x).a_data + (x).a_trsize+(x).a_drsize)
#define N_STROFF(x) \
    (N_SYMOFF(x) + (x).a_syms)
```

The file has five sections: a header, the program text and data, relocation information, a symbol table and a string table (in that order). The last three may be omitted if the program was loaded with the '-s' option of *ld* or if the symbols and relocation have been removed by *strip(1)*.

In the header the sizes of each section are given in bytes. The size of the header is not included in any of the other sizes.

When an *a.out* file is executed, three logical segments are set up: the text segment, the data segment (with uninitialized data, which starts off as all 0, following initialized), and a stack. The text segment begins at 0 in the core image; the header is not loaded. If the magic number in the header is OMAGIC (0407), 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. This is the

oldest kind of executable program and is rarely used. If the magic number is NMAGIC (0410) or ZMAGIC (0413), the data segment begins at the first 0 mod 1024 byte 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. For ZMAGIC format, the text segment begins at a 0 mod 1024 byte boundary in the *a.out* file, the remaining bytes after the header in the first block are reserved and should be zero. In this case the text and data sizes must both be multiples of 1024 bytes, and the pages of the file will be brought into the running image as needed, and not pre-loaded as with the other formats. This is especially suitable for very large programs and is the default format produced by *ld*(1).

The stack will occupy the highest possible locations in the core image: growing downwards from 0x7ffff000. The stack is automatically extended as required. The data segment is only extended as requested by *brk*(2).

After the header in the file follow the text, data, text relocation data relocation, symbol table and string table in that order. The text begins at the byte 1024 in the file for ZMAGIC format or just after the header for the other formats. The N_TXTOFF macro returns this absolute file position when given the name of an exec structure as argument. The data segment is contiguous with the text and immediately followed by the text relocation and then the data relocation information. The symbol table follows all this; its position is computed by the N_SYMOFF macro. Finally, the string table immediately follows the symbol table at a position which can be gotten easily using N_STROFF. The first 4 bytes of the string table are not used for string storage, but rather contain the size of the string table; this size INCLUDES the 4 bytes, the minimum string table size is thus 4.

The layout of a symbol table entry and the principal flag values that distinguish symbol types are given in the include file as follows:

```
/*
 * Format of a symbol table entry.
 */
struct nlist {
    union {
        char    *n_name; /* for use when in-core */
        long    n_strx;  /* index into file string table */
    } n_un;
    unsigned char n_type; /* type flag, i.e. N_TEXT etc; see below */
    char          n_other;
    short         n_desc; /* see <stab.h> */
    unsigned      n_value; /* value of this symbol (or offset) */
};
#define n_hash    n_desc /* used internally by ld */

/*
 * Simple values for n_type.
 */
#define N_UNDF    0x0    /* undefined */
#define N_ABS    0x2    /* absolute */
#define N_TEXT   0x4    /* text */
#define N_DATA   0x6    /* data */
#define N_BSS    0x8    /* bss */
#define N_COMM   0x12   /* common (internal to ld) */
#define N_FN     0x1f   /* file name symbol */

#define N_EXT    01     /* external bit, or'ed in */
```

```

#define N_TYPE      0x1e      /* mask for all the type bits */

/*
 * Other permanent symbol table entries have some of the N_STAB bits set.
 * These are given in <stab.h>
 */
#define N_STAB      0xe0      /* if any of these bits set, don't discard */

```

```

/*
 * Format for namelist values.
 */
#define N_FORMAT    "%08x"

```

In the *a.out* file a symbol's `n_un.n_strx` field gives an index into the string table. A `n_strx` value of 0 indicates that no name is associated with a particular symbol table entry. The field `n_un.n_name` can be used to refer to the symbol name only if the program sets this up using `n_strx` and appropriate data from the string table.

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 byte in the text or data which is not a portion of a reference to an undefined external symbol is exactly that value which will appear in memory when the file is executed. If a byte in the text or data involves a reference to an undefined external symbol, as indicated by the relocation information, then the value 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 to the bytes in the file.

If relocation information is present, it amounts to eight bytes per relocatable datum as in the following structure:

```

/*
 * Format of a relocation datum.
 */
struct relocation_info {
    int      r_address;      /* address which is relocated */
    unsigned r_symbolnum:24, /* local symbol ordinal */
           r_pcrel:1,      /* was relocated pc relative already */
           r_length:2,     /* 0=byte, 1=word, 2=long */
           r_extern:1,     /* does not include value of sym referenced */
           :4;             /* nothing, yet */
};

```

There is no relocation information if `a_trsize+a_drsz==0`. If `r_extern` is 0, then `r_symbolnum` is actually a `n_type` for the relocation (i.e. `N_TEXT` meaning relative to segment text origin.)

SEE ALSO

`adb(1)`, `as(1)`, `ld(1)`, `nm(1)`, `dbx(1)`, `stab(5)`, `strip(1)`

BUGS

Not having the size of the string table in the header is a loss, but expanding the header size would have meant stripped executable file incompatibility, and we couldn't hack this just now.

NAME

`ar` - archive (library) file format

SYNOPSIS

```
#include <ar.h>
```

DESCRIPTION

The archive command `ar` combines several files into one. Archives are used mainly as libraries to be searched by the link-editor `ld`.

A file produced by `ar` has a magic string 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 ARMAG      0177545
#define SARMAG 8

#define ARFMAG "\n"

struct ar_hdr {
    char    ar_name[14];
    long    ar_date;
    char    ar_uid;
    char    ar_gid;
    short   ar_mode;
    long    ar_size;
};
```

The name is a blank-padded string. The `ar_fmag` field contains ARFMAG to help verify the presence of a header. 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 a 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

`ar(1)`, `ld(1)`, `nm(1)`

BUGS

File names lose trailing blanks. Most software dealing with archives takes even an included blank as a name terminator.


```

#define N_TYPE      0xe     /* mask for all the type bits */

/*
 * Other permanent symbol table entries have some of the N_STAB bits set.
 * These are given in <stab.h>
 */
#define N_STAB      0xe0     /* if any of these bits set, don't discard */

/*
 * Format for namelist values.
 */
#define N_FORMAT    "%08x"

```

In the *a.out* file a symbol's `n_un.n_strx` field gives an index into the string table. A `n_strx` value of 0 indicates that no name is associated with a particular symbol table entry. The field `n_un.n_name` can be used to refer to the symbol name only if the program sets this up using `n_strx` and appropriate data from the string table.

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 byte in the text or data which is not a portion of a reference to an undefined external symbol is exactly that value which will appear in memory when the file is executed. If a byte in the text or data involves a reference to an undefined external symbol, as indicated by the relocation information, then the value 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 to the bytes in the file.

If relocation information is present, it amounts to eight bytes per relocatable datum as in the following structure:

```

/*
 * Format of a relocation datum.
 */
struct relocation_info {
    int      r_address;      /* address which is relocated */
    unsigned r_symbolnum:24, /* local symbol ordinal */
            r_pcrel:1,      /* was relocated pc relative already */
            r_length:2,     /* 0=byte, 1=word, 2=long */
            r_extern:1,     /* does not include value of sym referenced */
            :4;             /* nothing, yet */
};

```

There is no relocation information if `a_trsize+a_drsize==0`. If `r_extern` is 0, then `r_symbolnum` is actually a `n_type` for the relocation (i.e. `N_TEXT` meaning relative to segment text origin.)

SEE ALSO

`adb(1)`, `as(1)`, `ld(1)`, `nm(1)`, `dbx(1)`, `stab(5)`, `strip(1)`

BUGS

Not having the size of the string table in the header is a loss, but expanding the header size would have meant stripped executable file incompatibility, and we couldn't hack this just now.

NAME

ar - archive (library) file format

SYNOPSIS

```
#include <ar.h>
```

DESCRIPTION

The archive command *ar* combines several files into one. Archives are used mainly as libraries to be searched by the link-editor *ld*.

A file produced by *ar* has a magic string 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 ARMAG      0177545
#define SARMAG 8

#define ARFMAG "\n"

struct ar_hdr {
    char  ar_name[14];
    long  ar_date;
    char  ar_uid;
    char  ar_gid;
    short ar_mode;
    long  ar_size;
};
```

The name is a blank-padded string. The *ar_fm* field contains ARFMAG to help verify the presence of a header. 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 a 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

ar(1), *ld*(1), *nm*(1)

BUGS

File names lose trailing blanks. Most software dealing with archives takes even an included blank as a name terminator.

NAME

core — format of memory image file

SYNOPSIS

```
#include <machine/param.h>
```

DESCRIPTION

The UNIX System writes out a memory image of a terminated process when any of various errors occur. See *sigvec(2)* for the list of reasons; the most common are memory violations, illegal instructions, bus errors, and user-generated quit signals. The memory image is called 'core' and is written in the process's working directory (provided it can be: normal access controls apply).

The maximum size of a *core* file is limited by *setrlimit(2)*. Files which would be larger than the limit are not created.

The core file consists of the *u*. area, whose size (in pages) is defined by the UPAGES manifest in the *<machine/param.h>* file. The *u*. area starts with a *user* structure as given in *<sys/user.h>*. The remainder of the core file consists first of the data pages and then the stack pages of the process image. The amount of data space image in the core file is given (in pages) by the variable *u_dsize* in the *u*. area. The amount of stack image in the core file is given (in pages) by the variable *u_ssize* in the *u*. area.

In general the debugger *adb(1)* is sufficient to deal with core images.

SEE ALSO

adb(1), *dbx(1)*, *sigvec(2)*, *setrlimit(2)*

NAME

dir — format of directories

SYNOPSIS

```
#include <sys/types.h>
#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(5)*. The structure of a directory entry as given in the include file is:

```
/*
 * A directory consists of some number of blocks of DIRBLKSIZ
 * bytes, where DIRBLKSIZ is chosen such that it can be transferred
 * to disk in a single atomic operation (e.g. 512 bytes on most machines).
 *
 * Each DIRBLKSIZ byte block contains some number of directory entry
 * structures, which are of variable length. Each directory entry has
 * a struct direct at the front of it, containing its inode number,
 * the length of the entry, and the length of the name contained in
 * the entry. These are followed by the name padded to a 4 byte boundary
 * with null bytes. All names are guaranteed null terminated.
 * The maximum length of a name in a directory is MAXNAMLEN.
 *
 * The macro DIRSIZ(dp) gives the amount of space required to represent
 * a directory entry. Free space in a directory is represented by
 * entries which have dp->d_reclen > DIRSIZ(dp). All DIRBLKSIZ bytes
 * in a directory block are claimed by the directory entries. This
 * usually results in the last entry in a directory having a large
 * dp->d_reclen. When entries are deleted from a directory, the
 * space is returned to the previous entry in the same directory
 * block by increasing its dp->d_reclen. If the first entry of
 * a directory block is free, then its dp->d_ino is set to 0.
 * Entries other than the first in a directory do not normally have
 * dp->d_ino set to 0.
 */
#ifdef KERNEL
#define DIRBLKSIZ DEV_BSIZE
#else
#define DIRBLKSIZ 512
#endif

#define MAXNAMLEN 255

/*
 * The DIRSIZ macro gives the minimum record length which will hold
 * the directory entry. This requires the amount of space in struct direct
 * without the d_name field, plus enough space for the name with a terminating
 * null byte (dp->d_namlen + 1), rounded up to a 4 byte boundary.
 */
#undef DIRSIZ
#define DIRSIZ(dp) \
    ((sizeof (struct direct) - (MAXNAMLEN + 1)) + (((dp)->d_namlen + 1 + 3) &~ 3))
```

```
struct direct {
    u_long    d_ino;
    short     d_reclen;
    short     d_namlen;
    char      d_name[MAXNAMLEN + 1];
    /* typically shorter */
};

struct _dirdesc {
    int       dd_fd;
    long      dd_loc;
    long      dd_size;
    char      dd_buf[DIRBLKSIZ];
};
```

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 ("/"), where '..' has the same meaning as '.'.

SEE ALSO
fs(5)

NAME

disktab — disk description file

SYNOPSIS

```
#include <disktab.h>
```

DESCRIPTION

Disktab is a simple data base which describes disk geometries and disk partition characteristics. The format is patterned after the *termcap(5)* terminal data base. Entries in *disktab* consist of a number of ':' separated fields. The first entry for each disk gives the names which are known for the disk, separated by '|' characters. The last name given should be a long name fully identifying the disk.

The following list indicates the normal values stored for each disk entry.

Name	Type	Description
ns	num	Number of sectors per track
nt	num	Number of tracks per cylinder
nc	num	Total number of cylinders on the disk
ba	num	Block size for partition 'a' (bytes)
bd	num	Block size for partition 'd' (bytes)
be	num	Block size for partition 'e' (bytes)
bf	num	Block size for partition 'f' (bytes)
bg	num	Block size for partition 'g' (bytes)
bh	num	Block size for partition 'h' (bytes)
fa	num	Fragment size for partition 'a' (bytes)
fd	num	Fragment size for partition 'd' (bytes)
fe	num	Fragment size for partition 'e' (bytes)
ff	num	Fragment size for partition 'f' (bytes)
fg	num	Fragment size for partition 'g' (bytes)
fh	num	Fragment size for partition 'h' (bytes)
pa	num	Size of partition 'a' in sectors
pb	num	Size of partition 'b' in sectors
pc	num	Size of partition 'c' in sectors
pd	num	Size of partition 'd' in sectors
pe	num	Size of partition 'e' in sectors
pf	num	Size of partition 'f' in sectors
pg	num	Size of partition 'g' in sectors
ph	num	Size of partition 'h' in sectors
se	num	Sector size in bytes
ty	str	Type of disk (e.g. removable, winchester)

Disktab entries may be automatically generated with the *diskpart* program.

FILES

/etc/disktab

SEE ALSO

newfs(8), diskpart(8)

BUGS

This file shouldn't exist, the information should be stored on each disk pack.

NAME

dump, dumpdates — incremental dump format

SYNOPSIS

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

DESCRIPTION

Tapes used by *dump* and *restore(8)* contain:

- a header record
- two groups of bit map records
- a group of records describing directories
- a group of records describing files

The format of the header record and of the first record of each description as given in the include file *<dumprest.h>* is:

```
#define NTREC      10
#define MLEN       16
#define MSIZ       4096

#define TS_TAPE    1
#define TS_INODE   2
#define TS_BITS    3
#define TS_ADDR    4
#define TS_END     5
#define TS_CLRI    6
#define MAGIC      (int) 60011
#define CHECKSUM   (int) 84446
```

```
struct spcl {
    int          c_type;
    time_t      c_date;
    time_t      c_ddate;
    int         c_volume;
    daddr_t     c_tapea;
    ino_t       c_inumber;
    int         c_magic;
    int         c_checksum;
    struct      dinode      c_dinode;
    int         c_count;
    char        c_addr[BSIZE];
} spcl;
```

```
struct idates {
    char        id_name[16];
    char        id_incno;
    time_t      id_ddate;
};
```

```
#define DUMPOUTFMT      "%-16s %c %s"      /* for printf */
/* name, incno, ctime(date) */
#define DUMPINFMT      "%16s %c %[\n]\n"    /* inverse for scanf */
```

NTREC is the number of 1024 byte records in a physical tape block. MLEN is the number of bits in a bit map word. MSIZ is the number of bit map words.

The TS_ entries are used in the *c_type* field to indicate what sort of header this is. The types and their meanings are as follows:

TS_TAPE	Tape volume label
TS_INODE	A file or directory follows. The <i>c_dinode</i> field is a copy of the disk inode and contains bits telling what sort of file this is.
TS_BITS	A bit map follows. This bit map has a one bit for each inode that was dumped.
TS_ADDR	A subrecord of a file description. See <i>c_addr</i> below.
TS_END	End of tape record.
TS_CLRI	A bit map follows. This bit map contains a zero bit for all inodes that were empty on the file system when dumped.
MAGIC	All header records have this number in <i>c_magic</i> .
CHECKSUM	Header records checksum to this value.

The fields of the header structure are as follows:

<i>c_type</i>	The type of the header.
<i>c_date</i>	The date the dump was taken.
<i>c_ddate</i>	The date the file system was dumped from.
<i>c_volume</i>	The current volume number of the dump.
<i>c_tapea</i>	The current number of this (1024-byte) record.
<i>c_inumber</i>	The number of the inode being dumped if this is of type TS_INODE.
<i>c_magic</i>	This contains the value MAGIC above, truncated as needed.
<i>c_checksum</i>	This contains whatever value is needed to make the record sum to CHECKSUM.
<i>c_dinode</i>	This is a copy of the inode as it appears on the file system; see <i>fs(5)</i> .
<i>c_count</i>	The count of characters in <i>c_addr</i> .
<i>c_addr</i>	An array of characters describing the blocks of the dumped file. A character is zero if the block associated with that character was not present on the file system, otherwise the character is non-zero. If the block was not present on the file system, no block was dumped; the block will be restored as a hole in the file. If there is not sufficient space in this record to describe all of the blocks in a file, TS_ADDR records will be scattered through the file, each one picking up where the last left off.

Each volume except the last ends with a tapemark (read as an end of file). The last volume ends with a TS_END record and then the tapemark.

The structure *idates* describes an entry in the file */etc/dumpdates* where dump history is kept. The fields of the structure are:

<i>id_name</i>	The dumped filesystem is <i>'/dev/id_nam'</i> .
<i>id_incno</i>	The level number of the dump tape; see <i>dump(8)</i> .
<i>id_ddate</i>	The date of the incremental dump in system format see <i>types(5)</i> .

FILES

/etc/dumpdates

SEE ALSO

dump(8), *restore(8)*, *fs(5)*, *types(5)*

NAME

fs, inode — format of file system volume

SYNOPSIS

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

DESCRIPTION

Every file system storage volume (disk, nine-track tape, for instance) has a common format for certain vital information. Every such volume is divided into a certain number of blocks. The block size is a parameter of the file system. Sectors 0 to 15 on a file system are used to contain primary and secondary bootstrapping programs.

The actual file system begins at sector 16 with the *super block*. The layout of the super block as defined by the include file `<sys/fs.h>` is:

```
#define FS_MAGIC    0x011954
struct fs {
    struct fs *fs_link;           /* linked list of file systems */
    struct fs *fs_rlink;         /* used for incore super blocks */
    daddr_t fs_sblkno;           /* addr of super-block in filesys */
    daddr_t fs_cblkno;           /* offset of cyl-block in filesys */
    daddr_t fs_iblkn;           /* offset of inode-blocks in filesys */
    daddr_t fs_dblkno;           /* offset of first data after cg */
    long fs_cgoffset;           /* cylinder group offset in cylinder */
    long fs_cgmask;             /* used to calc mod fs_ntrak */
    time_t fs_time;             /* last time written */
    long fs_size;               /* number of blocks in fs */
    long fs_dsize;              /* number of data blocks in fs */
    long fs_ncg;                /* number of cylinder groups */
    long fs_bsize;              /* size of basic blocks in fs */
    long fs_fsize;              /* size of frag blocks in fs */
    long fs_frag;              /* number of frags in a block in fs */

    /* these are configuration parameters */
    long fs_minfree;            /* minimum percentage of free blocks */
    long fs_rotdelay;           /* num of ms for optimal next block */
    long fs_rps;                /* disk revolutions per second */

    /* these fields can be computed from the others */
    long fs_bmask;              /* "blkoff" calc of blk offsets */
    long fs_fmask;              /* "fragoff" calc of frag offsets */
    long fs_bshift;             /* "lblkno" calc of logical blkno */
    long fs_fshift;             /* "numfrags" calc number of frags */

    /* these are configuration parameters */
    long fs_maxcontig;           /* max number of contiguous blks */
    long fs_maxbpg;             /* max number of blks per cyl group */

    /* these fields can be computed from the others */
    long fs_fragshift;          /* block to frag shift */
    long fs_fsbtodb;            /* fsbtodb and dbtofsb shift constant */
    long fs_sbsize;             /* actual size of super block */
    long fs_csum;               /* csum block offset */
    long fs_cshift;             /* csum block number */
    long fs_nindir;             /* value of NINDIR */
    long fs_inopb;              /* value of INOPB */
    long fs_nspf;               /* value of NSPF */
};
```

```

    long   fs_sparecon[6];      /* reserved for future constants */
/* sizes determined by number of cylinder groups and their sizes */
    daddr_t fs_csaddr;         /* blk addr of cyl grp summary area */
    long   fs_cssize;          /* size of cyl grp summary area */
    long   fs_cgsize;          /* cylinder group size */
/* these fields should be derived from the hardware */
    long   fs_ntrak;           /* tracks per cylinder */
    long   fs_nsect;           /* sectors per track */
    long   fs_spc;             /* sectors per cylinder */
/* this comes from the disk driver partitioning */
    long   fs_ncyl;            /* cylinders in file system */
/* these fields can be computed from the others */
    long   fs_cpg;             /* cylinders per group */
    long   fs_ipg;             /* inodes per group */
    long   fs_fpg;             /* blocks per group * fs_frag */
/* this data must be re-computed after crashes */
    struct csum fs_cstotal; /* cylinder summary information */
/* these fields are cleared at mount time */
    char   fs_fmod;            /* super block modified flag */
    char   fs_clean;           /* file system is clean flag */
    char   fs_ronly;           /* mounted read-only flag */
    char   fs_flags;           /* currently unused flag */
    char   fs_fsmnt[MAXMNTLEN]; /* name mounted on */
/* these fields retain the current block allocation info */
    long   fs_cgrotor;         /* last cg searched */
    struct csum *fs_csp[MAXCSBUFS]; /* list of fs_cs info buffers */
    long   fs_cpc;             /* cyl per cycle in postbl */
    short  fs_postbl[MAXCPG][NRPOS]; /* head of blocks for each rotation */
    long   fs_magic;           /* magic number */
    u_char fs_rotbl[1];        /* list of blocks for each rotation */
/* actually longer */
};

```

Each disk drive contains some number of file systems. A file system consists of a number of cylinder groups. Each cylinder group has inodes and data.

A file system is described by its super-block, which in turn describes the cylinder groups. The super-block is critical data and is replicated in each cylinder group to protect against catastrophic loss. This is done at file system creation time and the critical super-block data does not change, so the copies need not be referenced further unless disaster strikes.

Addresses stored in inodes are capable of addressing fragments of 'blocks'. File system blocks of at most size MAXBSIZE can be optionally broken into 2, 4, or 8 pieces, each of which is addressable; these pieces may be DEV_BSIZE, or some multiple of a DEV_BSIZE unit.

Large files consist of exclusively large data blocks. To avoid undue wasted disk space, the last data block of a small file is allocated as only as many fragments of a large block as are necessary. The file system format retains only a single pointer to such a fragment, which is a piece of a single large block that has been divided. The size of such a fragment is determinable from information in the inode, using the "blksize(fs, ip, lbn)" macro.

The file system records space availability at the fragment level; to determine block availability, aligned fragments are examined.

The root inode is the root of the file system. Inode 0 can't be used for normal purposes and historically bad blocks were linked to inode 1, thus the root inode is 2 (inode 1 is no longer used for this purpose, however numerous dump tapes make this assumption, so we are stuck with it). The *lost+found* directory is given the next available inode when it is initially created by *mkfs*.

fs_minfree gives the minimum acceptable percentage of file system blocks which may be free. If the freelist drops below this level only the super-user may continue to allocate blocks. This may be set to 0 if no reserve of free blocks is deemed necessary, however severe performance degradations will be observed if the file system is run at greater than 90% full; thus the default value of *fs_minfree* is 10%.

Empirically the best trade-off between block fragmentation and overall disk utilization at a loading of 90% comes with a fragmentation of 4, thus the default fragment size is a fourth of the block size.

Cylinder group related limits: Each cylinder keeps track of the availability of blocks at different rotational positions, so that sequential blocks can be laid out with minimum rotational latency. NRPOS is the number of rotational positions which are distinguished. With NRPOS 8 the resolution of the summary information is 2ms for a typical 3600 rpm drive.

fs_rotdelay gives the minimum number of milliseconds to initiate another disk transfer on the same cylinder. It is used in determining the rotationally optimal layout for disk blocks within a file; the default value for *fs_rotdelay* is 2ms.

Each file system has a statically allocated number of inodes. An inode is allocated for each NBPI bytes of disk space. The inode allocation strategy is extremely conservative.

MAXIPG bounds the number of inodes per cylinder group, and is needed only to keep the structure simpler by having the only a single variable size element (the free bit map).

N.B.: MAXIPG must be a multiple of INOPB(fs).

MINBSIZE is the smallest allowable block size. With a MINBSIZE of 4096 it is possible to create files of size 2^{32} with only two levels of indirection. MINBSIZE must be big enough to hold a cylinder group block, thus changes to (struct cg) must keep its size within MINBSIZE. MAXCPG is limited only to dimension an array in (struct cg); it can be made larger as long as that structure's size remains within the bounds dictated by MINBSIZE. Note that super blocks are never more than size SBSIZE.

The path name on which the file system is mounted is maintained in *fs_fsmnt*. MAXMNTLEN defines the amount of space allocated in the super block for this name. The limit on the amount of summary information per file system is defined by MAXCSBUFS. It is currently parameterized for a maximum of two million cylinders.

Per cylinder group information is summarized in blocks allocated from the first cylinder group's data blocks. These blocks are read in from *fs_csaddr* (size *fs_cssize*) in addition to the super block.

N.B.: sizeof (struct csum) must be a power of two in order for the "fs_cs" macro to work.

Super block for a file system: MAXBPC bounds the size of the rotational layout tables and is limited by the fact that the super block is of size SBSIZE. The size of these tables is inversely proportional to the block size of the file system. The size of the tables is increased when sector sizes are not powers of two, as this increases the number of cylinders included before the rotational pattern repeats (*fs_cpc*). The size of the rotational layout tables is derived from the number of bytes remaining in (struct fs).

MAXBPG bounds the number of blocks of data per cylinder group, and is limited by the fact that cylinder groups are at most one block. The size of the free block table is derived from the size of blocks and the number of remaining bytes in the cylinder group structure (struct cg).

Inode: The inode is the focus of all file activity in the UNIX file system. There is a unique inode allocated for each active file, each current directory, each mounted-on file, text file, and the root. An inode is 'named' by its device/i-number pair. For further information, see the include file `<sys/inode.h>`.

NAME

`fstab` — static information about the filesystems

SYNOPSIS

```
#include <fstab.h>
```

DESCRIPTION

The file */etc/fstab* contains descriptive information about the various file systems. */etc/fstab* is only read by programs, and not written; it is the duty of the system administrator to properly create and maintain this file. The order of records in */etc/fstab* is important because *fsck*, *mount*, and *umount* sequentially iterate through */etc/fstab* doing their thing.

The special file name is the block special file name, and not the character special file name. If a program needs the character special file name, the program must create it by appending a "r" after the last "/" in the special file name.

If *fs_type* is "rw" or "ro" then the file system whose name is given in the *fs_file* field is normally mounted read-write or read-only on the specified special file. If *fs_type* is "rq", then the file system is normally mounted read-write with disk quotas enabled. The *fs_freq* field is used for these file systems by the *dump*(8) command to determine which file systems need to be dumped. The *fs_passno* field is used by the *fsck*(8) program to determine the order in which file system checks are done at reboot time. The root file system should be specified with a *fs_passno* of 1, and other file systems should have larger numbers. File systems within a drive should have distinct numbers, but file systems on different drives can be checked on the same pass to utilize parallelism available in the hardware.

If *fs_type* is "sw" then the special file is made available as a piece of swap space by the *swapon*(8) command at the end of the system reboot procedure. The fields other than *fs_spec* and *fs_type* are not used in this case.

If *fs_type* is "rq" then at boot time the file system is automatically processed by the *quota-check*(8) command and disk quotas are then enabled with *quotaon*(8). File system quotas are maintained in a file "quotas", which is located at the root of the associated file system.

If *fs_type* is specified as "xx" the entry is ignored. This is useful to show disk partitions which are currently not used.

```
#define FSTAB_RW    "rw"    /* read-write device */
#define FSTAB_RO    "ro"    /* read-only device */
#define FSTAB_RQ    "rq"    /* read-write with quotas */
#define FSTAB_SW    "sw"    /* swap device */
#define FSTAB_XX    "xx"    /* ignore totally */

struct fstab {
    char *fs_spec; /* block special device name */
    char *fs_file; /* file system path prefix */
    char *fs_type; /* rw,ro,sw or xx */
    int fs_freq; /* dump frequency, in days */
    int fs_passno; /* pass number on parallel dump */
};
```

The proper way to read records from */etc/fstab* is to use the routines *getfsent*(), *getfsspec*(), *getfstype*(), and *getfsfile*() .

FILES

/etc/fstab

SEE ALSO
getfsent(3X)

NAME

group — group file

DESCRIPTION

Group contains for each group the following information:

group name
encrypted password
numerical group ID
a comma separated list of all users 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

setgroups(2), initgroups(3X), crypt(3), passwd(1), passwd(5)

BUGS

The *passwd*(1) command won't change the passwords.

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

mtab — mounted file system table

SYNOPSIS

```
#include <fstab.h>
#include <mtab.h>
```

DESCRIPTION

Mtab resides in directory */etc* and contains a table of devices mounted by the *mount* command. *Umount* removes entries.

The table is a series of *mtab* structures, as defined in *<mtab.h>*. Each entry contains the null-padded name of the place where the special file is mounted, the null-padded name of the special file, and a type field, one of those defined in *<fstab.h>*. The special file has all its directories stripped away; that is, everything through the last '/' is thrown away. The type field indicates if the file system is mounted read-only, read-write, or read-write with disk quotas enabled.

This table is present only so people can look at it. It does not matter to *mount* if there are duplicated entries nor to *umount* if a name cannot be found.

FILES

/etc/mtab

SEE ALSO

mount(8)

NAME

passwd — password file

DESCRIPTION

Passwd contains for each user the following information:

name (login name, contains no upper case)

encrypted password

numerical user ID

numerical group ID

user's real name, office, extension, home phone.

initial working directory

program to use as Shell

The name may contain '&', meaning insert the login name. This information is set by the *chfn*(1) command and used by the *finger*(1) command.

This is an ASCII file. Each field within each user's entry is separated from the next by a colon. 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, then */bin/sh* 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 ID's to names.

Appropriate precautions must be taken to lock the file against changes if it is to be edited with a text editor; *vipw*(8) does the necessary locking.

FILES

/etc/passwd

SEE ALSO

getpwent(3), *login*(1), *crypt*(3), *passwd*(1), *group*(5), *chfn*(1), *finger*(1), *vipw*(8), *adduser*(8)

BUGS

A binary indexed file format should be available for fast access.

User information (name, office, etc.) should be stored elsewhere.

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

`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

stab — symbol table types

SYNOPSIS

```
#include <stab.h>
```

DESCRIPTION

Stab.h defines some values of the *n_type* field of the symbol table of *a.out* files. These are the types for permanent symbols (i.e. not local labels, etc.) used by the old debugger *sdb* and the Berkeley Pascal compiler *pc(1)*. Symbol table entries can be produced by the *.stabs* assembler directive. This allows one to specify a double-quote delimited name, a symbol type, one char and one short of information about the symbol, and an unsigned long (usually an address). To avoid having to produce an explicit label for the address field, the *.stabd* directive can be used to implicitly address the current location. If no name is needed, symbol table entries can be generated using the *.stabsn* directive. The loader promises to preserve the order of symbol table entries produced by *.stab* directives. As described in *a.out(5)*, an element of the symbol table consists of the following structure:

```
/*
 * Format of a symbol table entry.
 */
struct nlist {
    union {
        char *n_name; /* for use when in-core */
        long n_strx; /* index into file string table */
    } n_un;
    unsigned char n_type; /* type flag */
    char n_other; /* unused */
    short n_desc; /* see struct desc, below */
    unsigned n_value; /* address or offset or line */
};
```

The low bits of the *n_type* field are used to place a symbol into at most one segment, according to the following masks, defined in *<a.out.h>*. A symbol can be in none of these segments by having none of these segment bits set.

```
/*
 * Simple values for n_type.
 */
#define N_UNDF 0x0 /* undefined */
#define N_ABS 0x2 /* absolute */
#define N_TEXT 0x4 /* text */
#define N_DATA 0x6 /* data */
#define N_BSS 0x8 /* bss */

#define N_EXT 01 /* external bit, or'ed in */
```

The *n_value* field of a symbol is relocated by the linker, *ld(1)* as an address within the appropriate segment. *N_value* fields of symbols not in any segment are unchanged by the linker. In addition, the linker will discard certain symbols, according to rules of its own, unless the *n_type* field has one of the following bits set:

```
/*
 * Other permanent symbol table entries have some of the N_STAB bits set.
 * These are given in <stab.h>
 */
#define N_STAB 0xe0 /* if any of these bits set, don't discard */
```

This allows up to 112 ($7 \cdot 16$) symbol types, split between the various segments. Some of these have already been claimed. The old symbolic debugger, *sdb*, uses the following *n_type* values:

```
#define N_GSYM 0x20 /* global symbol: name,,0,type,0 */
#define N_FNAME 0x22 /* procedure name (f77 kludge): name,,0 */
#define N_FUN 0x24 /* procedure: name,,0,linenumber,address */
#define N_STSYM 0x26 /* static symbol: name,,0,type,address */
#define N_LCSYM 0x28 /* .lcomm symbol: name,,0,type,address */
#define N_RSYM 0x40 /* register sym: name,,0,type,register */
#define N_SLINE 0x44 /* src line: 0,,0,linenumber,address */
#define N_SSYM 0x60 /* structure elt: name,,0,type,struct_offset */
#define N_SO 0x64 /* source file name: name,,0,0,address */
#define N_LSYM 0x80 /* local sym: name,,0,type,offset */
#define N_SOL 0x84 /* #included file name: name,,0,0,address */
#define N_PSYM 0xa0 /* parameter: name,,0,type,offset */
#define N_ENTRY 0xa4 /* alternate entry: name,linenumber,address */
#define N_LBRAC 0xc0 /* left bracket: 0,,0,nesting level,address */
#define N_RBRAC 0xe0 /* right bracket: 0,,0,nesting level,address */
#define N_BCOMM 0xe2 /* begin common: name,, */
#define N_ECOMM 0xe4 /* end common: name,, */
#define N_ECOML 0xe8 /* end common (local name): ,,address */
#define N_LENG 0xfe /* second stab entry with length information */
```

where the comments give *sdb* conventional use for *.stabs* and the *n_name*, *n_other*, *n_desc*, and *n_value* fields of the given *n_type*. *Sdb* uses the *n_desc* field to hold a type specifier in the form used by the Portable C Compiler, *cc(1)*, in which a base type is qualified in the following structure:

```
struct desc {
    short q6:2,
          q5:2,
          q4:2,
          q3:2,
          q2:2,
          q1:2,
          basic:4;
};
```

There are four qualifications, with *q1* the most significant and *q6* the least significant:

```
0    none
1    pointer
2    function
3    array
```

The sixteen basic types are assigned as follows:

```
0    undefined
1    function argument
2    character
3    short
4    int
5    long
6    float
7    double
8    structure
9    union
```

10	enumeration
11	member of enumeration
12	unsigned character
13	unsigned short
14	unsigned int
15	unsigned long

The Berkeley Pascal compiler, *pc*(1), uses the following *n_type* value:

```
#define N_PC 0x30 /* global pascal symbol: name,,0,subtype,line */
```

and uses the following subtypes to do type checking across separately compiled files:

1	source file name
2	included file name
3	global label
4	global constant
5	global type
6	global variable
7	global function
8	global procedure
9	external function
10	external procedure
11	library variable
12	library routine

SEE ALSO

as(1), *ld*(1), *dbx*(1), *a.out*(5)

BUGS

Sdb assumes that a symbol of type *N_GSYM* with name *name* is located at address *_name*.

More basic types are needed.

NAME

tar — tape archive file format

DESCRIPTION

Tar, (the tape archive command) dumps several files into one, in a medium suitable for transportation.

A “tar tape” or file is a series of blocks. Each block is of size `TBLOCK`. A file on the tape is represented by a header block which describes the file, followed by zero or more blocks which give the contents of the file. At the end of the tape are two blocks filled with binary zeros, as an end-of-file indicator.

The blocks are grouped for physical I/O operations. Each group of n blocks (where n is set by the `b` keyletter on the `tar(1)` command line — default is 20 blocks) is written with a single system call; on nine-track tapes, the result of this write is a single tape record. The last group is always written at the full size, so blocks after the two zero blocks contain random data. On reading, the specified or default group size is used for the first read, but if that read returns less than a full tape block, the reduced block size is used for further reads.

The header block looks like:

```
#define TBLOCK      512
#define NAMSIZ      100

union hblock {
    char dummy[TBLOCK];
    struct header {
        char name[NAMSIZ];
        char mode[8];
        char uid[8];
        char gid[8];
        char size[12];
        char mtime[12];
        char chksum[8];
        char linkflag;
        char linkname[NAMSIZ];
    } dbuf;
};
```

Name is a null-terminated string. The other fields are zero-filled octal numbers in ASCII. Each field (of width w) contains $w-2$ digits, a space, and a null, except *size* and *mtime*, which do not contain the trailing null. *Name* is the name of the file, as specified on the `tar` command line. Files dumped because they were in a directory which was named in the command line have the directory name as prefix and *filename* as suffix. *Mode* is the file mode, with the top bit masked off. *Uid* and *gid* are the user and group numbers which own the file. *Size* is the size of the file in bytes. Links and symbolic links are dumped with this field specified as zero. *Mtime* is the modification time of the file at the time it was dumped. *Chksum* is a decimal ASCII value which represents the sum of all the bytes in the header block. When calculating the checksum, the *chksum* field is treated as if it were all blanks. *Linkflag* is ASCII '0' if the file is “normal” or a special file, ASCII '1' if it is an hard link, and ASCII '2' if it is a symbolic link. The name linked-to, if any, is in *linkname*, with a trailing null. Unused fields of the header are binary zeros (and are included in the checksum).

The first time a given i-node number is dumped, it is dumped as a regular file. The second and subsequent times, it is dumped as a link instead. Upon retrieval, if a link entry is retrieved, but not the file it was linked to, an error message is printed and the tape must be manually re-scanned to retrieve the linked-to file.

The encoding of the header is designed to be portable across machines.

SEE ALSO

tar(1)

BUGS

Names or linknames longer than NAMSIZ produce error reports and cannot be dumped.

NAME

termcap — terminal capability data base

SYNOPSIS

/etc/termcap

DESCRIPTION

Termcap is a data base describing terminals, used, e.g., by *vi(1)* and *curses(3X)*. 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 systemwide 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 millisec of bs delay needed
db	bool		Display may be retained below
dC	num		Number of millisec of cr delay needed
dc	str	(P*)	Delete character
dF	num		Number of millisec of ff delay needed
dl	str	(P*)	Delete line
dm	str		Delete mode (enter)
dN	num		Number of millisec of nl delay needed
do	str		Down one line
dT	num		Number of millisec 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	bool	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-l9	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 264?)
xt	bool	Tabs are destructive, magic so char (Telaray 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=\EU\Ef\E7\E5\E8\EI\ENH\EK\E\200\Eo&\200:\
:al=3*\E^R:am:bs:cd=16*\E^C:ce=16\E^S:cl=2*\L:cm=\Ea%+ %+ :co#80:\
:dc=16\E^A:dl=3*\E^B:ei=\E\200:eo:im=\E^P:in:ip=16*:li#24:mi:nd=\E=: \
:se=\Ed\Ee:so=\ED\EE:ta=8\t:ul:up=\E;:vb=\Ek\EK: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 newline, 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

```
t3|33|tty33:co#72:os
```

while the Lear Siegler ADM-3 is described as

```
cl|adm33|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=cl,ll,sf,sb:"`, 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 stand-out (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 *lusr/lib/tabset/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 *term*lib 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 canceled 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), curses(3X), termcap(3X), tset(1), vi(1), ul(1), more(1)

AUTHOR

William Joy

Mark Horton added underlining and keypad support

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.

NAME

ttys — terminal initialization data

DESCRIPTION

The *ttys* file is read by the *init* program and specifies which terminal special files are to have a process created for them so that people can log in. There is one line in the *ttys* file per special file.

The first character of a line in the *ttys* file is either '0' or '1'. If the first character on the line is a '0', the *init* program ignores that line. If the first character on the line is a '1', the *init* program creates a login process for that line. The second character on each line is used as an argument to *getty*(8), which performs such tasks as baud-rate recognition, reading the login name, and calling *login*. For normal lines, the character is '0'; other characters can be used, for example, with hard-wired terminals where speed recognition is unnecessary or which have special characteristics. (*Getty* will have to be fixed in such cases.) The remainder of the line is the terminal's entry in the device directory, */dev*.

FILES

/etc/ttys

SEE ALSO

gettytab(5), *init*(8), *getty*(8), *login*(1)

NAME

`ttytype` — data base of terminal types by port

SYNOPSIS

`/etc/ttytype`

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 variable at login time.

SEE ALSO

`tset`(1), `login`(1)

BUGS

Some lines are merely known as “dialup” or “plugboard”.

NAME

types — primitive system data types

SYNOPSIS

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

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:

```
/*      types.h   6.1   83/07/29*/

/*
 * Basic system types and major/minor device constructing/busting macros.
 */

/* major part of a device */
#define major(x) ((int)((((unsigned)(x)>>8)&0377))

/* minor part of a device */
#define minor(x) ((int)((x)&0377))

/* make a device number */
#define makedev(x,y) ((dev_t)(((x)<<8) | (y)))

typedef unsigned char    u_char;
typedef unsigned short   u_short;
typedef unsigned int     u_int;
typedef unsigned long    u_long;
typedef unsigned short   ushort; /* sys III compat */

#ifdef vax
typedef struct    _physadr { int r[1]; } *physadr;
typedef struct    label_t {
    int            val[14];
} label_t;
#endif
typedef struct    _quad { long val[2]; } quad;
typedef long      daddr_t;
typedef char *    caddr_t;
typedef u_long    ino_t;
typedef long      swblk_t;
typedef int       size_t;
typedef int       time_t;
typedef short     dev_t;
typedef int       off_t;

typedef struct    fd_set { int fds_bits[1]; } fd_set;
```

The form *daddr_t* is used for disk addresses except in an i-node on disk, see *fs(5)*. 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(5), time(3), lseek(2), adb(1)

NAME

utmp, wtmp — login records

SYNOPSIS

```
#include <utmp.h>
```

DESCRIPTION

The *utmp* file records information about who is currently using the system. The file is a sequence of entries with the following structure declared in the include file:

```
/*      utmp.h 4.2      83/05/22      */

/*
 * Structure of utmp and wtmp files.
 *
 * Assuming the number 8 is unwise.
 */
struct utmp {
    char    ut_line[8];          /* tty name */
    char    ut_name[8];         /* user id */
    char    ut_host[16];        /* host name, if remote */
    long    ut_time;            /* time on */
};
```

This structure gives the name of the special file associated with the user's terminal, the user's login name, and the time of the login in the form of *time(3C)*.

The *wtmp* file records all logins and logouts. A null user name indicates a logout on the associated terminal. Furthermore, the terminal name "" indicates that the system was rebooted at the indicated time; the adjacent pair of entries with terminal names '|' and '|' indicate the system-maintained time just before and just after a *date* command has changed the system's idea of the time.

Wtmp is maintained by *login(1)* and *init(8)*. Neither of these programs creates the file, so if it is removed record-keeping is turned off. It is summarized by *ac(8)*.

FILES

```
/etc/utmp
/usr/adm/wtmp
```

SEE ALSO

login(1), *init(8)*, *who(1)*, *ac(8)*

NAME

`vfont` — font formats for the Benson-Varian or Versatec

SYNOPSIS

`/usr/lib/vfont/*`

DESCRIPTION

The fonts for the printer/plotters have the following format. Each file contains a header, an array of 256 character description structures, and then the bit maps for the characters themselves. The header has the following format:

```
struct header {
    short          magic;
    unsigned short size;
    short          maxx;
    short          maxy;
    short          xtnd;
} header;
```

The *magic* number is 0436 (octal). The *maxx*, *maxy*, and *xtnd* fields are not used at the current time. *Maxx* and *maxy* are intended to be the maximum horizontal and vertical size of any glyph in the font, in raster lines. The *size* is the size of the bit maps for the characters in bytes. Before the maps for the characters is an array of 256 structures for each of the possible characters in the font. Each element of the array has the form:

```
struct dispatch {
    unsigned short addr;
    short          nbytes;
    char           up;
    char           down;
    char           left;
    char           right;
    short          width;
};
```

The *nbytes* field is nonzero for characters which actually exist. For such characters, the *addr* field is an offset into the rest of the file where the data for that character begins. There are *up+down* rows of data for each character, each of which has *left+right* bits, rounded up to a number of bytes. The *width* field is not used by `vcat`, although it is to make width tables for `troff`. It represents the logical width of the glyph, in raster lines, and shows where the base point of the next glyph would be.

FILES

`/usr/lib/vfont/*`

SEE ALSO

`troff(1)`, `pti(1)`, `vpr(1)`, `vtroff(1)`, `vfontinfo(1)`

NAME

miscellaneous — miscellaneous useful information pages

DESCRIPTION

This section contains miscellaneous documentation, mostly in the area of text processing macro packages for *troff*(1).

ascii	map of ASCII character set
environ	user environment
eqnchar	special character definitions for eqn
hier	file system hierarchy
mailaddr	mail addressing description
man	macros to typeset manual pages
me	macros for formatting papers
ms	macros for formatting manuscripts
term	conventional names for terminals

NAME

ascii — map of ASCII character set

SYNOPSIS

cat /usr/pub/ascii

DESCRIPTION*Ascii* is a map of the ASCII character set, 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

environ — user environment

SYNOPSIS

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

DESCRIPTION

An array of strings called the 'environment' is made available by *execve(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=:/usr/ucb:/bin:/usr/bin`.
- HOME** A user's login directory, set by *login(1)* from the password file *passwd(5)*.
- TERM** The kind of terminal for which output is to be prepared. This information is used by commands, such as *nroff* or *plot(1G)*, which may exploit special terminal capabilities. See *letctermcap (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)*, *termcap(3X)*.
- EXINIT** A startup list of commands read by *ex(1)*, *edit(1)*, and *vi(1)*.
- USER** The login name of the user.
- PRINTER** The name of the default printer to be used by *lpr(1)*, *lpq(1)*, and *lprm(1)*.

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 *execve(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)*, *execve(2)*, *system(3)*, *termcap(3X)*, *termcap(5)*

NAME

eqnchar — special character definitions for eqn

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 Graphic Systems typesetter. These definitions are primarily intended for use with *eqn* and *neqn*. It contains definitions for the following characters

<i>ciplus</i>	\oplus	\parallel	\parallel	<i>square</i>	\square
<i>citimes</i>	\otimes	<i>langle</i>	\langle	<i>circle</i>	\circ
<i>wig</i>	\sim	<i>rangle</i>	\rangle	<i>blot</i>	\blacksquare
<i>-wig</i>	\equiv	<i>hbar</i>	\hbar	<i>bullet</i>	\bullet
<i>> wig</i>	\succsim	<i>ppd</i>	\perp	<i>prop</i>	\propto
<i>< wig</i>	\precsim	<i><-></i>	\longleftrightarrow	<i>empty</i>	\emptyset
<i>=wig</i>	\equiv	<i><=></i>	\longleftrightarrow	<i>member</i>	\in
<i>star</i>	$*$	<i> <</i>	\leftarrow	<i>nomem</i>	\notin
<i>bigstar</i>	$*$	<i> ></i>	\rightarrow	<i>cup</i>	\cup
<i>=dot</i>	$\dot{=}$	<i>ang</i>	\angle	<i>cap</i>	\cap
<i>orsign</i>	\vee	<i>rang</i>	\sphericalangle	<i>incl</i>	\sqsubset
<i>andsign</i>	\wedge	<i>3dot</i>	\vdots	<i>subset</i>	\subset
<i>=del</i>	\triangle	<i>thf</i>	\ddots	<i>supset</i>	\supset
<i>oppA</i>	∇	<i>quarter</i>	$\frac{1}{4}$	<i>!subset</i>	$\not\subset$
<i>oppE</i>	\equiv	<i>3quarter</i>	$\frac{3}{4}$	<i>!supset</i>	$\not\supset$
<i>angstrom</i>	\AA	<i>degree</i>	$^{\circ}$		

FILES

/usr/pub/eqnchar

SEE ALSO

troff(1), eqn(1)

NAME

hier — file system hierarchy

DESCRIPTION

The following outline gives a quick tour through a representative directory hierarchy.

```

/      root
/vmunix
      the kernel binary (UNIX itself)
/lost+found
      directory for connecting detached files for fsck(8)
/dev/  devices (4)
      MAKEDEV
          shell script to create special files
      MAKEDEV.local
          site specific part of MAKEDEV
      console
          main console, ty(4)
      tty*  terminals, ty(4)
      hp*   disks, hp(4)
      rhp*  raw disks, hp(4)
      up*   UNIBUS disks up(4)
      ...
/bin/  utility programs, cf /usr/bin/ (1)
      as    assembler
      cc    C compiler executive, cf /lib/ccom, /lib/cpp, /lib/c2
      csh   C shell
      ...
/lib/  object libraries and other stuff, cf /usr/lib/
      libc.a  system calls, standard I/O, etc. (2,3,3S)
      ...
      ccom  C compiler proper
      cpp   C preprocessor
      c2    C code improver
      ...
/etc/  essential data and maintenance utilities; sect (8)
      dump  dump program dump(8)
      passwd password file, passwd(5)
      group  group file, group(5)
      motd  message of the day, login(1)
      termcap
          description of terminal capabilities, termcap(5)
      ttytype table of what kind of terminal is on each port, ttytype(5)
      mtab  mounted file table, mtab(5)
      dumpdates
          dump history, dump(8)
      fstab  file system configuration table fstab(5)
      disktab disk characteristics and partition tables, disktab(5)
      hosts  host name to network address mapping file, hosts(5)
      networks
          network name to network number mapping file, networks(5)
      protocols
          protocol name to protocol number mapping file, protocols(5)
      services

```

network services definition file, *services*(5)
 remote names and description of remote hosts for *tip*(1C), *remote*(5)
 phones private phone numbers for remote hosts, as described in *phones*(5)
 ttys properties of terminals, *tty*(5)
 getty part of *login*, *getty*(8)
 init the parent of all processes, *init*(8)
 rc shell program to bring the system up
 rc.local site dependent portion of *rc*
 cron the clock daemon, *cron*(8)
 mount *mount*(8)
 ...
 /sys/ system source
 h/ header (include) files
 acct.h *acct*(5)
 stat.h *stat*(2)
 ...
 sys/ machine independent system source
 init_main.c
 uipc_socket.c
 ufs_syscalls.c
 ...
 conf/ site configuration files
 GENERIC
 ...
 net/ general network source
 netinet/ DARPA Internet network source
 netimp/ network code related to use of an IMP
 if_imp.c
 if_imphost.c
 if_imphost.h
 ...
 vax/ source specific to the VAX
 locore.s
 machdep.c
 ...
 vaxuba/
 device drivers for hardware which resides on the UNIBUS
 uba.c
 dh.c
 up.c
 ...
 vaxmba/
 device drivers for hardware which resides on the MASBUS
 mba.c
 hp.c
 ht.c
 ...
 vaxif network interface drivers for the VAX
 if_en.c
 if_ec.c

```

        if_vv.c
        ...
/tmp/ temporary files, usually on a fast device, cf /usr/tmp/
e*   used by ed(1)
ctm* used by cc(1)
...
/usr/ general-purpose directory, usually a mounted file system
adm/  administrative information
      wtmp  login history, utmp(5)
      messages
            hardware error messages
      tracct phototypesetter accounting, troff(1)
      lpacct line printer accounting lpr(1)
      vaacct, vpacct
            varian and versatec accounting vpr(1), vtroff(1), pac(8)
/usr  /bin
      utility programs, to keep /bin/ small
      tmp/ temporaries, to keep /tmp/ small
            stm* used by sort(1)
            raster used by plot(1G)
      dict/ word lists, etc.
            words principal word list, used by look(1)
            spellhist
                  history file for spell(1)
      games/
            hangman
      lib/  library of stuff for the games
            quiz.k/ what quiz(6) knows
                  index category index
                  africa countries and capitals
            ...
      ...
include/
      standard #include files
      a.out.h object file layout, a.out(5)
      stdio.h standard I/O, intro(3S)
      math.h (3M)
      ...
      sys/  system-defined layouts, cf /sys/h
      net/  symbolic link to sys/net
      machine/
            symbolic link to sys/machine
      ...
lib/  object libraries and stuff, to keep /lib/ small
      atrun scheduler for at(1)
      lint/ utility files for lint
            lint[12]
                  subprocesses for lint(1)
      llib-lc dummy declarations for /lib/libc.a, used by lint(1)
      llib-lm dummy declarations for /lib/libc.m
      ...

```

```

struct/ passes of struct(1)
...
tmac/ macros for troff(1)
      tmac.an
           macros for man(7)
      tmac.s macros for ms(7)
...
font/ fonts for troff(1)
      ftR    Times Roman
      ftB    Times Bold
...
uucp/ programs and data for uucp(1C)
      L.sys  remote system names and numbers
      uucico the real copy program
...
units conversion tables for units(1)
eign  list of English words to be ignored by ptx(1)
/usr/ man/
      volume 1 of this manual, man(1)
      man0/ general
            intro  introduction to volume 1, ms(7) format
            xx     template for manual page
      man1/ chapter 1
            as.1
            mount.lm
            ...
      ...
      cat1/ preformatted pages for section 1
      ...
msgs/ messages, cf msgs(1)
      bounds highest and lowest message
new/  binaries of new versions of programs
preserve/
      editor temporaries preserved here after crashes/hangups
public/ binaries of user programs - write permission to everyone
spool/ delayed execution files
      at/    used by at(1)
      lpd/   used by lpr(1)
            lock   present when line printer is active
            cf*   copy of file to be printed, if necessary
            df*   daemon control file, lpd(8)
            tf*   transient control file, while lpr is working
uucp/  work files and staging area for uucp(1C)
      LOGFILE
            summary log
      LOG.* log file for one transaction
mail/  mailboxes for mail(1)
      name  mail file for user name
      name.lock
            lock file while name is receiving mail
secretmail/
      like mail/

```


uucp/ work files and staging area for *uucp*(1C)
 LOGFILE
 summary log
 LOG.* log file for one transaction
 mqueue/
 mail queue for *sendmail*(8)

wd initial working directory of a user, typically *wd* is the user's login name
 .profile set environment for *sh*(1), *environ*(7)
 .project
 what you are doing (used by (*finger*(1))
 .cshrc startup file for *csh*(1)
 .exrc startup file for *ex*(1)
 .plan what your short-term plans are (used by *finger*(1))
 .netrc startup file for various network programs
 .msgsrc
 startup file for *msgs*(1)
 .mailrc startup file for *mail*(1)
 calendar
 user's datebook for *calendar*(1)

doc/ papers, mostly in volume 2 of this manual, typically in *ms*(7) format
 as/ assembler manual
 c C manual
 ...

/usr/ *src/*
 source programs for utilities, etc.
 bin/ source of commands in */bin*
 as/ assembler
 ar.c source for *ar*(1)
 ...

usr.bin/
 source for commands in */usr/bin*
 troff/ source for *nroff* and *troff*(1)
 font/ source for font tables, */usr/lib/font/*
 ftR.c Roman
 ...
 term/ terminal characteristics tables, */usr/lib/term/*
 tab300.c
 DASI 300
 ...
 ...

 ucb source for programs in */usr/ucb*
 games/ source for */usr/games*
 lib/ source for programs and archives in */lib*
 libc/ C runtime library
 csu/ startup and wrapup routines needed with every C program
 crt0.s regular startup
 mcrt0.s modified startup for *cc -p*
 sys/ system calls (2)
 access.s
 brk.s
 ...

 stdio/ standard I/O functions (3S)

fgets.c
fopen.c
...
gen/ other functions in (3)
abs.c
...
net/ network functions in (3N)
gethostbyname.c
...
local/ source which isn't normally distributed
new/ source for new versions of commands and library routines
old/ source for old versions of commands and library routines
ucb/ binaries of programs developed at UCB
...
edit editor for beginners
ex command editor for experienced users
...
mail mail reading/sending subsystem
man on line documentation
...
pi Pascal translator
px Pascal interpreter
...
vi visual editor

SEE ALSO

ls(1), apropos(1), whatis(1), whereis(1), finger(1), which(1), ncheck(8), find(1), grep(1)

BUGS

The position of files is subject to change without notice.

NAME

mailaddr — mail addressing description

DESCRIPTION

Mail addresses are based on the ARPANET protocol listed at the end of this manual page. These addresses are in the general format

user@domain

where a domain is a hierarchical dot separated list of subdomains. For example, the address

eric@monet.Berkeley.ARPA

is normally interpreted from right to left: the message should go to the ARPA name tables (which do not correspond exactly to the physical ARPANET), then to the Berkeley gateway, after which it should go to the local host monet. When the message reaches monet it is delivered to the user "eric".

Unlike some other forms of addressing, this does not imply any routing. Thus, although this address is specified as an ARPA address, it might travel by an alternate route if that was more convenient or efficient. For example, at Berkeley the associated message would probably go directly to monet over the Ethernet rather than going via the Berkeley ARPANET gateway.

Abbreviation. Under certain circumstances it may not be necessary to type the entire domain name. In general anything following the first dot may be omitted if it is the same as the domain from which you are sending the message. For example, a user on "calder.Berkeley.ARPA" could send to "eric@monet" without adding the ".Berkeley.ARPA" since it is the same on both sending and receiving hosts.

Certain other abbreviations may be permitted as special cases. For example, at Berkeley ARPANET hosts can be referenced without adding the ".ARPA" as long as their names do not conflict with a local host name.

Compatibility. Certain old address formats are converted to the new format to provide compatibility with the previous mail system. In particular,

host:user

is converted to

user@host

to be consistent with the *rcp*(1C) command.

Also, the syntax:

host!user

is converted to:

user@host.UUCP

This is normally converted back to the "host!user" form before being sent on for compatibility with older UUCP hosts.

The current implementation is not able to route messages automatically through the UUCP network. Until that time you must explicitly tell the mail system which hosts to send your message through to get to your final destination.

Case Distinctions. Domain names (i.e., anything after the "@" sign) may be given in any mixture of upper and lower case with the exception of UUCP hostnames. Most hosts accept any mixture of case in user names, with the notable exception of MULTICS sites.

Differences with ARPA Protocols. Although the UNIX addressing scheme is based on the ARPA mail addressing protocols, there are some significant differences.

At the time of this writing the only "top level" domain defined by ARPA is the ".ARPA" domain itself. This is further restricted to having only one level of host specifier. That is, the only addresses that ARPA accepts at this time must be in the format "user@host.ARPA" (where "host" is one word). In particular, addresses such as:

eric@monet.Berkeley.ARPA

are not currently legal under the ARPA protocols. For this reason, these addresses are converted to a different format on output to the ARPANET, typically:

eric%monet@Berkeley.ARPA

Route-addr. Under some circumstances it may be necessary to route a message through several hosts to get it to the final destination. Normally this routing is done automatically, but sometimes it is desirable to route the message manually. An address that shows these relays are termed "route-addr." These use the syntax:

<@hosta,@hostb:user@hostc>

This specifies that the message should be sent to hosta, from there to hostb, and finally to hostc. This path is forced even if there is a more efficient path to hostc.

Route-addr occur frequently on return addresses, since these are generally augmented by the software at each host. It is generally possible to ignore all but the "user@host" part of the address to determine the actual sender.

Postmaster. Every site is required to have a user or user alias designated "postmaster" to which problems with the mail system may be addressed.

CSNET. Messages to CSNET sites can be sent to "user.host@UDel-Relay".

BERKELEY

The following comments apply only to the Berkeley environment.

Host Names. Many of the old familiar host names are being phased out. In particular, single character names as used in Berknet are incompatible with the larger world of which Berkeley is now a member. For this reason the following names are being obsoleted. You should notify any correspondents of your new address as soon as possible.

OLD	NEW	j ingvax	ucbingres
p	ucbcad	r arpavax	ucbarpa
v csvax	ucbernie		
n	ucbkim	y	ucbcory

The old addresses will be rejected as unknown hosts sometime in the near future.

What's My Address? If you are on a local machine, say monet, your address is

yourname@monet.Berkeley.ARPA

However, since most of the world does not have the new software in place yet, you will have to give correspondents slightly different addresses. From the ARPANET, your address would be:

yourname%monet@Berkeley.ARPA

From UUCP, your address would be:

ucbvax!yourname%monet

Computer Center. The Berkeley Computer Center is in a subdomain of Berkeley. Messages to the computer center should be addressed to:

user%host.CC@Berkeley.ARPA

The alternate syntax:

user@host.CC

may be used if the message is sent from inside Berkeley.

For the time being Computer Center hosts are known within the Berkeley domain, i.e., the ".CC" is optional. However, it is likely that this situation will change with time as both the Computer Science department and the Computer Center grow.

Bitnet. Hosts on bitnet may be accessed using:

user@host.BITNET

SEE ALSO

mail(1), sendmail(8); Crocker, D. H., *Standard for the Format of Arpa Internet Text Messages*, RFC822.

NAME

man — macros to typeset manual

SYNOPSIS

nroff **-man** file ...

troff **-man** file ...

DESCRIPTION

These macros are used to lay out pages of this manual. A skeleton page may be found in the file /usr/man/man0/xx.

Any text argument *t* may be zero to six words. Quotes may be used to include blanks in a 'word'. If *text* is empty, the special treatment is applied to the next input line with *text* to be printed. In this way .I may be used to italicize a whole line, or .SM followed by .B to make small bold letters.

A prevailing indent distance is remembered between successive indented paragraphs, and is reset to default value upon reaching a non-indented paragraph. Default units for indents *i* are ens.

Type font and size are reset to default values before each paragraph, and after processing font and size setting macros.

These strings are predefined by **-man**:

*R '•', '(Reg)' in *nroff*.

*S Change to default type size.

FILES

/usr/lib/tmac/tmac.an

/usr/man/man0/xx

SEE ALSO

troff(1), man(1)

BUGS

Relative indents don't nest.

REQUESTS

Request	Cause	If no	Explanation
	Break	Argument	
.B <i>t</i>	no	<i>t</i> =n.t.l.*	Text <i>t</i> is bold.
.BI <i>t</i>	no	<i>t</i> =n.t.l.	Join words of <i>t</i> alternating bold and italic.
.BR <i>t</i>	no	<i>t</i> =n.t.l.	Join words of <i>t</i> alternating bold and Roman.
.DT	no	.Si li...	Restore default tabs.
.HP <i>i</i>	yes	<i>i</i> =p.i.*	Set prevailing indent to <i>i</i> . Begin paragraph with hanging indent.
.I <i>t</i>	no	<i>t</i> =n.t.l.	Text <i>t</i> is italic.
.IB <i>t</i>	no	<i>t</i> =n.t.l.	Join words of <i>t</i> alternating italic and bold.
.IP <i>x i</i>	yes	<i>x</i> =""	Same as .TP with tag <i>x</i> .
.IR <i>t</i>	no	<i>t</i> =n.t.l.	Join words of <i>t</i> alternating italic and Roman.
.LP	yes	-	Same as .PP.
.PD <i>d</i>	no	<i>d</i> =.4v	Interparagraph distance is <i>d</i> .
.PP	yes	-	Begin paragraph. Set prevailing indent to .Si.
.RE	yes	-	End of relative indent. Set prevailing indent to amount of starting .RS.
.RB <i>t</i>	no	<i>t</i> =n.t.l.	Join words of <i>t</i> alternating Roman and bold.
.RI <i>t</i>	no	<i>t</i> =n.t.l.	Join words of <i>t</i> alternating Roman and italic.
.RS <i>i</i>	yes	<i>i</i> =p.i.	Start relative indent, move left margin in distance <i>i</i> . Set prevailing indent to .Si for nested indents.
.SH <i>t</i>	yes	<i>t</i> =n.t.l.	Subhead.

.SM *t* no *t*=n.t.l. Text *t* is small.
.TH *n c x v m* yes - Begin page named *n* of chapter *c*; *x* is extra commentary, e.g. 'local', for page foot center; *v* alters page foot left, e.g. '4th Berkeley Distribution'; *m* alters page head center, e.g. 'Brand X Programmer's Manual'. Set prevailing indent and tabs to .5i.
.TP *i* yes *i*=p.i. Set prevailing indent to *i*. Begin indented paragraph with hanging tag given by next text line. If tag doesn't fit, place it on separate line.

* n.t.l. = next text line; p.i. = prevailing indent

NAME

me — macros for formatting papers

SYNOPSIS

nroff **-me** [options] file ...

troff **-me** [options] file ...

DESCRIPTION

This package of *nroff* and *troff* macro definitions provides a canned formatting facility for technical papers in various formats. When producing 2-column output on a terminal, filter the output through *col(1)*.

The macro requests are defined below. Many *nroff* and *troff* requests are unsafe in conjunction with this package, however these requests may be used with impunity after the first *.pp*:

```
.bp      begin new page
.br      break output line here
.sp n    insert n spacing lines
.ls n    (line spacing) n=1 single, n=2 double space
.na      no alignment of right margin
.ce n    center next n lines
.ul n    underline next n lines
.sz +n   add n to point size
```

Output of the *eqn*, *neqn*, *refer*, and *tbl(1)* preprocessors for equations and tables is acceptable as input.

FILES

/usr/lib/tmac/tmac.e

/usr/lib/me/*

SEE ALSO

eqn(1), *troff(1)*, *refer(1)*, *tbl(1)*

-me Reference Manual, Eric P. Allman

Writing Papers with Nroff Using *-me*

REQUESTS

In the following list, "initialization" refers to the first *.pp*, *.lp*, *.ip*, *.np*, *.sh*, or *.uh* macro. This list is incomplete; see *The -me Reference Manual* for interesting details.

Request	Initial Value	Cause	Explanation
.(c	-	yes	Begin centered block
.(d	-	no	Begin delayed text
.(f	-	no	Begin footnote
.(l	-	yes	Begin list
.(q	-	yes	Begin major quote
.(x x	-	no	Begin indexed item in index x
.(z	-	no	Begin floating keep
.)c	-	yes	End centered block
.)d	-	yes	End delayed text
.)f	-	yes	End footnote
.)l	-	yes	End list
.)q	-	yes	End major quote
.)x	-	yes	End index item
.)z	-	yes	End floating keep
.)+ + m H -	-	no	Define paper section. <i>m</i> defines the part of the paper, and can be C (chapter), A (appendix), P (preliminary, e.g., abstract, table of contents, etc.), B

			(bibliography), RC (chapters renumbered from page one each chapter), or RA (appendix renumbered from page one).
.+c <i>T</i>	-	yes	Begin chapter (or appendix, etc., as set by .++). <i>T</i> is the chapter title.
.1c	1	yes	One column format on a new page.
.2c	1	yes	Two column format.
.EN	-	yes	Space after equation produced by <i>eqn</i> or <i>neqn</i> .
.EQ <i>x y</i>	-	yes	Precede equation; break out and add space. Equation number is <i>y</i> . The optional argument <i>x</i> may be <i>I</i> to indent equation (default), <i>L</i> to left-adjust the equation, or <i>C</i> to center the equation.
.TE	-	yes	End table.
.TH	-	yes	End heading section of table.
.TS <i>x</i>	-	yes	Begin table; if <i>x</i> is <i>H</i> table has repeated heading.
.ac <i>A N</i>	-	no	Set up for ACM style output. <i>A</i> is the Author's name(s), <i>N</i> is the total number of pages. Must be given before the first initialization.
.b <i>x</i>	no	no	Print <i>x</i> in boldface; if no argument switch to boldface.
.ba <i>+n</i>	0	yes	Augments the base indent by <i>n</i> . This indent is used to set the indent on regular text (like paragraphs).
.bc	no	yes	Begin new column
.bi <i>x</i>	no	no	Print <i>x</i> in bold italics (nofill only)
.bx <i>x</i>	no	no	Print <i>x</i> in a box (nofill only).
.ef ' <i>x'y'z</i> '	****	no	Set even footer to <i>x y z</i>
.eh ' <i>x'y'z</i> '	****	no	Set even header to <i>x y z</i>
.fo ' <i>x'y'z</i> '	****	no	Set footer to <i>x y z</i>
.hx	-	no	Suppress headers and footers on next page.
.he ' <i>x'y'z</i> '	****	no	Set header to <i>x y z</i>
.hl	-	yes	Draw a horizontal line
.i <i>x</i>	no	no	Italicize <i>x</i> ; if <i>x</i> missing, italic text follows.
.ip <i>x y</i>	no	yes	Start indented paragraph, with hanging tag <i>x</i> . Indentation is <i>y</i> ens (default 5).
.lp	yes	yes	Start left-blocked paragraph.
.lo	-	no	Read in a file of local macros of the form <i>.x</i> . Must be given before initialization.
.np	1	yes	Start numbered paragraph.
.of ' <i>x'y'z</i> '	****	no	Set odd footer to <i>x y z</i>
.oh ' <i>x'y'z</i> '	****	no	Set odd header to <i>x y z</i>
.pd	-	yes	Print delayed text.
.pp	no	yes	Begin paragraph. First line indented.
.r	yes	no	Roman text follows.
.re	-	no	Reset tabs to default values.
.sc	no	no	Read in a file of special characters and diacritical marks. Must be given before initialization.
.sh <i>n x</i>	-	yes	Section head follows, font automatically bold. <i>n</i> is level of section, <i>x</i> is title of section.
.sk	no	no	Leave the next page blank. Only one page is remembered ahead.
.sz <i>+n</i>	10p	no	Augment the point size by <i>n</i> points.
.th	no	no	Produce the paper in thesis format. Must be given before initialization.
.tp	no	yes	Begin title page.
.u <i>x</i>	-	no	Underline argument (even in <i>troff</i>). (Nofill only).
.uh	-	yes	Like <i>.sh</i> but unnumbered.
.xp <i>x</i>	-	no	Print index <i>x</i> .

NAME

ms — text formatting macros

SYNOPSIS

nroff **-ms** [options] file ...
troff **-ms** [options] file ...

DESCRIPTION

This package of *nroff* and *troff* macro definitions provides a formatting facility for various styles of articles, theses, and books. When producing 2-column output on a terminal or lineprinter, or when reverse line motions are needed, filter the output through *col(1)*. All external **-ms** macros are defined below. Many *nroff* and *troff* requests are unsafe in conjunction with this package. However, the first four requests below may be used with impunity after initialization, and the last two may be used even before initialization:

.bp begin new page
.br break output line
.sp n insert n spacing lines
.ce n center next n lines
.ls n line spacing: n=1 single, n=2 double space
.na no alignment of right margin

Font and point size changes with **\f** and **\s** are also allowed; for example, "**\fIword\fR**" will italicize *word*. Output of the *tbl*, *eqn*, and *refer(1)* preprocessors for equations, tables, and references is acceptable as input.

FILES

/usr/lib/tmac/tmac.x
/usr/lib/ms/x.???

SEE ALSO

eqn(1), *refer(1)*, *tbl(1)*, *troff(1)*

REQUESTS

Macro Name	Initial Value	Break? Reset?	Explanation
.AB <i>x</i>	—	y	begin abstract; if <i>x</i> = no don't label abstract
.AE	—	y	end abstract
.AI	—	y	author's institution
.AM	—	n	better accent mark definitions
.AU	—	y	author's name
.B <i>x</i>	—	n	embolden <i>x</i> ; if no <i>x</i> , switch to boldface
.B1	—	y	begin text to be enclosed in a box
.B2	—	y	end boxed text and print it
.BT	date	n	bottom title, printed at foot of page
.BX <i>x</i>	—	n	print word <i>x</i> in a box
.CM	if t	n	cut mark between pages
.CT	—	y,y	chapter title: page number moved to CF (TM only)
.DA <i>x</i>	if n	n	force date <i>x</i> at bottom of page; today if no <i>x</i>
.DE	—	y	end display (unfilled text) of any kind
.DS <i>x y</i>	I	y	begin display with keep; <i>x</i> = I,L,C,B; <i>y</i> = indent
.ID <i>y</i>	8n,.5i	y	indented display with no keep; <i>y</i> = indent
.LD	—	y	left display with no keep
.CD	—	y	centered display with no keep
.BD	—	y	block display; center entire block
.EF <i>x</i>	—	n	even page footer <i>x</i> (3 part as for .tl)
.EH <i>x</i>	—	n	even page header <i>x</i> (3 part as for .tl)

.EN	-	y	end displayed equation produced by <i>eqn</i>
.EQ <i>x y</i>	-	y	break out equation; <i>x</i> =L,I,C; <i>y</i> =equation number
.FE	-	n	end footnote to be placed at bottom of page
.FP	-	n	numbered footnote paragraph; may be redefined
.FS <i>x</i>	-	n	start footnote; <i>x</i> is optional footnote label
.HD	undef	n	optional page header below header margin
.I <i>x</i>	-	n	italicize <i>x</i> ; if no <i>x</i> , switch to italics
.IP <i>x y</i>	-	y,y	indented paragraph, with hanging tag <i>x</i> ; <i>y</i> =indent
.IX <i>x y</i>	-	y	index words <i>x y</i> and so on (up to 5 levels)
.KE	-	n	end keep of any kind
.KF	-	n	begin floating keep; text fills remainder of page
.KS	-	y	begin keep; unit kept together on a single page
.LG	-	n	larger; increase point size by 2
.LP	-	y,y	left (block) paragraph.
.MC <i>x</i>	-	y,y	multiple columns; <i>x</i> =column width
.ND <i>x</i>	if t	n	no date in page footer; <i>x</i> is date on cover
.NH <i>x y</i>	-	y,y	numbered header; <i>x</i> =level, <i>x</i> =0 resets, <i>x</i> =S sets to <i>y</i>
.NL	10p	n	set point size back to normal
.OF <i>x</i>	-	n	odd page footer <i>x</i> (3 part as for .tl)
.OH <i>x</i>	-	n	odd page header <i>x</i> (3 part as for .tl)
.P1	if TM	n	print header on 1st page
.PP	-	y,y	paragraph with first line indented
.PT	- % -	n	page title, printed at head of page
.PX <i>x</i>	-	y	print index (table of contents); <i>x</i> =no suppresses title
.QP	-	y,y	quote paragraph (indented and shorter)
.R	on	n	return to Roman font
.RE	5n	y,y	retreat: end level of relative indentation
.RP <i>x</i>	-	n	released paper format; <i>x</i> =no stops title on 1st page
.RS	5n	y,y	right shift: start level of relative indentation
.SH	-	y,y	section header, in boldface
.SM	-	n	smaller; decrease point size by 2
.TA	8n,5n	n	set tabs to 8n 16n ... (nroff) 5n 10n ... (troff)
.TC <i>x</i>	-	y	print table of contents at end; <i>x</i> =no suppresses title
.TE	-	y	end of table processed by <i>tbl</i>
.TH	-	y	end multi-page header of table
.TL	-	y	title in boldface and two points larger
.TM	off	n	UC Berkeley thesis mode
.TS <i>x</i>	-	y,y	begin table; if <i>x</i> =H table has multi-page header
.UL <i>x</i>	-	n	underline <i>x</i> , even in <i>troff</i>
.UX <i>x</i>	-	n	UNIX; trademark message first time; <i>x</i> appended
.XA <i>x y</i>	-	y	another index entry; <i>x</i> =page or no for none; <i>y</i> =indent
.XE	-	y	end index entry (or series of .IX entries)
.XP	-	y,y	paragraph with first line exdented, others indented
.XS <i>x y</i>	-	y	begin index entry; <i>x</i> =page or no for none; <i>y</i> =indent
.1C	on	y,y	one column format, on a new page
.2C	-	y,y	begin two column format
.]-	-	n	beginning of <i>refer</i> reference
.[0	-	n	end of unclassifiable type of reference
.[N	-	n	N= 1:journal-article, 2:book, 3:book-article, 4:report

REGISTERS

Formatting distances can be controlled in `-ms` by means of built-in number registers. For example, this sets the line length to 6.5 inches:

```
.nr LL 6.5i
```

Here is a table of number registers and their default values:

Name	Register Controls	Takes Effect	Default
PS	point size	paragraph	10
VS	vertical spacing	paragraph	12
LL	line length	paragraph	6i
LT	title length	next page	same as LL
FL	footnote length	next .FS	5.5i
PD	paragraph distance	paragraph	1v (if n), .3v (if t)
DD	display distance	displays	1v (if n), .5v (if t)
PI	paragraph indent	paragraph	5n
QI	quote indent	next .QP	5n
FI	footnote indent	next .FS	2n
PO	page offset	next page	0 (if n), ~1i (if t)
HM	header margin	next page	1i
FM	footer margin	next page	1i
FF	footnote format	next .FS	0 (1, 2, 3 available)

When resetting these values, make sure to specify the appropriate units. Setting the line length to 7, for example, will result in output with one character per line. Setting FF to 1 suppresses footnote superscripting; setting it to 2 also suppresses indentation of the first line; and setting it to 3 produces an .IP-like footnote paragraph.

Here is a list of string registers available in `-ms`; they may be used anywhere in the text:

Name	String's Function
*Q	quote (" in <i>nroff</i> , " in <i>troff</i>)
*U	unquote (" in <i>nroff</i> , " in <i>troff</i>)
*-	dash (-- in <i>nroff</i> , - in <i>troff</i>)
*(MO	month (month of the year)
*(DY	day (current date)
**	automatically numbered footnote
*'	acute accent (before letter)
*`	grave accent (before letter)
*^	circumflex (before letter)
*,	cedilla (before letter)
*:	umlaut (before letter)
*~	tilde (before letter)

When using the extended accent mark definitions available with `.AM`, these strings should come after, rather than before, the letter to be accented.

BUGS

Floating keeps and regular keeps are diverted to the same space, so they cannot be mixed together with predictable results.

NAME

term — conventional names for terminals

DESCRIPTION

Certain commands use these terminal names. They are maintained as part of the shell environment (see *sh(1)*, *environ(7)*).

adm3a	Lear Seigler Adm-3a
2621	Hewlett-Packard HP262? series terminals
hp	Hewlett-Packard HP264? series terminals
c100	Human Designed Systems Concept 100
h19	Heathkit H19
mime	Microterm mime in enhanced ACT IV mode
1620	DIABLO 1620 (and others using HyType II)
300	DASI/DTC/GSI 300 (and others using HyType I)
33	TELETYPE® Model 33
37	TELETYPE Model 37
43	TELETYPE Model 43
735	Texas Instruments TI735 (and TI725)
745	Texas Instruments TI745
dumb	terminals with no special features
dialup	a terminal on a phone line with no known characteristics
network	a terminal on a network connection with no known characteristics
4014	Tektronix 4014
vt52	Digital Equipment Corp. VT52

The list goes on and on. Consult */etc/termcap* (see *termcap(5)*) for an up-to-date and locally correct list.

Commands whose behavior may depend on the terminal either consult **TERM** in the environment, or accept arguments of the form **-Tterm**, where *term* is one of the names given above.

SEE ALSO

stty(1), *tabs(1)*, *plot(1G)*, *sh(1)*, *environ(7)* *ex(1)*, *clear(1)*, *more(1)*, *ul(1)*, *tset(1)*, *termcap(5)*, *termcap(3X)*, *ttytype(5)*
troff(1) for *nroff*

BUGS

The programs that ought to adhere to this nomenclature do so only fitfully.

NAME

intro — introduction to system maintenance and operation commands

DESCRIPTION

This section contains information related to system operation and maintenance. In particular, commands used to create new file systems, *newfs*, *mkfs*, and verify the integrity of the file systems, *fsck*, *icheck*, *dcheck*, and *ncheck* are described here. The section *format* should be consulted when formatting disk packs. The section *crash* should be consulted in understanding how to interpret system crash dumps.

LIST OF PROGRAMS

<i>Program</i>	<i>Appears on Page</i>	<i>Description</i>
ac	ac.8	login accounting
accton	sa.8	system accounting
adduser	adduser.8	procedure for adding new users
analyze	analyze.8	Virtual UNIX postmortem crash analyzer
arcv	arcv.8	convert archives to new format
arff	arff.8v	archiver and copier for floppy
bad144	bad144.8	read/write dec standard 144 bad sector information
badsect	badsect.8	create files to contain bad sectors
bugfiler	bugfiler.8	file bug reports in folders automatically
catman	catman.8	create the cat files for the manual
chown	chown.8	change owner
clri	clri.8	clear i-node
comsat	comsat.8c	biff server
config	config.8	build system configuration files
crash	crash.8v	what happens when the system crashes
cron	cron.8	clock daemon
dcheck	dcheck.8	file system directory consistency check
diskpart	diskpart.8	calculate default disk partition sizes
dmesg	dmesg.8	collect system diagnostic messages to form error log
drtest	drtest.8	standalone disk test program
dump	dump.8	incremental file system dump
dumpfs	dumpfs.8	dump file system information
edquota	edquota.8	edit user quotas
fastboot	fastboot.8	reboot/halt the system without checking the disks
fasthalt	fastboot.8	reboot/halt the system without checking the disks
flcopy	arff.8v	archiver and copier for floppy
format	format.8v	how to format disk packs
fsck	fsck.8	file system consistency check and interactive repair
ftpd	ftpd.8c	DARPA Internet File Transfer Protocol server
gettable	gettable.8c	get NIC format host tables from a host
getty	getty.8	set terminal mode
halt	halt.8	stop the processor
htable	htable.8	convert NIC standard format host tables
icheck	icheck.8	file system storage consistency check
ifconfig	ifconfig.8c	configure network interface parameters
implog	implog.8c	IMP log interpreter
implogd	implogd.8c	IMP logger process
init	init.8	process control initialization
kgmon	kgmon.8	generate a dump of the operating systems profile buffers
lpc	lpc.8	line printer control program
lpd	lpd.8	line printer daemon

makedev	makedev.8	make system special files
makekey	makekey.8	generate encryption key
mkfs	mkfs.8	construct a file system
mklost+found	mklost+found.8	make a lost+found directory for fsck
mknod	mknod.8	build special file
mkproto	mkproto.8	construct a prototype file system
mount	mount.8	mount and dismount file system
ncheck	ncheck.8	generate names from i-numbers
newfs	newfs.8	construct a new file system
pac	pac.8	printer/ploter accounting information
pstat	pstat.8	print system facts
quot	quot.8	summarize file system ownership
quotacheck	quotacheck.8	file system quota consistency checker
quotaoff	quotaon.8	turn file system quotas on and off
quotaon	quotaon.8	turn file system quotas on and off
rc	rc.8	command script for auto-reboot and daemons
rdump	rdump.8c	file system dump across the network
reboot	reboot.8	UNIX bootstrapping procedures
renice	renice.8	alter priority of running processes
repquota	repquota.8	summarize quotas for a file system
restore	restore.8	incremental file system restore
rexecd	rexecd.8c	remote execution server
rlogind	rlogind.8c	remote login server
rmt	rmt.8c	remote magtape protocol module
route	route.8c	manually manipulate the routing tables
routed	routed.8c	network routing daemon
rrestore	rrestore.8c	restore a file system dump across the network
rshd	rshd.8c	remote shell server
rwhod	rwhod.8c	system status server
rxformat	rxformat.8v	format floppy disks
sa	sa.8	system accounting
savecore	savecore.8	save a core dump of the operating system
sendmail	sendmail.8	send mail over the internet
shutdown	shutdown.8	close down the system at a given time
sticky	sticky.8	executable files with persistent text
swapon	swapon.8	specify additional device for paging and swapping
sync	sync.8	update the super block
syslog	syslog.8	log systems messages
telnetd	telnetd.8c	DARPA TELNET protocol server
tftpd	tftpd.8c	DARPA Trivial File Transfer Protocol server
trpt	trpt.8c	transliterate protocol trace
tunefs	tunefs.8	tune up an existing file system
umount	mount.8	mount and dismount file system
update	update.8	periodically update the super block
uuclean	uuclean.8c	uucp spool directory clean-up
uusnap	uusnap.8c	show snapshot of the UUCP system
vipw	vipw.8	edit the password file

NAME

ac — login accounting

SYNOPSIS

/etc/ac [**-w** *wtmp*] [**-p**] [**-d**] [*people*] ...

DESCRIPTION

Ac produces a printout giving connect time for each user who has logged in during the life of the current *wtmp* file. A total is also produced. **-w** is used to specify an alternate *wtmp* file. **-p** prints individual totals; without this option, only totals are printed. **-d** causes a printout for each midnight to midnight period. Any *people* will limit the printout to only the specified login names. If no *wtmp* file is given, */usr/adm/wtmp* is used.

The accounting file */usr/adm/wtmp* is maintained by *init* and *login*. Neither of these programs creates the file, so if it does not exist no connect-time accounting is done. To start accounting, it should be created with length 0. On the other hand if the file is left undisturbed it will grow without bound, so periodically any information desired should be collected and the file truncated.

FILES

/usr/adm/wtmp

SEE ALSO

init(8), *sa*(8), *login*(1), *utmp*(5).

NAME

backup - make backup tapes

SYNOPSIS

backup

DESCRIPTION

Backup is an interactive front-end to the UNIX dump program, *dump(8)*. Its user is prompted as to the type of dump (daily/weekly/ monthly), whether to dump the root file system (this usually doesn't change -- no need to make frequent backups), and whether or not to invoke *fsck(8)* to check the file systems before dumping. The file systems will only be checked if the system is running single user. (Note: races can develop between users of files and dumpers; to ensure that all files are backed up when they should be, we recommend that *backup* be run from the single user shell). *Backup* proceeds by calling *getfs(8)* to produce a list of all file systems in the system; it calls *dump* on each one of these in turn, prompting the user to change tapes as required.

FILES

/etc/fstab - file system data base

/etc/dump - actual dump program

DIAGNOSTICS

Are intended to be self-explanatory; mostly about invalid typed input.

NAME

`chkhosts` - update `/etc/hosts`, `/etc/.rhosts`, and `/etc/hosts.equiv` data bases.

SYNOPSIS

```
chkhosts  
chkhosts -clean  
chkhosts -add <hostname> <hostaddress>  
chkhosts -delete <hostname> ...
```

DESCRIPTION

Chkhosts maintains the data base files used by the Berkeley TCP/IP based facilities *rlogin(1)* and *rsh(1)*. The default invocation (no argument) is to have `/etc/hosts` add to the growing list of address-to-names by querying the incore connection table. New address/name pairs can be explicitly added using the **-add** option. Address/name pairs can be deleted using the **-delete** option. The **-clean** option clears out all the data bases and inserts a dummy entry for the local machine in each file.

Currently, `/etc/hosts.equiv` and `/.rhosts` reflect the current state of `/etc/hosts` {i.e., each time *chkhosts* is run, these two files are rebuilt from the same information that is used to build the `/etc/hosts` file. With the `/.rhosts` file, the name list associated with each host, if applicable, is saved.

FILES

```
/etc/hosts  
/etc/hosts.equiv  
/.rhosts  
/tmp/hosts,tmp  
/tmp/hosts.equiv,tmp  
/tmp/rhosts,tmp
```

SEE ALSO

`rlogin(1)`, `rsh(1)`, `hosts(5)`, `rlogind(8)`, `rshd(8)`

BUGS

The program should handle internetting.

The program should allow the user to manipulate the `/.rhosts` name lists.

The program should allow the user to specify if the `/etc/hosts.equiv` file should be an image of `/etc/hosts`.

NAME

chown - change owner

SYNOPSIS

/etc/chown [-f] owner file ...

DESCRIPTION

Chown changes the owner of the *files* to *owner*. The owner may be either a decimal UID or a login name found in the password file.

Only the super-user can change owner, in order to simplify accounting procedures. No errors are reported when the -f (force) option is given.

FILES

/etc/passwd

SEE ALSO

chgrp(1), chown(2), passwd(5), group(5)

NAME

chuid - changes uid/gid's on directory trees according to command line arguments.

SYNOPSIS

chuid [- **d** directory tree] [- **D**] [- **o** old password file] [- **n** new password file] [- **l** logfile]
[- **r** restart file] [- **u** user olduid]

DESCRIPTION

Chuid executes a wholesale *chown(2)* on the entire directory tree. It does this by constructing a map of old to new uid and gid's, following the directory tree, logging nodes touched, and calling *chown(2)* for each accessible file that should be changed. The user has the option of either changing the ownerships of only one user or of the entire password domain. *Chuid* is not for the faint of heart.

Chuid has the following command line arguments:

-d *directory tree*

Use *directory tree* as the root for *chuid*. The user must specify a directory tree with each invocation of *chuid*.

-D Open a debugging file which will contain the names of every changed file and its new uid and gid. The name of the file will be of the form:

/tmp/chdbXXXXXX

where XXXXXX is a unique 6 digit number.

-o *old passwd file*

Use *old password file* as the name of the previous password file. Only the super-user is allowed to use this option. Default is /etc/passwd.old.

-n *new passwd file*

Use *new password file* as the name of the new password file. Only the super-user is allowed to use this option. Default is /etc/passwd.

-l *logfile*

Use *logfile* as the name of the logging file. The default is a file of the form:

/tmp/chlgXXXXXX

where XXXXXX is a unique 6 digit number.

-u *user olduid*

Change only those files with a uid of *olduid* to the uid for *user*

-r *restart file*

Restart using the previous log file *restart file*. Note that *logfile* and *restart file* should not be the same file. *Chuid* allows the user to restart from a previous log file in case of program or system failure. In this case, the user should give as a command line argument, the name of the restart file and all other arguments from the first execution. The command line information is **not** saved in the log file.

If *chuid* is run with privileged (i.e. *root*) permission, *chuid* prints an entry in the file "/etc/defunct.log" for each file it encounters which does not have an entry in the new password file. This entry is of the form

<name of file> <user id> <group id>

Chuid does not change the ownerships of these files nor does it create a defunct user entry in the password file.

SEE ALSO

chown(2), *chown(8)*, *chgrp(8)*

DIAGNOSTICS

Chuid uses extensively the facilities of *perror(3)* in reporting errors.

Unable to open logfile:

the file used as the logfile was inaccessible for writing. If it cannot open another file for logging (see next diagnostic), the error is fatal.

Using logfile as the log file:

a new file was created for logging because the first one was inaccessible.

Cannot open log file for reading:

chuid was unable to open the log file for reading. This is a fatal error.

Only super user allowed to use option:

the user was not the super-user. Occurs with **-n** and **-o**.

Error in loading password file:

chuid was unable to load either the new or the old password file for mapping. This is a fatal error.

Creating password entry:

chuid has found a file with an unknown owner. It has created a password entry in the *new password file* for the defunct owner.

User not in password file:

there is a user who is not in both the old and the new password file. This is done mostly as a consistency check.

RESTRICTIONS

When changing the entire password domain (i.e. not using the **-u** option), *chuid* should be run single user. The result of password file manipulation or file system manipulation with a concurrent execution of *chuid* is undefined.

BUGS

Argument handling is a real pain.

The defunct file is over written with each execution.

AUTHOR

Scott Schoenthal

NAME

clri — clear i-node

SYNOPSIS

/etc/clri filesystem i-number ...

DESCRIPTION

N.B.: *Clri* is obsoleted for normal file system repair work by *fsck(8)*.

Clri writes zeros on the i-nodes with the decimal *i-numbers* on the *filesystem*. After *clri*, any blocks in the affected file will show up as 'missing' in an *icheck(8)* of the *filesystem*.

Read and write permission is required on the specified file system device. The i-node becomes allocatable.

The primary purpose of this routine is to remove a file which for some reason appears in no directory. If it is used to zap an i-node which does appear in a directory, care should be taken to track down the entry and remove it. Otherwise, when the i-node is reallocated to some new file, the old entry will still point to that file. At that point removing the old entry will destroy the new file. The new entry will again point to an unallocated i-node, so the whole cycle is likely to be repeated again and again.

SEE ALSO

icheck(8)

BUGS

If the file is open, *clri* is likely to be ineffective.

NAME

`conn` - Valid network management program

SYNOPSIS

`conn` [`disable`] [`enable`] [`listen`] [`nolsten`] [`show`] [`shutdown period`] [`remove nodeName`]

DESCRIPTION

Conn is used to maintain the network connection for the Extended File System. *Disable* stops the periodic broadcast of "I am alive" message. *Enable* starts the periodic broadcast of "I am alive" message. *Listen* starts the connection manager listening to the net. *Nolisten* makes our system stop listening to the net. *Show* displays a list of known nodes and their status. *Shutdown period* lets node keep any impending I/O active for the specified period, allowing a graceful shutdown. *Remove nodename* removes those nodes who have either crashed or are shutdown.

The connection manager maintains the incore host table. If a node does not broadcast "I am alive" message for 90 seconds, it is marked inactive. If a node is inactive for 90 seconds, it is considered to be down.

A node should check that no other active node has the same name as itself before announcing its presence to everyone on the net. For this reason a node should listen to the net for at least 30 seconds before broadcasting the first "I am alive" message.

FILES

`conn.h`

DIAGNOSTICS

`/net` can not be opened

SEE ALSO

`chkhosts(8V)`

BUGS

None

NAME

crash — what happens when the system crashes

DESCRIPTION

This section explains what happens when the system crashes and how you can analyze crash dumps.

When the system crashes voluntarily it prints a message of the form

panic: why i gave up the ghost

on the console, takes a dump on a mass storage peripheral, and then invokes an automatic reboot procedure as described in *reboot(8)*. (If auto-reboot is disabled on the front panel of the machine the system will simply halt at this point.) Unless some unexpected inconsistency is encountered in the state of the file systems due to hardware or software failure the system will then resume multi-user operations.

The system has a large number of internal consistency checks; if one of these fails, then it will panic with a very short message indicating which one failed.

The most common cause of system failures is hardware failure, which can reflect itself in different ways. Here are the messages which you are likely to encounter, with some hints as to causes. Left unstated in all cases is the possibility that hardware or software error produced the message in some unexpected way.

IO err in push**hard IO err in swap**

The system encountered an error trying to write to the paging device or an error in reading critical information from a disk drive. You should fix your disk if it is broken or unreliable.

timeout table overflow

This really shouldn't be a panic, but until we fix up the data structure involved, running out of entries causes a crash. If this happens, you should make the timeout table bigger.

KSP not valid**SBI fault****CHM? in kernel**

These indicate either a serious bug in the system or, more often, a glitch or failing hardware. If SBI faults recur, check out the hardware or call field service. If the other faults recur, there is likely a bug somewhere in the system, although these can be caused by a flakey processor. Run processor microdiagnostics.

machine check %x:

description

machine dependent machine-check information

We should describe machine checks, and will someday. For now, ask someone who knows (like your friendly field service people).

trap type %d, code=%d, pc=%x

A unexpected trap has occurred within the system; the trap types are:

0	reserved addressing fault
1	privileged instruction fault
2	reserved operand fault
3	bpt instruction fault
4	xfc instruction fault
5	system call trap

6	arithmetic trap
7	ast delivery trap
8	segmentation fault
9	protection fault
10	trace trap
11	compatibility mode fault
12	page fault
13	page table fault

The favorite trap types in system crashes are trap types 8 and 9, indicating a wild reference. The code is the referenced address, and the pc at the time of the fault is printed. These problems tend to be easy to track down if they are kernel bugs since the processor stops cold, but random flakiness seems to cause this sometimes.

init died

The system initialization process has exited. This is bad news, as no new users will then be able to log in. Rebooting is the only fix, so the system just does it right away.

That completes the list of panic types you are likely to see.

When the system crashes it writes (or at least attempts to write) an image of memory into the back end of the primary swap area. After the system is rebooted, the program *savecore*(8) runs and preserves a copy of this core image and the current system in a specified directory for later perusal. See *savecore*(8) for details.

To analyze a dump you should begin by running *adb*(1) with the `-k` flag on the core dump. Normally the command `*(intstack-4)$c` will provide a stack trace from the point of the crash and this will provide a clue as to what went wrong. A more complete discussion of system debugging is impossible here. See, however, "Using ADB to Debug the UNIX Kernel".

SEE ALSO

adb(1), *analyze*(8), *reboot*(8)

VAX 11/780 System Maintenance Guide for more information about machine checks.

Using ADB to Debug the UNIX Kernel

NAME

`cron` - clock daemon

SYNOPSIS

`/etc/cron`

DESCRIPTION

Cron executes commands at specified dates and times according to the instructions in the files `/usr/lib/crontab` and `/usr/lib/crontab.local`. Any site dependent commands should be added to this latter file.

Since *cron* never exits, it should only be executed once. This is best done by running *cron* from the initialization process through the file `/etc/rc`; see *init*(8).

`Crontab` and `crontab.local` consist of lines of six fields each. The fields are separated by spaces or tabs. The first five are integer patterns to specify the minute (0-59), hour (0-23), day of the month (1-31), month of the year (1-12), and day of the week (1-7 with 1=Monday). Each of these patterns may contain a number in the range above; two numbers separated by a minus meaning a range inclusive; a list of numbers separated by commas meaning any of the numbers; or an asterisk meaning all legal values. The sixth field is a string that is executed by the Shell at the specified times. A percent character in this field is translated to a new-line character. Only the first line (up to a % or end of line) of the command field is executed by the Shell. The other lines are made available to the command as standard input.

`Crontab` and `crontab.local` are examined by *cron* every minute.

FILES

`/usr/lib/crontab`

`/usr/lib/crontab.local`

NAME

`dcheck` — file system directory consistency check

SYNOPSIS

`/etc/dcheck [-i numbers] [filesystem]`

DESCRIPTION

N.B.: *Dcheck* is obsoleted for normal consistency checking by *fsck*(8).

Dcheck reads the directories in a file system and compares the link-count in each i-node with the number of directory entries by which it is referenced. If the file system is not specified, a set of default file systems is checked.

The `-i` flag is followed by a list of i-numbers; when one of those i-numbers turns up in a directory, the number, the i-number of the directory, and the name of the entry are reported.

The program is fastest if the raw version of the special file is used, since the i-list is read in large chunks.

FILES

Default file systems vary with installation.

SEE ALSO

fsck(8), *icheck*(8), *fs*(5), *clri*(8), *ncheck*(8)

DIAGNOSTICS

When a file turns up for which the link-count and the number of directory entries disagree, the relevant facts are reported. Allocated files which have 0 link-count and no entries are also listed. The only dangerous situation occurs when there are more entries than links; if entries are removed, so the link-count drops to 0, the remaining entries point to thin air. They should be removed. When there are more links than entries, or there is an allocated file with neither links nor entries, some disk space may be lost but the situation will not degenerate.

BUGS

Since *dcheck* is inherently two-pass in nature, extraneous diagnostics may be produced if applied to active file systems.

Dcheck is obsoleted by *fsck* and remains for historical reasons.

NAME

`dmesg` — collect system diagnostic messages to form error log

SYNOPSIS

`/etc/dmesg [-]`

DESCRIPTION

Dmesg looks in a system buffer for recently printed diagnostic messages and prints them on the standard output. The messages are those printed by the system when device (hardware) errors occur and (occasionally) when system tables overflow non-fatally. If the `-` flag is given, then *dmesg* computes (incrementally) the new messages since the last time it was run and places these on the standard output. This is typically used with *cron*(8) to produce the error log */usr/adm/messages* by running the command

```
/etc/dmesg - >> /usr/adm/messages
```

every 10 minutes.

FILES

<i>/usr/adm/messages</i>	error log (conventional location)
<i>/usr/adm/msgbuf</i>	scratch file for memory of <code>-</code> option

BUGS

The system error message buffer is of small finite size. As *dmesg* is run only every few minutes, not all error messages are guaranteed to be logged. This can be construed as a blessing rather than a curse.

Error diagnostics generated immediately before a system crash will never get logged.

NAME

`dump` — incremental file system dump

SYNOPSIS

`/etc/dump [key [argument ...] filesystem]`

DESCRIPTION

Dump copies to magnetic tape all files changed after a certain date in the *filesystem*. The *key* specifies the date and other options about the dump. *Key* consists of characters from the set 0123456789fusdWn.

- 0-9 This number is the 'dump level'. All files modified since the last date stored in the file */etc/dumpdates* for the same filesystem at lesser levels will be dumped. If no date is determined by the level, the beginning of time is assumed; thus the option 0 causes the entire filesystem to be dumped.
- f Place the dump on the next *argument* file instead of the tape. If the name of the file is "--", *dump* writes to standard output.
- u If the dump completes successfully, write the date of the beginning of the dump on file */etc/dumpdates*. This file records a separate date for each filesystem and each dump level. The format of */etc/dumpdates* is readable by people, consisting of one free format record per line: filesystem name, increment level and *ctime(3)* format dump date. */etc/dumpdates* may be edited to change any of the fields, if necessary.
- s The size of the dump tape is specified in feet. The number of feet is taken from the next *argument*. When the specified size is reached, *dump* will wait for reels to be changed. The default tape size is 2300 feet.
- d The density of the tape, expressed in BPI, is taken from the next *argument*. This is used in calculating the amount of tape used per reel. The default is 1600.
- W *Dump* tells the operator what file systems need to be dumped. This information is gleaned from the files */etc/dumpdates* and */etc/fstab*. The W option causes *dump* to print out, for each file system in */etc/dumpdates* the most recent dump date and level, and highlights those file systems that should be dumped. If the W option is set, all other options are ignored, and *dump* exits immediately.
- w Is like W, but prints only those filesystems which need to be dumped.
- n Whenever *dump* requires operator attention, notify by means similar to a *wall(1)* all of the operators in the group "operator".

If no arguments are given, the *key* is assumed to be 9u and a default file system is dumped to the default tape.

Dump requires operator intervention on these conditions: end of tape, end of dump, tape write error, tape open error or disk read error (if there are more than a threshold of 32). In addition to alerting all operators implied by the n key, *dump* interacts with the operator on *dump's* control terminal at times when *dump* can no longer proceed, or if something is grossly wrong. All questions *dump* poses must be answered by typing "yes" or "no", appropriately.

Since making a dump involves a lot of time and effort for full dumps, *dump* checkpoints itself at the start of each tape volume. If writing that volume fails for some reason, *dump* will, with operator permission, restart itself from the checkpoint after the old tape has been rewound and removed, and a new tape has been mounted.

Dump tells the operator what is going on at periodic intervals, including usually low estimates of the number of blocks to write, the number of tapes it will take, the time to completion, and the time to the tape change. The output is verbose, so that others know that the terminal controlling *dump* is busy, and will be for some time.

Now a short suggestion on how to perform dumps. Start with a full level 0 dump
 dump 0un

Next, dumps of active file systems are taken on a daily basis, using a modified Tower of Hanoi algorithm, with this sequence of dump levels:

3 2 5 4 7 6 9 8 9 9 ...

For the daily dumps, a set of 10 tapes per dumped file system is used on a cyclical basis. Each week, a level 1 dump is taken, and the daily Hanoi sequence repeats with 3. For weekly dumps, a set of 5 tapes per dumped file system is used, also on a cyclical basis. Each month, a level 0 dump is taken on a set of fresh tapes that is saved forever.

FILES

/dev/rrplg	default filesystem to dump from
/dev/rmt8	default tape unit to dump to
/etc/ddate	old format dump date record (obsolete after <code>-J</code> option)
/etc/dumpdates	new format dump date record
/etc/fstab	dump table: file systems and frequency
/etc/group	to find group <i>operator</i>

SEE ALSO

restor (8), dump(5), fstab(5)

DIAGNOSTICS

Many, and verbose.

BUGS

Sizes are based on 1600 BPI blocked tape; the raw magtape device has to be used to approach these densities. Fewer than 32 read errors on the filesystem are ignored. Each reel requires a new process, so parent processes for reels already written just hang around until the entire tape is written.

It would be nice if *dump* knew about the dump sequence, kept track of the tapes scribbled on, told the operator which tape to mount when, and provided more assistance for the operator running *restor*.

NAME

efsioctl - EFS Superuser Access Control

SYNOPSIS

efsioctl {suid <uid> |sufew |sumost}

DESCRIPTION

When a superuser accesses a remote system's files through the Extended File System (EFS), the access rights are defined on the remote system. These rights are settable on each system by using the program *efsioctl*. Each system can determine how all other system's superuser (uid of 0) is interpreted during file access. The superuser's uid is either used as is (0) or is changed to a locally defined value.

efsioctl sumost

This command specifies that all superusers on other machines are to be treated as the local superuser during access control checks.

efsioctl sufew

This command specifies that when a superuser on another machine accesses the local machine the uid will be changed before access control checks. The default uid for this transformation is 11.

efsioctl suuid <uid>

This command specifies the uid used for the transformation when **sufew** is in effect.

FILES

/net - EFS network interface device file

DIAGNOSTICS

efsioctl - cannot open /net

efsioctl - ioctl failed

NAME

ether - 3Com 3C400 driver control program

SYNOPSIS

ether <verb> [<object> [<arguments>]]

DESCRIPTION

This program allows a user to control and observe the actions of the 3Com Ethernet driver and controller. Each command consists of an action verb and an object of the driver. Some commands also take parameter arguments after the object to qualify the action performed. Keywords in all commands can be abbreviated and are not case-sensitive.

The verbs supported by *ether* are **clear**, **disable**, **enable**, **help**, **set**, and **show**. The complete command syntax is:

```
clear {counters | distribution | multicast | statistics}
      {disable | enable}
      {debug | multicast <address> |
       mode {distribution | normal | promiscuous | trailer}}
```

help

```
set address <address>
```

```
show {address | counters | debug |
      distribution | mode | multicast | version}
```

COUNTERS

The counters maintained by the driver show how received and transmitted packets have been processed. Total counts of interrupts, packets successfully processed, and packets discarded are shown. The verbs **show** and **clear** are permitted on the counters object.

DISTRIBUTION

When the **distribution** mode of the driver is **enabled** the driver collects packet size information. This is maintained based on Ethertype and packet size. Separate information is collected for received and transmitted packets. The verbs **show** and **clear** are permitted on the distribution object. In addition the **enable/disable** verb on the driver mode (see mode below) is used to control whether this information is collected.

STATISTICS

The statistics object is the union of the counter and distribution objects. The only verb allowed on this object is **clear**. It is used to synchronize the data contained in the counters and packet distribution data bases.

ADDRESS

The physical Ethernet object of the controller can be displayed and set using the verbs **show** and **set address** respectively. For the set verb, the address is specified as six (6) hexadecimal 2 digit values separated by hypens (e.g., 02-60-8c-12-34-56). When the address is set, the controller is reset and the address will propagate through all software using the controller.

MULTICAST

The driver supports multicast receive filtering of an arbitrary set of multicast addresses. The verbs supported by the multicast object are **enable**, **disable**, **clear**, and **show**. The enable and disable verbs operate on single entries in the multicast filter. A multicast address is specified as six (6) hexadecimal 2 digit values separated by hypens (e.g., FF-01-02-65-43-21). The clear verb disables all entries in the multicast filter instead of having to perform individual disable operations for each.

MODE

The mode object controls actions of the driver during packet processing. The three verbs that the mode object support are: **enable**, **disable**, and **show**. The enable and disable verbs take an extra parameter after the mode object to qualify what part of the driver mode should be affected. The **distribution** qualifier controls whether the driver collects packet size distribution statistics. The **normal** qualifier controls whether the driver proceses normal packets or just discards them. The **trailer** qualifier controls whether the LAN trailer protocol is used on IP packets. The **promiscuous** qualifier controls whether all packets on the net are received. The only packets that are passed to higher protocol levels for processing are those that pass address match tests. In promiscuous mode, all packets received contribute to counters and distribution data.

VERSION

The version object only supports the **show** verb. The information shown consists of the versions of the 3Com driver in the kernel and the version of the ether program itself.

DEBUG

The debug object controls the debugging flags internal to the 3Com driver. Currently the verbs supported are **enable**, **disable**, and **show**.

NAME

fsck — file system consistency check and interactive repair

SYNOPSIS

```
/etc/fsck -p [ filesystem ... ]
/etc/fsck [ -b block# ] [ -y ] [ -n ] [ filesystem ] ...
```

DESCRIPTION

The first form of *fsck* preens a standard set of filesystems or the specified file systems. It is normally used in the script */etc/rc* during automatic reboot. In this case *fsck* reads the table */etc/fstab* to determine which file systems to check. It uses the information there to inspect groups of disks in parallel taking maximum advantage of i/o overlap to check the file systems as quickly as possible. Normally, the root file system will be checked on pass 1, other "root" ("a" partition) file systems on pass 2, other small file systems on separate passes (e.g. the "d" file systems on pass 3 and the "e" file systems on pass 4), and finally the large user file systems on the last pass, e.g. pass 5. A pass number of 0 in *fstab* causes a disk to not be checked; similarly partitions which are not shown as to be mounted "rw" or "ro" are not checked.

The system takes care that only a restricted class of innocuous inconsistencies can happen unless hardware or software failures intervene. These are limited to the following:

- Unreferenced inodes
- Link counts in inodes too large
- Missing blocks in the free list
- Blocks in the free list also in files
- Counts in the super-block wrong

These are the only inconsistencies which *fsck* with the *-p* option will correct; if it encounters other inconsistencies, it exits with an abnormal return status and an automatic reboot will then fail. For each corrected inconsistency one or more lines will be printed identifying the file system on which the correction will take place, and the nature of the correction. After successfully correcting a file system, *fsck* will print the number of files on that file system and the number of used and free blocks.

Without the *-p* option, *fsck* audits and interactively repairs inconsistent conditions for file systems. If the file system is inconsistent the operator is prompted for concurrence before each correction is attempted. It should be noted that a number of the corrective actions which are not fixable under the *-p* option will result in some loss of data. The amount and severity of data lost may be determined from the diagnostic output. The default action for each consistency correction is to wait for the operator to respond *yes* or *no*. If the operator does not have write permission *fsck* will default to a *-n* action.

Fsck has more consistency checks than its predecessors *check*, *dcheck*, *fcheck*, and *ichck* combined.

The following flags are interpreted by *fsck*.

- b** Use the block specified immediately after the flag as the super block for the file system. Block 32 is always an alternate super block.
- y** Assume a yes response to all questions asked by *fsck*; this should be used with great caution as this is a free license to continue after essentially unlimited trouble has been encountered.
- n** Assume a no response to all questions asked by *fsck*; do not open the file system for writing.

If no filesystems are given to *fsck* then a default list of file systems is read from the file */etc/fstab*.

Inconsistencies checked are as follows:

1. Blocks claimed by more than one inode or the free list.
2. Blocks claimed by an inode or the free list outside the range of the file system.
3. Incorrect link counts.
4. Size checks:
 - Directory size not of proper format.
5. Bad inode format.
6. Blocks not accounted for anywhere.
7. Directory checks:
 - File pointing to unallocated inode.
 - Inode number out of range.
8. Super Block checks:
 - More blocks for inodes than there are in the file system.
9. Bad free block list format.
10. Total free block and/or free inode count incorrect.

Orphaned files and directories (allocated but unreferenced) are, with the operator's concurrence, reconnected by placing them in the *lost+found* directory. The name assigned is the inode number. The only restriction is that the directory *lost+found* must preexist in the root of the filesystem being checked and must have empty slots in which entries can be made. This is accomplished by making *lost+found*, copying a number of files to the directory, and then removing them (before *fsck* is executed).

Checking the raw device is almost always faster.

FILES

/etc/fstab contains default list of file systems to check.

DIAGNOSTICS

The diagnostics produced by *fsck* are intended to be self-explanatory.

SEE ALSO

fstab(5), *fs*(5), *newfs*(8), *mkfs*(8), *crash*(8V), *reboot*(8)

BUGS

Inode numbers for *.* and *..* in each directory should be checked for validity.

There should be some way to start a *fsck -p* at pass *n*.

NAME

getfs - print list of file systems and directories

SYNOPSIS

/etc/getfs

DESCRIPTION

Getfs reads the file system data base *fstab*(5) and writes on its standard output the device name (minus its starting /dev/) and mount directory of each file system that it finds there. The chief use of this program is in maintenance scripts such as *backup*(8V).

FILES

/etc/fstab - file system data base

DIAGNOSTICS

Returns non-zero if it cannot read *fstab*.

NAME

getty - set terminal mode

SYNOPSIS

/etc/getty [char]

DESCRIPTION

Getty is invoked by *init*(8) immediately after a terminal is opened. It reads the user's login name and calls *login*(1) with the name as argument. While reading the name *getty* attempts to adapt the system to the speed and type of terminal being used.

The *char* argument, which *init* reads from the *ttys* file, */etc/ttys*, specifies the name of a circular list, internal to *getty*, of data rates and initial characteristics of the user's terminal. Normally, *getty* sets the speed of the interface to 9600 baud, specifies that raw mode is to be used (break on every character), that echo is to be suppressed, and either parity allowed. If a *loginbanner* file exists, *getty* reads it and prints a login banner (see *loginbanner*(5)), otherwise *getty* types the 'login:' message. Special banner requests such as clearing the screen, inverse video, and the like are handled by obtaining the terminal capabilities from the */etc/termcap* file. The user's name is read and echoed, a character at a time. If a null character is received, it is assumed to be the result of the user pushing the 'break' ('interrupt') key, and *getty* sets the terminal to another speed and re-types the 'login:' message. Further breaks cause *getty* to cycle through its circular list of speeds, attempting to match the data rate of the user's terminal.

The user's name is terminated by a new-line or carriage-return character. The latter results in the system being set to treat carriage returns appropriately (see *ioctl*(2)).

The user's name is scanned to see if it contains any lower-case alphabetic characters; if not, and if the name is nonempty, the system is told to map any future upper-case characters into the corresponding lower-case characters.

The following arguments from the *ttys* file are recognized by *getty* as valid speeds.

- 0 Cycles through 300-1200-150-110 baud.
- 110 baud.
- 1 150 baud.
- 2 Intended for an on-line 9600 baud terminal.
- 3 Starts at 1200 baud, cycles to 300 and back.
- 4 Intended for 300 baud DECwriter console
- 5 Same as '3' but starts at 300 baud.
- 6 2400 baud.
- 7 4800 baud.
- 8 Starts at 9600 baud, cycles to 300 and back.
- 9 Same as '8' but starts at 300 baud.
- p Cycles through 9600-300-1200 baud.
- q Same as 'p' but starts at 300 baud.
- r Same as 'p' but starts at 1200 baud.

SEE ALSO

init(8), *login*(1), *ioctl*(2), *ttys*(5), *loginbanner*(5), *termcap*(5)

NAME

halt - stop the processor

SYNOPSIS

/etc/halt

DESCRIPTION

Halt syncs the disks and then stops the processor. The machine does not reboot, even if the auto-reboot switch is set on the console.

SEE ALSO

reboot(8), shutdown(8)

BUGS

It is very difficult to halt a VAX, as the machine wants to then reboot itself. A rather tight loop suffices.

NAME

`icheck` — file system storage consistency check

SYNOPSIS

`/etc/icheck [-s] [-b numbers] [filesystem]`

DESCRIPTION

N.B.: *Icheck* is obsoleted for normal consistency checking by *fsck*(8).

Icheck examines a file system, builds a bit map of used blocks, and compares this bit map against the free list maintained on the file system. If the file system is not specified, a set of default file systems is checked. The normal output of *icheck* includes a report of

The total number of files and the numbers of regular, directory, block special and character special files.

The total number of blocks in use and the numbers of single-, double-, and triple-indirect blocks and directory blocks.

The number of free blocks.

The number of blocks missing; i.e. not in any file nor in the free list.

The `-s` option causes *icheck* to ignore the actual free list and reconstruct a new one by rewriting the super-block of the file system. The file system should be dismounted while this is done; if this is not possible (for example if the root file system has to be salvaged) care should be taken that the system is quiescent and that it is rebooted immediately afterwards so that the old, bad in-core copy of the super-block will not continue to be used. Notice also that the words in the super-block which indicate the size of the free list and of the i-list are believed. If the super-block has been curdled these words will have to be patched. The `-s` option causes the normal output reports to be suppressed.

Following the `-b` option is a list of block numbers; whenever any of the named blocks turns up in a file, a diagnostic is produced.

Icheck is faster if the raw version of the special file is used, since it reads the i-list many blocks at a time.

FILES

Default file systems vary with installation.

SEE ALSO

`fsck`(8), `dcheck`(8), `ncheck`(8), `fs`(5), `clri`(8)

DIAGNOSTICS

For duplicate blocks and bad blocks (which lie outside the file system) *icheck* announces the difficulty, the i-number, and the kind of block involved. If a read error is encountered, the block number of the bad block is printed and *icheck* considers it to contain 0. 'Bad freeblock' means that a block number outside the available space was encountered in the free list. '*n* dups in free' means that *n* blocks were found in the free list which duplicate blocks either in some file or in the earlier part of the free list.

BUGS

Since *icheck* is inherently two-pass in nature, extraneous diagnostics may be produced if applied to active file systems.

It believes even preposterous super-blocks and consequently can get core images.

The system should be fixed so that the reboot after fixing the root file system is not necessary.

NAME

`init` — process control initialization

SYNOPSIS

`/etc/init`

DESCRIPTION

Init is invoked inside UNIX as the last step in the boot procedure. It normally then runs the automatic reboot sequence as described in *reboot(8)*, and if this succeeds, begins multi-user operation. If the reboot fails, it commences single user operation by giving the super-user a shell on the console. It is possible to pass parameters from the boot program to *init* so that single user operation is commenced immediately. When such single user operation is terminated by killing the single-user shell (i.e. by hitting ^D), *init* runs *etc/rc* without the reboot parameter. This command file performs housekeeping operations such as removing temporary files, mounting file systems, and starting daemons.

In multi-user operation, *init*'s role is to create a process for each terminal port on which a user may log in. To begin such operations, it reads the file *etc/ttys* and forks several times to create a process for each terminal specified in the file. Each of these processes opens the appropriate terminal for reading and writing. These channels thus receive file descriptors 0, 1 and 2, the standard input and output and the diagnostic output. Opening the terminal will usually involve a delay, since the *open* is not completed until someone is dialed up and carrier established on the channel. If a terminal exists but an error occurs when trying to open the terminal *init* complains by writing a message to the system console; the message is repeated every 10 minutes for each such terminal until the terminal is shut off in *etc/ttys* and *init* notified (by a hangup, as described below), or the terminal becomes accessible (*init* checks again every minute). After an open succeeds, *etc/getty* is called with argument as specified by the second character of the *ttys* file line. *Getty* reads the user's name and invokes *login* to log in the user and execute the Shell.

Ultimately the Shell will terminate because of an end-of-file either typed explicitly or generated as a result of hanging up. The main path of *init*, which has been waiting for such an event, wakes up and removes the appropriate entry from the file *utmp*, which records current users, and makes an entry in *usr/adm/wtmp*, which maintains a history of logins and logouts. The *wtmp* entry is made only if a user logged in successfully on the line. Then the appropriate terminal is reopened and *getty* is reinvoked.

Init catches the *hangup* signal (signal SIGHUP) and interprets it to mean that the file *etc/ttys* should be read again. The Shell process on each line which used to be active in *ttys* but is no longer there is terminated; a new process is created for each added line; lines unchanged in the file are undisturbed. Thus it is possible to drop or add phone lines without rebooting the system by changing the *ttys* file and sending a *hangup* signal to the *init* process: use 'kill -HUP 1.'

Init will terminate multi-user operations and resume single-user mode if sent a terminate (TERM) signal, i.e. "kill -TERM 1". If there are processes outstanding which are deadlocked (due to hardware or software failure), *init* will not wait for them all to die (which might take forever), but will time out after 30 seconds and print a warning message.

Init will cease creating new *getty*'s and allow the system to slowly die away, if it is sent a terminal stop (TSTP) signal, i.e. "kill -TSTP 1". A later hangup will resume full multi-user operations, or a terminate will initiate a single user shell. This hook is used by *reboot(8)* and *halt(8)*.

Init's role is so critical that if it dies, the system will reboot itself automatically. If, at bootstrap time, the *init* process cannot be located, the system will loop in user mode at location 0x13.

DIAGNOSTICS

init: *ty*: cannot open. A terminal which is turned on in the *rc* file cannot be opened, likely because the requisite lines are either not configured into the system or the associated device

was not attached during boot-time system configuration.

WARNING: Something is hung (wont die); ps axl advised. A process is hung and could not be killed when the system was shutting down. This is usually caused by a process which is stuck in a device driver due to a persistent device error condition.

FILES

/dev/console, /dev/tty*, /etc/utmp, /usr/adm/wtmp, /etc/ttys, /etc/rc

SEE ALSO

login(1), kill(1), sh(1), ttys(5), crash(8V), getty(8), rc(8), reboot(8), halt(8), shutdown(8)

NAME

lpd — line printer daemon

SYNOPSIS

```
/usr/lib/lpd [ -l ] [ -L logfile ] [ port # ]
```

DESCRIPTION

Lpd is the line printer daemon (spool area handler) and is normally invoked at boot time from the *rc(8)* file. It makes a single pass through the *printcap(5)* file to find out about the existing printers and prints any files left after a crash. It then uses the system calls *listen(2)* and *accept(2)* to receive requests to print files in the queue, transfer files to the spooling area, display the queue, or remove jobs from the queue. In each case, it forks a child to handle the request so the parent can continue to listen for more requests. The Internet port number used to rendezvous with other processes is normally obtained with *getservbyname(3)* but can be changed with the *port#* argument. The *-L* option changes the file used for writing error conditions from the system console to *logfile*. The *-l* flag causes *lpd* to log valid requests received from the network. This can be useful for debugging purposes.

Access control is provided by two means. First, All requests must come from one of the machines listed in the file */etc/hosts.equiv*. Second, if the "rs" capability is specified in the *printcap* entry for the printer being accessed, *lpr* requests will only be honored for those users with accounts on the machine with the printer.

The file *lock* in each spool directory is used to prevent multiple daemons from becoming active simultaneously, and to store information about the daemon process for *lpr(1)*, *lpq(1)*, and *lprm(1)*. After the daemon has successfully set the lock, it scans the directory for files beginning with *cf*. Lines in each *cf* file specify files to be printed or non-printing actions to be performed. Each such line begins with a key character to specify what to do with the remainder of the line.

- J Job Name. String to be used for the job name on the burst page.
- C Classification. String to be used for the classification line on the burst page.
- L Literal. The line contains identification info from the password file and causes the banner page to be printed.
- T Title. String to be used as the title for *pr(1)*.
- H Host Name. Name of the machine where *lpr* was invoked.
- P Person. Login name of the person who invoked *lpr*. This is used to verify ownership by *lprm*.
- M Send mail to the specified user when the current print job completes.
- f Formatted File. Name of a file to print which is already formatted.
- l Like "f" but passes control characters and does not make page breaks.
- p Name of a file to print using *pr(1)* as a filter.
- t Troff File. The file contains *troff(1)* output (cat phototypesetter commands).
- d DVI File. The file contains *Tex(1)* output (DVI format from Stanford).
- g Graph File. The file contains data produced by *plot(3X)*.
- c Cifplot File. The file contains data produced by *cifplot*.
- v The file contains a raster image.
- r The file contains text data with FORTRAN carriage control characters.
- l Troff Font R. Name of the font file to use instead of the default.

- 2 Troff Font I. Name of the font file to use instead of the default.
- 3 Troff Font B. Name of the font file to use instead of the default.
- 4 Troff Font S. Name of the font file to use instead of the default.
- W Width. Changes the page width (in characters) used by *pr*(1) and the text filters.
- I Indent. The number of characters to indent the output by (in ascii).
- U Unlink. Name of file to remove upon completion of printing.
- N File name. The name of the file which is being printed, or a blank for the standard input (when *lpr* is invoked in a pipeline).

If a file can not be opened, a message will be placed in the log file (normally the console). *Lpd* will try up to 20 times to reopen a file it expects to be there, after which it will skip the file to be printed.

Lpd uses *flock*(2) to provide exclusive access to the lock file and to prevent multiple daemons from becoming active simultaneously. If the daemon should be killed or die unexpectedly, the lock file need not be removed. The lock file is kept in a readable ASCII form and contains two lines. The first is the process id of the daemon and the second is the control file name of the current job being printed. The second line is updated to reflect the current status of *lpd* for the programs *lpq*(1) and *lprm*(1).

FILES

<i>/etc/printcap</i>	printer description file
<i>/usr/spool/*</i>	spool directories
<i>/dev/lp*</i>	line printer devices
<i>/dev/printer</i>	socket for local requests
<i>/etc/hosts.equiv</i>	lists machine names allowed printer access

SEE ALSO

lpc(8), *pac*(1), *lpr*(1), *lpq*(1), *lprm*(1), *printcap*(5)
4.2BSD Line Printer Spooler Manual

NAME

`makekey` — generate encryption key

SYNOPSIS

`/usr/lib/makekey`

DESCRIPTION

Makekey improves the usefulness of encryption schemes depending on a key by increasing the amount of time required to search the key space. It reads 10 bytes from its standard input, and writes 13 bytes on its standard output. The output depends on the input in a way intended to be difficult to compute (that is, to require a substantial fraction of a second).

The first eight input bytes (the *input key*) can be arbitrary ASCII characters. The last two (the *salt*) are best chosen from the set of digits, upper- and lower-case letters, and '.' and '/'. The salt characters are repeated as the first two characters of the output. The remaining 11 output characters are chosen from the same set as the salt and constitute the *output key*.

The transformation performed is essentially the following: the salt is used to select one of 4096 cryptographic machines all based on the National Bureau of Standards DES algorithm, but modified in 4096 different ways. Using the input key as key, a constant string is fed into the machine and recirculated a number of times. The 64 bits that come out are distributed into the 66 useful key bits in the result.

Makekey is intended for programs that perform encryption (for instance, *ed* and *crypt(1)*). Usually *makekey*'s input and output will be pipes.

SEE ALSO

crypt(1), *ed(1)*

NAME

mkfs — construct a file system

SYNOPSIS

/etc/mkfs special size [nsect] [ntrack] [blksize] [fragsize] [ncpng] [minfree] [rps]

DESCRIPTION

N.B.: file system are normally created with the *newfs(8)* command.

Mkfs constructs a file system by writing on the special file *special*. The numeric size specifies the number of sectors in the file system. *Mkfs* builds a file system with a root directory and a *lost+found* directory. (see *fsck(8)*) The number of i-nodes is calculated as a function of the file system size. No boot program is initialized by *mkfs* (see *newfs(8)*.)

The optional arguments allow fine tune control over the parameters of the file system. **Nsect** specify the number of sectors per track on the disk. **Ntrack** specify the number of tracks per cylinder on the disk. **Blksize** gives the primary block size for files on the file system. It must be a power of two, currently selected from 4096 or 8192. **Fragsize** gives the fragment size for files on the file system. The **fragsize** represents the smallest amount of disk space that will be allocated to a file. It must be a power of two currently selected from the range 512 to 8192. **Ncpng** specifies the number of disk cylinders per cylinder group. This number must be in the range 1 to 32. **Minfree** specifies the minimum percentage of free disk space allowed. Once the file system capacity reaches this threshold, only the super-user is allowed to allocate disk blocks. The default value is 10%. If a disk does not revolve at 60 revolutions per second, the **rps** parameter may be specified. Users with special demands for their file systems are referred to the paper cited below for a discussion of the tradeoffs in using different configurations.

SEE ALSO

fs(5), *dir(5)*, *fsck(8)*, *newfs(8)*, *tunefs(8)*

McKusick, Joy, Leffler; "A Fast File System for Unix", Computer Systems Research Group, Dept of EECS, Berkeley, CA 94720; TR #7, September 1982.

BUGS

There should be some way to specify bad blocks.

NAME

mklost+found — make a lost+found directory for fsck

SYNOPSIS

/etc/mklost+found

DESCRIPTION

A directory *lost+found* is created in the current directory and a number of empty files are created therein and then removed so that there will be empty slots for *fsck(8)*. This command should not normally be needed since *mkfs(8)* automatically creates the *lost+found* directory when a new file system is created.

SEE ALSO

fsck(8), *mkfs(8)*

NAME

mknod — build special file

SYNOPSIS

/etc/mknod name [c] [b] major minor

DESCRIPTION

Mknod makes a special file. The first argument is the *name* of the entry. The second is **b** if the special file is block-type (disks, tape) or **c** if it is character-type (other devices). The last two arguments are numbers specifying the *major* device type and the *minor* device (e.g. unit, drive, or line number).

The assignment of major device numbers is specific to each system. They have to be dug out of the system source file *conf.c*.

SEE ALSO

mknod(2)

NAME

mkproto — construct a prototype file system

SYNOPSIS

/etc/mkproto special proto

DESCRIPTION

Mkproto is used to bootstrap a new file system. First a new file system is created using *newfs(8)*. *Mkproto* is then used to copy files from the old file system into the new file system according to the directions found in the prototype file *proto*. The prototype file contains tokens separated by spaces or new lines. The first tokens comprise the specification for the root directory. File specifications consist of tokens giving the mode, the user-id, the group id, and the initial contents of the file. The syntax of the contents field depends on the mode.

The mode token for a file is a 6 character string. The first character specifies the type of the file. (The characters *-bcd* specify regular, block special, character special and directory files respectively.) The second character of the type is either *u* or *-* to specify set-user-id mode or not. The third is *g* or *-* for the set-group-id mode. The rest of the mode is a three digit octal number giving the owner, group, and other read, write, execute permissions, see *chmod(1)*.

Two decimal number tokens come after the mode; they specify the user and group ID's of the owner of the file.

If the file is a regular file, the next token is a pathname whence the contents and size are copied.

If the file is a block or character special file, two decimal number tokens follow which give the major and minor device numbers.

If the file is a directory, *mkproto* makes the entries *.* and *..* and then reads a list of names and (recursively) file specifications for the entries in the directory. The scan is terminated with the token *\$*.

A sample prototype specification follows:

```

d--777 3 1
usr   d--777 3 1
      sh   ---755 3 1 /bin/sh
      ken  d--755 6 1
      $
      b0   b--644 3 1 0 0
      c0   c--644 3 1 0 0
      $
$

```

SEE ALSO

fs(5), *dir(5)*, *fsck(8)*, *newfs(8)*

BUGS

There should be some way to specify links.

There should be some way to specify bad blocks.

Mkproto can only be run on virgin file systems. It should be possible to copy files into existent file systems.

NAME

mkusr - make a new user

SYNOPSIS

mkusr

DESCRIPTION

Mkusr creates a new user on the system by prompting for information about the new user and then updating the */etc/passwd* and */etc/group* files. *Mkusr* then calls the script */etc/mkscaldusr* to build the new user's home directory and set up default files needed by the Scald software, the C shell and the Bourne shell.

FILES

<i>/etc/mknewpentry</i>	creates password and group entries
<i>/etc/mkscaldusr</i>	script to set up user directory
<i>/u0/user/*</i>	directory for default startup files

DIAGNOSTICS

Error messages are meant to be self explanatory.

NAME

mount, umount — mount and dismount file system

SYNOPSIS

/etc/mount [special name [-r]]

/etc/mount -a

/etc/umount special

/etc/umount -a

DESCRIPTION

Mount announces to the system that a removable file system is present on the device *special*. The file *name* must exist already; it must be a directory (unless the root of the mounted file system is not a directory). It becomes the name of the newly mounted root. The optional argument -r indicates that the file system is to be mounted read-only.

Umount announces to the system that the removable file system previously mounted on device *special* is to be removed.

If the -a option is present for either *mount* or *umount*, all of the file systems described in */etc/fstab* are attempted to be mounted or unmounted. In this case, *special* and *name* are taken from */etc/fstab*. The *special* file name from */etc/fstab* is the block special name.

These commands maintain a table of mounted devices in */etc/mtab*. If invoked without an argument, *mount* prints the table.

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.

FILES

<i>/etc/mtab</i>	mount table
<i>/etc/fstab</i>	file system table

SEE ALSO

mount(2), mtab(5), fstab(5)

BUGS

Mounting file systems full of garbage will crash the system.

Mounting a root directory on a non-directory makes some apparently good pathnames invalid.

NAME

`ncheck` — generate names from i-numbers

SYNOPSIS

`/etc/ncheck` [`-i` numbers] [`-a`] [`-s`] [filesystem]

DESCRIPTION

N.B.: For most normal file system maintenance, the function of *ncheck* is subsumed by *fsck*(8).

Ncheck with no argument generates a pathname vs. i-number list of all files on a set of default file systems. Names of directory files are followed by '/.'. The `-i` option reduces the report to only those files whose i-numbers follow. The `-a` option allows printing of the names '.' and '..', which are ordinarily suppressed. The `-s` option reduces the report to special files and files with set-user-ID mode; it is intended to discover concealed violations of security policy.

A file system may be specified.

The report is in no useful order, and probably should be sorted.

SEE ALSO

`sort`(1), `dcheck`(8), `fsck`(8), `icheck`(8)

DIAGNOSTICS

When the filesystem structure is improper, '??' denotes the 'parent' of a parentless file and a pathname beginning with '...' denotes a loop.

NAME

`newfs` — construct a new file system

SYNOPSIS

`/etc/newfs [-v] [-n] [mkfs-options] special disk-type`

DESCRIPTION

Newfs is a “friendly” front-end to the *mkfs*(8) program. *Newfs* will look up the type of disk a file system is being created on in the disk description file *etcdisktab*, calculate the appropriate parameters to use in calling *mkfs*, then build the file system by forking *mkfs* and, if the file system is a root partition, install the necessary bootstrap programs in the initial 8 sectors of the device. The `-n` option prevents the bootstrap programs from being installed.

If the `-v` option is supplied, *newfs* will print out its actions, including the parameters passed to *mkfs*.

Options which may be used to override default parameters passed to *mkfs* are:

- `-s size` The size of the file system in sectors.
- `-b block-size`
The block size of the file system in bytes.
- `-f frag-size`
The fragment size of the file system in bytes.
- `-t #tracks/cylinder`
- `-c #cylinders/group`
The number of cylinders per cylinder group in a file system. The default value used is 16.
- `-m free space %`
The percentage of space reserved from normal users; the minimum free space threshold. The default value used is 10%.
- `-r revolutions/minute`
The speed of the disk in revolutions per minute (normally 3600).
- `-S sector-size`
The size of a sector in bytes (almost never anything but 512).
- `-i number of bytes per inode`
This specifies the density of inodes in the file system. The default is to create an inode for each 2048 bytes of data space. If fewer inodes are desired, a larger number should be used; to create more inodes a smaller number should be given.

FILES

<code>/etc/disktab</code>	for disk geometry and file system partition information
<code>/etc/mkfs</code>	to actually build the file system
<code>/usr/mdec</code>	for boot strapping programs

SEE ALSO

`disktab(5)`, `fs(5)`, `diskpart(8)`, `fsck(8)`, `format(8)`, `mkfs(8)`, `tunefs(8)`

McKusick, Joy, Leffler; “A Fast File System for Unix”, Computer Systems Research Group, Dept of EECS, Berkeley, CA 94720; TR #7, September 1982.

BUGS

Should figure out the type of the disk without the user's help.

NAME

pstat — print system facts

SYNOPSIS

/etc/pstat -aixptufT [suboptions] [system] [corefile]

DESCRIPTION

Pstat interprets the contents of certain system tables. If *corefile* is given, the tables are sought there, otherwise in */dev/kmem*. The required namelist is taken from *lvmunix* unless *system* is specified. Options are

-a Under -p, describe all process slots rather than just active ones.

-l Print the inode table with the these headings:

LOC The core location of this table entry.

FLAGS Miscellaneous state variables encoded thus:

L locked
 U update time (*/s(5)*) must be corrected
 A access time must be corrected
 M file system is mounted here
 W wanted by another process (L flag is on)
 T contains a text file
 C changed time must be corrected
 S shared lock applied
 E exclusive lock applied
 Z someone waiting for an exclusive lock

CNT Number of open file table entries for this inode.

DEV Major and minor device number of file system in which this inode resides.

RDC Reference count of shared locks on the inode.

WRC Reference count of exclusive locks on the inode (this may be > 1 if, for example, a file descriptor is inherited across a fork).

INO I-number within the device.

MODE Mode bits, see *chmod(2)*.

NLK Number of links to this inode.

UID User ID of owner.

SIZ/DEV

Number of bytes in an ordinary file, or major and minor device of special file.

-x Print the text table with these headings:

LOC The core location of this table entry.

FLAGS Miscellaneous state variables encoded thus:

T *ptrace(2)* in effect
 W text not yet written on swap device
 L loading in progress
 K locked
 w wanted (L flag is on)
 P resulted from demand-page-from-inode exec format (see *execve(2)*)

DADDR Disk address in swap, measured in multiples of 512 bytes.

CADDR Head of a linked list of loaded processes using this text segment.

SIZE Size of text segment, measured in multiples of 512 bytes.

IPTR Core location of corresponding inode.

CNT Number of processes using this text segment.

CCNT Number of processes in core using this text segment.

-p Print process table for active processes with these headings:

LOC The core location of this table entry.

S Run state encoded thus:

0	no process
1	waiting for some event
3	runnable
4	being created
5	being terminated
6	stopped under trace

F Miscellaneous state variables, or-ed together (hexadecimal):

000001	loaded
000002	the scheduler process
000004	locked for swap out
000008	swapped out
000010	traced
000020	used in tracing
000080	in page-wait
000100	prevented from swapping during <i>fork(2)</i>
000200	gathering pages for raw i/o
000400	exiting
001000	process resulted from a <i>vfork(2)</i> which is not yet complete
002000	another flag for <i>vfork(2)</i>
004000	process has no virtual memory, as it is a parent in the context of <i>vfork(2)</i>
008000	process is demand paging data pages from its text inode.
010000	process has advised of anomalous behavior with <i>vadvise(2)</i> .
020000	process has advised of sequential behavior with <i>vadvise(2)</i> .
040000	process is in a sleep which will timeout.
080000	a parent of this process has exited and this process is now considered detached.
100000	process used 4.1BSD compatibility mode signal primitives, no system calls will restart.
200000	process is owed a profiling tick.

POIP number of pages currently being pushed out from this process.

PRI Scheduling priority, see *setpriority(2)*.

SIGNAL Signals received (signals 1-32 coded in bits 0-31),

UID Real user ID.

SLP Amount of time process has been blocked.

TIM Time resident in seconds; times over 127 coded as 127.

CPU Weighted integral of CPU time, for scheduler.

NI Nice level, see *setpriority(2)*.

PGRP Process number of root of process group (the opener of the controlling terminal).

PID The process ID number.

PPID The process ID of parent process.

ADDR If in core, the page frame number of the first page of the 'u-area' of the process. If swapped out, the position in the swap area measured in multiples of 512 bytes.

RSS Resident set size — the number of physical page frames allocated to this process.

SRSS RSS at last swap (0 if never swapped).

SIZE Virtual size of process image (data+stack) in multiples of 512 bytes.

WCHAN Wait channel number of a waiting process.

LINK Link pointer in list of runnable processes.

TEXTP If text is pure, pointer to location of text table entry.

CLKT Countdown for real interval timer, *setitimer(2)* measured in clock ticks (10

- milliseconds).
- t** Print table for terminals with these headings:
 - RAW** Number of characters in raw input queue.
 - CAN** Number of characters in canonicalized input queue.
 - OUT** Number of characters in putput queue.
 - MODE** See *ty(4)*.
 - ADDR** Physical device address.
 - DEL** Number of delimiters (newlines) in canonicalized input queue.
 - COL** Calculated column position of terminal.
 - STATE** Miscellaneous state variables encoded thus:
 - W** waiting for open to complete
 - O** open
 - S** has special (output) start routine
 - C** carrier is on
 - B** busy doing output
 - A** process is awaiting output
 - X** open for exclusive use
 - H** hangup on close
 - PGRP** Process group for which this is controlling terminal.
 - DISC** Line discipline; blank is old tty OTTYDISC or "new tty" for NTTYDISC or "net" for NETLDISC (see *bk(4)*).
 - u** print information about a user process; the next argument is its address as given by *ps(1)*. The process must be in main memory, or the file used can be a core image and the address 0.
 - f** Print the open file table with these headings:
 - LOC** The core location of this table entry.
 - TYPE** The type of object the file table entry points to.
 - FLG** Miscellaneous state variables encoded thus:
 - R** open for reading
 - W** open for writing
 - A** open for appending
 - CNT** Number of processes that know this open file.
 - INO** The location of the inode table entry for this file.
 - OFFS/SOCK**
 - The file offset (see *lseek(2)*), or the core address of the associated socket structure.
 - s** print information about swap space usage: the number of (1k byte) pages used and free is given as well as the number of used pages which belong to text images.
 - T** prints the number of used and free slots in the several system tables and is useful for checking to see how full system tables have become if the system is under heavy load.

FILES

/vmunix namelist
/dev/kmem default source of tables

SEE ALSO

ps(1), *stat(2)*, *fs(5)*
 K. Thompson, *UNIX Implementation*

BUGS

It would be very useful if the system recorded "maximum occupancy" on the tables reported by **-T**; even more useful if these tables were dynamically allocated.

NAME

rc — command script for auto-reboot and daemons

SYNOPSIS

/etc/rc
/etc/rc.local

DESCRIPTION

Rc is the command script which controls the automatic reboot and *rc.local* is the script holding commands which are pertinent only to a specific site.

When an automatic reboot is in progress, *rc* is invoked with the argument *autoboot* and runs a *fsck* with option *-p* to “preen” all the disks of minor inconsistencies resulting from the last system shutdown and to check for serious inconsistencies caused by hardware or software failure. If this auto-check and repair succeeds, then the second part of *rc* is run.

The second part of *rc*, which is run after a auto-reboot succeeds and also if *rc* is invoked when a single user shell terminates (see *init(8)*), starts all the daemons on the system, preserves editor files and clears the scratch directory */tmp*. *Rc.local* is executed immediately before any other commands after a successful *fsck*. Normally, the first commands placed in the *rc.local* file define the machine's name, using *hostname(1)*, and save any possible core image that might have been generated as a result of a system crash, *savecore(8)*. The latter command is included in the *rc.local* file because the directory in which core dumps are saved is usually site specific.

SEE ALSO

init(8), *reboot(8)*, *savecore(8)*

BUGS

NAME

reboot - reboot or halt system

SYNOPSIS

`/etc/reboot [-h] [-s]`

DESCRIPTION

Reboot reboots, or with the `-h` option, halts the system without syncing. If `-s` is specified the disks are synced before reboot or halt.

SEE ALSO

reboot(2), halt(8)

NAME

restor - incremental file system restore

SYNOPSIS

/etc/restor key [name ...]

DESCRIPTION

Restor is used to read tapes dumped with the *dump*(8) command. Its actions are controlled by the *key* argument. The *key* is a string of characters containing at most one function letter and possibly one or more function modifiers. Other arguments to the command are file or directory names specifying which files are to be restored. Unless the *-h* flag is specified (see below), the appearance of a directory name refers to the files and (recursively) subdirectories of that directory.

The function portion of the key is specified by one of the following letters:

- x** If file names are specified, the named files are extracted from the tape. If the named file matches a directory whose contents had been written onto the tape, this directory is (recursively) extracted. The owner, modification time, and mode are restored (if possible). If no file argument is given, the entire content of the tape is extracted.
- t** The names of the specified files are listed if they occur on the tape. If no file argument is given, all of the names on the tape are listed. Note that this key replaces the function of *dumpdir*(8).

The following characters may be used in addition to the letter which selects the function desired.

- v** Normally *restor* does its work silently. The **v** (verbose) option causes it to type the name of each file it treats preceded by the function letter. [With the **t** function, **v** gives more information about the tape entries than just the name.]
- f** causes *restor* to use the next argument as the name of the archive instead of */dev/rmt?*. [If the name of the file is *'-'*, *restor* reads from standard input. Thus, *dump*(8) and *restor* can be used in a pipeline to dump and restore a file system with the command


```
dump 0f - /usr |(cd /mnt; restor xf -) ]
```
- y** tells *restor* not to complain if gets a tape error, but simply to skip over the bad tape blocks and continue as best it can.
- m** causes *restor* to extract by inode numbers rather than by file name.
- h** causes *restor* to extract the actual directory, rather than the files that it references.
- s** causes *restor* to use the next argument as the number of the dump on the tape to skip to. This option is used for a tape with multiple dumps on it.

SEE ALSO

dump(8), *mkfs*(8)

FILES

/dev/rmt? the default tape drive
*rst** the temporary file used by *restor*

DIAGNOSTICS

Complaints about bad key characters.

Complaints if it gets a read error. If *-y* has been specified, or the user responds "y", *restor* will attempt to continue the restore.

If the dump extends over more than one tape, *restor* will ask the user to change tapes.

NAME

`rimioctl` – send ioctl commands to the Rimfire 44/45 controller

SYNOPSIS

`rimioctl {device} option [{arguments}]`

DESCRIPTION

The *rimioctl* utility is used to perform disk or tape operations for devices that are attached to the Ciprico Rimfire 44 or 45 disk/tape controller. The options supported by *rimioctl* are **map**, **rmbad**, **tracktype**, **badtracks**, **fsize**, **diskinfo**, and **errorcount**. The complete command syntax is:

`rimioctl /dev/rrim06 map {cyl} {head}`

`rimioctl /dev/rrim06 rmbad {inode}`

`rimioctl /dev/rrim06 tracktype {cyl} {head}`

`rimioctl /dev/rrim06 badtracks`

`rimioctl /dev/rrim06 fsize`

`rimioctl /dev/rrim06 diskinfo`

`rimioctl /dev/rct0 errorcount`

In the above, the files `/dev/rrim06` and `/dev/rct0` are representative examples of the special character devices that would be used during a typical *rimioctl* operation. Each of these options is described below.

Map

Option *map* remaps a bad track, specified by *cyl* and *head* , to another track.

Rmbad

Option *rmbad* removes bad block markings associated with inode *inode* .

Tracktype

Option *tracktype* prints out information about the specified track.

Badtracks

Option *badtracks* prints out a list of bad tracks on the given device.

Fsize

Option *fsize* prints information about the given filesystem including its size in blocks.

Diskinfo

Option *diskinfo* prints out information about each filesystem on the disk where *device* is located.

Errorcount

Option *errorcount* dumps all error counters associated with the given tape device. *NOTE*: This option only works with cartridge tape drives.

SEE, ALSO

`conn(5)`

DIAGNOSTICS

`/net` can not be opened

NAME

rlogind — remote login server

SYNOPSIS

/etc/rlogind [-d]

DESCRIPTION

Rlogind is the server for the *rlogin*(1C) program. The server provides a remote login facility with authentication based on privileged port numbers.

Rlogind listens for service requests at the port indicated in the “login” service specification; see *services*(5). When a service request is received the following protocol is initiated:

- 1) The server checks the client's source port. If the port is not in the range 0-1023, the server aborts the connection.
- 2) The server checks the client's source address. If the address is associated with a host for which no corresponding entry exists in the host name data base (see *hosts*(5)), the server aborts the connection.

Once the source port and address have been checked, *rlogind* allocates a pseudo terminal (see *pty*(4)), and manipulates file descriptors so that the slave half of the pseudo terminal becomes the *stdin*, *stdout*, and *stderr* for a login process. The login process is an instance of the *login*(1) program, invoked with the *-r* option. The login process then proceeds with the authentication process as described in *rshd*(8C), but if automatic authentication fails, it reprompts the user to login as one finds on a standard terminal line.

The parent of the login process manipulates the master side of the pseudo terminal, operating as an intermediary between the login process and the client instance of the *rlogin* program. In normal operation, the packet protocol described in *pty*(4) is invoked to provide *^S/^Q* type facilities and propagate interrupt signals to the remote programs. The login process propagates the client terminal's baud rate and terminal type, as found in the environment variable, “*TERM*”; see *environ*(7).

DIAGNOSTICS

All diagnostic messages are returned on the connection associated with the *stderr*, after which any network connections are closed. An error is indicated by a leading byte with a value of 1.

“**Hostname for your address unknown.**”

No entry in the host name database existed for the client's machine.

“**Try again.**”

A *fork* by the server failed.

“**/bin/sh: ...**”

The user's login shell could not be started.

BUGS

The authentication procedure used here assumes the integrity of each client machine and the connecting medium. This is insecure, but is useful in an “open” environment.

A facility to allow all data exchanges to be encrypted should be present.

NAME

rshd — remote shell server

SYNOPSIS

/etc/rshd

DESCRIPTION

Rshd is the server for the *rcmd*(3X) routine and, consequently, for the *rsh*(1C) program. The server provides remote execution facilities with authentication based on privileged port numbers.

Rshd listens for service requests at the port indicated in the “cmd” service specification; see *services*(5). When a service request is received the following protocol is initiated:

- 1) The server checks the client's source port. If the port is not in the range 0-1023, the server aborts the connection.
- 2) The server reads characters from the socket up to a null ('\0') byte. The resultant string is interpreted as an ASCII number, base 10.
- 3) If the number received in step 1 is non-zero, it is interpreted as the port number of a secondary stream to be used for the *stderr*. A second connection is then created to the specified port on the client's machine. The source port of this second connection is also in the range 0-1023.
- 4) The server checks the client's source address. If the address is associated with a host for which no corresponding entry exists in the host name data base (see *hosts*(5)), the server aborts the connection.
- 5) A null terminated user name of at most 16 characters is retrieved on the initial socket. This user name is interpreted as a user identity to use on the server's machine.
- 6) A null terminated user name of at most 16 characters is retrieved on the initial socket. This user name is interpreted as the user identity on the client's machine.
- 7) A null terminated command to be passed to a shell is retrieved on the initial socket. The length of the command is limited by the upper bound on the size of the system's argument list.
- 8) *Rshd* then validates the user according to the following steps. The remote user name is looked up in the password file and a *chdir* is performed to the user's home directory. If either the lookup or *chdir* fail, the connection is terminated. If the user is not the super-user, (user id 0), the file */etc/hosts.equiv* is consulted for a list of hosts considered “equivalent”. If the client's host name is present in this file, the authentication is considered successful. If the lookup fails, or the user is the super-user, then the file *.rhosts* in the home directory of the remote user is checked for the machine name and identity of the user on the client's machine. If this lookup fails, the connection is terminated.
- 9) A null byte is returned on the connection associated with the *stderr* and the command line is passed to the normal login shell of the user. The shell inherits the network connections established by *rshd*.

DIAGNOSTICS

All diagnostic messages are returned on the connection associated with the *stderr*, after which any network connections are closed. An error is indicated by a leading byte with a value of 1 (0 is returned in step 9 above upon successful completion of all the steps prior to the command execution).

“locuser too long”

The name of the user on the client's machine is longer than 16 characters.

“remuser too long”

The name of the user on the remote machine is longer than 16 characters.

“command too long”

The command line passed exceeds the size of the argument list (as configured into the system).

“Hostname for your address unknown.”

No entry in the host name database existed for the client's machine.

“Login incorrect.”

No password file entry for the user name existed.

“No remote directory.”

The *chdir* command to the home directory failed.

“Permission denied.”

The authentication procedure described above failed.

“Can't make pipe.”

The pipe needed for the *stderr*, wasn't created.

“Try again.”

A *fork* by the server failed.

“/bin/sh: ...”

The user's login shell could not be started.

SEE ALSO

rsh(1C), rcmd(3X)

BUGS

The authentication procedure used here assumes the integrity of each client machine and the connecting medium. This is insecure, but is useful in an “open” environment.

A facility to allow all data exchanges to be encrypted should be present.

NAME

rwhod — system status server

SYNOPSIS

/etc/rwhod

DESCRIPTION

Rwhod is the server which maintains the database used by the *rwho*(1C) and *runtime*(1C) programs. Its operation is predicated on the ability to *broadcast* messages on a network.

Rwhod operates as both a producer and consumer of status information. As a producer of information it periodically queries the state of the system and constructs status messages which are broadcast on a network. As a consumer of information, it listens for other *rwhod* servers' status messages, validating them, then recording them in a collection of files located in the directory */usr/spool/rwho*.

The *rwho* server transmits and receives messages at the port indicated in the "rwho" service specification, see *services*(5). The messages sent and received, are of the form:

```
struct outmp {
    char   out_line[8]; /* tty name */
    char   out_name[8]; /* user id */
    long   out_time; /* time on */
};

struct whod {
    char   wd_vers;
    char   wd_type;
    char   wd_fill[2];
    int    wd_sendtime;
    int    wd_recvtime;
    char   wd_hostname[32];
    int    wd_loadav[3];
    int    wd_boottime;
    struct whoent {
        struct outmp we_utmp;
        int    we_idle;
    } wd_we[1024 / sizeof (struct whoent)];
};
```

All fields are converted to network byte order prior to transmission. The load averages are as calculated by the *w*(1) program, and represent load averages over the 5, 10, and 15 minute intervals prior to a server's transmission. The host name included is that returned by the *gethostname*(2) system call. The array at the end of the message contains information about the users logged in to the sending machine. This information includes the contents of the *utmp*(5) entry for each non-idle terminal line and a value indicating the time since a character was last received on the terminal line.

Messages received by the *rwho* server are discarded unless they originated at a *rwho* server's port. In addition, if the host's name, as specified in the message, contains any unprintable ASCII characters, the message is discarded. Valid messages received by *rwhod* are placed in files named *whod.hostname* in the directory */usr/spool/rwho*. These files contain only the most recent message, in the format described above.

Status messages are generated approximately once every 60 seconds. *Rwhod* performs an *nlist*(3) on */vmunix* every 10 minutes to guard against the possibility that this file is not the system image currently operating.

SEE ALSO

`rwho(1C)`, `ruptime(1C)`

BUGS

Should relay status information between networks. People often interpret the server dieing as a machine going down.

NAME

savecore — save a core dump of the operating system

SYNOPSIS

/etc/savecore dirname [system]

DESCRIPTION

Savecore is meant to be called near the end of the */etc/rc* file. Its function is to save the core dump of the system (assuming one was made) and to write a reboot message in the shutdown log.

Savecore checks the core dump to be certain it corresponds with the current running unix. If it does it saves the core image in the file *dirname/vmcore.n* and it's brother, the namelist, *dirname/vmunix.n*. The trailing ".n" in the pathnames is replaced by a number which grows every time *savecore* is run in that directory.

Before *savecore* writes out a core image, it reads a number from the file *dirname/minfree*. If there are fewer free blocks on the filesystem which contains *dirname* than the number obtained from the *minfree* file, the core dump is not done. If the *minfree* file does not exist, *savecore* always writes out the core file (assuming that a core dump was taken).

Savecore also writes a reboot message in the shut down log. If the system crashed as a result of a panic, *savecore* records the panic string in the shut down log too.

If the core dump was from a system other than */vmunix*, the name of that system must be supplied as *sysname*.

FILES

/usr/adm/shutdownlog shut down log
/vmunix current UNIX

BUGS

Can be fooled into thinking a core dump is the wrong size.

NAME

shutdown — close down the system at a given time

SYNOPSIS

```
/etc/shutdown [ -k ] [ -r ] [ -h ] time [ warning-message ... ]
```

DESCRIPTION

Shutdown provides an automated shutdown procedure which a super-user can use to notify users nicely when the system is shutting down, saving them from system administrators, hackers, and gurus, who would otherwise not bother with niceties.

Time is the time at which *shutdown* will bring the system down and may be the word **now** (indicating an immediate shutdown) or specify a future time in one of two formats: **+number** and **hour:min**. The first form brings the system down in *number* minutes and the second brings the system down at the time of day indicated (as a 24-hour clock).

At intervals which get closer together as apocalypse approaches, warning messages are displayed at the terminals of all users on the system. Five minutes before shutdown, or immediately if shutdown is in less than 5 minutes, logins are disabled by creating */etc/nologin* and writing a message there. If this file exists when a user attempts to log in, *login(1)* prints its contents and exits. The file is removed just before *shutdown* exits.

At shutdown time a message is written in the file */usr/adm/shutdownlog*, containing the time of shutdown, who ran shutdown and the reason. Then a terminate signal is sent at *init* to bring the system down to single-user state. Alternatively, if **-r**, **-h**, or **-k** was used, then *shutdown* will exec *reboot(8)*, *halt(8)*, or avoid shutting the system down (respectively). (If it isn't obvious, **-k** is to make people *think* the system is going down!)

The time of the shutdown and the warning message are placed in */etc/nologin* and should be used to inform the users about when the system will be back up and why it is going down (or anything else).

FILES

/etc/nologin tells login not to let anyone log in
/usr/adm/shutdownlog log file for successful shutdowns.

SEE ALSO

login(1), *reboot(8)*

BUGS

Only allows you to kill the system between now and 23:59 if you use the absolute time for shutdown.

NAME

`sync` — update the super block

SYNOPSIS

`/etc/sync`

DESCRIPTION

Sync executes the *sync* system primitive. *Sync* can be called to insure all disk writes have been completed before the processor is halted in a way not suitably done by *reboot(8)* or *halt(8)*.

See *sync(2)* for details on the system primitive.

SEE ALSO

sync(2), *fsync(2)*, *halt(8)*, *reboot(8)*, *update(8)*

NAME

tunefs — tune up an existing file system

SYNOPSIS

/etc/tunefs tuneup-options special/filesys

DESCRIPTION

Tunefs is designed to change the dynamic parameters of a file system which affect the layout policies. The parameters which are to be changed are indicated by the flags given below:

-a maxcontig

This specifies the maximum number of contiguous blocks that will be laid out before forcing a rotational delay (see **-d** below). The default value is one, since most device drivers require an interrupt per disk transfer. Device drivers that can chain several buffers together in a single transfer should set this to the maximum chain length.

-d rotdelay

This specifies the expected time (in milliseconds) to service a transfer completion interrupt and initiate a new transfer on the same disk. It is used to decide how much rotational spacing to place between successive blocks in a file.

-e maxbpg

This indicates the maximum number of blocks any single file can allocate out of a cylinder group before it is forced to begin allocating blocks from another cylinder group. Typically this value is set to about one quarter of the total blocks in a cylinder group. The intent is to prevent any single file from using up all the blocks in a single cylinder group, thus degrading access times for all files subsequently allocated in that cylinder group. The effect of this limit is to cause big files to do long seeks more frequently than if they were allowed to allocate all the blocks in a cylinder group before seeking elsewhere. For file systems with exclusively large files, this parameter should be set higher.

-m minfree

This value specifies the percentage of space held back from normal users: the minimum free space threshold. The default value used is 10%. This value can be set to zero, however up to a factor of three in throughput will be lost over the performance obtained at a 10% threshold. Note that if the value is raised above the current usage level, users will be unable to allocate files until enough files have been deleted to get under the higher threshold.

SEE ALSO

fs(5), newfs(8), mkfs(8)

McKusick, Joy, Leffler; "A Fast File System for Unix". Computer Systems Research Group, Dept of EECS, Berkeley, CA 94720; TR #7, September 1982.

BUGS

This program should work on mounted and active file systems. Because the super-block is not kept in the buffer cache, the program will only take effect if it is run on dismounted file systems. (if run on the root file system, the system must be rebooted)

You can tune a file system, but you can't tune a fish.

NAME

`update` — periodically update the super block

SYNOPSIS

`/etc/update`

DESCRIPTION

Update is a program that executes the *sync*(2) primitive every 30 seconds. This insures that the file system is fairly up to date in case of a crash. This command should not be executed directly, but should be executed out of the initialization shell command file.

SEE ALSO

sync(2), *sync*(8), *init*(8), *rc*(8)

BUGS

With *update* running, if the CPU is halted just as the *sync* is executed, a file system can be damaged. This is partially due to DEC hardware that writes zeros when NPR requests fail. A fix would be to have *sync*(8) temporarily increment the system time by at least 30 seconds to trigger the execution of *update*. This would give 30 seconds grace to halt the CPU.

An introduction to the C shell

William Joy

Computer Science Division
Department of Electrical Engineering and Computer Science
University of California, Berkeley
Berkeley, California 94720

ABSTRACT

Csh is a new command language interpreter for UNIX† systems. It incorporates good features of other shells and a *history* mechanism similar to the *redo* of INTERLISP. While incorporating many features of other shells which make writing shell programs (shell scripts) easier, most of the features unique to *csh* are designed more for the interactive UNIX user.

UNIX users who have read a general introduction to the system will find a valuable basic explanation of the shell here. Simple terminal interaction with *csh* is possible after reading just the first section of this document. The second section describes the shells capabilities which you can explore after you have begun to become acquainted with the shell. Later sections introduce features which are useful, but not necessary for all users of the shell.

Back matter includes an appendix listing special characters of the shell and a glossary of terms and commands introduced in this manual.

November 8, 1980

†UNIX is a Trademark of Bell Laboratories.

An introduction to the C shell

William Joy

Computer Science Division
Department of Electrical Engineering and Computer Science
University of California, Berkeley
Berkeley, California 94720

Introduction

A *shell* is a command language interpreter. *Csh* is the name of one particular command interpreter on UNIX. The primary purpose of *csh* is to translate command lines typed at a terminal into system actions, such as invocation of other programs. *Csh* is a user program just like any you might write. Hopefully, *csh* will be a very useful program for you in interacting with the UNIX system.

In addition to this document, you will want to refer to a copy of the UNIX programmer's manual. The *csh* documentation in the manual provides a full description of all features of the shell and is a final reference for questions about the shell.

Many words in this document are shown in *italics*. These are important words: names of commands, and words which have special meaning in discussing the shell and UNIX. Many of the words are defined in a glossary at the end of this document. If you don't know what is meant by a word, you should look for it in the glossary.

Acknowledgements

Numerous people have provided good input about previous versions of *csh* and aided in its debugging and in the debugging of its documentation. I would especially like to thank Michael Ubell who made the crucial observation that history commands could be done well over the word structure of input text, and implemented a prototype history mechanism in an older version of the shell. Eric Allman has also provided a large number of useful comments on the shell, helping to unify those concepts which are present and to identify and eliminate useless and marginally useful features. Mike O'Brien suggested the pathname hashing mechanism which speeds command execution. Jim Kulp added the job control and directory stack primitives and added their documentation to this introduction.

1. Terminal usage of the shell

1.1. The basic notion of commands

A *shell* in UNIX acts mostly as a medium through which other *programs* are invoked. While it has a set of *builtin* functions which it performs directly, most commands cause execution of programs that are, in fact, external to the shell. The shell is thus distinguished from the command interpreters of other systems both by the fact that it is just a user program, and by the fact that it is used almost exclusively as a mechanism for invoking other programs.

Commands in the UNIX system consist of a list of strings or *words* interpreted as a *command name* followed by *arguments*. Thus the command

```
mail bill
```

consists of two words. The first word *mail* names the command to be executed, in this case the mail program which sends messages to other users. The shell uses the name of the command in attempting to execute it for you. It will look in a number of *directories* for a file with the name *mail* which is expected to contain the mail program.

The rest of the words of the command are given as *arguments* to the command itself when it is executed. In this case we specified also the argument *bill* which is interpreted by the *mail* program to be the name of a user to whom mail is to be sent. In normal terminal usage we might use the *mail* command as follows.

```
% mail bill
I have a question about the csh documentation.
My document seems to be missing page 5.
Does a page five exist?
Bill
EOT
%
```

Here we typed a message to send to *bill* and ended this message with a \uparrow D which sent an end-of-file to the mail program. (Here and throughout this document, the notation " \uparrow x" is to be read "control-x" and represents the striking of the x key while the control key is held down.) The mail program then echoed the characters 'EOT' and transmitted our message. The characters '%' were printed before and after the mail command by the shell to indicate that input was needed.

After typing the '%' prompt the shell was reading command input from our terminal. We typed a complete command 'mail bill'. The shell then executed the *mail* program with argument *bill* and went dormant waiting for it to complete. The mail program then read input from our terminal until we signalled an end-of-file via typing a \uparrow D after which the shell noticed that mail had completed and signaled us that it was ready to read from the terminal again by printing another '%' prompt.

This is the essential pattern of all interaction with UNIX through the shell. A complete command is typed at the terminal, the shell executes the command and when this execution completes, it prompts for a new command. If you run the editor for an hour, the shell will patiently wait for you to finish editing and obediently prompt you again whenever you finish editing.

An example of a useful command you can execute now is the *reset* command, which sets the default *erase* and *kill* characters on your terminal — the erase character erases the last character you typed and the kill character erases the entire line you have entered so far. By default, the erase character is '#' and the kill character is '@'. Most people who use CRT displays prefer to use the backspace (\uparrow H) character as their erase character since it is then easier to see what you have typed so far. You can make this be true by typing

```
tset -e
```

which tells the program *tset* to set the erase character, and its default setting for this character is a backspace.

1.2. Flag arguments

A useful notion in UNIX is that of a *flag* argument. While many arguments to commands specify file names or user names some arguments rather specify an optional capability of the command which you wish to invoke. By convention, such arguments begin with the character '-' (hyphen). Thus the command

```
ls
```

will produce a list of the files in the current *working directory*. The option `-s` is the size option, and

```
ls -s
```

causes *ls* to also give, for each file the size of the file in blocks of 512 characters. The manual section for each command in the UNIX reference manual gives the available options for each command. The *ls* command has a large number of useful and interesting options. Most other commands have either no options or only one or two options. It is hard to remember options of commands which are not used very frequently, so most UNIX utilities perform only one or two functions rather than having a large number of hard to remember options.

1.3. Output to files

Commands that normally read input or write output on the terminal can also be executed with this input and/or output done to a file.

Thus suppose we wish to save the current date in a file called 'now'. The command

```
date
```

will print the current date on our terminal. This is because our terminal is the default *standard output* for the *date* command and the *date* command prints the date on its standard output. The shell lets us *redirect* the *standard output* of a command through a notation using the *metacharacter* '>' and the name of the file where output is to be placed. Thus the command

```
date > now
```

runs the *date* command such that its standard output is the file 'now' rather than the terminal. Thus this command places the current date and time into the file 'now'. It is important to know that the *date* command was unaware that its output was going to a file rather than to the terminal. The shell performed this *redirection* before the command began executing.

One other thing to note here is that the file 'now' need not have existed before the *date* command was executed; the shell would have created the file if it did not exist. And if the file did exist? If it had existed previously these previous contents would have been discarded! A shell option *noclobber* exists to prevent this from happening accidentally; it is discussed in section 2.2.

The system normally keeps files which you create with '>' and all other files. Thus the default is for files to be permanent. If you wish to create a file which will be removed automatically, you can begin its name with a '#' character, this 'scratch' character denotes the fact that the file will be a scratch file.* The system will remove such files after a couple of days, or

*Note that if your erase character is a '#', you will have to precede the '#' with a '\'. The fact that the '#' character is the old (pre-CRT) standard erase character means that it seldom appears in a file name, and allows this convention to be used for scratch files. If you are using a CRT, your erase character should be a [H, as we demonstrated in section 1.1 how this could be set up.

sooner if file space becomes very tight. Thus, in running the *date* command above, we don't really want to save the output forever, so we would more likely do

```
date > #now
```

1.4. Metacharacters in the shell

The shell has a large number of special characters (like '>') which indicate special functions. We say that these notations have *syntactic* and *semantic* meaning to the shell. In general, most characters which are neither letters nor digits have special meaning to the shell. We shall shortly learn a means of *quotation* which allows us to use *metacharacters* without the shell treating them in any special way.

Metacharacters normally have effect only when the shell is reading our input. We need not worry about placing shell metacharacters in a letter we are sending via *mail*, or when we are typing in text or data to some other program. Note that the shell is only reading input when it has prompted with '% '.

1.5. Input from files; pipelines

We learned above how to *redirect* the *standard output* of a command to a file. It is also possible to redirect the *standard input* of a command from a file. This is not often necessary since most commands will read from a file whose name is given as an argument. We can give the command

```
sort < data
```

to run the *sort* command with standard input, where the command normally reads its input, from the file 'data'. We would more likely say

```
sort data
```

letting the *sort* command open the file 'data' for input itself since this is less to type.

We should note that if we just typed

```
sort
```

then the *sort* program would sort lines from its *standard input*. Since we did not *redirect* the standard input, it would sort lines as we typed them on the terminal until we typed a \uparrow D to indicate an end-of-file.

A most useful capability is the ability to combine the standard output of one command with the standard input of another, i.e. to run the commands in a sequence known as a *pipeline*. For instance the command

```
ls -s
```

normally produces a list of the files in our directory with the size of each in blocks of 512 characters. If we are interested in learning which of our files is largest we may wish to have this sorted by size rather than by name, which is the default way in which *ls* sorts. We could look at the many options of *ls* to see if there was an option to do this but would eventually discover that there is not. Instead we can use a couple of simple options of the *sort* command, combining it with *ls* to get what we want.

The *-n* option of *sort* specifies a numeric sort rather than an alphabetic sort. Thus

```
ls -s | sort -n
```

specifies that the output of the *ls* command run with the option *-s* is to be *piped* to the command *sort* run with the numeric sort option. This would give us a sorted list of our files by size, but with the smallest first. We could then use the *-r* reverse sort option and the *head* command in combination with the previous command doing

```
ls -s | sort -n -r | head -5
```

Here we have taken a list of our files sorted alphabetically, each with the size in blocks. We have run this to the standard input of the *sort* command asking it to sort numerically in reverse order (largest first). This output has then been run into the command *head* which gives us the first few lines. In this case we have asked *head* for the first 5 lines. Thus this command gives us the names and sizes of our 5 largest files.

The notation introduced above is called the *pipe* mechanism. Commands separated by '|' characters are connected together by the shell and the standard output of each is run into the standard input of the next. The leftmost command in a pipeline will normally take its standard input from the terminal and the rightmost will place its standard output on the terminal. Other examples of pipelines will be given later when we discuss the history mechanism; one important use of pipes which is illustrated there is in the routing of information to the line printer.

1.6. Filenames

Many commands to be executed will need the names of files as arguments. UNIX *pathnames* consist of a number of *components* separated by '/'. Each component except the last names a directory in which the next component resides, in effect specifying the *path* of directories to follow to reach the file. Thus the pathname

```
/etc/motd
```

specifies a file in the directory 'etc' which is a subdirectory of the *root* directory '/'. Within this directory the file named is 'motd' which stands for 'message of the day'. A *pathname* that begins with a slash is said to be an *absolute* pathname since it is specified from the absolute top of the entire directory hierarchy of the system (the *root*). *Pathnames* which do not begin with '/' are interpreted as starting in the current *working directory*, which is, by default, your *home* directory and can be changed dynamically by the *cd* change directory command. Such pathnames are said to be *relative* to the working directory since they are found by starting in the working directory and descending to lower levels of directories for each *component* of the pathname. If the pathname contains no slashes at all then the file is contained in the working directory itself and the pathname is merely the name of the file in this directory. Absolute pathnames have no relation to the working directory.

Most filenames consist of a number of alphanumeric characters and '.'s (periods). In fact, all printing characters except '/' (slash) may appear in filenames. It is inconvenient to have most non-alphabetic characters in filenames because many of these have special meaning to the shell. The character '.' (period) is not a shell-metacharacter and is often used to separate the *extension* of a file name from the base of the name. Thus

```
prog.c prog.o prog.errs prog.output
```

are four related files. They share a *base* portion of a name (a base portion being that part of the name that is left when a trailing '.' and following characters which are not '.' are stripped off). The file 'prog.c' might be the source for a C program, the file 'prog.o' the corresponding object file, the file 'prog.errs' the errors resulting from a compilation of the program and the file 'prog.output' the output of a run of the program.

If we wished to refer to all four of these files in a command, we could use the notation

```
prog.*
```

This word is expanded by the shell, before the command to which it is an argument is executed, into a list of names which begin with 'prog.'. The character '*' here matches any sequence (including the empty sequence) of characters in a file name. The names which match are alphabetically sorted and placed in the *argument list* of the command. Thus the command

```
echo prog.*
```

will echo the names

```
prog.c prog.errs prog.o prog.output
```

Note that the names are in sorted order here, and a different order than we listed them above. The *echo* command receives four words as arguments, even though we only typed one word as argument directly. The four words were generated by *filename expansion* of the one input word.

Other notations for *filename expansion* are also available. The character '?' matches any single character in a filename. Thus

```
echo ? ?? ???
```

will echo a line of filenames; first those with one character names, then those with two character names, and finally those with three character names. The names of each length will be independently sorted.

Another mechanism consists of a sequence of characters between '[' and ']'. This metasequence matches any single character from the enclosed set. Thus

```
prog.[co]
```

will match

```
prog.c prog.o
```

in the example above. We can also place two characters around a '-' in this notation to denote a range. Thus

```
chap.[1-5]
```

might match files

```
chap.1 chap.2 chap.3 chap.4 chap.5
```

if they existed. This is shorthand for

```
chap.[12345]
```

and otherwise equivalent.

An important point to note is that if a list of argument words to a command (an *argument list*) contains filename expansion syntax, and if this filename expansion syntax fails to match any existing file names, then the shell considers this to be an error and prints a diagnostic

```
No match.
```

and does not execute the command.

Another very important point is that files with the character '.' at the beginning are treated specially. Neither '*' or '?' or the '[' ']' mechanism will match it. This prevents accidental matching of the filenames '.' and '..' in the working directory which have special meaning to the system, as well as other files such as *.cshrc* which are not normally visible. We will discuss the special role of the file *.cshrc* later.

Another filename expansion mechanism gives access to the pathname of the *home* directory of other users. This notation consists of the character '~' (tilde) followed by another users' login name. For instance the word '~bill' would map to the pathname '/usr/bill' if the home directory for 'bill' was '/usr/bill'. Since, on large systems, users may have login directories scattered over many different disk volumes with different prefix directory names, this notation provides a reliable way of accessing the files of other users.

A special case of this notation consists of a '~' alone, e.g. '~/mbox'. This notation is expanded by the shell into the file 'mbox' in your *home* directory, i.e. into '/usr/bill/mbox' for me on Ernie Co-vax, the UCB Computer Science Department VAX machine, where this document was prepared. This can be very useful if you have used *cd* to change to another directory and have found a file you wish to copy using *cp*. If I give the command

```
cp thatfile *
```

the shell will expand this command to

```
cp thatfile /usr/bill
```

since my home directory is /usr/bill.

There also exists a mechanism using the characters '[' and ']' for abbreviating a set of words which have common parts but cannot be abbreviated by the above mechanisms because they are not files, are the names of files which do not yet exist, are not thus conveniently described. This mechanism will be described much later, in section 4.2, as it is used less frequently.

1.7. Quotation

We have already seen a number of metacharacters used by the shell. These metacharacters pose a problem in that we cannot use them directly as parts of words. Thus the command

```
echo *
```

will not echo the character '*'. It will either echo an sorted list of filenames in the current *working directory*, or print the message 'No match' if there are no files in the working directory.

The recommended mechanism for placing characters which are neither numbers, digits, '/', '.', or '-' in an argument word to a command is to enclose it with single quotation characters "'", i.e.

```
echo '*'
```

There is one special character '!' which is used by the *history* mechanism of the shell and which cannot be *escaped* by placing it within "'" characters. It and the character "'" itself can be preceded by a single '\ ' to prevent their special meaning. Thus

```
echo '\!'
```

prints

```
!'
```

These two mechanisms suffice to place any printing character into a word which is an argument to a shell command. They can be combined, as in

```
echo '\''
```

which prints

```
''
```

since the first '\ ' escaped the first "'" and the "*" was enclosed between "'" characters.

1.8. Terminating commands

When you are executing a command and the shell is waiting for it to complete there are several ways to force it to stop. For instance if you type the command

```
cat /etc/passwd
```

the system will print a copy of a list of all users of the system on your terminal. This is likely to continue for several minutes unless you stop it. You can send an INTERRUPT *signal* to the *cat* command by typing the DEL or RUBOUT key on your terminal.* Since *cat* does not take any precautions to avoid or otherwise handle this signal the INTERRUPT will cause it to terminate. The shell notices that *cat* has terminated and prompts you again with '% '. If you hit INTERRUPT

*Many users use *stty(1)* to change the interrupt character to |C.

again, the shell will just repeat its prompt since it handles INTERRUPT signals and chooses to continue to execute commands rather than terminating like *cat* did, which would have the effect of logging you out.

Another way in which many programs terminate is when they get an end-of-file from their standard input. Thus the *mail* program in the first example above was terminated when we typed a `^D` which generates an end-of-file from the standard input. The shell also terminates when it gets an end-of-file printing 'logout'; UNIX then logs you off the system. Since this means that typing too many `^D`'s can accidentally log us off, the shell has a mechanism for preventing this. This *ignoreeof* option will be discussed in section 2.2.

If a command has its standard input redirected from a file, then it will normally terminate when it reaches the end of this file. Thus if we execute

```
mail bill < prepared.text
```

the *mail* command will terminate without our typing a `^D`. This is because it read to the end-of-file of our file 'prepared.text' in which we placed a message for 'bill' with an editor program. We could also have done

```
cat prepared.text | mail bill
```

since the *cat* command would then have written the text through the pipe to the standard input of the *mail* command. When the *cat* command completed it would have terminated, closing down the pipeline and the *mail* command would have received an end-of-file from it and terminated. Using a pipe here is more complicated than redirecting input so we would more likely use the first form. These commands could also have been stopped by sending an INTERRUPT.

Another possibility for stopping a command is to suspend its execution temporarily, with the possibility of continuing execution later. This is done by sending a STOP signal via typing a `^Z`. This signal causes all commands running on the terminal (usually one but more if a pipeline is executing) to become suspended. The shell notices that the command(s) have been suspended, types 'Stopped' and then prompts for a new command. The previously executing command has been suspended, but otherwise unaffected by the STOP signal. Any other commands can be executed while the original command remains suspended. The suspended command can be continued using the *fg* command with no arguments. The shell will then retype the command to remind you which command is being continued, and cause the command to resume execution. Unless any input files in use by the suspended command have been changed in the meantime, the suspension has no effect whatsoever on the execution of the command. This feature can be very useful during editing, when you need to look at another file before continuing. An example of command suspension follows.

```
% mail harold
Someone just copied a big file into my directory and its name is
^Z
Stopped
% ls
funnyfile
prog.c
prog.o
% jobs
[1] + Stopped          mail harold
% fg
mail harold
funnyfile. Do you know who did it?
EOT
%
```

In this example someone was sending a message to Harold and forgot the name of the file he wanted to mention. The *mail* command was suspended by typing `^Z`. When the shell noticed

that the mail program was suspended, it typed 'Stopped' and prompted for a new command. Then the *ls* command was typed to find out the name of the file. The *jobs* command was run to find out which command was suspended. At this time the *fg* command was typed to continue execution of the mail program. Input to the mail program was then continued and ended with a `^D` which indicated the end of the message at which time the mail program typed EOT. The *jobs* command will show which commands are suspended. The `^Z` should only be typed at the beginning of a line since everything typed on the current line is discarded when a signal is sent from the keyboard. This also happens on INTERRUPT, and QUIT signals. More information on suspending jobs and controlling them is given in section 2.6.

If you write or run programs which are not fully debugged then it may be necessary to stop them somewhat ungracefully. This can be done by sending them a QUIT signal, sent by typing a `^C`. This will usually provoke the shell to produce a message like:

```
Quit (Core dumped)
```

indicating that a file 'core' has been created containing information about the program 'a.out's state when it terminated due to the QUIT signal. You can examine this file yourself, or forward information to the maintainer of the program telling him/her where the *core file* is.

If you run background commands (as explained in section 2.6) then these commands will ignore INTERRUPT and QUIT signals at the terminal. To stop them you must use the *kill* command. See section 2.6 for an example.

If you want to examine the output of a command without having it move off the screen as the output of the

```
cat /etc/passwd
```

command will, you can use the command

```
more /etc/passwd
```

The *more* program pauses after each complete screenful and types '--More--' at which point you can hit a space to get another screenful, a return to get another line, or a 'q' to end the *more* program. You can also use *more* as a filter, i.e.

```
cat /etc/passwd | more
```

works just like the more simple *more* command above.

For stopping output of commands not involving *more* you can use the `^S` key to stop the typeout. The typeout will resume when you hit `^Q` or any other key, but `^Q` is normally used because it only restarts the output and does not become input to the program which is running. This works well on low-speed terminals, but at 9600 baud it is hard to type `^S` and `^Q` fast enough to paginate the output nicely, and a program like *more* is usually used.

An additional possibility is to use the `^O` flush output character; when this character is typed, all output from the current command is thrown away (quickly) until the next input read occurs or until the next shell prompt. This can be used to allow a command to complete without having to suffer through the output on a slow terminal; `^O` is a toggle, so flushing can be turned off by typing `^O` again while output is being flushed.

1.9. What now?

We have so far seen a number of mechanisms of the shell and learned a lot about the way in which it operates. The remaining sections will go yet further into the internals of the shell, but you will surely want to try using the shell before you go any further. To try it you can log in to UNIX and type the following command to the system:

```
chsh myname /bin/csh
```

Here 'myname' should be replaced by the name you typed to the system prompt of 'login:' to get onto the system. Thus I would use 'chsh bill /bin/csh'. You only have to do this once: it

takes effect at next login. You are now ready to try using *cs/*.

Before you do the 'chsh' command, the shell you are using when you log into the system is '/bin/sh'. In fact, much of the above discussion is applicable to '/bin/sh'. The next section will introduce many features particular to *cs/* so you should change your shell to *cs/* before you begin reading it.

2. Details on the shell for terminal users

2.1. Shell startup and termination

When you login, the shell is started by the system in your *home* directory and begins by reading commands from a file *.cshrc* in this directory. All shells which you may start during your terminal session will read from this file. We will later see what kinds of commands are usefully placed there. For now we need not have this file and the shell does not complain about its absence.

A *login shell*, executed after you login to the system, will, after it reads commands from *.cshrc*, read commands from a file *.login* also in your home directory. This file contains commands which you wish to do each time you login to the UNIX system. My *.login* file looks something like:

```
set ignoreeof
set mail = (/usr/spool/mail/bill)
echo "$promptusers" ; users
alias ts \
    'set noglob ; eval `tset -s -m dialup:c100rv4pna -m plugboard:?hp2621nl *`;
ts; stty intr ↑C kill ↑U crt
set time=15 history=10
msgs -f
if (-e $mail) then
    echo "$promptmail"
    mail
end
```

This file contains several commands to be executed by UNIX each time I login. The first is a *set* command which is interpreted directly by the shell. It sets the shell variable *ignoreeof* which causes the shell to not log me off if I hit ↑D. Rather, I use the *logout* command to log off of the system. By setting the *mail* variable, I ask the shell to watch for incoming mail to me. Every 5 minutes the shell looks for this file and tells me if more mail has arrived there. An alternative to this is to put the command

```
biff y
```

in place of this *set*; this will cause me to be notified immediately when mail arrives, and to be shown the first few lines of the new message.

Next I set the shell variable 'time' to '15' causing the shell to automatically print out statistics lines for commands which execute for at least 15 seconds of CPU time. The variable 'history' is set to 10 indicating that I want the shell to remember the last 10 commands I type in its *history list*, (described later).

I create an *alias* "ts" which executes a *tset*(1) command setting up the modes of the terminal. The parameters to *tset* indicate the kinds of terminal which I usually use when not on a hardwired port. I then execute "ts" and also use the *stty* command to change the interrupt character to ↑C and the line kill character to ↑U.

I then run the 'msgs' program, which provides me with any system messages which I have not seen before; the '-f' option here prevents it from telling me anything if there are no new messages. Finally, if my mailbox file exists, then I run the 'mail' program to process my mail.

When the 'mail' and 'msgs' programs finish, the shell will finish processing my *.login* file and begin reading commands from the terminal, prompting for each with '%'. When I log off (by giving the *logout* command) the shell will print 'logout' and execute commands from the file *.logout* if it exists in my home directory. After that the shell will terminate and UNIX will log me off the system. If the system is not going down, I will receive a new login message. In

any case, after the 'logout' message the shell is committed to terminating and will take no further input from my terminal.

2.2. Shell variables

The shell maintains a set of *variables*. We saw above the variables *history* and *time* which had values '10' and '15'. In fact, each shell variable has as value an array of zero or more *strings*. Shell variables may be assigned values by the *set* command. It has several forms, the most useful of which was given above and is

```
set name=value
```

Shell variables may be used to store values which are to be used in commands later through a substitution mechanism. The shell variables most commonly referenced are, however, those which the shell itself refers to. By changing the values of these variables one can directly affect the behavior of the shell.

One of the most important variables is the variable *path*. This variable contains a sequence of directory names where the shell searches for commands. The *set* command with no arguments shows the value of all variables currently defined (we usually say *set*) in the shell. The default value for *path* will be shown by *set* to be

```
% set
argv      ()
cwd       /usr/bill
home      /usr/bill
path      (. /usr/ucb /bin /usr/bin)
prompt    %
shell     /bin/csh
status    0
term      c100rv4pna
user      bill
%
```

This output indicates that the variable *path* points to the current directory '.' and then '/usr/ucb', '/bin' and '/usr/bin'. Commands which you may write might be in '.' (usually one of your directories). Commands developed at Berkeley, live in '/usr/ucb' while commands developed at Bell Laboratories live in '/bin' and '/usr/bin'.

A number of locally developed programs on the system live in the directory '/usr/local'. If we wish that all shells which we invoke to have access to these new programs we can place the command

```
set path=(. /usr/ucb /bin /usr/bin /usr/local)
```

in our file *.cshrc* in our home directory. Try doing this and then logging out and back in and do

```
set
```

again to see that the value assigned to *path* has changed.

One thing you should be aware of is that the shell examines each directory which you insert into your *path* and determines which commands are contained there. Except for the current directory '.', which the shell treats specially, this means that if commands are added to a directory in your search *path* after you have started the shell, they will not necessarily be found by the shell. If you wish to use a command which has been added in this way, you should give the command

```
rehash
```

to the shell, which will cause it to recompute its internal table of command locations, so that it will find the newly added command. Since the shell has to look in the current directory '.' on

each command, placing it at the end of the path specification usually works equivalently and reduces overhead.

Other useful built in variables are the variable *home* which shows your home directory, *cwd* which contains your current working directory, the variable *ignoreeof* which can be set in your *.login* file to tell the shell not to exit when it receives an end-of-file from a terminal (as described above). The variable 'ignoreeof' is one of several variables which the shell does not care about the value of, only whether they are *set* or *unset*. Thus to set this variable you simply do

```
set ignoreeof
```

and to unset it do

```
unset ignoreeof
```

These give the variable 'ignoreeof' no value, but none is desired or required.

Finally, some other built-in shell variables of use are the variables *noclobber* and *mail*. The metasyntax

```
> filename
```

which redirects the standard output of a command will overwrite and destroy the previous contents of the named file. In this way you may accidentally overwrite a file which is valuable. If you would prefer that the shell not overwrite files in this way you can

```
set noclobber
```

in your *.login* file. Then trying to do

```
date > now
```

would cause a diagnostic if 'now' existed already. You could type

```
date >! now
```

if you really wanted to overwrite the contents of 'now'. The '>!' is a special metasyntax indicating that clobbering the file is ok.†

2.3. The shell's history list

The shell can maintain a *history list* into which it places the words of previous commands. It is possible to use a notation to reuse commands or words from commands in forming new commands. This mechanism can be used to repeat previous commands or to correct minor typing mistakes in commands.

The following figure gives a sample session involving typical usage of the history mechanism of the shell. In this example we have a very simple C program which has a bug (or two) in it in the file 'bug.c', which we 'cat' out on our terminal. We then try to run the C compiler on it, referring to the file again as '!S', meaning the last argument to the previous command. Here the '!' is the history mechanism invocation metacharacter, and the 'S' stands for the last argument, by analogy to '\$' in the editor which stands for the end of the line. The shell echoed the command, as it would have been typed without use of the history mechanism, and then executed it. The compilation yielded error diagnostics so we now run the editor on the file we were trying to compile, fix the bug, and run the C compiler again, this time referring to this command simply as '!c', which repeats the last command which started with the letter 'c'. If there were other commands starting with 'c' done recently we could have said '!cc' or even '!cc:p' which would have printed the last command starting with 'cc' without executing it.

†The space between the '!' and the word 'now' is critical here, as '!now' would be an invocation of the *history* mechanism, and have a totally different effect.

```
% cat bug.c
main()

{
    printf("hello");
}
% cc !S
cc bug.c
"bug.c", line 4: newline in string or char constant
"bug.c", line 5: syntax error
% ed !S
ed bug.c
29
4s/):/"&/p
    printf("hello");

w
30
q
% !c
cc bug.c
% a.out
hello% !e
ed bug.c
30
4s/lo/lo\\n/p
    printf("hello\n");

w
32
q
% !c -o bug
cc bug.c -o bug
% size a.out bug
a.out: 2784+364+1028 = 4176b = 0x1050b
bug: 2784+364+1028 = 4176b = 0x1050b
% ls -l !*
ls -l a.out bug
-rwxr-xr-x 1 bill    3932 Dec 19 09:41 a.out
-rwxr-xr-x 1 bill    3932 Dec 19 09:42 bug
% bug
hello
% num bug.c | spp
spp: Command not found.
% |spp|ssp
num bug.c | ssp
  1 main()
  3 {
  4     printf("hello\n");
  5 }
% !! | lpr
num bug.c | ssp | lpr
%
```

After this recompilation, we ran the resulting 'a.out' file, and then noting that there still was a bug, ran the editor again. After fixing the program we ran the C compiler again, but tacked onto the command an extra '-o bug' telling the compiler to place the resultant binary in the file 'bug' rather than 'a.out'. In general, the history mechanisms may be used anywhere in the formation of new commands and other characters may be placed before and after the substituted commands.

We then ran the 'size' command to see how large the binary program images we have created were, and then an 'ls -l' command with the same argument list, denoting the argument list '*'. Finally we ran the program 'bug' to see that its output is indeed correct.

To make a numbered listing of the program we ran the 'num' command on the file 'bug.c'. In order to compress out blank lines in the output of 'num' we ran the output through the filter 'ssp', but misspelled it as spp. To correct this we used a shell substitute, placing the old text and new text between '[' characters. This is similar to the substitute command in the editor. Finally, we repeated the same command with '!!', but sent its output to the line printer.

There are other mechanisms available for repeating commands. The *history* command prints out a number of previous commands with numbers by which they can be referenced. There is a way to refer to a previous command by searching for a string which appeared in it, and there are other, less useful, ways to select arguments to include in a new command. A complete description of all these mechanisms is given in the C shell manual pages in the UNIX Programmers Manual.

2.4. Aliases

The shell has an *alias* mechanism which can be used to make transformations on input commands. This mechanism can be used to simplify the commands you type, to supply default arguments to commands, or to perform transformations on commands and their arguments. The alias facility is similar to a macro facility. Some of the features obtained by aliasing can be obtained also using shell command files, but these take place in another instance of the shell and cannot directly affect the current shells environment or involve commands such as *cd* which must be done in the current shell.

As an example, suppose that there is a new version of the mail program on the system called 'newmail' you wish to use, rather than the standard mail program which is called 'mail'. If you place the shell command

```
alias mail newmail
```

in your *.cshrc* file, the shell will transform an input line of the form

```
mail bill
```

into a call on 'newmail'. More generally, suppose we wish the command 'ls' to always show sizes of files, that is to always do '-s'. We can do

```
alias ls ls -s
```

or even

```
alias dir ls -s
```

creating a new command syntax 'dir' which does an 'ls -s'. If we say

```
dir ~bill
```

then the shell will translate this to

```
ls -s /mnt/bill
```

Thus the *alias* mechanism can be used to provide short names for commands, to provide default arguments, and to define new short commands in terms of other commands. It is also possible to define aliases which contain multiple commands or pipelines, showing where the

arguments to the original command are to be substituted using the facilities of the history mechanism. Thus the definition

```
alias cd 'cd \!* ; ls '
```

would do an *ls* command after each change directory *cd* command. We enclosed the entire alias definition in "" characters to prevent most substitutions from occurring and the character "." from being recognized as a metacharacter. The "!" here is escaped with a "\" to prevent it from being interpreted when the alias command is typed in. The "\!" here substitutes the entire argument list to the pre-aliasing *cd* command, without giving an error if there were no arguments. The ";" separating commands is used here to indicate that one command is to be done and then the next. Similarly the definition

```
alias whois 'grep \!| /etc/passwd'
```

defines a command which looks up its first argument in the password file.

Warning: The shell currently reads the *.cshrc* file each time it starts up. If you place a large number of commands there, shells will tend to start slowly. A mechanism for saving the shell environment after reading the *.cshrc* file and quickly restoring it is under development, but for now you should try to limit the number of aliases you have to a reasonable number... 10 or 15 is reasonable, 50 or 60 will cause a noticeable delay in starting up shells, and make the system seem sluggish when you execute commands from within the editor and other programs.

2.5. More redirection; >> and >&

There are a few more notations useful to the terminal user which have not been introduced yet.

In addition to the standard output, commands also have a *diagnostic output* which is normally directed to the terminal even when the standard output is redirected to a file or a pipe. It is occasionally desirable to direct the diagnostic output along with the standard output. For instance if you want to redirect the output of a long running command into a file and wish to have a record of any error diagnostic it produces you can do

```
command >& file
```

The '>&' here tells the shell to route both the diagnostic output and the standard output into 'file'. Similarly you can give the command

```
command |& lpr
```

to route both standard and diagnostic output through the pipe to the line printer daemon *lpr.#*

Finally, it is possible to use the form

```
command >> file
```

to place output at the end of an existing file.†

#A command form

```
command >&! file
```

exists, and is used when *noclobber* is set and *file* already exists.

†If *noclobber* is set, then an error will result if *file* does not exist, otherwise the shell will create *file* if it doesn't exist. A form

```
command >>! file
```

makes it not be an error for file to not exist when *noclobber* is set.

2.6. Jobs: Background, Foreground, or Suspended

When one or more commands are typed together as a pipeline or as a sequence of commands separated by semicolons, a single *job* is created by the shell consisting of these commands together as a unit. Single commands without pipes or semicolons create the simplest jobs. Usually, every line typed to the shell creates a job. Some lines that create jobs (one per line) are

```
sort < data
ls -s | sort -n | head -5
mail harold
```

If the metacharacter '&' is typed at the end of the commands, then the job is started as a *background* job. This means that the shell does not wait for it to complete but immediately prompts and is ready for another command. The job runs *in the background* at the same time that normal jobs, called *foreground* jobs, continue to be read and executed by the shell one at a time. Thus

```
du > usage &
```

would run the *du* program, which reports on the disk usage of your working directory (as well as any directories below it), put the output into the file 'usage' and return immediately with a prompt for the next command without waiting for *du* to finish. The *du* program would continue executing in the background until it finished, even though you can type and execute more commands in the mean time. When a background job terminates, a message is typed by the shell just before the next prompt telling you that the job has completed. In the following example the *du* job finishes sometime during the execution of the *mail* command and its completion is reported just before the prompt after the *mail* job is finished.

```
% du > usage &
[1] 503
% mail bill
How do you know when a background job is finished?
EOT
[1] - Done          du > usage
%
```

If the job did not terminate normally the 'Done' message might say something else like 'Killed'. If you want the terminations of background jobs to be reported at the time they occur (possibly interrupting the output of other foreground jobs), you can set the *notify* variable. In the previous example this would mean that the 'Done' message might have come right in the middle of the message to Bill. Background jobs are unaffected by any signals from the keyboard like the STOP, INTERRUPT, or QUIT signals mentioned earlier.

Jobs are recorded in a table inside the shell until they terminate. In this table, the shell remembers the command names, arguments and the *process numbers* of all commands in the job as well as the working directory where the job was started. Each job in the table is either running *in the foreground* with the shell waiting for it to terminate, running *in the background*, or *suspended*. Only one job can be running in the foreground at one time, but several jobs can be suspended or running in the background at once. As each job is started, it is assigned a small identifying number called the *job number* which can be used later to refer to the job in the commands described below. Job numbers remain the same until the job terminates and then are re-used.

When a job is started in the background using '&', its number, as well as the process numbers of all its (top level) commands, is typed by the shell before prompting you for another command. For example,

```
% ls -s | sort -n > usage &  
[2] 2034 2035  
%
```

runs the 'ls' program with the '-s' options, pipes this output into the 'sort' program with the '-n' option which puts its output into the file 'usage'. Since the '&' was at the end of the line, these two programs were started together as a background job. After starting the job, the shell prints the job number in brackets (2 in this case) followed by the process number of each program started in the job. Then the shell immediately prompts for a new command, leaving the job running simultaneously.

As mentioned in section 1.8, foreground jobs become *suspended* by typing ↑Z which sends a STOP signal to the currently running foreground job. A background job can become suspended by using the *stop* command described below. When jobs are suspended they merely stop any further progress until started again, either in the foreground or the background. The shell notices when a job becomes stopped and reports this fact, much like it reports the termination of background jobs. For foreground jobs this looks like

```
% du > usage  
↑Z  
Stopped  
%
```

'Stopped' message is typed by the shell when it notices that the *du* program stopped. For background jobs, using the *stop* command, it is

```
% sort usage &  
[1] 2345  
% stop %1  
[1] + Stopped (signal)    sort usage  
%
```

Suspending foreground jobs can be very useful when you need to temporarily change what you are doing (execute other commands) and then return to the suspended job. Also, foreground jobs can be suspended and then continued as background jobs using the *bg* command, allowing you to continue other work and stop waiting for the foreground job to finish. Thus

```
% du > usage  
↑Z  
Stopped  
% bg  
[1] du > usage &  
%
```

starts 'du' in the foreground, stops it before it finishes, then continues it in the background allowing more foreground commands to be executed. This is especially helpful when a foreground job ends up taking longer than you expected and you wish you had started it in the background in the beginning.

All *job control* commands can take an argument that identifies a particular job. All job name arguments begin with the character '%', since some of the job control commands also accept process numbers (printed by the *ps* command.) The default job (when no argument is given) is called the *current* job and is identified by a '+' in the output of the *jobs* command, which shows you which jobs you have. When only one job is stopped or running in the background (the usual case) it is always the current job thus no argument is needed. If a job is stopped while running in the foreground it becomes the *current* job and the existing current job becomes the *previous* job — identified by a '-' in the output of *jobs*. When the current job terminates, the previous job becomes the current job. When given, the argument is either '%-' (indicating the previous job); '%#', where # is the job number; '%pref' where pref is some

unique prefix of the command name and arguments of one of the jobs; or '%?' followed by some string found in only one of the jobs.

The *jobs* command types the table of jobs, giving the job number, commands and status ('Stopped' or 'Running') of each background or suspended job. With the '-l' option the process numbers are also typed.

```
% du > usage &
[1] 3398
% ls -s | sort -n > myfile &
[2] 3405
% mail bill
↑Z
Stopped
% jobs
[1] - Running      du > usage
[2]  Running      ls -s | sort -n > myfile
[3] + Stopped     mail bill
% fg %ls
ls -s | sort -n > myfile
% more myfile
```

The *fg* command runs a suspended or background job in the foreground. It is used to restart a previously suspended job or change a background job to run in the foreground (allowing signals or input from the terminal). In the above example we used *fg* to change the 'ls' job from the background to the foreground since we wanted to wait for it to finish before looking at its output file. The *bg* command runs a suspended job in the background. It is usually used after stopping the currently running foreground job with the STOP signal. The combination of the STOP signal and the *bg* command changes a foreground job into a background job. The *stop* command suspends a background job.

The *kill* command terminates a background or suspended job immediately. In addition to jobs, it may be given process numbers as arguments, as printed by *ps*. Thus, in the example above, the running *du* command could have been terminated by the command

```
% kill %1
[1] Terminated      du > usage
%
```

The *notify* command (not the variable mentioned earlier) indicates that the termination of a specific job should be reported at the time it finishes instead of waiting for the next prompt.

If a job running in the background tries to read input from the terminal it is automatically stopped. When such a job is then run in the foreground, input can be given to the job. If desired, the job can be run in the background again until it requests input again. This is illustrated in the following sequence where the 's' command in the text editor might take a long time.

```
% ed bigfile
120000
1,$s/thisword/thatword/
↑Z
Stopped
% bg
[1] ed bigfile &
%
... some foreground commands
[1] Stopped (tty input)  ed bigfile
% fg
```

```
ed bigfile
w
120000
q
%
```

So after the 's' command was issued, the 'ed' job was stopped with ↑Z and then put in the background using *bg*. Some time later when the 's' command was finished, *ed* tried to read another command and was stopped because jobs in the background cannot read from the terminal. The *fg* command returned the 'ed' job to the foreground where it could once again accept commands from the terminal.

The command

```
stty tostop
```

causes all background jobs run on your terminal to stop when they are about to write output to the terminal. This prevents messages from background jobs from interrupting foreground job output and allows you to run a job in the background without losing terminal output. It also can be used for interactive programs that sometimes have long periods without interaction. Thus each time it outputs a prompt for more input it will stop before the prompt. It can then be run in the foreground using *fg*, more input can be given and, if necessary stopped and returned to the background. This *stty* command might be a good thing to put in your *.login* file if you do not like output from background jobs interrupting your work. It also can reduce the need for redirecting the output of background jobs if the output is not very big:

```
% stty tostop
% wc hugefile &
[1] 10387
% ed text
... some time later
q
[1] Stopped (tty output)   wc hugefile
% fg wc
wc hugefile
13371 30123 302577
% stty -tostop
```

Thus after some time the 'wc' command, which counts the lines, words and characters in a file, had one line of output. When it tried to write this to the terminal it stopped. By restarting it in the foreground we allowed it to write on the terminal exactly when we were ready to look at its output. Programs which attempt to change the mode of the terminal will also block, whether or not *istop* is set, when they are not in the foreground, as it would be very unpleasant to have a background job change the state of the terminal.

Since the *jobs* command only prints jobs started in the currently executing shell, it knows nothing about background jobs started in other login sessions or within shell files. The *ps* can be used in this case to find out about background jobs not started in the current shell.

2.7. Working Directories

As mentioned in section 1.6, the shell is always in a particular *working directory*. The 'change directory' command *chdir* (its short form *cd* may also be used) changes the working directory of the shell, that is, changes the directory you are located in.

It is useful to make a directory for each project you wish to work on and to place all files related to that project in that directory. The 'make directory' command, *mkdir*, creates a new directory. The *pwd* ('print working directory') command reports the absolute pathname of the working directory of the shell, that is, the directory you are located in. Thus in the example below:

```
% pwd
/usr/bill
% mkdir newspaper
% chdir newspaper
% pwd
/usr/bill/newspaper
%
```

the user has created and moved to the directory *newspaper*. where, for example, he might place a group of related files.

No matter where you have moved to in a directory hierarchy, you can return to your 'home' login directory by doing just

```
cd
```

with no arguments. The name `..` always means the directory above the current one in the hierarchy, thus

```
cd ..
```

changes the shell's working directory to the one directly above the current one. The name `..` can be used in any pathname, thus,

```
cd ../programs
```

means change to the directory 'programs' contained in the directory above the current one. If you have several directories for different projects under, say, your home directory, this shorthand notation permits you to switch easily between them.

The shell always remembers the pathname of its current working directory in the variable *cwd*. The shell can also be requested to remember the previous directory when you change to a new working directory. If the 'push directory' command *pushd* is used in place of the *cd* command, the shell saves the name of the current working directory on a *directory stack* before changing to the new one. You can see this list at any time by typing the 'directories' command *dirs*.

```
% pushd newspaper/references
~/newspaper/references ~
% pushd /usr/lib/tmac
/usr/lib/tmac ~/newspaper/references ~
% dirs
/usr/lib/tmac ~/newspaper/references ~
% popd
~/newspaper/references ~
% popd
-
%
```

The list is printed in a horizontal line, reading left to right, with a tilde (~) as shorthand for your home directory—in this case `~/usr/bill`. The directory stack is printed whenever there is more than one entry on it and it changes. It is also printed by a *dirs* command. *Dirs* is usually faster and more informative than *pwd* since it shows the current working directory as well as any other directories remembered in the stack.

The *pushd* command with no argument alternates the current directory with the first directory in the list. The 'pop directory' *popd* command without an argument returns you to the directory you were in prior to the current one, discarding the previous current directory from the stack (forgetting it). Typing *popd* several times in a series takes you backward through the directories you had been in (changed to) by *pushd* command. There are other options to *pushd* and *popd* to manipulate the contents of the directory stack and to change to directories not at the top of the stack; see the *csh* manual page for details.

Since the shell remembers the working directory in which each job was started, it warns you when you might be confused by restarting a job in the foreground which has a different working directory than the current working directory of the shell. Thus if you start a background job, then change the shell's working directory and then cause the background job to run in the foreground, the shell warns you that the working directory of the currently running foreground job is different from that of the shell.

```
% dirs -l
/mnt/bill
% cd myproject
% dirs
~/myproject
% ed prog.c
1143
|Z
Stopped
% cd ..
% ls
myproject
textfile
% fg
ed prog.c (wd: ~/myproject)
```

This way the shell warns you when there is an implied change of working directory, even though no `cd` command was issued. In the above example the 'ed' job was still in `~/mnt/bill/project` even though the shell had changed to `~/mnt/bill`. A similar warning is given when such a foreground job terminates or is suspended (using the STOP signal) since the return to the shell again implies a change of working directory.

```
% fg
ed prog.c (wd: ~/myproject)
... after some editing
q
(wd now: ~)
%
```

These messages are sometimes confusing if you use programs that change their own working directories, since the shell only remembers which directory a job is started in, and assumes it stays there. The `-l` option of `jobs` will type the working directory of suspended or background jobs when it is different from the current working directory of the shell.

2.8. Useful built-in commands

We now give a few of the useful built-in commands of the shell describing how they are used.

The `alias` command described above is used to assign new aliases and to show the existing aliases. With no arguments it prints the current aliases. It may also be given only one argument such as

```
alias ls
```

to show the current alias for, e.g., 'ls'.

The `echo` command prints its arguments. It is often used in *shell scripts* or as an interactive command to see what filename expansions will produce.

The `history` command will show the contents of the history list. The numbers given with the history events can be used to reference previous events which are difficult to reference using the contextual mechanisms introduced above. There is also a shell variable called *prompt*.

By placing a `'!` character in its value the shell will there substitute the number of the current command in the history list. You can use this number to refer to this command in a history substitution. Thus you could

```
set prompt = \! % '
```

Note that the `'!` character had to be *escaped* here even within `'` characters.

The `limit` command is used to restrict use of resources. With no arguments it prints the current limitations:

```
cpulimit      unlimited
filesize      unlimited
datasize      5616 kbytes
stacksize     512 kbytes
coredumpsize  unlimited
```

Limits can be set, e.g.:

```
limit coredumpsize 128k
```

Most reasonable units abbreviations will work; see the `cs` manual page for more details.

The `logout` command can be used to terminate a login shell which has *ignoreeof* set.

The `rehash` command causes the shell to recompute a table of where commands are located. This is necessary if you add a command to a directory in the current shell's search path and wish the shell to find it, since otherwise the hashing algorithm may tell the shell that the command wasn't in that directory when the hash table was computed.

The `repeat` command can be used to repeat a command several times. Thus to make 5 copies of the file `one` in the file `five` you could do

```
repeat 5 cat one >> five
```

The `setenv` command can be used to set variables in the environment. Thus

```
setenv TERM adm3a
```

will set the value of the environment variable `TERM` to `'adm3a'`. A user program `printenv` exists which will print out the environment. It might then show:

```
% printenv
HOME = /usr/bill
SHELL = /bin/csh
PATH = :/usr/ucb:/bin:/usr/bin:/usr/local
TERM = adm3a
USER = bill
%
```

The `source` command can be used to force the current shell to read commands from a file. Thus

```
source .cshrc
```

can be used after editing in a change to the `.cshrc` file which you wish to take effect before the next time you login.

The `time` command can be used to cause a command to be timed no matter how much CPU time it takes. Thus


```
% time cp /etc/rc /usr/bill/rc
0.0u 0.1s 0:01 8% 2+1k 3+2io 1pf+0w
% time wc /etc/rc /usr/bill/rc
  52  178 1347 /etc/rc
  52  178 1347 /usr/bill/rc
 104  356 2694 total
0.1u 0.1s 0:00 13% 3+3k 5+3io 7pf+0w
%
```

indicates that the *cp* command used a negligible amount of user time (u) and about 1/10th of a system time (s); the elapsed time was 1 second (0:01), there was an average memory usage of 2k bytes of program space and 1k bytes of data space over the cpu time involved (2+1k); the program did three disk reads and two disk writes (3+2io), and took one page fault and was not swapped (1pf+0w). The word count command *wc* on the other hand used 0.1 seconds of user time and 0.1 seconds of system time in less than a second of elapsed time. The percentage '13%' indicates that over the period when it was active the command 'wc' used an average of 13 percent of the available CPU cycles of the machine.

The *unalias* and *unset* commands can be used to remove aliases and variable definitions from the shell, and *unsetenv* removes variables from the environment.

2.9. What else?

This concludes the basic discussion of the shell for terminal users. There are more features of the shell to be discussed here, and all features of the shell are discussed in its manual pages. One useful feature which is discussed later is the *foreach* built-in command which can be used to run the same command sequence with a number of different arguments.

If you intend to use UNIX a lot you should look through the rest of this document and the shell manual pages to become familiar with the other facilities which are available to you.

3. Shell control structures and command scripts

3.1. Introduction

It is possible to place commands in files and to cause shells to be invoked to read and execute commands from these files, which are called *shell scripts*. We here detail those features of the shell useful to the writers of such scripts.

3.2. Make

It is important to first note what shell scripts are *not* useful for. There is a program called *make* which is very useful for maintaining a group of related files or performing sets of operations on related files. For instance a large program consisting of one or more files can have its dependencies described in a *makefile* which contains definitions of the commands used to create these different files when changes occur. Definitions of the means for printing listings, cleaning up the directory in which the files reside, and installing the resultant programs are easily, and most appropriately placed in this *makefile*. This format is superior and preferable to maintaining a group of shell procedures to maintain these files.

Similarly when working on a document a *makefile* may be created which defines how different versions of the document are to be created and which options of *troff* or *nrff* are appropriate.

3.3. Invocation and the argv variable

A *cs*h command script may be interpreted by saying

```
% csh script ...
```

where *script* is the name of the file containing a group of *cs*h commands and '...' is replaced by a sequence of arguments. The shell places these arguments in the variable *argv* and then begins to read commands from the script. These parameters are then available through the same mechanisms which are used to reference any other shell variables.

If you make the file 'script' executable by doing

```
chmod 755 script
```

and place a shell comment at the beginning of the shell script (i.e. begin the file with a '#' character) then a '/bin/csh' will automatically be invoked to execute 'script' when you type

```
script
```

If the file does not begin with a '#' then the standard shell '/bin/sh' will be used to execute it. This allows you to convert your older shell scripts to use *cs*h at your convenience.

3.4. Variable substitution

After each input line is broken into words and history substitutions are done on it, the input line is parsed into distinct commands. Before each command is executed a mechanism know as *variable substitution* is done on these words. Keyed by the character 'S' this substitution replaces the names of variables by their values. Thus

```
echo $argv
```

when placed in a command script would cause the current value of the variable *argv* to be echoed to the output of the shell script. It is an error for *argv* to be unset at this point.

A number of notations are provided for accessing components and attributes of variables. The notation

```
$?name
```

expands to '1' if name is *set* or to '0' if name is not *set*. It is the fundamental mechanism used

for checking whether particular variables have been assigned values. All other forms of reference to undefined variables cause errors.

The notation

`S#name`

expands to the number of elements in the variable *name*. Thus

```
% set argv=(a b c)
% echo S?argv
1
% echo S#argv
3
% unset argv
% echo S?argv
0
% echo Sargv
Undefined variable: argv.
%
```

It is also possible to access the components of a variable which has several values. Thus

`Sargv[1]`

gives the first component of *argv* or in the example above 'a'. Similarly

`Sargv[S#argv]`

would give 'c', and

`Sargv[1-2]`

would give 'a b'. Other notations useful in shell scripts are

`Sn`

where *n* is an integer as a shorthand for

`Sargv[n]`

the *n*th parameter and

`S*`

which is a shorthand for

`Sargv`

The form

`$$`

expands to the process number of the current shell. Since this process number is unique in the system it can be used in generation of unique temporary file names. The form

`S<`

is quite special and is replaced by the next line of input read from the shell's standard input (not the script it is reading). This is useful for writing shell scripts that are interactive, reading commands from the terminal, or even writing a shell script that acts as a filter, reading lines from its input file. Thus the sequence

```
echo `yes or no?\c`
set a=(S<)
```

would write out the prompt 'yes or no?' without a newline and then read the answer into the

variable 'a'. In this case 'S#a' would be '0' if either a blank line or end-of-file (␣) was typed.

One minor difference between 'S*n*' and 'Sargv[*n*]' should be noted here. The form 'Sargv[*n*]' will yield an error if *n* is not in the range '1-S#argv' while 'S*n*' will never yield an out of range subscript error. This is for compatibility with the way older shells handled parameters.

Another important point is that it is never an error to give a subrange of the form 'm-n': if there are less than *n* components of the given variable then no words are substituted. A range of the form 'm-n' likewise returns an empty vector without giving an error when *m* exceeds the number of elements of the given variable, provided the subscript *n* is in range.

3.5. Expressions

In order for interesting shell scripts to be constructed it must be possible to evaluate expressions in the shell based on the values of variables. In fact, all the arithmetic operations of the language C are available in the shell with the same precedence that they have in C. In particular, the operations '=' and '!=' compare strings and the operators '&&' and '||' implement the boolean and/or operations. The special operators '=~' and '!~' are similar to '=' and '!=' except that the string on the right side can have pattern matching characters (like '.', '? or []) and the test is whether the string on the left matches the pattern on the right.

The shell also allows file enquiries of the form

-? filename

where '?' is replace by a number of single characters. For instance the expression primitive

-e filename

tell whether the file 'filename' exists. Other primitives test for read, write and execute access to the file, whether it is a directory, or has non-zero length.

It is possible to test whether a command terminates normally, by a primitive of the form '{ command }' which returns true, i.e. '1' if the command succeeds exiting normally with exit status 0, or '0' if the command terminates abnormally or with exit status non-zero. If more detailed information about the execution status of a command is required, it can be executed and the variable 'Sstatus' examined in the next command. Since 'Sstatus' is set by every command, it is very transient. It can be saved if it is inconvenient to use it only in the single immediately following command.

For a full list of expression components available see the manual section for the shell.

3.6. Sample shell script

A sample shell script which makes use of the expression mechanism of the shell and some of its control structure follows:

```
% cat copyc
#
# Copyc copies those C programs in the specified list
# to the directory ~/backup if they differ from the files
# already in ~/backup
#
set noglob
foreach i (Sargv)

    if ($i != *.c) continue # not a .c file so do nothing

    if (! -r ~/backup/$i:t) then
        echo $i:t not in backup... not cp'ed
        continue
    endif

    cmp -s $i ~/backup/$i:t # to set Sstatus

    if (Sstatus != 0) then
        echo new backup of $i
        cp $i ~/backup/$i:t
    endif
end
```

This script makes use of the *foreach* command, which causes the shell to execute the commands between the *foreach* and the matching *end* for each of the values given between '(' and ')' with the named variable, in this case 'i' set to successive values in the list. Within this loop we may use the command *break* to stop executing the loop and *continue* to prematurely terminate one iteration and begin the next. After the *foreach* loop the iteration variable (i in this case) has the value at the last iteration.

We set the variable *noglob* here to prevent filename expansion of the members of *argv*. This is a good idea, in general, if the arguments to a shell script are filenames which have already been expanded or if the arguments may contain filename expansion metacharacters. It is also possible to quote each use of a '\$' variable expansion, but this is harder and less reliable.

The other control construct used here is a statement of the form

```
if ( expression ) then
    command
    ...
endif
```

The placement of the keywords here is not flexible due to the current implementation of the shell.†

†The following two formats are not currently acceptable to the shell:

```
if ( expression )          # Won't work!
then
    command
    ...
endif
```

and

```
if ( expression ) then command endif          # Won't work
```

The shell does have another form of the if statement of the form

```
if ( expression ) command
```

which can be written

```
if ( expression ) \  
    command
```

Here we have escaped the newline for the sake of appearance. The command must not involve '|', '&' or ';' and must not be another control command. The second form requires the final '\ ' to immediately precede the end-of-line.

The more general *if* statements above also admit a sequence of *else-if* pairs followed by a single *else* and an *endif*, e.g.:

```
if ( expression ) then  
    commands  
else if (expression ) then  
    commands  
...  
  
else  
    commands  
endif
```

Another important mechanism used in shell scripts is the ':' modifier. We can use the modifier ':r' here to extract a root of a filename or ':e' to extract the *extension*. Thus if the variable *i* has the value '/mnt/foo.bar' then

```
% echo $i $i:r $i:e  
/mnt/foo.bar /mnt/foo bar  
%
```

shows how the ':r' modifier strips off the trailing '.bar' and the ':e' modifier leaves only the 'bar'. Other modifiers will take off the last component of a pathname leaving the head ':h' or all but the last component of a pathname leaving the tail ':t'. These modifiers are fully described in the *cs*h manual pages in the programmers manual. It is also possible to use the *command substitution* mechanism described in the next major section to perform modifications on strings to then reenter the shells environment. Since each usage of this mechanism involves the creation of a new process, it is much more expensive to use than the ':' modification mechanism. # Finally, we note that the character '#' lexically introduces a shell comment in shell scripts (but not from the terminal). All subsequent characters on the input line after a '#' are discarded by the shell. This character can be quoted using '"' or '\ ' to place it in an argument word.

#It is also important to note that the current implementation of the shell limits the number of ':' modifiers on a '\$' substitution to 1. Thus

```
% echo $i $i:h:t  
/a/b/c /a/b:t  
%
```

does not do what one would expect.

3.7. Other control structures

The shell also has control structures *while* and *switch* similar to those of C. These take the forms

```
while ( expression )
    commands
end
```

and

```
switch ( word )

case str1:
    commands
    breaksw
```

...

```
case strn:
    commands
    breaksw
```

```
default:
    commands
    breaksw
```

```
endsw
```

For details see the manual section for *cs**h*. C programmers should note that we use *breaksw* to exit from a *switch* while *break* exits a *while* or *foreach* loop. A common mistake to make in *cs**h* scripts is to use *break* rather than *breaksw* in switches.

Finally, *cs**h* allows a *goto* statement, with labels looking like they do in C, i.e.:

```
loop:
    commands
    goto loop
```

3.8. Supplying input to commands

Commands run from shell scripts receive by default the standard input of the shell which is running the script. This is different from previous shells running under UNIX. It allows shell scripts to fully participate in pipelines, but mandates extra notation for commands which are to take inline data.

Thus we need a metanotation for supplying inline data to commands in shell scripts. As an example, consider this script which runs the editor to delete leading blanks from the lines in each argument file

```
% cat deblank
# deblank -- remove leading blanks
foreach i (Sargv)
ed - Si << 'EOF'
1.\$s/[ ]*//
w
q
'EOF'
end
%
```

The notation '<< 'EOF'' means that the standard input for the *ed* command is to come from the text in the shell script file up to the next line consisting of exactly 'EOF'. The fact that the 'EOF' is enclosed in "" characters, i.e. quoted, causes the shell to not perform variable substitution on the intervening lines. In general, if any part of the word following the '<<' which the shell uses to terminate the text to be given to the command is quoted then these substitutions will not be performed. In this case since we used the form '1,S' in our editor script we needed to insure that this 'S' was not variable substituted. We could also have insured this by preceding the 'S' here with a '\', i.e.:

```
1.\$s/[ ]*//
```

but quoting the 'EOF' terminator is a more reliable way of achieving the same thing.

3.9. Catching interrupts

If our shell script creates temporary files, we may wish to catch interruptions of the shell script so that we can clean up these files. We can then do

```
onintr label
```

where *label* is a label in our program. If an interrupt is received the shell will do a 'goto label' and we can remove the temporary files and then do an *exit* command (which is built in to the shell) to exit from the shell script. If we wish to exit with a non-zero status we can do

```
exit(1)
```

e.g. to exit with status '1'.

3.10. What else?

There are other features of the shell useful to writers of shell procedures. The *verbose* and *echo* options and the related *-v* and *-x* command line options can be used to help trace the actions of the shell. The *-n* option causes the shell only to read commands and not to execute them and may sometimes be of use.

One other thing to note is that *csh* will not execute shell scripts which do not begin with the character '#', that is shell scripts that do not begin with a comment. Similarly, the '/bin/sh' on your system may well defer to 'csh' to interpret shell scripts which begin with '#'. This allows shell scripts for both shells to live in harmony.

There is also another quotation mechanism using "" which allows only some of the expansion mechanisms we have so far discussed to occur on the quoted string and serves to make this string into a single word as "" does.

4. Other, less commonly used, shell features

4.1. Loops at the terminal; variables as vectors

It is occasionally useful to use the *foreach* control structure at the terminal to aid in performing a number of similar commands. For instance, there were at one point three shells in use on the Cory UNIX system at Cory Hall, `"/bin/sh"`, `"/bin/nsh"`, and `"/bin/csh"`. To count the number of persons using each shell one could have issued the commands

```
% grep -c csh$ /etc/passwd
27
% grep -c nsh$ /etc/passwd
128
% grep -c -v sh$ /etc/passwd
430
%
```

Since these commands are very similar we can use *foreach* to do this more easily.

```
% foreach i ('sh$' 'csh$' '-v sh$')
? grep -c $i /etc/passwd
? end
27
128
430
%
```

Note here that the shell prompts for input with `'?'` when reading the body of the loop.

Very useful with loops are variables which contain lists of filenames or other words. You can, for example, do

```
% set a=('ls')
% echo $a
csh.n csh.rm
% ls
csh.n
csh.rm
% echo $#a
2
%
```

The *set* command here gave the variable *a* a list of all the filenames in the current directory as value. We can then iterate over these names to perform any chosen function.

The output of a command within `""` characters is converted by the shell to a list of words. You can also place the `""` quoted string within `""` characters to take each (non-empty) line as a component of the variable; preventing the lines from being split into words at blanks and tabs. A modifier `'x'` exists which can be used later to expand each component of the variable into another variable splitting it into separate words at embedded blanks and tabs.

4.2. Braces { ... } in argument expansion

Another form of filename expansion, alluded to before involves the characters `'{'` and `'}'`. These characters specify that the contained strings, separated by `','` are to be consecutively substituted into the containing characters and the results expanded left to right. Thus

```
A(str1,str2,....strn)B
```

expands to

Astr1B Astr2B ... AstrnB

This expansion occurs before the other filename expansions, and may be applied recursively (i.e. nested). The results of each expanded string are sorted separately, left to right order being preserved. The resulting filenames are not required to exist if no other expansion mechanisms are used. This means that this mechanism can be used to generate arguments which are not filenames, but which have common parts.

A typical use of this would be

```
mkdir ~/ {hdrs,retrofit,csH}
```

to make subdirectories 'hdrs', 'retrofit' and 'csH' in your home directory. This mechanism is most useful when the common prefix is longer than in this example, i.e.

```
chown root /usr/{ucb/{ex,edit},lib/{ex??.*,how_ex}}
```

4.3. Command substitution

A command enclosed in `` characters is replaced, just before filenames are expanded, by the output from that command. Thus it is possible to do

```
set pwd=`pwd`
```

to save the current directory in the variable *pwd* or to do

```
ex `grep -l TRACE *.c`
```

to run the editor *ex* supplying as arguments those files whose names end in '.c' which have the string 'TRACE' in them.*

4.4. Other details not covered here

In particular circumstances it may be necessary to know the exact nature and order of different substitutions performed by the shell. The exact meaning of certain combinations of quotations is also occasionally important. These are detailed fully in its manual section.

The shell has a number of command line option flags mostly of use in writing UNIX programs, and debugging shell scripts. See the shells manual section for a list of these options.

*Command expansion also occurs in input redirected with '<<' and within `` quotations. Refer to the shell manual section for full details.

Appendix — Special characters

The following table lists the special characters of *cs*h and the UNIX system, giving for each the section(s) in which it is discussed. A number of these characters also have special meaning in expressions. See the *cs*h manual section for a complete list.

Syntactic metacharacters

;	2.4	separates commands to be executed sequentially
	1.5	separates commands in a pipeline
()	2.2.3.6	brackets expressions and variable values
&	2.5	follows commands to be executed without waiting for completion

Filename metacharacters

/	1.6	separates components of a file's pathname
?	1.6	expansion character matching any single character
*	1.6	expansion character matching any sequence of characters
[]	1.6	expansion sequence matching any single character from a set
~	1.6	used at the beginning of a filename to indicate home directories
{ }	4.2	used to specify groups of arguments with common parts

Quotation metacharacters

\	1.7	prevents meta-meaning of following single character
'	1.7	prevents meta-meaning of a group of characters
"	4.3	like ', but allows variable and command expansion

Input/output metacharacters

<	1.5	indicates redirected input
>	1.3	indicates redirected output

Expansion/substitution metacharacters

\$	3.4	indicates variable substitution
!	2.3	indicates history substitution
:	3.6	precedes substitution modifiers
↑	2.3	used in special forms of history substitution
`	4.3	indicates command substitution

Other metacharacters

#	1.3.3.6	begins scratch file names; indicates shell comments
-	1.2	prefixes option (flag) arguments to commands
%	2.6	prefixes job name specifications

Glossary

This glossary lists the most important terms introduced in the introduction to the shell and gives references to sections of the shell document for further information about them. References of the form 'pr (1)' indicate that the command *pr* is in the UNIX programmer's manual in section 1. You can get an online copy of its manual page by doing

```
man 1 pr
```

References of the form (2.5) indicate that more information can be found in section 2.5 of this manual.

- . Your current directory has the name '.' as well as the name printed by the command *pwd*; see also *dirs*. The current directory '.' is usually the first *component* of the search path contained in the variable *path*, thus commands which are in '.' are found first (2.2). The character '.' is also used in separating *components* of filenames (1.6). The character '.' at the beginning of a *component* of a *pathname* is treated specially and not matched by the *filename expansion* meta-characters '?', '*', and '[' ']' pairs (1.6).
- .. Each directory has a file '..' in it which is a reference to its parent directory. After changing into the directory with *chdir*, i.e.

```
chdir paper
```

you can return to the parent directory by doing

```
chdir ..
```

The current directory is printed by *pwd* (2.7).
- a.out Compilers which create executable images create them, by default, in the file *a.out*. for historical reasons (2.3).
- absolute pathname A *pathname* which begins with a '/' is *absolute* since it specifies the *path* of directories from the beginning of the entire directory system — called the *root* directory. *Pathnames* which are not *absolute* are called *relative* (see definition of *relative pathname*) (1.6).
- alias An *alias* specifies a shorter or different name for a UNIX command, or a transformation on a command to be performed in the shell. The shell has a command *alias* which establishes *aliases* and can print their current values. The command *unalias* is used to remove *aliases* (2.4).
- argument Commands in UNIX receive a list of *argument* words. Thus the command

```
echo a b c
```

consists of the *command name* 'echo' and three *argument* words 'a', 'b' and 'c'. The set of *arguments* after the *command name* is said to be the *argument list* of the command (1.1).
- argv The list of arguments to a command written in the shell language (a shell script or shell procedure) is stored in a variable called *argv* within the shell. This name is taken from the conventional name in the C programming language (3.4).
- background Commands started without waiting for them to complete are called *background* commands (2.6).
- base A filename is sometimes thought of as consisting of a *base* part, before any '.' character, and an *extension* — the part after the '.'. See *filename* and *extension* (1.6)

- bg** The *bg* command causes a *suspended* job to continue execution in the *background* (2.6).
- bin** A directory containing binaries of programs and shell scripts to be executed is typically called a *bin* directory. The standard system *bin* directories are */bin* containing the most heavily used commands and */usr/bin* which contains most other user programs. Programs developed at UC Berkeley live in */usr/ucb*, while locally written programs live in */usr/local*. Games are kept in the directory */usr/games*. You can place binaries in any directory. If you wish to execute them often, the name of the directories should be a *component* of the variable *path*.
- break** *Break* is a builtin command used to exit from loops within the control structure of the shell (3.7).
- breaksw** The *breaksw* builtin command is used to exit from a *switch* control structure. like a *break* exits from loops (3.7).
- builtin** A command executed directly by the shell is called a *builtin* command. Most commands in UNIX are not built into the shell, but rather exist as files in *bin* directories. These commands are accessible because the directories in which they reside are named in the *path* variable.
- case** A *case* command is used as a label in a *switch* statement in the shell's control structure, similar to that of the language C. Details are given in the shell documentation *'csh(1)'* (3.7).
- cat** The *cat* program catenates a list of specified files on the *standard output*. It is usually used to look at the contents of a single file on the terminal, to *'cat a file'* (1.8, 2.3).
- cd** The *cd* command is used to change the *working directory*. With no arguments, *cd* changes your *working directory* to be your *home* directory (2.4, 2.7).
- chdir** The *chdir* command is a synonym for *cd*. *Cd* is usually used because it is easier to type.
- chsh** The *chsh* command is used to change the shell which you use on UNIX. By default, you use an different version of the shell which resides in */bin/sh*. You can change your shell to */bin/csh* by doing
- ```
chsh your-login-name /bin/csh
```
- Thus I would do
- ```
chsh bill /bin/csh
```
- It is only necessary to do this once. The next time you log in to UNIX after doing this command, you will be using *csh* rather than the shell in */bin/sh* (1.9).
- cmp** *Cmp* is a program which compares files. It is usually used on binary files, or to see if two files are identical (3.6). For comparing text files the program *diff*, described in *'diff (1)'* is used.
- command** A function performed by the system, either by the shell (a builtin *command*) or by a program residing in a file in a directory within the UNIX system, is called a *command* (1.1).
- command name** When a command is issued, it consists of a *command name*, which is the first word of the command, followed by arguments. The convention on UNIX is that the first word of a command names the function to be performed (1.1).

- command substitution** The replacement of a command enclosed in "" characters by the text output by that command is called *command substitution* (4.3).
- component** A part of a *pathname* between '/' characters is called a *component* of that *pathname*. A variable which has multiple strings as value is said to have several *components*; each string is a *component* of the variable.
- continue** A builtin command which causes execution of the enclosing *foreach* or *while* loop to cycle prematurely. Similar to the *continue* command in the programming language C (3.6).
- control-** Certain special characters, called *control* characters, are produced by holding down the CONTROL key on your terminal and simultaneously pressing another character, much like the SHIFT key is used to produce upper case characters. Thus *control-c* is produced by holding down the CONTROL key while pressing the 'c' key. Usually UNIX prints an up-arrow (↑) followed by the corresponding letter when you type a *control* character (e.g. '↑C' for *control-c* (1.8).
- core dump** When a program terminates abnormally, the system places an image of its current state in a file named 'core'. This *core dump* can be examined with the system debugger 'adb(1)' or 'sdb(1)' in order to determine what went wrong with the program (1.8). If the shell produces a message of the form
- Illegal instruction (core dumped)
- (where 'Illegal instruction' is only one of several possible messages), you should report this to the author of the program or a system administrator, saving the 'core' file.
- cp** The *cp* (copy) program is used to copy the contents of one file into another file. It is one of the most commonly used UNIX commands (1.6).
- cs** The name of the shell program that this document describes.
- .cshrc** The file *.cshrc* in your *home* directory is read by each shell as it begins execution. It is usually used to change the setting of the variable *path* and to set *alias* parameters which are to take effect globally (2.1).
- cwd** The *cwd* variable in the shell holds the *absolute pathname* of the current *working directory*. It is changed by the shell whenever your current *working directory* changes and should not be changed otherwise (2.2).
- date** The *date* command prints the current date and time (1.3).
- debugging** *Debugging* is the process of correcting mistakes in programs and shell scripts. The shell has several options and variables which may be used to aid in shell *debugging* (4.4).
- default:** The label *default:* is used within shell *switch* statements, as it is in the C language to label the code to be executed if none of the *case* labels matches the value switched on (3.7).
- DELETE** The DELETE or RUBOUT key on the terminal normally causes an interrupt to be sent to the current job. Many users change the interrupt character to be ↑C.
- detached** A command that continues running in the *background* after you logout is said to be *detached*.
- diagnostic** An error message produced by a program is often referred to as a *diagnostic*. Most error messages are not written to the *standard output*, since that is often directed away from the terminal (1.3, 1.5). Error messages are instead written to the *diagnostic output* which may be directed away from the terminal, but usually is not. Thus *diagnostics* will usually appear on the terminal (2.5).

- directory** A structure which contains files. At any time you are in one particular *directory* whose names can be printed by the command *pwd*. The *chdir* command will change you to another *directory*, and make the files in that *directory* visible. The *directory* in which you are when you first login is your *home* directory (1.1, 2.7).
- directory stack** The shell saves the names of previous *working directories* in the *directory stack* when you change your current *working directory* via the *pushd* command. The *directory stack* can be printed by using the *dirc* command, which includes your current *working directory* as the first directory name on the left (2.7).
- dirc** The *dirc* command prints the shell's *directory stack* (2.7).
- du** The *du* command is a program (described in 'du(1)') which prints the number of disk blocks is all directories below and including your current *working directory* (2.6).
- echo** The *echo* command prints its arguments (1.6, 3.6).
- else** The *else* command is part of the 'if-then-else-endif' control command construct (3.6).
- endif** If an *if* statement is ended with the word *then*, all lines following the *if* up to a line starting with the word *endif* or *else* are executed if the condition between parentheses after the *if* is true (3.6).
- EOF** An *end-of-file* is generated by the terminal by a control-d, and whenever a command reads to the end of a file which it has been given as input. Commands receiving input from a *pipe* receive an *end-of-file* when the command sending them input completes. Most commands terminate when they receive an *end-of-file*. The shell has an option to ignore *end-of-file* from a terminal input which may help you keep from logging out accidentally by typing too many control-d's (1.1, 1.8, 3.8).
- escape** A character '\ ' used to prevent the special meaning of a metacharacter is said to *escape* the character from its special meaning. Thus
- ```
echo *
```
- will echo the character '\*' while just
- ```
echo *
```
- will echo the names of the file in the current directory. In this example, \ *escapes* '*' (1.7). There is also a non-printing character called *escape*, usually labelled ESC or ALTMODE on terminal keyboards. Some older UNIX systems use this character to indicate that output is to be *suspended*. Most systems use control-s to stop the output and control-q to start it.
- /etc/passwd** This file contains information about the accounts currently on the system. It consists of a line for each account with fields separated by ':' characters (1.8). You can look at this file by saying
- ```
cat /etc/passwd
```
- The commands *finger* and *grep* are often used to search for information in this file. See 'finger(1)', 'passwd(5)', and 'grep(1)' for more details.
- exit** The *exit* command is used to force termination of a shell script, and is built into the shell (3.9).
- exit status** A command which discovers a problem may reflect this back to the command (such as a shell) which invoked (executed) it. It does this by returning a non-zero number as its *exit status*, a status of zero being considered 'normal termination'. The *exit* command can be used to force a shell command script

- to give a non-zero *exit status* (3.6).
- expansion** The replacement of strings in the shell input which contain metacharacters by other strings is referred to as the process of *expansion*. Thus the replacement of the word "\*" by a sorted list of files in the current directory is a 'filename expansion'. Similarly the replacement of the characters "!" by the text of the last command is a 'history expansion'. *Expansions* are also referred to as *substitutions* (1.6, 3.4, 4.2).
- expressions** *Expressions* are used in the shell to control the conditional structures used in the writing of shell scripts and in calculating values for these scripts. The operators available in shell *expressions* are those of the language C (3.5).
- extension** Filenames often consist of a *base* name and an *extension* separated by the character '.'. By convention, groups of related files often share the same *root* name. Thus if 'prog.c' were a C program, then the object file for this program would be stored in 'prog.o'. Similarly a paper written with the '-me' nroff macro package might be stored in 'paper.me' while a formatted version of this paper might be kept in 'paper.out' and a list of spelling errors in 'paper.errs' (1.6).
- fg** The *job control* command *fg* is used to run a *background* or *suspended* job in the *foreground* (1.8, 2.6).
- filename** Each file in UNIX has a name consisting of up to 14 characters and not including the character '/' which is used in *pathname* building. Most *filenames* do not begin with the character '.', and contain only letters and digits with perhaps a '.' separating the *base* portion of the *filename* from an *extension* (1.6).
- filename expansion** *Filename expansion* uses the metacharacters "\*", "?" and "[" and "]" to provide a convenient mechanism for naming files. Using *filename expansion* it is easy to name all the files in the current directory, or all files which have a common *root* name. Other *filename expansion* mechanisms use the metacharacter "" and allow files in other users' directories to be named easily (1.6, 4.2).
- flag** Many UNIX commands accept arguments which are not the names of files or other users but are used to modify the action of the commands. These are referred to as *flag* options, and by convention consist of one or more letters preceded by the character '-'. Thus the *ls* (list files) command has an option '-s' to list the sizes of files. This is specified
- ```
ls -s
```
- foreach** The *foreach* command is used in shell scripts and at the terminal to specify repetition of a sequence of commands while the value of a certain shell variable ranges through a specified list (3.6, 4.1).
- foreground** When commands are executing in the normal way such that the shell is waiting for them to finish before prompting for another command they are said to be *foreground jobs* or *running in the foreground*. This is as opposed to *background*. *Foreground* jobs can be stopped by signals from the terminal caused by typing different control characters at the keyboard (1.8, 2.6).
- goto** The shell has a command *goto* used in shell scripts to transfer control to a given label (3.7).
- grep** The *grep* command searches through a list of argument files for a specified string. Thus
- ```
grep bill /etc/passwd
```
- will print each line in the file */etc/passwd* which contains the string 'bill'.



- Actually, *grep* scans for *regular expressions* in the sense of the editors 'ed(1)' and 'ex(1)'. *Grep* stands for 'globally find *regular expression* and print' (2.4).
- head** The *head* command prints the first few lines of one or more files. If you have a bunch of files containing text which you are wondering about it is sometimes useful to run *head* with these files as arguments. This will usually show enough of what is in these files to let you decide which you are interested in (1.5).  
*Head* is also used to describe the part of a *pathname* before and including the last '/' character. The *tail* of a *pathname* is the part after the last '/'. The ':h' and ':t' modifiers allow the *head* or *tail* of a *pathname* stored in a shell variable to be used (3.6).
- history** The *history* mechanism of the shell allows previous commands to be repeated, possibly after modification to correct typing mistakes or to change the meaning of the command. The shell has a *history list* where these commands are kept, and a *history* variable which controls how large this list is (2.3).
- home directory** Each user has a *home directory*, which is given in your entry in the password file, */etc/passwd*. This is the directory which you are placed in when you first login. The *cd* or *chdir* command with no arguments takes you back to this directory, whose name is recorded in the shell variable *home*. You can also access the *home directories* of other users in forming filenames using a *filename expansion* notation and the character '~' (1.6).
- if** A conditional command within the shell, the *if* command is used in shell command scripts to make decisions about what course of action to take next (3.6).
- ignoreeof** Normally, your shell will exit, printing 'logout' if you type a control-d at a prompt of '% '. This is the way you usually log off the system. You can set the *ignoreeof* variable if you wish in your *.login* file and then use the command *logout* to logout. This is useful if you sometimes accidentally type too many control-d characters, logging yourself off (2.2).
- input** Many commands on UNIX take information from the terminal or from files which they then act on. This information is called *input*. Commands normally read for *input* from their *standard input* which is, by default, the terminal. This *standard input* can be redirected from a file using a shell metanotation with the character '<'. Many commands will also read from a file specified as argument. Commands placed in *pipelines* will read from the output of the previous command in the *pipeline*. The leftmost command in a *pipeline* reads from the terminal if you neither redirect its *input* nor give it a filename to use as *standard input*. Special mechanisms exist for supplying input to commands in shell scripts (1.5, 3.8).
- interrupt** An *interrupt* is a signal to a program that is generated by hitting the RUBOUT or DELETE key (although users can and often do change the interrupt character, usually to ↑C). It causes most programs to stop execution. Certain programs, such as the shell and the editors, handle an *interrupt* in special ways, usually by stopping what they are doing and prompting for another command. While the shell is executing another command and waiting for it to finish, the shell does not listen to *interrupts*. The shell often wakes up when you hit *interrupt* because many commands die when they receive an *interrupt* (1.8, 3.9).
- job** One or more commands typed on the same input line separated by '|' or ':' characters are run together and are called a *job*. Simple commands run by themselves without any '|' or ':' characters are the simplest *jobs*. *Jobs* are classified as *foreground*, *background*, or *suspended* (2.6).

- job control** The builtin functions that control the execution of jobs are called *job control* commands. These are *bg, fg, stop, kill* (2.6).
- job number** When each job is started it is assigned a small number called a *job number* which is printed next to the job in the output of the *jobs* command. This number, preceded by a '%' character, can be used as an argument to *job control* commands to indicate a specific job (2.6).
- jobs** The *jobs* command prints a table showing jobs that are either running in the *background* or are *suspended* (2.6).
- kill** A command which sends a signal to a job causing it to terminate (2.6).
- .login** The file *.login* in your *home* directory is read by the shell each time you login to UNIX and the commands there are executed. There are a number of commands which are usefully placed here, especially *set* commands to the shell itself (2.1).
- login shell** The shell that is started on your terminal when you login is called your *login shell*. It is different from other shells which you may run (e.g. on shell scripts) in that it reads the *.login* file before reading commands from the terminal and it reads the *.logout* file after you logout (2.1).
- logout** The *logout* command causes a login shell to exit. Normally, a login shell will exit when you hit control-d generating an *end-of-file*, but if you have set *ignoreeof* in your *.login* file then this will not work and you must use *logout* to log off the UNIX system (2.8).
- .logout** When you log off of UNIX the shell will execute commands from the file *.logout* in your *home* directory after it prints 'logout'.
- lpr** The command *lpr* is the line printer daemon. The standard input of *lpr* spooled and printed on the UNIX line printer. You can also give *lpr* a list of filenames as arguments to be printed. It is most common to use *lpr* as the last component of a *pipeline* (2.3).
- ls** The *ls* (list files) command is one of the most commonly used UNIX commands. With no argument filenames it prints the names of the files in the current directory. It has a number of useful *flag* arguments, and can also be given the names of directories as arguments, in which case it lists the names of the files in these directories (1.2).
- mail** The *mail* program is used to send and receive messages from other UNIX users (1.1, 2.1).
- make** The *make* command is used to maintain one or more related files and to organize functions to be performed on these files. In many ways *make* is easier to use, and more helpful than shell command scripts (3.2).
- makefile** The file containing commands for *make* is called *makefile* (3.2).
- manual** The *manual* often referred to is the 'UNIX programmer's manual'. It contains a number of sections and a description of each UNIX program. An online version of the *manual* is accessible through the *man* command. Its documentation can be obtained online via
- man man
- metacharacter** Many characters which are neither letters nor digits have special meaning either to the shell or to UNIX. These characters are called *metacharacters*. If it is necessary to place these characters in arguments to commands without them having their special meaning then they must be *quoted*. An example of a *metacharacter* is the character '>' which is used to indicate placement of output

into a file. For the purposes of the *history* mechanism, most unquoted *meta-characters* form separate words (1.4). The appendix to this user's manual lists the *metacharacters* in groups by their function.

- mkdir** The *mkdir* command is used to create a new directory.
- modifier** Substitutions with the *history* mechanism, keyed by the character '*!*' or of variables using the metacharacter '*\$*', are often subjected to modifications, indicated by placing the character '*:*' after the substitution and following this with the *modifier* itself. The *command substitution* mechanism can also be used to perform modification in a similar way, but this notation is less clear (3.6).
- more** The program *more* writes a file on your terminal allowing you to control how much text is displayed at a time. *More* can move through the file screenful by screenful, line by line, search forward for a string, or start again at the beginning of the file. It is generally the easiest way of viewing a file (1.8).
- noclobber** The shell has a variable *noclobber* which may be set in the file *.login* to prevent accidental destruction of files by the '*>*' output redirection metasyntax of the shell (2.2, 2.5).
- noglob** The shell variable *noglob* is set to suppress the *filename expansion* of arguments containing the metacharacters '*~*', '*~\**', '*?*', '*[*' and '*]*' (3.6).
- notify** The *notify* command tells the shell to report on the termination of a specific *background job* at the exact time it occurs as opposed to waiting until just before the next prompt to report the termination. The *notify* variable, if set, causes the shell to always report the termination of *background jobs* exactly when they occur (2.6).
- onintr** The *onintr* command is built into the shell and is used to control the action of a shell command script when an *interrupt* signal is received (3.9).
- output** Many commands in UNIX result in some lines of text which are called their *output*. This *output* is usually placed on what is known as the *standard output* which is normally connected to the user's terminal. The shell has a syntax using the metacharacter '*>*' for redirecting the *standard output* of a command to a file (1.3). Using the *pipe* mechanism and the metacharacter '*|*' it is also possible for the *standard output* of one command to become the *standard input* of another command (1.5). Certain commands such as the line printer daemon *p* do not place their results on the *standard output* but rather in more useful places such as on the line printer (2.3). Similarly the *write* command places its output on another user's terminal rather than its *standard output* (2.3). Commands also have a *diagnostic output* where they write their error messages. Normally these go to the terminal even if the *standard output* has been sent to a file or another command, but it is possible to direct error diagnostics along with *standard output* using a special metanotation (2.5).
- pushd** The *pushd* command, which means 'push directory', changes the shell's *working directory* and also remembers the current *working directory* before the change is made, allowing you to return to the same directory via the *popd* command later without retyping its name (2.7).
- path** The shell has a variable *path* which gives the names of the directories in which it searches for the commands which it is given. It always checks first to see if the command it is given is built into the shell. If it is, then it need not search for the command as it can do it internally. If the command is not builtin, then the shell searches for a file with the name given in each of the directories in the *path* variable, left to right. Since the normal definition of the *path* variable is

**path** (./usr/ucb/bin/usr/bin)

the shell normally looks in the current directory, and then in the standard system directories '/usr/ucb', '/bin' and '/usr/bin' for the named command (2.2). If the command cannot be found the shell will print an error diagnostic. Scripts of shell commands will be executed using another shell to interpret them if they have 'execute' permission set. This is normally true because a command of the form

`chmod 755 script`

was executed to turn this execute permission on (3.3). If you add new commands to a directory in the *path*, you should issue the command *rehash* (2.2).

- pathname** A list of names, separated by '/' characters, forms a *pathname*. Each *component*, between successive '/' characters, names a directory in which the next *component* file resides. *Pathnames* which begin with the character '/' are interpreted relative to the *root* directory in the filesystem. Other *pathnames* are interpreted relative to the current directory as reported by *pwd*. The last component of a *pathname* may name a directory, but usually names a file.
- pipeline** A group of commands which are connected together, the *standard output* of each connected to the *standard input* of the next, is called a *pipeline*. The *pipe* mechanism used to connect these commands is indicated by the shell meta-character '|' (1.5, 2.3).
- popd** The *popd* command changes the shell's *working directory* to the directory you most recently left using the *pushd* command. It returns to the directory without having to type its name, forgetting the name of the current *working directory* before doing so (2.7).
- port** The part of a computer system to which each terminal is connected is called a *port*. Usually the system has a fixed number of *ports*, some of which are connected to telephone lines for dial-up access, and some of which are permanently wired directly to specific terminals.
- pr** The *pr* command is used to prepare listings of the contents of files with headers giving the name of the file and the date and time at which the file was last modified (2.3).
- printenv** The *printenv* command is used to print the current setting of variables in the environment (2.8).
- process** An instance of a running program is called a *process* (2.6). UNIX assigns each *process* a unique number when it is started — called the *process number*. *Process numbers* can be used to stop individual *processes* using the *kill* or *stop* commands when the *processes* are part of a detached *background job*.
- program** Usually synonymous with *command*, a binary file or shell command script which performs a useful function is often called a *program*.
- programmer's manual** Also referred to as the *manual*. See the glossary entry for 'manual'.
- prompt** Many programs will print a *prompt* on the terminal when they expect input. Thus the editor 'ex(1)' will print a ':' when it expects input. The shell *prompts* for input with '%' and occasionally with '?' when reading commands from the terminal (1.1). The shell has a variable *prompt* which may be set to a different value to change the shell's main *prompt*. This is mostly used when debugging the shell (2.8).

- ps** The *ps* command is used to show the processes you are currently running. Each process is shown with its unique process number, an indication of the terminal name it is attached to, an indication of the state of the process (whether it is running, stopped, awaiting some event (sleeping), and whether it is swapped out), and the amount of CPU time it has used so far. The command is identified by printing some of the words used when it was invoked (2.6). Shells, such as the *cs*h you use to run the *ps* command, are not normally shown in the output.
- pwd** The *pwd* command prints the full *pathname* of the current *working directory*. The *dirc*s builtin command is usually a better and faster choice.
- quit** The *quit* signal, generated by a control-\, is used to terminate programs which are behaving unreasonably. It normally produces a core image file (1.8).
- quotation** The process by which metacharacters are prevented their special meaning, usually by using the character " in pairs, or by using the character \, is referred to as *quotation* (1.7).
- redirection** The routing of input or output from or to a file is known as *redirection* of input or output (1.3).
- rehash** The *rehash* command tells the shell to rebuild its internal table of which commands are found in which directories in your *path*. This is necessary when a new program is installed in one of these directories (2.8).
- relative pathname** A *pathname* which does not begin with a '/' is called a *relative pathname* since it is interpreted *relative* to the current *working directory*. The first *component* of such a *pathname* refers to some file or directory in the *working directory*, and subsequent *components* between '/' characters refer to directories below the *working directory*. *Pathnames* that are not *relative* are called *absolute pathnames* (1.6).
- repeat** The *repeat* command iterates another command a specified number of times.
- root** The directory that is at the top of the entire directory structure is called the *root* directory since it is the 'root' of the entire tree structure of directories. The name used in *pathnames* to indicate the *root* is '/'. *Pathnames* starting with '/' are said to be *absolute* since they start at the *root* directory. *Root* is also used as the part of a *pathname* that is left after removing the *extension*. See *filename* for a further explanation (1.6).
- RUBOUT** The RUBOUT or DELETE key sends an interrupt to the current job. Most interactive commands return to their command level upon receipt of an interrupt, while non-interactive commands usually terminate, returning control to the shell. Users often change interrupt to be generated by ^C rather than DELETE by using the *str*y command.
- scratch file** Files whose names begin with a '#' are referred to as *scratch files*, since they are automatically removed by the system after a couple of days of non-use, or more frequently if disk space becomes tight (1.3).
- script** Sequences of shell commands placed in a file are called shell command *scripts*. It is often possible to perform simple tasks using these *scripts* without writing a program in a language such as C, by using the shell to selectively run other programs (3.3, 3.10).
- set** The builtin *set* command is used to assign new values to shell variables and to show the values of the current variables. Many shell variables have special meaning to the shell itself. Thus by using the *set* command the behavior of the shell can be affected (2.1).

- setenv** Variables in the environment 'environ(5)' can be changed by using the *setenv* builtin command (2.8). The *printenv* command can be used to print the value of the variables in the environment.
- shell** A *shell* is a command language interpreter. It is possible to write and run your own *shell*, as *shells* are no different than any other programs as far as the system is concerned. This manual deals with the details of one particular *shell*, called *csh*.
- shell script** See *script* (3.3, 3.10).
- signal** A *signal* in UNIX is a short message that is sent to a running program which causes something to happen to that process. *Signals* are sent either by typing special *control* characters on the keyboard or by using the *kill* or *stop* commands (1.8, 2.6).
- sort** The *sort* program sorts a sequence of lines in ways that can be controlled by argument *flags* (1.5).
- source** The *source* command causes the shell to read commands from a specified file. It is most useful for reading files such as *.cshrc* after changing them (2.8).
- special character** See *metacharacters* and the appendix to this manual.
- standard** We refer often to the *standard input* and *standard output* of commands. See *input* and *output* (1.3, 3.8).
- status** A command normally returns a *status* when it finishes. By convention a *status* of zero indicates that the command succeeded. Commands may return non-zero *status* to indicate that some abnormal event has occurred. The shell variable *status* is set to the *status* returned by the last command. It is most useful in shell command scripts (3.6).
- stop** The *stop* command causes a *background* job to become *suspended* (2.6).
- string** A sequential group of characters taken together is called a *string*. *Strings* can contain any printable characters (2.2).
- stty** The *stty* program changes certain parameters inside UNIX which determine how your terminal is handled. See 'stty(1)' for a complete description (2.6).
- substitution** The shell implements a number of *substitutions* where sequences indicated by metacharacters are replaced by other sequences. Notable examples of this are history *substitution* keyed by the metacharacter '!' and variable *substitution* indicated by '\$'. We also refer to *substitutions* as *expansions* (3.4).
- suspended** A job becomes *suspended* after a STOP signal is sent to it, either by typing a *control-z* at the terminal (for *foreground* jobs) or by using the *stop* command (for *background* jobs). When *suspended*, a job temporarily stops running until it is restarted by either the *fg* or *bg* command (2.6).
- switch** The *switch* command of the shell allows the shell to select one of a number of sequences of commands based on an argument string. It is similar to the *switch* statement in the language C (3.7).
- termination** When a command which is being executed finishes we say it undergoes *termination* or *terminates*. Commands normally terminate when they read an *end-of-file* from their *standard input*. It is also possible to terminate commands by sending them an *interrupt* or *quit* signal (1.8). The *kill* program terminates specified jobs (2.6).
- then** The *then* command is part of the shell's 'if-then-else-endif' control construct used in command scripts (3.6).

- time** The *time* command can be used to measure the amount of CPU and real time consumed by a specified command as well as the amount of disk i/o, memory utilized, and number of page faults and swaps taken by the command (2.1, 2.8).
- tset** The *tset* program is used to set standard erase and kill characters and to tell the system what kind of terminal you are using. It is often invoked in a *.login* file (2.1).
- tty** The word *tty* is a historical abbreviation for 'teletype' which is frequently used in UNIX to indicate the *port* to which a given terminal is connected. The *tty* command will print the name of the *tty* or *port* to which your terminal is presently connected.
- unalias** The *unalias* command removes aliases (2.8).
- UNIX** UNIX is an operating system on which *cs*h runs. UNIX provides facilities which allow *cs*h to invoke other programs such as editors and text formatters which you may wish to use.
- unset** The *unset* command removes the definitions of shell variables (2.2, 2.8).
- variable expansion**  
See *variables* and *expansion* (2.2, 3.4).
- variables** *Variables* in *cs*h hold one or more strings as value. The most common use of *variables* is in controlling the behavior of the shell. See *path*, *noclobber*, and *ignoreeof* for examples. *Variables* such as *argv* are also used in writing shell programs (shell command scripts) (2.2).
- verbose** The *verbose* shell variable can be set to cause commands to be echoed after they are history expanded. This is often useful in debugging shell scripts. The *verbose* variable is set by the shell's *-v* command line option (3.10).
- wc** The *wc* program calculates the number of characters, words, and lines in the files whose names are given as arguments (2.6).
- while** The *while* builtin control construct is used in shell command scripts (3.7).
- word** A sequence of characters which forms an argument to a command is called a *word*. Many characters which are neither letters, digits, '-', '.', nor '/' form *words* all by themselves even if they are not surrounded by blanks. Any sequence of characters may be made into a *word* by surrounding it with "" characters except for the characters "" and '!' which require special treatment (1.1). This process of placing special characters in *words* without their special meaning is called *quoting*.
- working directory**  
At any given time you are in one particular directory, called your *working directory*. This directory's name is printed by the *pwd* command and the files listed by *ls* are the ones in this directory. You can change *working directories* using *chdir*.
- write** The *write* command is used to communicate with other users who are logged in to UNIX.





# An Introduction to Display Editing with Vi

*William Joy*

*Revised for versions 3.5/2.13 by*

*Mark Horton*

Computer Science Division  
Department of Electrical Engineering and Computer Science  
University of California, Berkeley  
Berkeley, Ca. 94720

## *ABSTRACT*

*Vi* (visual) is a display oriented interactive text editor. When using *vi* the screen of your terminal acts as a window into the file which you are editing. Changes which you make to the file are reflected in what you see.

Using *vi* you can insert new text any place in the file quite easily. Most of the commands to *vi* move the cursor around in the file. There are commands to move the cursor forward and backward in units of characters, words, sentences and paragraphs. A small set of operators, like *d* for delete and *c* for change, are combined with the motion commands to form operations such as delete word or change paragraph, in a simple and natural way. This regularity and the mnemonic assignment of commands to keys makes the editor command set easy to remember and to use.

*Vi* will work on a large number of display terminals, and new terminals are easily driven after editing a terminal description file. While it is advantageous to have an intelligent terminal which can locally insert and delete lines and characters from the display, the editor will function quite well on dumb terminals over slow phone lines. The editor makes allowance for the low bandwidth in these situations and uses smaller window sizes and different display updating algorithms to make best use of the limited speed available.

It is also possible to use the command set of *vi* on hardcopy terminals, storage tubes and "glass tty's" using a one line editing window; thus *vi*'s command set is available on all terminals. The full command set of the more traditional, line oriented editor *ex* is available within *vi*; it is quite simple to switch between the two modes of editing.

September 16, 1980

# An Introduction to Display Editing with Vi

*William Joy*

*Revised for versions 3.512.13 by  
Mark Horton*

Computer Science Division  
Department of Electrical Engineering and Computer Science  
University of California, Berkeley  
Berkeley, Ca. 94720

## 1. Getting started

This document provides a quick introduction to *vi*. (Pronounced *vee-eye*.) You should be running *vi* on a file you are familiar with while you are reading this. The first part of this document (sections 1 through 5) describes the basics of using *vi*. Some topics of special interest are presented in section 6, and some nitty-gritty details of how the editor functions are saved for section 7 to avoid cluttering the presentation here.

There is also a short appendix here, which gives for each character the special meanings which this character has in *vi*. Attached to this document should be a quick reference card. This card summarizes the commands of *vi* in a very compact format. You should have the card handy while you are learning *vi*.

### 1.1. Specifying terminal type

Before you can start *vi* you must tell the system what kind of terminal you are using. Here is a (necessarily incomplete) list of terminal type codes. If your terminal does not appear here, you should consult with one of the staff members on your system to find out the code for your terminal. If your terminal does not have a code, one can be assigned and a description for the terminal can be created.

| Code   | Full name                | Type        |
|--------|--------------------------|-------------|
| 2621   | Hewlett-Packard 2621A/P  | Intelligent |
| 2645   | Hewlett-Packard 264x     | Intelligent |
| act4   | Microterm ACT-IV         | Dumb        |
| act5   | Microterm ACT-V          | Dumb        |
| adm3a  | Lear Siegler ADM-3a      | Dumb        |
| adm31  | Lear Siegler ADM-31      | Intelligent |
| c100   | Human Design Concept 100 | Intelligent |
| dm1520 | Datamedia 1520           | Dumb        |
| dm2500 | Datamedia 2500           | Intelligent |
| dm3025 | Datamedia 3025           | Intelligent |
| fox    | Perkin-Elmer Fox         | Dumb        |
| h1500  | Hazeltine 1500           | Intelligent |
| h19    | Heathkit h19             | Intelligent |
| i100   | Infoton 100              | Intelligent |
| mime   | Imitating a smart act4   | Intelligent |

---

The financial support of an IBM Graduate Fellowship and the National Science Foundation under grants MCS74-07644-A03 and MCS78-07291 is gratefully acknowledged.

t1061  
vt52

Teleray 1061  
Dec VT-52

Intelligent  
Dumb

Suppose for example that you have a Hewlett-Packard HP2621A terminal. The code used by the system for this terminal is '2621'. In this case you can use one of the following commands to tell the system the type of your terminal:

```
% setenv TERM 2621
```

This command works with the shell *cs*h on both version 6 and 7 systems. If you are using the standard version 7 shell then you should give the commands

```
$ TERM=2621
$ export TERM
```

If you want to arrange to have your terminal type set up automatically when you log in, you can use the *tset* program. If you dial in on a *mime*, but often use hardwired ports, a typical line for your *.login* file (if you use *cs*h) would be

```
setenv TERM `tset - -d mime`
```

or for your *.profile* file (if you use *sh*)

```
TERM=`tset - -d mime`
```

*Tset* knows which terminals are hardwired to each port and needs only to be told that when you dial in you are probably on a *mime*. *Tset* is usually used to change the erase and kill characters, too.

## 1.2. Editing a file

After telling the system which kind of terminal you have, you should make a copy of a file you are familiar with, and run *vi* on this file, giving the command

```
% vi name
```

replacing *name* with the name of the copy file you just created. The screen should clear and the text of your file should appear on the screen. If something else happens refer to the footnote.†

## 1.3. The editor's copy: the buffer

The editor does not directly modify the file which you are editing. Rather, the editor makes a copy of this file, in a place called the *buffer*, and remembers the file's name. You do not affect the contents of the file unless and until you write the changes you make back into the original file.

---

† If you gave the system an incorrect terminal type code then the editor may have just made a mess out of your screen. This happens when it sends control codes for one kind of terminal to some other kind of terminal. In this case hit the keys :q (colon and the q key) and then hit the RETURN key. This should get you back to the command level interpreter. Figure out what you did wrong (ask someone else if necessary) and try again.

Another thing which can go wrong is that you typed the wrong file name and the editor just printed an error diagnostic. In this case you should follow the above procedure for getting out of the editor, and try again this time spelling the file name correctly.

If the editor doesn't seem to respond to the commands which you type here, try sending an interrupt to it by hitting the DEL or RUB key on your terminal, and then hitting the :q command again followed by a carriage return.

#### 1.4. Notational conventions

In our examples, input which must be typed as is will be presented in **bold face**. Text which should be replaced with appropriate input will be given in *italics*. We will represent special characters in SMALL CAPITALS.

#### 1.5. Arrow keys

The editor command set is independent of the terminal you are using. On most terminals with cursor positioning keys, these keys will also work within the editor. If you don't have cursor positioning keys, or even if you do, you can use the **h j k** and **l** keys as cursor positioning keys (these are labelled with arrows on an *adm3a*).\*

(Particular note for the HP2621: on this terminal the function keys must be *shifted* (ick) to send to the machine, otherwise they only act locally. Unshifted use will leave the cursor positioned incorrectly.)

#### 1.6. Special characters: ESC, CR and DEL

Several of these special characters are very important, so be sure to find them right now. Look on your keyboard for a key labelled ESC or ALT. It should be near the upper left corner of your terminal. Try hitting this key a few times. The editor will ring the bell to indicate that it is in a quiescent state.‡ Partially formed commands are cancelled by ESC, and when you insert text in the file you end the text insertion with ESC. This key is a fairly harmless one to hit, so you can just hit it if you don't know what is going on until the editor rings the bell.

The CR or RETURN key is important because it is used to terminate certain commands. It is usually at the right side of the keyboard, and is the same command used at the end of each shell command.

Another very useful key is the DEL or RUB key, which generates an interrupt, telling the editor to stop what it is doing. It is a forceful way of making the editor listen to you, or to return it to the quiescent state if you don't know or don't like what is going on. Try hitting the **'/'** key on your terminal. This key is used when you want to specify a string to be searched for. The cursor should now be positioned at the bottom line of the terminal after a **'/'** printed as a prompt. You can get the cursor back to the current position by hitting the DEL or RUB key; try this now.\* From now on we will simply refer to hitting the DEL or RUB key as "sending an interrupt."\*\*

The editor often echoes your commands on the last line of the terminal. If the cursor is on the first position of this last line, then the editor is performing a computation, such as computing a new position in the file after a search or running a command to reformat part of the buffer. When this is happening you can stop the editor by sending an interrupt.

#### 1.7. Getting out of the editor

After you have worked with this introduction for a while, and you wish to do something else, you can give the command **ZZ** to the editor. This will write the contents of the editor's buffer back into the file you are editing, if you made any changes, and then quit from the editor. You can also end an editor session by giving the command **:q!CR;**† this is a dangerous but occasionally essential command which ends the editor session and discards all your changes. You need to know about this command in case you change the editor's copy of a file you wish

\* As we will see later, **h** moves back to the left (like control-h which is a backspace), **j** moves down (in the same column), **k** moves up (in the same column), and **l** moves to the right.

‡ On smart terminals where it is possible, the editor will quietly flash the screen rather than ringing the bell.

\* Backspacing over the **'/'** will also cancel the search.

\*\* On some systems, this interruptibility comes at a price: you cannot type ahead when the editor is computing with the cursor on the bottom line.

† All commands which read from the last display line can also be terminated with a ESC as well as an CR.

only to look at. Be very careful not to give this command when you really want to save the changes you have made.

## 2. Moving around in the file

### 2.1. Scrolling and paging

The editor has a number of commands for moving around in the file. The most useful of these is generated by hitting the control and D keys at the same time, a control-D or '^D'. We will use this two character notation for referring to these control keys from now on. You may have a key labelled '^' on your terminal. This key will be represented as '^' in this document; '^' is exclusively used as part of the '^x' notation for control characters.‡

As you know now if you tried hitting '^D', this command scrolls down in the file. The D thus stands for down. Many editor commands are mnemonic and this makes them much easier to remember. For instance the command to scroll up is '^U'. Many dumb terminals can't scroll up at all, in which case hitting '^U' clears the screen and refreshes it with a line which is farther back in the file at the top.

If you want to see more of the file below where you are, you can hit '^E' to expose one more line at the bottom of the screen, leaving the cursor where it is. †† The command '^Y' (which is hopelessly non-mnemonic, but next to '^U' on the keyboard) exposes one more line at the top of the screen.

There are other ways to move around in the file; the keys '^F' and '^B' ‡ move forward and backward a page, keeping a couple of lines of continuity between screens so that it is possible to read through a file using these rather than '^D' and '^U' if you wish.

Notice the difference between scrolling and paging. If you are trying to read the text in a file, hitting '^F' to move forward a page will leave you only a little context to look back at. Scrolling on the other hand leaves more context, and happens more smoothly. You can continue to read the text as scrolling is taking place.

### 2.2. Searching, goto, and previous context

Another way to position yourself in the file is by giving the editor a string to search for. Type the character '/' followed by a string of characters terminated by CR. The editor will position the cursor at the next occurrence of this string. Try hitting 'n' to then go to the next occurrence of this string. The character '?' will search backwards from where you are, and is otherwise like '/.†

If the search string you give the editor is not present in the file the editor will print a diagnostic on the last line of the screen, and the cursor will be returned to its initial position.

If you wish the search to match only at the beginning of a line, begin the search string with an '^'. To match only at the end of a line, end the search string with a '\$'. Thus '^/searchCR will search for the word 'search' at the beginning of a line, and '/last\$CR searches for the word 'last' at the end of a line.\*

---

‡ If you don't have a '^' key on your terminal then there is probably a key labelled '^'; in any case these characters are one and the same.

†† Version 3 only.

‡ Not available in all v2 editors due to memory constraints.

† These searches will normally wrap around the end of the file, and thus find the string even if it is not on a line in the direction you search provided it is anywhere else in the file. You can disable this wraparound in scans by giving the command `:se nowrapscanCR`, or more briefly `:se nowrapCR`.

\*Actually, the string you give to search for here can be a *regular expression* in the sense of the editors *ex(1)* and *ed(1)*. If you don't wish to learn about this yet, you can disable this more general facility by doing `:se nomagicCR`: by putting this command in EXINIT in your environment, you can have this always be in effect (more about *EXINIT* later.)

The command G, when preceded by a number will position the cursor at that line in the file. Thus 1G will move the cursor to the first line of the file. If you give G no count, then it moves to the end of the file.

If you are near the end of the file, and the last line is not at the bottom of the screen, the editor will place only the character "" on each remaining line. This indicates that the last line in the file is on the screen; that is, the "" lines are past the end of the file.

You can find out the state of the file you are editing by typing a ^G. The editor will show you the name of the file you are editing, the number of the current line, the number of lines in the buffer, and the percentage of the way through the buffer which you are. Try doing this now, and remember the number of the line you are on. Give a G command to get to the end and then another G command to get back where you were.

You can also get back to a previous position by using the command "" (two back quotes). This is often more convenient than G because it requires no advance preparation. Try giving a G or a search with / or ? and then a "" to get back to where you were. If you accidentally hit n or any command which moves you far away from a context of interest, you can quickly get back by hitting "".

### 2.3. Moving around on the screen

Now try just moving the cursor around on the screen. If your terminal has arrow keys (4 or 5 keys with arrows going in each direction) try them and convince yourself that they work. (On certain terminals using v2 editors, they won't.) If you don't have working arrow keys, you can always use h, j, k, and l. Experienced users of vi prefer these keys to arrow keys, because they are usually right underneath their fingers.

Hit the + key. Each time you do, notice that the cursor advances to the next line in the file, at the first non-white position on the line. The - key is like + but goes the other way.

These are very common keys for moving up and down lines in the file. Notice that if you go off the bottom or top with these keys then the screen will scroll down (and up if possible) to bring a line at a time into view. The RETURN key has the same effect as the + key.

Vi also has commands to take you to the top, middle and bottom of the screen. H will take you to the top (home) line on the screen. Try preceding it with a number as in 3H. This will take you to the third line on the screen. Many vi commands take preceding numbers and do interesting things with them. Try M, which takes you to the middle line on the screen, and L, which takes you to the last line on the screen. L also takes counts, thus 5L will take you to the fifth line from the bottom.

### 2.4. Moving within a line

Now try picking a word on some line on the screen, not the first word on the line. move the cursor using RETURN and - to be on the line where the word is. Try hitting the w key. This will advance the cursor to the next word on the line. Try hitting the b key to back up words in the line. Also try the e key which advances you to the end of the current word rather than to the beginning of the next word. Also try SPACE (the space bar) which moves right one character and the BS (backspace or ^H) key which moves left one character. The key h works as ^H does and is useful if you don't have a BS key. (Also, as noted just above, l will move to the right.)

If the line had punctuation in it you may have noticed that that the w and b keys stopped at each group of punctuation. You can also go back and forwards words without stopping at punctuation by using W and B rather than the lower case equivalents. Think of these as bigger words. Try these on a few lines with punctuation to see how they differ from the lower case w and b.

The word keys wrap around the end of line, rather than stopping at the end. Try moving to a word on a line below where you are by repeatedly hitting w.

## 2.5. Summary

|       |                                          |
|-------|------------------------------------------|
| SPACE | advance the cursor one position          |
| ^B    | backwards to previous page               |
| ^D    | scrolls down in the file                 |
| ^E    | exposes another line at the bottom (v3)  |
| ^F    | forward to next page                     |
| ^G    | tell what is going on                    |
| ^H    | backspace the cursor                     |
| ^N    | next line, same column                   |
| ^P    | previous line, same column               |
| ^U    | scrolls up in the file                   |
| ^Y    | exposes another line at the top (v3)     |
| +     | next line, at the beginning              |
| -     | previous line, at the beginning          |
| /     | scan for a following string forwards     |
| ?     | scan backwards                           |
| B     | back a word, ignoring punctuation        |
| G     | go to specified line, last default       |
| H     | home screen line                         |
| M     | middle screen line                       |
| L     | last screen line                         |
| W     | forward a word, ignoring punctuation     |
| b     | back a word                              |
| e     | end of current word                      |
| n     | scan for next instance of / or ? pattern |
| w     | word after this word                     |

## 2.6. View †

If you want to use the editor to look at a file, rather than to make changes, invoke it as *view* instead of *vi*. This will set the *readonly* option which will prevent you from accidentally overwriting the file.

## 3. Making simple changes

### 3.1. Inserting

One of the most useful commands is the *i* (insert) command. After you type *i*, everything you type until you hit ESC is inserted into the file. Try this now; position yourself to some word in the file and try inserting text before this word. If you are on a dumb terminal it will seem, for a minute, that some of the characters in your line have been overwritten, but they will reappear when you hit ESC.

Now try finding a word which can, but does not, end in an 's'. Position yourself at this word and type *e* (move to end of word), then *a* for append and then *^ESC* to terminate the textual insert. This sequence of commands can be used to easily pluralize a word.

Try inserting and appending a few times to make sure you understand how this works; *i* placing text to the left of the cursor, *a* to the right.

It is often the case that you want to add new lines to the file you are editing, before or after some specific line in the file. Find a line where this makes sense and then give the command *o* to create a new line after the line you are on, or the command *O* to create a new line before the line you are on. After you create a new line in this way, text you type up to an ESC

---

† Not available in all v2 editors due to memory constraints.

is inserted on the new line.

Many related editor commands are invoked by the same letter key and differ only in that one is given by a lower case key and the other is given by an upper case key. In these cases, the upper case key often differs from the lower case key in its sense of direction, with the upper case key working backward and/or up, while the lower case key moves forward and/or down.

Whenever you are typing in text, you can give many lines of input or just a few characters. To type in more than one line of text, hit a RETURN at the middle of your input. A new line will be created for text, and you can continue to type. If you are on a slow and dumb terminal the editor may choose to wait to redraw the tail of the screen, and will let you type over the existing screen lines. This avoids the lengthy delay which would occur if the editor attempted to keep the tail of the screen always up to date. The tail of the screen will be fixed up, and the missing lines will reappear, when you hit ESC.

While you are inserting new text, you can use the characters you normally use at the system command level (usually `^H` or `#`) to backspace over the last character which you typed, and the character which you use to kill input lines (usually `@`, `^X`, or `^U`) to erase the input you have typed on the current line.† The character `^W` will erase a whole word and leave you after the space after the previous word; it is useful for quickly backing up in an insert.

Notice that when you backspace during an insertion the characters you backspace over are not erased; the cursor moves backwards, and the characters remain on the display. This is often useful if you are planning to type in something similar. In any case the characters disappear when when you hit ESC; if you want to get rid of them immediately, hit an ESC and then a again.

Notice also that you can't erase characters which you didn't insert, and that you can't backspace around the end of a line. If you need to back up to the previous line to make a correction, just hit ESC and move the cursor back to the previous line. After making the correction you can return to where you were and use the insert or append command again.

### 3.2. Making small corrections

You can make small corrections in existing text quite easily. Find a single character which is wrong or just pick any character. Use the arrow keys to find the character, or get near the character with the word motion keys and then either backspace (hit the BS key or `^H` or even just `h`) or SPACE (using the space bar) until the cursor is on the character which is wrong. If the character is not needed then hit the `x` key; this deletes the character from the file. It is analogous to the way you `x` out characters when you make mistakes on a typewriter (except it's not as messy).

If the character is incorrect, you can replace it with the correct character by giving the command `rc`, where `c` is replaced by the correct character. Finally if the character which is incorrect should be replaced by more than one character, give the command `s` which substitutes a string of characters, ending with ESC, for it. If there are a small number of characters which are wrong you can precede `s` with a count of the number of characters to be replaced. Counts are also useful with `x` to specify the number of characters to be deleted.

### 3.3. More corrections: operators

You already know almost enough to make changes at a higher level. All you need to know now is that the `d` key acts as a delete operator. Try the command `dw` to delete a word. Try hitting `.` a few times. Notice that this repeats the effect of the `dw`. The command `.` repeats the last command which made a change. You can remember it by analogy with an ellipsis `'...'`.

---

† In fact, the character `^H` (backspace) always works to erase the last input character here, regardless of what your erase character is.



Now try **db**. This deletes a word backwards, namely the preceding word. Try **dSPACE**. This deletes a single character, and is equivalent to the **x** command.

Another very useful operator is **c** or change. The command **cw** thus changes the text of a single word. You follow it by the replacement text ending with an **ESC**. Find a word which you can change to another, and try this now. Notice that the end of the text to be changed was marked with the character '\$' so that you can see this as you are typing in the new material.

### 3.4. Operating on lines

It is often the case that you want to operate on lines. Find a line which you want to delete, and type **dd**, the **d** operator twice. This will delete the line. If you are on a dumb terminal, the editor may just erase the line on the screen, replacing it with a line with only an **@** on it. This line does not correspond to any line in your file, but only acts as a place holder. It helps to avoid a lengthy redraw of the rest of the screen which would be necessary to close up the hole created by the deletion on a terminal without a delete line capability.

Try repeating the **c** operator twice; this will change a whole line, erasing its previous contents and replacing them with text you type up to an **ESC**.†

You can delete or change more than one line by preceding the **dd** or **cc** with a count, i.e. **5dd** deletes 5 lines. You can also give a command like **dL** to delete all the lines up to and including the last line on the screen, or **d3L** to delete through the third from the bottom line. Try some commands like this now.\* Notice that the editor lets you know when you change a large number of lines so that you can see the extent of the change. The editor will also always tell you when a change you make affects text which you cannot see.

### 3.5. Undoing

Now suppose that the last change which you made was incorrect; you could use the insert, delete and append commands to put the correct material back. However, since it is often the case that we regret a change or make a change incorrectly, the editor provides a **u** (undo) command to reverse the last change which you made. Try this a few times, and give it twice in a row to notice that an **u** also undoes a **u**.

The undo command lets you reverse only a single change. After you make a number of changes to a line, you may decide that you would rather have the original state of the line back. The **U** command restores the current line to the state before you started changing it.

You can recover text which you delete, even if undo will not bring it back; see the section on recovering lost text below.

### 3.6. Summary

|              |                                                                                         |
|--------------|-----------------------------------------------------------------------------------------|
| <b>SPACE</b> | advance the cursor one position                                                         |
| <b>^H</b>    | backspace the cursor                                                                    |
| <b>^W</b>    | erase a word during an insert                                                           |
| <b>erase</b> | your erase (usually <b>^H</b> or <b>#</b> ), erases a character during an insert        |
| <b>kill</b>  | your kill (usually <b>@</b> , <b>^X</b> , or <b>^U</b> ), kills the insert on this line |
| <b>.</b>     | repeats the changing command                                                            |
| <b>O</b>     | opens and inputs new lines, above the current                                           |
| <b>U</b>     | undoes the changes you made to the current line                                         |
| <b>a</b>     | appends text after the cursor                                                           |
| <b>c</b>     | changes the object you specify to the following text                                    |

† The command **S** is a convenient synonym for **cc**, by analogy with **s**. Think of **S** as a substitute on lines, while **s** is a substitute on characters.

\* One subtle point here involves using the **/** search after a **d**. This will normally delete characters from the current position to the point of the match. If what is desired is to delete whole lines including the two points, give the pattern as **/pat/+0**, a line address.

|   |                                               |
|---|-----------------------------------------------|
| d | deletes the object you specify                |
| i | inserts text before the cursor                |
| o | opens and inputs new lines, below the current |
| u | undoes the last change                        |

#### 4. Moving about; rearranging and duplicating text

##### 4.1. Low level character motions

Now move the cursor to a line where there is a punctuation or a bracketing character such as a parenthesis or a comma or period. Try the command `fx` where `x` is this character. This command finds the next `x` character to the right of the cursor in the current line. Try then hitting a `;`, which finds the next instance of the same character. By using the `f` command and then a sequence of `;`'s you can often get to a particular place in a line much faster than with a sequence of word motions or `SPACES`. There is also a `F` command, which is like `f`, but searches backward. The `;` command repeats `F` also.

When you are operating on the text in a line it is often desirable to deal with the characters up to, but not including, the first instance of a character. Try `dfx` for some `x` now and notice that the `x` character is deleted. Undo this with `u` and then try `dtx`; the `t` here stands for to, i.e. delete up to the next `x`, but not the `x`. The command `T` is the reverse of `t`.

When working with the text of a single line, an `↑` moves the cursor to the first non-white position on the line, and a `$` moves it to the end of the line. Thus `$a` will append new text at the end of the current line.

Your file may have tab (`^I`) characters in it. These characters are represented as a number of spaces expanding to a tab stop, where tab stops are every 8 positions.\* When the cursor is at a tab, it sits on the last of the several spaces which represent that tab. Try moving the cursor back and forth over tabs so you understand how this works.

On rare occasions, your file may have nonprinting characters in it. These characters are displayed in the same way they are represented in this document, that is with a two character code, the first character of which is `^^`. On the screen non-printing characters resemble a `^^` character adjacent to another, but spacing or backspacing over the character will reveal that the two characters are, like the spaces representing a tab character, a single character.

The editor sometimes discards control characters, depending on the character and the setting of the *beautify* option, if you attempt to insert them in your file. You can get a control character in the file by beginning an insert and then typing a `^V` before the control character. The `^V` quotes the following character, causing it to be inserted directly into the file.

##### 4.2. Higher level text objects

In working with a document it is often advantageous to work in terms of sentences, paragraphs, and sections. The operations `(` and `)` move to the beginning of the previous and next sentences respectively. Thus the command `d)` will delete the rest of the current sentence; likewise `d(` will delete the previous sentence if you are at the beginning of the current sentence, or the current sentence up to where you are if you are not at the beginning of the current sentence.

A sentence is defined to end at a `.`, `!` or `?` which is followed by either the end of a line, or by two spaces. Any number of closing `)`, `]`, `""` and `'''` characters may appear after the `.`, `!` or `?` before the spaces or end of line.

The operations `{` and `}` move over paragraphs and the operations `[[` and `]]` move over sections.†

\* This is settable by a command of the form `:set ts=xCR`, where *x* is 4 to set tabstops every four columns. This has effect on the screen representation within the editor.

† The `[[` and `]]` operations require the operation character to be doubled because they can move the cursor far

A paragraph begins after each empty line, and also at each of a set of paragraph macros, specified by the pairs of characters in the definition of the string valued option *paragraphs*. The default setting for this option defines the paragraph macros of the *-ms* and *-mm* macro packages, i.e. the `.IP`, `.LP`, `.PP` and `.QP`, `.P` and `.LI` macros.† Each paragraph boundary is also a sentence boundary. The sentence and paragraph commands can be given counts to operate over groups of sentences and paragraphs.

Sections in the editor begin after each macro in the *sections* option, normally `.NH`, `.SH`, `.H` and `.HU`, and each line with a formfeed `^L` in the first column. Section boundaries are always line and paragraph boundaries also.

Try experimenting with the sentence and paragraph commands until you are sure how they work. If you have a large document, try looking through it using the section commands. The section commands interpret a preceding count as a different window size in which to redraw the screen at the new location, and this window size is the base size for newly drawn windows until another size is specified. This is very useful if you are on a slow terminal and are looking for a particular section. You can give the first section command a small count to then see each successive section heading in a small window.

### 4.3. Rearranging and duplicating text

The editor has a single unnamed buffer where the last deleted or changed away text is saved, and a set of named buffers `a-z` which you can use to save copies of text and to move text around in your file and between files.

The operator `y` yanks a copy of the object which follows into the unnamed buffer. If preceded by a buffer name, `"xy`, where `x` here is replaced by a letter `a-z`, it places the text in the named buffer. The text can then be put back in the file with the commands `p` and `P`; `p` puts the text after or below the cursor, while `P` puts the text before or above the cursor.

If the text which you yank forms a part of a line, or is an object such as a sentence which partially spans more than one line, then when you put the text back, it will be placed after the cursor (or before if you use `P`). If the yanked text forms whole lines, they will be put back as whole lines, without changing the current line. In this case, the `put` acts much like a `o` or `O` command.

Try the command `YP`. This makes a copy of the current line and leaves you on this copy, which is placed before the current line. The command `Y` is a convenient abbreviation for `yy`. The command `Yp` will also make a copy of the current line, and place it after the current line. You can give `Y` a count of lines to yank, and thus duplicate several lines; try `3YP`.

To move text within the buffer, you need to delete it in one place, and put it back in another. You can precede a delete operation by the name of a buffer in which the text is to be stored as in `"a5dd` deleting 5 lines into the named buffer `a`. You can then move the cursor to the eventual resting place of these lines and do a `"ap` or `"aP` to put them back. In fact, you can switch and edit another file before you put the lines back, by giving a command of the form `:e nameCR` where `name` is the name of the other file you want to edit. You will have to write back the contents of the current editor buffer (or discard them) if you have made changes before the editor will let you switch to the other file. An ordinary delete command saves the text in the unnamed buffer, so that an ordinary `put` can move it elsewhere. However, the unnamed buffer is lost when you change files, so to move text from one file to another you should use an unnamed buffer.

---

from where it currently is. While it is easy to get back with the command `"`, these commands would still be frustrating if they were easy to hit accidentally.

† You can easily change or extend this set of macros by assigning a different string to the *paragraphs* option in your EXINIT. See section 6.2 for details. The `.bp` directive is also considered to start a paragraph.

#### 4.4. Summary.

|    |                                                    |
|----|----------------------------------------------------|
|    | first non-white on line                            |
| \$ | end of line                                        |
| )  | forward sentence                                   |
| }  | forward paragraph                                  |
|    | forward section                                    |
| (  | backward sentence                                  |
| {  | backward paragraph                                 |
| [[ | backward section                                   |
| fx | find x forward in line                             |
| p  | put text back, after cursor or below current line  |
| y  | yank operator, for copies and moves                |
| tx | up to x forward, for operators                     |
| Fx | f backward in line                                 |
| P  | put text back, before cursor or above current line |
| Tx | t backward in line                                 |

### 5. High level commands

#### 5.1. Writing, quitting, editing new files

So far we have seen how to enter *vi* and to write out our file using either **ZZ** or **:wCR**. The first exits from the editor, (writing if changes were made), the second writes and stays in the editor.

If you have changed the editor's copy of the file but do not wish to save your changes, either because you messed up the file or decided that the changes are not an improvement to the file, then you can give the command **:q!CR** to quit from the editor without writing the changes. You can also reedit the same file (starting over) by giving the command **:e!CR**. These commands should be used only rarely, and with caution, as it is not possible to recover the changes you have made after you discard them in this manner.

You can edit a different file without leaving the editor by giving the command **:e nameCR**. If you have not written out your file before you try to do this, then the editor will tell you this, and delay editing the other file. You can then give the command **:wCR** to save your work and then the **:e nameCR** command again, or carefully give the command **:e! nameCR**, which edits the other file discarding the changes you have made to the current file. To have the editor automatically save changes, include *set autowrite* in your EXINIT, and use **:n** instead of **:e**.

#### 5.2. Escaping to a shell

You can get to a shell to execute a single command by giving a *vi* command of the form **!:cmdCR**. The system will run the single command *cmd* and when the command finishes, the editor will ask you to hit a RETURN to continue. When you have finished looking at the output on the screen, you should hit RETURN and the editor will clear the screen and redraw it. You can then continue editing. You can also give another **:** command when it asks you for a RETURN; in this case the screen will not be redrawn.

If you wish to execute more than one command in the shell, then you can give the command **:shCR**. This will give you a new shell, and when you finish with the shell, ending it by typing a **^D**, the editor will clear the screen and continue.

On systems which support it, **^Z** will suspend the editor and return to the (top level) shell. When the editor is resumed, the screen will be redrawn.

### 5.3. Marking and returning

The command ```` returned to the previous place after a motion of the cursor by a command such as `/`, `?` or `G`. You can also mark lines in the file with single letter tags and return to these marks later by naming the tags. Try marking the current line with the command `m.x`, where you should pick some letter for `x`, say `'a'`. Then move the cursor to a different line (any way you like) and hit `'a`. The cursor will return to the place which you marked. Marks last only until you edit another file.

When using operators such as `d` and referring to marked lines, it is often desirable to delete whole lines rather than deleting to the exact position in the line marked by `m`. In this case you can use the form `'x` rather than ``x`. Used without an operator, `'x` will move to the first non-white character of the marked line; similarly ```` moves to the first non-white character of the line containing the previous context mark ````.

### 5.4. Adjusting the screen

If the screen image is messed up because of a transmission error to your terminal, or because some program other than the editor wrote output to your terminal, you can hit a `^L`, the ASCII form-feed character, to cause the screen to be refreshed.

On a dumb terminal, if there are `@` lines in the middle of the screen as a result of line deletion, you may get rid of these lines by typing `^R` to cause the editor to retype the screen, closing up these holes.

Finally, if you wish to place a certain line on the screen at the top middle or bottom of the screen, you can position the cursor to that line, and then give a `z` command. You should follow the `z` command with a RETURN if you want the line to appear at the top of the window, a `.` if you want it at the center, or a `-` if you want it at the bottom. (`z.`, `z-`, and `z+` are not available on all v2 editors.)

## 6. Special topics

### 6.1. Editing on slow terminals

When you are on a slow terminal, it is important to limit the amount of output which is generated to your screen so that you will not suffer long delays, waiting for the screen to be refreshed. We have already pointed out how the editor optimizes the updating of the screen during insertions on dumb terminals to limit the delays, and how the editor erases lines to `@` when they are deleted on dumb terminals.

The use of the slow terminal insertion mode is controlled by the *slowopen* option. You can force the editor to use this mode even on faster terminals by giving the command `:se slowCR`. If your system is sluggish this helps lessen the amount of output coming to your terminal. You can disable this option by `:se noslowCR`.

The editor can simulate an intelligent terminal on a dumb one. Try giving the command `:se redrawCR`. This simulation generates a great deal of output and is generally tolerable only on lightly loaded systems and fast terminals. You can disable this by giving the command `:se noredrawCR`.

The editor also makes editing more pleasant at low speed by starting editing in a small window, and letting the window expand as you edit. This works particularly well on intelligent terminals. The editor can expand the window easily when you insert in the middle of the screen on these terminals. If possible, try the editor on an intelligent terminal to see how this works.

You can control the size of the window which is redrawn each time the screen is cleared by giving window sizes as argument to the commands which cause large screen motions:

```
: / ? [] ` `
```

Thus if you are searching for a particular instance of a common string in a file you can precede

the first search command by a small number, say 3, and the editor will draw three line windows around each instance of the string which it locates.

You can easily expand or contract the window, placing the current line as you choose, by giving a number on a z command, after the z and before the following RETURN, . or -. Thus the command z5. redraws the screen with the current line in the center of a five line window.†

If the editor is redrawing or otherwise updating large portions of the display, you can interrupt this updating by hitting a DEL or RUB as usual. If you do this you may partially confuse the editor about what is displayed on the screen. You can still edit the text on the screen if you wish; clear up the confusion by hitting a ^L; or move or search again, ignoring the current state of the display.

See section 7.8 on open mode for another way to use the vi command set on slow terminals.

### 6.2. Options, set, and editor startup files

The editor has a set of options, some of which have been mentioned above. The most useful options are given in the following table.

| Name       | Default               | Description                                    |
|------------|-----------------------|------------------------------------------------|
| autoindent | noai                  | Supply indentation automatically               |
| autowrite  | noaw                  | Automatic write before :n, :ta, ^I, !          |
| ignorecase | noic                  | Ignore case in searching                       |
| lisp       | nolisp                | ( { ) commands deal with S-expressions         |
| list       | nolist                | Tabs print as ^I; end of lines marked with \$  |
| magic      | nomagic               | The characters . [ and * are special in scans  |
| number     | nonu                  | Lines are displayed prefixed with line numbers |
| paragraphs | para = IPLPPPQPbpP LI | Macro names which start paragraphs             |
| redraw     | nore                  | Simulate a smart terminal on a dumb one        |
| sections   | sect = NHSHH HU       | Macro names which start new sections           |
| shiftwidth | sw = 8                | Shift distance for <, > and input ^D and ^T    |
| showmatch  | nosm                  | Show matching ( or { as ) or } is typed        |
| slowopen   | slow                  | Postpone display updates during inserts        |
| term       | dumb                  | The kind of terminal you are using.            |

The options are of three kinds: numeric options, string options, and toggle options. You can set numeric and string options by a statement of the form

```
set opt=val
```

and toggle options can be set or unset by statements of one of the forms

```
set opt
set noopt
```

These statements can be placed in your EXINIT in your environment, or given while you are running vi by preceding them with a : and following them with a CR.

You can get a list of all options which you have changed by the command :setCR, or the value of a single option by the command :set opt?CR. A list of all possible options and their values is generated by :set allCR. Set can be abbreviated se. Multiple options can be placed on one line, e.g. :se ai aw nuCR.

Options set by the set command only last while you stay in the editor. It is common to want to have certain options set whenever you use the editor. This can be accomplished by creating a list of ex commands† which are to be run every time you start up ex, edit, or vi. A

† Note that the command Sz. has an entirely different effect, placing line S in the center of a new window.  
† All commands which start with : are ex commands.

typical list includes a `set` command, and possibly a few `map` commands (on v3 editors). Since it is advisable to get these commands on one line, they can be separated with the `|` character, for example:

```
set ai aw terse|map @ dd|map # x
```

which sets the options *autoindent*, *autowrite*, *terse*, (the `set` command), makes `@` delete a line, (the first `map`), and makes `#` delete a character, (the second `map`). (See section 6.9 for a description of the `map` command, which only works in version 3.) This string should be placed in the variable `EXINIT` in your environment. If you use *csh*, put this line in the file *.login* in your home directory:

```
setenv EXINIT 'set ai aw terse|map @ dd|map # x'
```

If you use the standard v7 shell, put these lines in the file *.profile* in your home directory:

```
EXINIT='set ai aw terse|map @ dd|map # x'
export EXINIT
```

On a version 6 system, the concept of environments is not present. In this case, put the line in the file *.exrc* in your home directory.

```
set ai aw terse|map @ dd|map # x
```

Of course, the particulars of the line would depend on which options you wanted to set.

### 6.3. Recovering lost lines

You might have a serious problem if you delete a number of lines and then regret that they were deleted. Despair not, the editor saves the last 9 deleted blocks of text in a set of numbered registers 1–9. You can get the *n*'th previous deleted text back in your file by the command "`n`p". The "`n`" here says that a buffer name is to follow, *n* is the number of the buffer you wish to try (use the number 1 for now), and `p` is the put command, which puts text in the buffer after the cursor. If this doesn't bring back the text you wanted, hit `u` to undo this and then `.` (period) to repeat the put command. In general the `.` command will repeat the last change you made. As a special case, when the last command refers to a numbered text buffer, the `.` command increments the number of the buffer before repeating the command. Thus a sequence of the form

```
"1pu.u.u.
```

will, if repeated long enough, show you all the deleted text which has been saved for you. You can omit the `u` commands here to gather up all this text in the buffer, or stop after any `.` command to keep just the then recovered text. The command `P` can also be used rather than `p` to put the recovered text before rather than after the cursor.

### 6.4. Recovering lost files

If the system crashes, you can recover the work you were doing to within a few changes. You will normally receive mail when you next login giving you the name of the file which has been saved for you. You should then change to the directory where you were when the system crashed and give a command of the form:

```
% vi -r name
```

replacing *name* with the name of the file which you were editing. This will recover your work to a point near where you left off.†

---

† In rare cases, some of the lines of the file may be lost. The editor will give you the numbers of these lines and the text of the lines will be replaced by the string 'LOST'. These lines will almost always be among the last few which you changed. You can either choose to discard the changes which you made (if they are easy to remake) or to replace the few lost lines by hand.

You can get a listing of the files which are saved for you by giving the command:

```
% vi -r
```

If there is more than one instance of a particular file saved, the editor gives you the newest instance each time you recover it. You can thus get an older saved copy back by first recovering the newer copies.

For this feature to work, *vi* must be correctly installed by a super user on your system, and the *mail* program must exist to receive mail. The invocation "*vi -r*" will not always list all saved files, but they can be recovered even if they are not listed.

### 6.5. Continuous text input

When you are typing in large amounts of text it is convenient to have lines broken near the right margin automatically. You can cause this to happen by giving the command `:se wm=10CR`. This causes all lines to be broken at a space at least 10 columns from the right hand edge of the screen.\*

If the editor breaks an input line and you wish to put it back together you can tell it to join the lines with `J`. You can give `J` a count of the number of lines to be joined as in `3J` to join 3 lines. The editor supplies white space, if appropriate, at the juncture of the joined lines, and leaves the cursor at this white space. You can kill the white space with `x` if you don't want it.

### 6.6. Features for editing programs

The editor has a number of commands for editing programs. The thing that most distinguishes editing of programs from editing of text is the desirability of maintaining an indented structure to the body of the program. The editor has a *autoindent* facility for helping you generate correctly indented programs.

To enable this facility you can give the command `:se aiCR`. Now try opening a new line with `o` and type some characters on the line after a few tabs. If you now start another line, notice that the editor supplies white space at the beginning of the line to line it up with the previous line. You cannot backspace over this indentation, but you can use `^D` key to backtab over the supplied indentation.

Each time you type `^D` you back up one position, normally to an 8 column boundary. This amount is settable; the editor has an option called *shiftwidth* which you can set to change this value. Try giving the command `:se sw=4CR` and then experimenting with autoindent again.

For shifting lines in the program left and right, there are operators `<` and `>`. These shift the lines you specify right or left by one *shiftwidth*. Try `<<` and `>>` which shift one line left or right, and `<L` and `>L` shifting the rest of the display left and right.

If you have a complicated expression and wish to see how the parentheses match, put the cursor at a left or right parenthesis and hit `%`. This will show you the matching parenthesis. This works also for braces `{` and `}`, and brackets `[` and `]`.

If you are editing C programs, you can use the `[[` and `]]` keys to advance or retreat to a line starting with a `{`, i.e. a function declaration at a time. When `]]` is used with an operator it stops after a line which starts with `};` this is sometimes useful with `y]]`.

---

\* This feature is not available on some v2 editors. In v2 editors where it is available, the break can only occur to the right of the specified boundary instead of to the left.



### 6.7. Filtering portions of the buffer

You can run system commands over portions of the buffer using the operator `!`. You can use this to sort lines in the buffer, or to reformat portions of the buffer with a pretty-printer. Try typing in a list of random words, one per line and ending them with a blank line. Back up to the beginning of the list, and then give the command `!|sortCR`. This says to sort the next paragraph of material, and the blank line ends a paragraph.

### 6.8. Commands for editing LISP†

If you are editing a LISP program you should set the option `lisp` by doing `:se lispCR`. This changes the `(` and `)` commands to move backward and forward over s-expressions. The `{` and `}` commands are like `(` and `)` but don't stop at atoms. These can be used to skip to the next list, or through a comment quickly.

The `autoindent` option works differently for LISP, supplying indent to align at the first argument to the last open list. If there is no such argument then the indent is two spaces more than the last level.

There is another option which is useful for typing in LISP, the `showmatch` option. Try setting it with `:se smCR` and then try typing a `'` some words and then a `'`. Notice that the cursor shows the position of the `'` which matches the `'` briefly. This happens only if the matching `'` is on the screen, and the cursor stays there for at most one second.

The editor also has an operator to realign existing lines as though they had been typed in with `lisp` and `autoindent` set. This is the `=` operator. Try the command `=%` at the beginning of a function. This will realign all the lines of the function declaration.

When you are editing LISP., the `[` and `]` advance and retreat to lines beginning with a `(`, and are useful for dealing with entire function definitions.

### 6.9. Macros‡

`Vi` has a parameterless macro facility, which lets you set it up so that when you hit a single keystroke, the editor will act as though you had hit some longer sequence of keys. You can set this up if you find yourself typing the same sequence of commands repeatedly.

Briefly, there are two flavors of macros:

- a) Ones where you put the macro body in a buffer register, say `x`. You can then type `@x` to invoke the macro. The `@` may be followed by another `@` to repeat the last macro.
- b) You can use the `map` command from `vi` (typically in your `EXINIT`) with a command of the form:

```
:map lhs rhsCR
```

mapping `lhs` into `rhs`. There are restrictions: `lhs` should be one keystroke (either 1 character or one function key) since it must be entered within one second (unless `notimeout` is set, in which case you can type it as slowly as you wish, and `vi` will wait for you to finish it before it echoes anything). The `lhs` can be no longer than 10 characters, the `rhs` no longer than 100. To get a space, tab or newline into `lhs` or `rhs` you should escape them with a `^V`. (It may be necessary to double the `^V` if the map command is given inside `vi`, rather than in `ex`.) Spaces and tabs inside the `rhs` need not be escaped.

Thus to make the `q` key write and exit the editor, you can give the command

```
:map q :wq^V^VCR CR
```

which means that whenever you type `q`, it will be as though you had typed the four characters `:wqCR`. A `^V`'s is needed because without it the `CR` would end the `:` command, rather than

† The LISP features are not available on some v2 editors due to memory constraints.

‡ The macro feature is available only in version 3 editors.

becoming part of the *map* definition. There are two `^V`'s because from within *vi*, two `^V`'s must be typed to get one. The first CR is part of the *rhs*, the second terminates the `:` command.

Macros can be deleted with

```
unmap lhs
```

If the *lhs* of a macro is `"#0"` through `"#9"`, this maps the particular function key instead of the 2 character `"#"` sequence. So that terminals without function keys can access such definitions, the form `"#x"` will mean function key *x* on all terminals (and need not be typed within one second.) The character `"#"` can be changed by using a macro in the usual way:

```
:map ^V^V^I #
```

to use `tab`, for example. (This won't affect the *map* command, which still uses `#`, but just the invocation from visual mode.)

The undo command reverses an entire macro call as a unit, if it made any changes.

Placing a `!` after the word `map` causes the mapping to apply to input mode, rather than command mode. Thus, to arrange for `^T` to be the same as 4 spaces in input mode, you can type:

```
:map ^T ^VBBBB
```

where `B` is a blank. The `^V` is necessary to prevent the blanks from being taken as white space between the *lhs* and *rhs*.

## 7. Word Abbreviations `##`

A feature similar to macros in input mode is word abbreviation. This allows you to type a short word and have it expanded into a longer word or words. The commands are `:abbreviate` and `:unabbreviate` (`:ab` and `:una`) and have the same syntax as `:map`. For example:

```
:ab eecs Electrical Engineering and Computer Sciences
```

causes the word 'eecs' to always be changed into the phrase 'Electrical Engineering and Computer Sciences'. Word abbreviation is different from macros in that only whole words are affected. If 'eecs' were typed as part of a larger word, it would be left alone. Also, the partial word is echoed as it is typed. There is no need for an abbreviation to be a single keystroke, as it should be with a macro.

### 7.1. Abbreviations

The editor has a number of short commands which abbreviate longer commands which we have introduced here. You can find these commands easily on the quick reference card. They often save a bit of typing and you can learn them as convenient.

## 8. Nitty-gritty details

### 8.1. Line representation in the display

The editor folds long logical lines onto many physical lines in the display. Commands which advance lines advance logical lines and will skip over all the segments of a line in one motion. The command `|` moves the cursor to a specific column, and may be useful for getting near the middle of a long line to split it in half. Try `80|` on a line which is more than 80 columns long.†

The editor only puts full lines on the display; if there is not enough room on the display to fit a logical line, the editor leaves the physical line empty, placing only an `@` on the line as a

---

## Version 3 only.

† You can make long lines very easily by using `J` to join together short lines.

place holder. When you delete lines on a dumb terminal, the editor will often just clear the lines to @ to save time (rather than rewriting the rest of the screen.) You can always maximize the information on the screen by giving the ^R command.

If you wish, you can have the editor place line numbers before each line on the display. Give the command :se nuCR to enable this, and the command :se nonuCR to turn it off. You can have tabs represented as ^I and the ends of lines indicated with ^S by giving the command :se listCR; :se nolistCR turns this off.

Finally, lines consisting of only the character "" are displayed when the last line in the file is in the middle of the screen. These represent physical lines which are past the logical end of file.

## 8.2. Counts

Most vi commands will use a preceding count to affect their behavior in some way. The following table gives the common ways in which the counts are used:

|                    |                  |
|--------------------|------------------|
| new window size    | : / ? [ [ ] ` `  |
| scroll amount      | ^D ^U            |
| line/column number | z G              |
| repeat effect      | most of the rest |

The editor maintains a notion of the current default window size. On terminals which run at speeds greater than 1200 baud the editor uses the full terminal screen. On terminals which are slower than 1200 baud (most dialup lines are in this group) the editor uses 8 lines as the default window size. At 1200 baud the default is 16 lines.

This size is the size used when the editor clears and refills the screen after a search or other motion moves far from the edge of the current window. The commands which take a new window size as count all often cause the screen to be redrawn. If you anticipate this, but do not need as large a window as you are currently using, you may wish to change the screen size by specifying the new size before these commands. In any case, the number of lines used on the screen will expand if you move off the top with a - or similar command or off the bottom with a command such as RETURN or ^D. The window will revert to the last specified size the next time it is cleared and refilled.†

The scroll commands ^D and ^U likewise remember the amount of scroll last specified, using half the basic window size initially. The simple insert commands use a count to specify a repetition of the inserted text. Thus 10a+----ESC will insert a grid-like string of text. A few commands also use a preceding count as a line or column number.

Except for a few commands which ignore any counts (such as ^R), the rest of the editor commands use a count to indicate a simple repetition of their effect. Thus 5w advances five words on the current line, while 5RETURN advances five lines. A very useful instance of a count as a repetition is a count given to the . command, which repeats the last changing command. If you do dw and then 3., you will delete first one and then three words. You can then delete two more words with 2..

## 8.3. More file manipulation commands

The following table lists the file manipulation commands which you can use when you are in vi. All of these commands are followed by a CR or ESC. The most basic commands are :w and :e. A normal editing session on a single file will end with a ZZ command. If you are editing for a long period of time you can give :w commands occasionally after major amounts of editing, and then finish with a ZZ. When you edit more than one file, you can finish with one

† But not by a ^L which just redraws the screen as it is.

```
:w write back changes
:wq write and quit
:x write (if necessary) and quit (same as ZZ).
:e name edit file name
:e! reedit, discarding changes
:e + name edit, starting at end
:e + n edit, starting at line n
:e # edit alternate file
:w name write file name
:w! name overwrite file name
:x,yw name write lines x through y to name
:r name read file name into buffer
:r !cmd read output of cmd into buffer
:n edit next file in argument list
:n! edit next file, discarding changes to current
:n args specify new argument list
:ta tag edit file containing tag tag, at tag
```

with a `:w` and start editing a new file by giving a `:e` command, or set `autowrite` and use `:n <file>`.

If you make changes to the editor's copy of a file, but do not wish to write them back, then you must give an `!` after the command you would otherwise use; this forces the editor to discard any changes you have made. Use this carefully.

The `:e` command can be given a `+` argument to start at the end of the file, or a `+n` argument to start at line `n`. In actuality, `n` may be any editor command not containing a space, usefully a scan like `+/pat` or `+?pat`. In forming new names to the `e` command, you can use the character `%` which is replaced by the current file name, or the character `#` which is replaced by the alternate file name. The alternate file name is generally the last name you typed other than the current file. Thus if you try to do a `:e` and get a diagnostic that you haven't written the file, you can give a `:w` command and then a `:e #` command to redo the previous `:e`.

You can write part of the buffer to a file by finding out the lines that bound the range to be written using `^G`, and giving these numbers after the `:` and before the `w`, separated by `,`'s. You can also mark these lines with `m` and then use an address of the form `'x,'y` on the `w` command here.

You can read another file into the buffer after the current line by using the `:r` command. You can similarly read in the output from a command, just use `!cmd` instead of a file name.

If you wish to edit a set of files in succession, you can give all the names on the command line, and then edit each one in turn using the command `:n`. It is also possible to respecify the list of files to be edited by giving the `:n` command a list of file names, or a pattern to be expanded as you would have given it on the initial `w` command.

If you are editing large programs, you will find the `:ta` command very useful. It utilizes a data base of function names and their locations, which can be created by programs such as `ctags`, to quickly find a function whose name you give. If the `:ta` command will require the editor to switch files, then you must `:w` or abandon any changes before switching. You can repeat the `:ta` command without any arguments to look for the same tag again. (The tag feature is not available in some v2 editors.)

#### 8.4. More about searching for strings

When you are searching for strings in the file with `/` and `?`, the editor normally places you at the next or previous occurrence of the string. If you are using an operator such as `d`, `c` or `y`, then you may well wish to affect lines up to the line before the line containing the pattern.

You can give a search of the form */pat/-n* to refer to the *n*'th line before the next line containing *pat*, or you can use *+* instead of *-* to refer to the lines after the one containing *pat*. If you don't give a line offset, then the editor will affect characters up to the match place, rather than whole lines; thus use *"+0"* to affect to the line which matches.

You can have the editor ignore the case of words in the searches it does by giving the command *:se icCR*. The command *:se noicCR* turns this off.

Strings given to searches may actually be regular expressions. If you do not want or need this facility, you should

```
set nomagic
```

in your EXINIT. In this case, only the characters *↑* and *\$* are special in patterns. The character *\* is also then special (as it is most everywhere in the system), and may be used to get at the an extended pattern matching facility. It is also necessary to use a *\* before a */* in a forward scan or a *?* in a backward scan, in any case. The following table gives the extended forms when *magic* is set.

|               |                                                     |
|---------------|-----------------------------------------------------|
| <i>↑</i>      | at beginning of pattern, matches beginning of line  |
| <i>\$</i>     | at end of pattern, matches end of line              |
| <i>.</i>      | matches any character                               |
| <i>\&lt;</i>  | matches the beginning of a word                     |
| <i>\&gt;</i>  | matches the end of a word                           |
| <i>[str]</i>  | matches any single character in <i>str</i>          |
| <i>[↑str]</i> | matches any single character not in <i>str</i>      |
| <i>[x-y]</i>  | matches any character between <i>x</i> and <i>y</i> |
| <i>*</i>      | matches any number of the preceding pattern         |

If you use *nomagic* mode, then the *.* *[* and *\** primitives are given with a preceding *\*.

### 8.5. More about input mode

There are a number of characters which you can use to make corrections during input mode. These are summarized in the following table.

|              |                                                       |
|--------------|-------------------------------------------------------|
| <i>^H</i>    | deletes the last input character                      |
| <i>^W</i>    | deletes the last input word, defined as by <i>b</i>   |
| <i>erase</i> | your erase character, same as <i>^H</i>               |
| <i>kill</i>  | your kill character, deletes the input on this line   |
| <i>\</i>     | escapes a following <i>^H</i> and your erase and kill |
| <i>ESC</i>   | ends an insertion                                     |
| <i>DEL</i>   | interrupts an insertion, terminating it abnormally    |
| <i>CR</i>    | starts a new line                                     |
| <i>^D</i>    | backtabs over <i>autoindent</i>                       |
| <i>0^D</i>   | kills all the <i>autoindent</i>                       |
| <i>↑^D</i>   | same as <i>0^D</i> , but restores indent next line    |
| <i>^V</i>    | quotes the next non-printing character into the file  |

The most usual way of making corrections to input is by typing *^H* to correct a single character, or by typing one or more *^W*'s to back over incorrect words. If you use *#* as your erase character in the normal system, it will work like *^H*.

Your system kill character, normally *@*, *^X* or *^U*, will erase all the input you have given on the current line. In general, you can neither erase input back around a line boundary nor can you erase characters which you did not insert with this insertion command. To make corrections on the previous line after a new line has been started you can hit *ESC* to end the insertion, move over and make the correction, and then return to where you were to continue.

The command `A` which appends at the end of the current line is often useful for continuing.

If you wish to type in your erase or kill character (say `#` or `@`) then you must precede it with a `\`, just as you would do at the normal system command level. A more general way of typing non-printing characters into the file is to precede them with a `^V`. The `^V` echoes as a `↑` character on which the cursor rests. This indicates that the editor expects you to type a control character. In fact you may type any character and it will be inserted into the file at that point.\*

If you are using *autoindent* you can backtab over the indent which it supplies by typing a `^D`. This backs up to a *shiftwidth* boundary. This only works immediately after the supplied *autoindent*.

When you are using *autoindent* you may wish to place a label at the left margin of a line. The way to do this easily is to type `↑` and then `^D`. The editor will move the cursor to the left margin for one line, and restore the previous indent on the next. You can also type a `0` followed immediately by a `^D` if you wish to kill all the indent and not have it come back on the next line.

### 8.6. Upper case only terminals

If your terminal has only upper case, you can still use *vi* by using the normal system convention for typing on such a terminal. Characters which you normally type are converted to lower case, and you can type upper case letters by preceding them with a `\`. The characters `{` `~` `|` `'` are not available on such terminals, but you can escape them as `\({\|}\) \! \\'`. These characters are represented on the display in the same way they are typed.† ‡

### 8.7. Vi and ex

*Vi* is actually one mode of editing within the editor *ex*. When you are running *vi* you can escape to the line oriented editor of *ex* by giving the command `Q`. All of the `:` commands which were introduced above are available in *ex*. Likewise, most *ex* commands can be invoked from *vi* using `:`. Just give them without the `:` and follow them with a CR.

In rare instances, an internal error may occur in *vi*. In this case you will get a diagnostic and be left in the command mode of *ex*. You can then save your work and quit if you wish by giving a command `x` after the `:` which *ex* prompts you with, or you can reenter *vi* by giving *ex* a *vi* command.

There are a number of things which you can do more easily in *ex* than in *vi*. Systematic changes in line oriented material are particularly easy. You can read the advanced editing documents for the editor *ed* to find out a lot more about this style of editing. Experienced users often mix their use of *ex* command mode and *vi* command mode to speed the work they are doing.

### 8.8. Open mode: vi on hardcopy terminals and "glass tty's" ‡

If you are on a hardcopy terminal or a terminal which does not have a cursor which can move off the bottom line, you can still use the command set of *vi*, but in a different mode. When you give a *vi* command, the editor will tell you that it is using *open* mode. This name comes from the *open* command in *ex*, which is used to get into the same mode.

The only difference between *visual* mode and *open* mode is the way in which the text is

---

\* This is not quite true. The implementation of the editor does not allow the NULL (`^@`) character to appear in files. Also the LF (linefeed or `^J`) character is used by the editor to separate lines in the file, so it cannot appear in the middle of a line. You can insert any other character, however, if you wait for the editor to echo the `↑` before you type the character. In fact, the editor will treat a following letter as a request for the corresponding control character. This is the only way to type `^S` or `^Q`, since the system normally uses them to suspend and resume output and never gives them to the editor to process.

† The `\` character you give will not echo until you type another key.

‡ Not available in all v2 editors due to memory constraints.

displayed.

In *open* mode the editor uses a single line window into the file, and moving backward and forward in the file causes new lines to be displayed, always below the current line. Two commands of *vi* work differently in *open*: *z* and *^R*. The *z* command does not take parameters, but rather draws a window of context around the current line and then returns you to the current line.

If you are on a hardcopy terminal, the *^R* command will retype the current line. On such terminals, the editor normally uses two lines to represent the current line. The first line is a copy of the line as you started to edit it, and you work on the line below this line. When you delete characters, the editor types a number of *\*'s to show you the characters which are deleted. The editor also reprints the current line soon after such changes so that you can see what the line looks like again.

It is sometimes useful to use this mode on very slow terminals which can support *vi* in the full screen mode. You can do this by entering *ex* and using an *open* command.

### Acknowledgements

Bruce Englar encouraged the early development of this display editor. Peter Kessler helped bring sanity to version 2's command layout. Bill Joy wrote versions 1 and 2.0 through 2.7, and created the framework that users see in the present editor. Mark Horton added macros and other features and made the editor work on a large number of terminals and Unix systems.

## Appendix: character functions

This appendix gives the uses the editor makes of each character. The characters are presented in their order in the ASCII character set: Control characters come first, then most special characters, then the digits, upper and then lower case characters.

For each character we tell a meaning it has as a command and any meaning it has during an insert. If it has only meaning as a command, then only this is discussed. Section numbers in parentheses indicate where the character is discussed; a 'f' after the section number means that the character is mentioned in a footnote.

- ^@** Not a command character. If typed as the first character of an insertion it is replaced with the last text inserted, and the insert terminates. Only 128 characters are saved from the last insert; if more characters were inserted the mechanism is not available. A ^@ cannot be part of the file due to the editor implementation (7.5f).
- ^A** Unused.
- ^B** Backward window. A count specifies repetition. Two lines of continuity are kept if possible (2.1, 6.1, 7.2).
- ^C** Unused.
- ^D** As a command, scrolls down a half-window of text. A count gives the number of (logical) lines to scroll, and is remembered for future ^D and ^U commands (2.1, 7.2). During an insert, backtabs over *autoindent* white space at the beginning of a line (6.6, 7.5); this white space cannot be backspaced over.
- ^E** Exposes one more line below the current screen in the file, leaving the cursor where it is if possible. (Version 3 only.)
- ^F** Forward window. A count specifies repetition. Two lines of continuity are kept if possible (2.1, 6.1, 7.2).
- ^G** Equivalent to :fCR, printing the current file, whether it has been modified, the current line number and the number of lines in the file, and the percentage of the way through the file that you are.
- ^H (BS)** Same as **left arrow**. (See h). During an insert, eliminates the last input character, backing over it but not erasing it; it remains so you can see what you typed if you wish to type something only slightly different (3.1, 7.5).
- ^I (TAB)** Not a command character. When inserted it prints as some number of spaces. When the cursor is at a tab character it rests at the last of the spaces which represent the tab. The spacing of tabstops is controlled by the *tabstop* option (4.1, 6.6).
- ^J (LF)** Same as **down arrow** (see j).
- ^K** Unused.
- ^L** The ASCII formfeed character, this causes the screen to be cleared and redrawn. This is useful after a transmission error, if characters typed by a program other than the editor scramble the screen, or after output is stopped by an interrupt (5.4, 7.2f).
- ^M (CR)** A carriage return advances to the next line, at the first non-white position in the line. Given a count, it advances that many lines (2.3). During an insert, a CR causes the insert to continue onto another line (3.1).
- ^N** Same as **down arrow** (see j).
- ^O** Unused.



- ^P** Same as **up arrow** (see **k**).
- ^Q** Not a command character. In input mode, **^Q** quotes the next character, the same as **^V**, except that some teletype drivers will eat the **^Q** so that the editor never sees it.
- ^R** Redraws the current screen, eliminating logical lines not corresponding to physical lines (lines with only a single **@** character on them). On hardcopy terminals in *open* mode, retypes the current line (5.4, 7.2, 7.8).
- ^S** Unused. Some teletype drivers use **^S** to suspend output until **^Q**is
- ^T** Not a command character. During an insert, with *autoindent* set and at the beginning of the line, inserts *shiftwidth* white space.
- ^U** Scrolls the screen up, inverting **^D** which scrolls down. Counts work as they do for **^D**, and the previous scroll amount is common to both. On a dumb terminal, **^U** will often necessitate clearing and redrawing the screen further back in the file (2.1, 7.2).
- ^V** Not a command character. In input mode, quotes the next character so that it is possible to insert non-printing and special characters into the file (4.2, 7.5).
- ^W** Not a command character. During an insert, backs up as **b** would in command mode; the deleted characters remain on the display (see **^H**) (7.5).
- ^X** Unused.
- ^Y** Exposes one more line above the current screen, leaving the cursor where it is if possible. (No mnemonic value for this key; however, it is next to **^U** which scrolls up a bunch.) (Version 3 only.)
- ^Z** If supported by the Unix system, stops the editor, exiting to the top level shell. Same as **:stopCR**. Otherwise, unused.
- ^[\_ (ESC)** Cancels a partially formed command, such as a **z** when no following character has yet been given; terminates inputs on the last line (read by commands such as **:/** and **?**); ends insertions of new text into the buffer. If an **ESC** is given when quiescent in command state, the editor rings the bell or flashes the screen. You can thus hit **ESC** if you don't know what is happening till the editor rings the bell. If you don't know if you are in insert mode you can type **ESCa**, and then material to be input; the material will be inserted correctly whether or not you were in insert mode when you started (1.5, 3.1, 7.5).
- ^`** Unused.
- ^]** Searches for the word which is after the cursor as a tag. Equivalent to typing **:ta**, this word, and then a **CR**. Mnemonically, this command is "go right to" (7.3).
- ^j** Equivalent to **:e #CR**, returning to the previous position in the last edited file, or editing a file which you specified if you got a 'No write since last change diagnostic' and do not want to have to type the file name again (7.3). (You have to do a **:w** before **^j** will work in this case. If you do not wish to write the file you should do **:e! #CR** instead.)
- ^\_** Unused. Reserved as the command character for the Tektronix 4025 and 4027 terminal.
- SPACE** Same as **right arrow** (see **l**).
- !** An operator, which processes lines from the buffer with reformatting commands. Follow **!** with the object to be processed, and then the command name terminated by **CR**. Doubling **!** and preceding it by a count causes count lines to be filtered; otherwise the count is passed on to the object after the **!**. Thus **2!)fmtCR** reformats the next two paragraphs by running them through the program *fmt*. If you are working on LISP, the command **!%grindCR,\*** given at the

---

\*Both *fmt* and *grind* are Berkeley programs and may not be present at all installations.

beginning of a function, will run the text of the function through the LISP grinder (6.7, 7.3). To read a file or the output of a command into the buffer use `:r` (7.3). To simply execute a command use `:! (7.3)`.

- " Precedes a named buffer specification. There are named buffers 1-9 used for saving deleted text and named buffers a-z into which you can place text (4.3, 6.3)
- # The macro character which, when followed by a number, will substitute for a function key on terminals without function keys (6.9). In input mode, if this is your erase character, it will delete the last character you typed in input mode, and must be preceded with a `\` to insert it, since it normally backs over the last input character you gave.
- \$ Moves to the end of the current line. If you `:se listCR`, then the end of each line will be shown by printing a \$ after the end of the displayed text in the line. Given a count, advances to the count'th following end of line; thus `2$` advances to the end of the following line.
- % Moves to the parenthesis or brace { } which balances the parenthesis or brace at the current cursor position.
- & A synonym for `:&CR`, by analogy with the `ex &` command.
- ' When followed by a ' returns to the previous context at the beginning of a line. The previous context is set whenever the current line is moved in a non-relative way. When followed by a letter a-z, returns to the line which was marked with this letter with a `m` command, at the first non-white character in the line. (2.2, 5.3). When used with an operator such as `d`, the operation takes place over complete lines; if you use ```, the operation takes place from the exact marked place to the current cursor position within the line.
- ( Retreats to the beginning of a sentence, or to the beginning of a LISP s-expression if the `lisp` option is set. A sentence ends at a `. !` or `?` which is followed by either the end of a line or by two spaces. Any number of closing `) ] " and ' characters may appear after the . ! or ?, and before the spaces or end of line. Sentences also begin at paragraph and section boundaries (see { and || below). A count advances that many sentences (4.2, 6.8).`
- ) Advances to the beginning of a sentence. A count repeats the effect. See ( above for the definition of a sentence (4.2, 6.8).
- \* Unused.
- + Same as `CR` when used as a command.
- , Reverse of the last `f F t` or `T` command, looking the other way in the current line. Especially useful after hitting too many `;` characters. A count repeats the search.
- Retreats to the previous line at the first non-white character. This is the inverse of `+` and `RETURN`. If the line moved to is not on the screen, the screen is scrolled, or cleared and redrawn if this is not possible. If a large amount of scrolling would be required the screen is also cleared and redrawn, with the current line at the center (2.3).
- . Repeats the last command which changed the buffer. Especially useful when deleting words or lines; you can delete some words/lines and then hit `.` to delete more and more words/lines. Given a count, it passes it on to the command being repeated. Thus after a `2dw`, `3.` deletes three words (3.3, 6.3, 7.2, 7.4).

- /** Reads a string from the last line on the screen, and scans forward for the next occurrence of this string. The normal input editing sequences may be used during the input on the bottom line; an returns to command state without ever searching. The search begins when you hit CR to terminate the pattern; the cursor moves to the beginning of the last line to indicate that the search is in progress; the search may then be terminated with a DEL or RUB, or by backspacing when at the beginning of the bottom line, returning the cursor to its initial position. Searches normally wrap end-around to find a string anywhere in the buffer.
- When used with an operator the enclosed region is normally affected. By mentioning an offset from the line matched by the pattern you can force whole lines to be affected. To do this give a pattern with a closing a closing / and then an offset  $+n$  or  $-n$ .
- To include the character / in the search string, you must escape it with a preceding \. A ↑ at the beginning of the pattern forces the match to occur at the beginning of a line only; this speeds the search. A \$ at the end of the pattern forces the match to occur at the end of a line only. More extended pattern matching is available, see section 7.4; unless you set **nomagic** in your .*exrc* file you will have to precede the characters . | \* and ~ in the search pattern with a \ to get them to work as you would naively expect (1.5, 2.2, 6.1, 7.2, 7.4).
- 0** Moves to the first character on the current line. Also used, in forming numbers, after an initial 1–9.
- 1–9** Used to form numeric arguments to commands (2.3, 7.2).
- :** A prefix to a set of commands for file and option manipulation and escapes to the system. Input is given on the bottom line and terminated with an CR, and the command then executed. You can return to where you were by hitting DEL or RUB if you hit : accidentally (see primarily 6.2 and 7.3).
- ;** Repeats the last single character find which used **f F t** or **T**. A count iterates the basic scan (4.1).
- <** An operator which shifts lines left one *shiftwidth*, normally 8 spaces. Like all operators, affects lines when repeated, as in <<. Counts are passed through to the basic object, thus 3<< shifts three lines (6.6, 7.2).
- =** Reindents line for LISP, as though they were typed in with *lisp* and *autoindent* set (6.8).
- >** An operator which shifts lines right one *shiftwidth*, normally 8 spaces. Affects lines when repeated as in >>. Counts repeat the basic object (6.6, 7.2).
- ?** Scans backwards, the opposite of /. See the / description above for details on scanning (2.2, 6.1, 7.4).
- @** A macro character (6.9). If this is your kill character, you must escape it with a \ to type it in during input mode, as it normally backs over the input you have given on the current line (3.1, 3.4, 7.5).
- A** Appends at the end of line, a synonym for **\$a** (7.2).
- B** Backs up a word, where words are composed of non-blank sequences, placing the cursor at the beginning of the word. A count repeats the effect (2.4).
- C** Changes the rest of the text on the current line; a synonym for **c\$**.
- D** Deletes the rest of the text on the current line; a synonym for **d\$**.

- E** Moves forward to the end of a word, defined as blanks and non-blanks, like **B** and **W**. A count repeats the effect.
- F** Finds a single following character, backwards in the current line. A count repeats this search that many times (4.1).
- G** Goes to the line number given as preceding argument, or the end of the file if no preceding count is given. The screen is redrawn with the new current line in the center if necessary (7.2).
- H** **Home arrow**. Homes the cursor to the top line on the screen. If a count is given, then the cursor is moved to the count'th line on the screen. In any case the cursor is moved to the first non-white character on the line. If used as the target of an operator, full lines are affected (2.3, 3.2).
- I** Inserts at the beginning of a line; a synonym for  $\uparrow$ i.
- J** Joins together lines, supplying appropriate white space: one space between words, two spaces after a ., and no spaces at all if the first character of the joined on line is ). A count causes that many lines to be joined rather than the default two (6.5, 7.1f).
- K** Unused.
- L** Moves the cursor to the first non-white character of the last line on the screen. With a count, to the first non-white of the count'th line from the bottom. Operators affect whole lines when used with **L** (2.3).
- M** Moves the cursor to the middle line on the screen, at the first non-white position on the line (2.3).
- N** Scans for the next match of the last pattern given to / or ?, but in the reverse direction; this is the reverse of **n**.
- O** Opens a new line above the current line and inputs text there up to an ESC. A count can be used on dumb terminals to specify a number of lines to be opened; this is generally obsolete, as the *slowopen* option works better (3.1).
- P** Puts the last deleted text back before/above the cursor. The text goes back as whole lines above the cursor if it was deleted as whole lines. Otherwise the text is inserted between the characters before and at the cursor. May be preceded by a named buffer specification "x to retrieve the contents of the buffer; buffers 1-9 contain deleted material, buffers a-z are available for general use (6.3).
- Q** Quits from *vi* to *ex* command mode. In this mode, whole lines form commands, ending with a RETURN. You can give all the : commands; the editor supplies the : as a prompt (7.7).
- R** Replaces characters on the screen with characters you type (overlay fashion). Terminates with an ESC.
- S** Changes whole lines, a synonym for **cc**. A count substitutes for that many lines. The lines are saved in the numeric buffers, and erased on the screen before the substitution begins.
- T** Takes a single following character, locates the character before the cursor in the current line, and places the cursor just after that character. A count repeats the effect. Most useful with operators such as **d** (4.1).
- U** Restores the current line to its state before you started changing it (3.5).
- V** Unused.

- W** Moves forward to the beginning of a word in the current line, where words are defined as sequences of blank/non-blank characters. A count repeats the effect (2.4).
- X** Deletes the character before the cursor. A count repeats the effect, but only characters on the current line are deleted.
- Y** Yanks a copy of the current line into the unnamed buffer, to be put back by a later **p** or **P**; a very useful synonym for **yy**. A count yanks that many lines. May be preceded by a buffer name to put lines in that buffer (7.4).
- ZZ** Exits the editor. (Same as **:xCR.**) If any changes have been made, the buffer is written out to the current file. Then the editor quits.
- ||** Backs up to the previous section boundary. A section begins at each macro in the *sections* option, normally a **‘.NH’** or **‘.SH’** and also at lines which which start with a formfeed **^L**. Lines beginning with **{** also stop **||**; this makes it useful for looking backwards, a function at a time, in C programs. If the option *lisp* is set, stops at each **(** at the beginning of a line, and is thus useful for moving backwards at the top level LISP objects. (4.2, 6.1, 6.6, 7.2).
- \** Unused.
- ||** Forward to a section boundary, see **||** for a definition (4.2, 6.1, 6.6, 7.2).
- |** Moves to the first non-white position on the current line (4.4).
- Unused.
- ‘** When followed by a **’** returns to the previous context. The previous context is set whenever the current line is moved in a non-relative way. When followed by a letter **a—z**, returns to the position which was marked with this letter with a **m** command. When used with an operator such as **d**, the operation takes place from the exact marked place to the current position within the line; if you use **’**, the operation takes place over complete lines (2.2, 5.3).
- a** Appends arbitrary text after the current cursor position; the insert can continue onto multiple lines by using **RETURN** within the insert. A count causes the inserted text to be replicated, but only if the inserted text is all on one line. The insertion terminates with an **ESC** (3.1, 7.2).
- b** Backs up to the beginning of a word in the current line. A word is a sequence of alphanumerics, or a sequence of special characters. A count repeats the effect (2.4).
- c** An operator which changes the following object, replacing it with the following input text up to an **ESC**. If more than part of a single line is affected, the text which is changed away is saved in the numeric named buffers. If only part of the current line is affected, then the last character to be changed away is marked with a **\$**. A count causes that many objects to be affected, thus both **3c** and **c3** change the following three sentences (7.4).
- d** An operator which deletes the following object. If more than part of a line is affected, the text is saved in the numeric buffers. A count causes that many objects to be affected; thus **3dw** is the same as **d3w** (3.3, 3.4, 4.1, 7.4).
- e** Advances to the end of the next word, defined as for **b** and **w**. A count repeats the effect (2.4, 3.1).
- f** Finds the first instance of the next character following the cursor on the current line. A count repeats the find (4.1).
- g** Unused.

Arrow keys **h**, **j**, **k**, **l**, and **H**.

- h** **Left arrow.** Moves the cursor one character to the left. Like the other arrow keys, either **h**, the **left arrow** key, or one of the synonyms (**^H**) has the same effect. On v2 editors, arrow keys on certain kinds of terminals (those which send escape sequences, such as vt52, c100, or hp) cannot be used. A count repeats the effect (3.1, 7.5).
- i** Inserts text before the cursor, otherwise like **a** (7.2).
- j** **Down arrow.** Moves the cursor one line down in the same column. If the position does not exist, *vi* comes as close as possible to the same column. Synonyms include **^J** (linefeed) and **^N**.
- k** **Up arrow.** Moves the cursor one line up. **^P** is a synonym.
- l** **Right arrow.** Moves the cursor one character to the right. **SPACE** is a synonym.
- m** Marks the current position of the cursor in the mark register which is specified by the next character **a-z**. Return to this position or use with an operator using **`** or **'** (5.3).
- n** Repeats the last / or ? scanning commands (2.2).
- o** Opens new lines below the current line; otherwise like **O** (3.1).
- p** Puts text after/below the cursor; otherwise like **P** (6.3).
- q** Unused.
- r** Replaces the single character at the cursor with a single character you type. The new character may be a **RETURN**; this is the easiest way to split lines. A count replaces each of the following count characters with the single character given; see **R** above which is the more usually useful iteration of **r** (3.2).
- s** Changes the single character under the cursor to the text which follows up to an **ESC**; given a count, that many characters from the current line are changed. The last character to be changed is marked with **\$** as in **c** (3.2).
- t** Advances the cursor upto the character before the next character typed. Most useful with operators such as **d** and **c** to delete the characters up to a following character. You can use **.** to delete more if this doesn't delete enough the first time (4.1).
- u** Undoes the last change made to the current buffer. If repeated, will alternate between these two states, thus is its own inverse. When used after an insert which inserted text on more than one line, the lines are saved in the numeric named buffers (3.5).
- v** Unused.
- w** Advances to the beginning of the next word, as defined by **b** (2.4).
- x** Deletes the single character under the cursor. With a count deletes deletes that many characters forward from the cursor position, but only on the current line (6.5).
- y** An operator, yanks the following object into the unnamed temporary buffer. If preceded by a named buffer specification, **"x**, the text is placed in that buffer also. Text can be recovered by a later **p** or **P** (7.4).
- z** Redraws the screen with the current line placed as specified by the following character: **RETURN** specifies the top of the screen, **.** the center of the screen, and **-** at the bottom of the screen. A count may be given after the **z** and before the following character to specify the new screen size for the redraw. A count before the **z** gives the number of the line to place in the center of the screen instead of the default current line. (5.4)

- { Retreats to the beginning of the beginning of the preceding paragraph. A paragraph begins at each macro in the *paragraphs* option, normally '.IP', '.LP', '.PP', '.QP' and '.bp'. A paragraph also begins after a completely empty line, and at each section boundary (see [[ above) (4.2, 6.8, 7.6).
- | Places the cursor on the character in the column specified by the count (7.1, 7.2).
- } Advances to the beginning of the next paragraph. See { for the definition of paragraph (4.2, 6.8, 7.6).
- Unused.
- ^? (DEL) Interrupts the editor, returning it to command accepting state (1.5, 7.5)





# NROFF/TROFF User's Manual

Joseph F. Ossanna

Bell Laboratories  
Murray Hill, New Jersey 07974

## Introduction

NROFF and TROFF are text processors under the PDP-11 UNIX Time-Sharing System<sup>1</sup> that format text for typewriter-like terminals and for a Graphic Systems phototypesetter, respectively. They accept lines of text interspersed with lines of format control information and format the text into a printable, paginated document having a user-designed style. NROFF and TROFF offer unusual freedom in document styling, including: arbitrary style headers and footers; arbitrary style footnotes; multiple automatic sequence numbering for paragraphs, sections, etc; multiple column output; dynamic font and point-size control; arbitrary horizontal and vertical local motions at any point; and a family of automatic overstriking, bracket construction, and line drawing functions.

NROFF and TROFF are highly compatible with each other and it is almost always possible to prepare input acceptable to both. Conditional input is provided that enables the user to embed input expressly destined for either program. NROFF can prepare output directly for a variety of terminal types and is capable of utilizing the full resolution of each terminal.

## Usage

The general form of invoking NROFF (or TROFF) at UNIX command level is

**nroff** *options files* (or **troff** *options files*)

where *options* represents any of a number of option arguments and *files* represents the list of files containing the document to be formatted. An argument consisting of a single minus (-) is taken to be a file name corresponding to the standard input. If no file names are given input is taken from the standard input. The options, which may appear in any order so long as they appear before the files, are:

| <i>Option</i> | <i>Effect</i>                                                                                                                                                                                                                                                                                                                         |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-olist</b> | Print only pages whose page numbers appear in <i>list</i> , which consists of comma-separated numbers and number ranges. A number range has the form $N-M$ and means pages $N$ through $M$ ; a initial $-N$ means from the beginning to page $N$ ; and a final $N-$ means from $N$ to the end.                                        |
| <b>-nN</b>    | Number first generated page $N$ .                                                                                                                                                                                                                                                                                                     |
| <b>-sN</b>    | Stop every $N$ pages. NROFF will halt prior to every $N$ pages (default $N=1$ ) to allow paper loading or changing, and will resume upon receipt of a newline. TROFF will stop the phototypesetter every $N$ pages, produce a trailer to allow changing cassettes, and will resume after the phototypesetter START button is pressed. |
| <b>-mname</b> | Prepends the macro file <code>/usr/lib/tmac.name</code> to the input <i>files</i> .                                                                                                                                                                                                                                                   |
| <b>-raN</b>   | Register <i>a</i> (one-character) is set to $N$ .                                                                                                                                                                                                                                                                                     |
| <b>-i</b>     | Read standard input after the input files are exhausted.                                                                                                                                                                                                                                                                              |
| <b>-q</b>     | Invoke the simultaneous input-output mode of the <code>rd</code> request.                                                                                                                                                                                                                                                             |

**NROFF Only**

- Tname Specifies the name of the output terminal type. Currently defined names are 37 for the (default) Model 37 Teletype®, tn300 for the GE TermiNet 300 (or any terminal without half-line capabilities), 300S for the DASI-300S, 300 for the DASI-300, and 450 for the DASI-450 (Diablo Hyterm).
- e Produce equally-spaced words in adjusted lines, using full terminal resolution.

**TROFF Only**

- t Direct output to the standard output instead of the phototypesetter.
- f Refrain from feeding out paper and stopping phototypesetter at the end of the run.
- w Wait until phototypesetter is available, if currently busy.
- b TROFF will report whether the phototypesetter is busy or available. No text processing is done.
- a Send a printable (ASCII) approximation of the results to the standard output.
- pN Print all characters in point size N while retaining all prescribed spacings and motions, to reduce phototypesetter elapsed time.
- g Prepare output for the Murray Hill Computation Center phototypesetter and direct it to the standard output.

Each option is invoked as a separate argument; for example,

```
nroff -o4,8-10 -T300S -mabc file1 file2
```

requests formatting of pages 4, 8, 9, and 10 of a document contained in the files named *file1* and *file2*, specifies the output terminal as a DASI-300S, and invokes the macro package *abc*.

Various pre- and post-processors are available for use with NROFF and TROFF. These include the equation preprocessors NEQN and EQN<sup>2</sup> (for NROFF and TROFF respectively), and the table-construction preprocessor TBL<sup>3</sup>. A reverse-line postprocessor COL<sup>4</sup> is available for multiple-column NROFF output on terminals without reverse-line ability; COL expects the Model 37 Teletype escape sequences that NROFF produces by default. TK<sup>4</sup> is a 37 Teletype simulator postprocessor for printing NROFF output on a Tektronix 4014. TCAT<sup>4</sup> is phototypesetter-simulator postprocessor for TROFF that produces an approximation of phototypesetter output on a Tektronix 4014. For example, in

```
tbl files | eqn | troff -t options | tcat
```

the first | indicates the piping of TBL's output to EQN's input; the second the piping of EQN's output to TROFF's input; and the third indicates the piping of TROFF's output to TCAT. GCAT<sup>4</sup> can be used to send TROFF (-g) output to the Murray Hill Computation Center.

The remainder of this manual consists of: a Summary and Index; a Reference Manual keyed to the index; and a set of Tutorial Examples. Another tutorial is [5].

Joseph F. Ossanna

**References**

- [1] K. Thompson, D. M. Ritchie, *UNIX Programmer's Manual*, Sixth Edition (May 1975).
- [2] B. W. Kernighan, L. L. Cherry, *Typesetting Mathematics - User's Guide (Second Edition)*, Bell Laboratories internal memorandum.
- [3] M. E. Lesk, *Tbl - A Program to Format Tables*, Bell Laboratories internal memorandum.
- [4] Internal on-line documentation, on UNIX.
- [5] B. W. Kernighan, *A TROFF Tutorial*, Bell Laboratories internal memorandum.

SUMMARY AND INDEX

| <i>Request Form</i>                                      | <i>Initial Value*</i> | <i>If No Argument</i> | <i>Notes#</i> | <i>Explanation</i>                                      |
|----------------------------------------------------------|-----------------------|-----------------------|---------------|---------------------------------------------------------|
| <b>1. General Explanation</b>                            |                       |                       |               |                                                         |
| <b>2. Font and Character Size Control</b>                |                       |                       |               |                                                         |
| .ps ±N                                                   | 10 point              | previous              | E             | Point size; also \s±N.†                                 |
| .ss N                                                    | 12/36 em              | ignored               | E             | Space-character size set to N/36 em.†                   |
| .cs FNM                                                  | off                   | -                     | P             | Constant character space (width) mode (font F).†        |
| .bd FN                                                   | off                   | -                     | P             | Embolden font F by N-1 units.†                          |
| .bd S FN                                                 | off                   | -                     | P             | Embolden Special Font when current font is F.†          |
| .ft F                                                    | Roman                 | previous              | E             | Change to font F = x, xx, or 1-4. Also \fx, \f(xx, \fN. |
| .fp NF                                                   | R,I,B,S               | ignored               | -             | Font named F mounted on physical position 1 ≤ N ≤ 4.    |
| <b>3. Page Control</b>                                   |                       |                       |               |                                                         |
| .pl ±N                                                   | 11 in                 | 11 in                 | v             | Page length.                                            |
| .bp ±N                                                   | N=1                   | -                     | B‡,v          | Eject current page; next page number N.                 |
| .pn ±N                                                   | N=1                   | ignored               | -             | Next page number N.                                     |
| .po ±N                                                   | 0; 26/27 in           | previous              | v             | Page offset.                                            |
| .ne N                                                    | -                     | N=1 V                 | D,v           | Need N vertical space (V = vertical spacing).           |
| .mk R                                                    | none                  | internal              | D             | Mark current vertical place in register R.              |
| .rt ±N                                                   | none                  | internal              | D,v           | Return ( <i>upward only</i> ) to marked vertical place. |
| <b>4. Text Filling, Adjusting, and Centering</b>         |                       |                       |               |                                                         |
| .br                                                      | -                     | -                     | B             | Break.                                                  |
| .fi                                                      | fill                  | -                     | B,E           | Fill output lines.                                      |
| .nf                                                      | fill                  | -                     | B,E           | No filling or adjusting of output lines.                |
| .ad c                                                    | adj,both              | adjust                | E             | Adjust output lines with mode c.                        |
| .na                                                      | adjust                | -                     | E             | No output line adjusting.                               |
| .ce N                                                    | off                   | N=1                   | B,E           | Center following N input text lines.                    |
| <b>5. Vertical Spacing</b>                               |                       |                       |               |                                                         |
| .vs N                                                    | 1/6in;12pts           | previous              | E,p           | Vertical base line spacing (V).                         |
| .ls N                                                    | N=1                   | previous              | E             | Output N-1 Vs after each text output line.              |
| .sp N                                                    | -                     | N=1 V                 | B,v           | Space vertical distance N in either direction.          |
| .sv N                                                    | -                     | N=1 V                 | v             | Save vertical distance N.                               |
| .os                                                      | -                     | -                     | -             | Output saved vertical distance.                         |
| .ns                                                      | space                 | -                     | D             | Turn no-space mode on.                                  |
| .rs                                                      | -                     | -                     | D             | Restore spacing; turn no-space mode off.                |
| <b>6. Line Length and Indenting</b>                      |                       |                       |               |                                                         |
| .ll ±N                                                   | 6.5 in                | previous              | E,m           | Line length.                                            |
| .in ±N                                                   | N=0                   | previous              | B,E,m         | Indent.                                                 |
| .ti ±N                                                   | -                     | ignored               | B,E,m         | Temporary indent.                                       |
| <b>7. Macros, Strings, Diversion, and Position Traps</b> |                       |                       |               |                                                         |
| .de xx yy                                                | -                     | .yy=..                | -             | Define or redefine macro xx; end at call of yy.         |
| .am xx yy                                                | -                     | .yy=..                | -             | Append to a macro.                                      |
| .ds xx string                                            | -                     | ignored               | -             | Define a string xx containing string.                   |
| .as xx string                                            | -                     | ignored               | -             | Append string to string xx.                             |

\*Values separated by ";" are for NROFF and TROFF respectively.

#Notes are explained at the end of this Summary and Index

†No effect in NROFF.

‡The use of " " as control character (instead of ".") suppresses the break function.

| <i>Request Form</i>                                                      | <i>Initial Value</i> | <i>If No Argument</i> | <i>Notes</i> | <i>Explanation</i>                                                                                    |
|--------------------------------------------------------------------------|----------------------|-----------------------|--------------|-------------------------------------------------------------------------------------------------------|
| .rm <i>xx</i>                                                            | -                    | ignored               | -            | Remove request, macro, or string.                                                                     |
| .rn <i>xx yy</i>                                                         | -                    | ignored               | -            | Rename request, macro, or string <i>xx</i> to <i>yy</i> .                                             |
| .di <i>xx</i>                                                            | -                    | end                   | D            | Divert output to macro <i>xx</i> .                                                                    |
| .da <i>xx</i>                                                            | -                    | end                   | D            | Divert and append to <i>xx</i> .                                                                      |
| .wh <i>Nxx</i>                                                           | -                    | -                     | v            | Set location trap; negative is w.r.t. page bottom.                                                    |
| .ch <i>xx N</i>                                                          | -                    | -                     | v            | Change trap location.                                                                                 |
| .dt <i>Nxx</i>                                                           | -                    | off                   | D,v          | Set a diversion trap.                                                                                 |
| .it <i>Nxx</i>                                                           | -                    | off                   | E            | Set an input-line count trap.                                                                         |
| .em <i>xx</i>                                                            | none                 | none                  | -            | End macro is <i>xx</i> .                                                                              |
| <b>8. Number Registers</b>                                               |                      |                       |              |                                                                                                       |
| .nr <i>R ±NM</i>                                                         | -                    | -                     | u            | Define and set number register <i>R</i> ; auto-increment by <i>M</i> .                                |
| .af <i>R c</i>                                                           | arabic               | -                     | -            | Assign format to register <i>R</i> ( <i>c</i> =1, i, I, a, A).                                        |
| .rr <i>R</i>                                                             | -                    | -                     | -            | Remove register <i>R</i> .                                                                            |
| <b>9. Tabs, Leaders, and Fields</b>                                      |                      |                       |              |                                                                                                       |
| .ta <i>Nt ...</i>                                                        | 0.8; 0.5in           | none                  | E,m          | Tab settings; <i>left</i> type, unless <i>t</i> =R(right), C(centered).                               |
| .tc <i>c</i>                                                             | none                 | none                  | E            | Tab repetition character.                                                                             |
| .lc <i>c</i>                                                             | .                    | none                  | E            | Leader repetition character.                                                                          |
| .fc <i>a b</i>                                                           | off                  | off                   | -            | Set field delimiter <i>a</i> and pad character <i>b</i> .                                             |
| <b>10. Input and Output Conventions and Character Translations</b>       |                      |                       |              |                                                                                                       |
| .ec <i>c</i>                                                             | \                    | \                     | -            | Set escape character.                                                                                 |
| .eo                                                                      | on                   | -                     | -            | Turn off escape character mechanism.                                                                  |
| .lg <i>N</i>                                                             | -; on                | on                    | -            | Ligature mode on if <i>N</i> >0.                                                                      |
| .ul <i>N</i>                                                             | off                  | <i>N</i> =1           | E            | Underline (italicize in TROFF) <i>N</i> input lines.                                                  |
| .cu <i>N</i>                                                             | off                  | <i>N</i> =1           | E            | Continuous underline in NROFF; like <i>ul</i> in TROFF.                                               |
| .uf <i>F</i>                                                             | Italic               | Italic                | -            | Underline font set to <i>F</i> (to be switched to by <i>ul</i> ).                                     |
| .cc <i>c</i>                                                             | :                    | :                     | E            | Set control character to <i>c</i> .                                                                   |
| .c2 <i>c</i>                                                             | :                    | :                     | E            | Set nobreak control character to <i>c</i> .                                                           |
| .tr <i>abcd....</i>                                                      | none                 | -                     | O            | Translate <i>a</i> to <i>b</i> , etc. on output.                                                      |
| <b>11. Local Horizontal and Vertical Motions, and the Width Function</b> |                      |                       |              |                                                                                                       |
| <b>12. Overstrike, Bracket, Line-drawing, and Zero-width Functions</b>   |                      |                       |              |                                                                                                       |
| <b>13. Hyphenation.</b>                                                  |                      |                       |              |                                                                                                       |
| .nh                                                                      | hyphenate            | -                     | E            | No hyphenation.                                                                                       |
| .hy <i>N</i>                                                             | hyphenate            | hyphenate             | E            | Hyphenate; <i>N</i> = mode.                                                                           |
| .hc <i>c</i>                                                             | \%                   | \%                    | E            | Hyphenation indicator character <i>c</i> .                                                            |
| .hw <i>word1 ...</i>                                                     |                      | ignored               | -            | Exception words.                                                                                      |
| <b>14. Three Part Titles.</b>                                            |                      |                       |              |                                                                                                       |
| .tl ' <i>left center right</i> '                                         |                      | -                     | -            | Three part title.                                                                                     |
| .pc <i>c</i>                                                             | %                    | off                   | -            | Page number character.                                                                                |
| .lt <i>±N</i>                                                            | 6.5 in               | previous              | E,m          | Length of title.                                                                                      |
| <b>15. Output Line Numbering.</b>                                        |                      |                       |              |                                                                                                       |
| .nm <i>±NMSI</i>                                                         |                      | off                   | E            | Number mode on or off, set parameters.                                                                |
| .nn <i>N</i>                                                             | -                    | <i>N</i> =1           | E            | Do not number next <i>N</i> lines.                                                                    |
| <b>16. Conditional Acceptance of Input</b>                               |                      |                       |              |                                                                                                       |
| .if <i>c anything</i>                                                    |                      | -                     | -            | If condition <i>c</i> true, accept <i>anything</i> as input, for multi-line use <i>\{anything\}</i> . |

| <i>Request Form</i>                             | <i>Initial Value</i> | <i>If No Argument</i> | <i>Notes</i> | <i>Explanation</i>                                                              |
|-------------------------------------------------|----------------------|-----------------------|--------------|---------------------------------------------------------------------------------|
| .if ! <i>c anything</i>                         |                      | -                     | -            | If condition <i>c</i> false, accept <i>anything</i> .                           |
| .if <i>N anything</i>                           |                      | -                     | u            | If expression <i>N</i> > 0, accept <i>anything</i> .                            |
| .if ! <i>N anything</i>                         |                      | -                     | u            | If expression <i>N</i> ≤ 0, accept <i>anything</i> .                            |
| .if ' <i>string1 string2</i> ' <i>anything</i>  |                      | -                     | -            | If <i>string1</i> identical to <i>string2</i> , accept <i>anything</i> .        |
| .if '! <i>string1 string2</i> ' <i>anything</i> |                      | -                     | -            | If <i>string1</i> not identical to <i>string2</i> , accept <i>anything</i> .    |
| .ie <i>c anything</i>                           |                      | -                     | u            | If portion of if-else; all above forms (like if).                               |
| .el <i>anything</i>                             |                      | -                     | -            | Else portion of if-else.                                                        |
| <b>17. Environment Switching.</b>               |                      |                       |              |                                                                                 |
| .ev <i>N</i>                                    | <i>N=0</i>           | previous              | -            | Environment switched ( <i>push down</i> ).                                      |
| <b>18. Insertions from the Standard Input</b>   |                      |                       |              |                                                                                 |
| .rd <i>prompt</i>                               | -                    | <i>prompt=BEL</i>     | -            | Read insertion.                                                                 |
| .ex                                             | -                    | -                     | -            | Exit from NROFF/TROFF.                                                          |
| <b>19. Input/Output File Switching</b>          |                      |                       |              |                                                                                 |
| .so <i>filename</i>                             |                      | -                     | -            | Switch source file ( <i>push down</i> ).                                        |
| .nx <i>filename</i>                             |                      | end-of-file           | -            | Next file.                                                                      |
| .pi <i>program</i>                              |                      | -                     | -            | Pipe output to <i>program</i> (NROFF only).                                     |
| <b>20. Miscellaneous</b>                        |                      |                       |              |                                                                                 |
| .mc <i>c N</i>                                  | -                    | off                   | E,m          | Set margin character <i>c</i> and separation <i>N</i> .                         |
| .tm <i>string</i>                               | -                    | newline               | -            | Print <i>string</i> on terminal (UNIX standard message output).                 |
| .lg <i>yy</i>                                   | -                    | .yy=..                | -            | Ignore till call of <i>yy</i> .                                                 |
| .pm <i>t</i>                                    | -                    | all                   | -            | Print macro names and sizes;<br>if <i>t</i> present, print only total of sizes. |
| .fl                                             | -                    | -                     | B            | Flush output buffer.                                                            |
| <b>21. Output and Error Messages</b>            |                      |                       |              |                                                                                 |

Notes-

- B Request normally causes a break.
- D Mode or relevant parameters associated with current diversion level.
- E Relevant parameters are a part of the current environment.
- O Must stay in effect until logical output.
- P Mode must be still or again in effect at the time of physical output.
- v,p,m,u Default scale indicator; if not specified, scale indicators are *ignored*.

Alphabetical Request and Section Number Cross Reference

|       |       |       |       |       |       |       |       |       |       |      |
|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|------|
| ad 4  | cc 10 | ds 7  | fc 9  | ie 16 | ll 6  | nh 13 | pi 19 | rn 7  | ta 9  | vs 5 |
| af 8  | ce 4  | dt 7  | fi 4  | if 16 | ls 5  | nm 15 | pl 3  | rr 8  | tc 9  | wh 7 |
| am 7  | ch 7  | ec 10 | fl 20 | ig 20 | lt 14 | nn 15 | pm 20 | rs 5  | ti 6  |      |
| as 7  | cs 2  | el 16 | fp 2  | in 6  | mc 20 | nr 8  | pn 3  | rt 3  | tl 14 |      |
| bd 2  | cu 10 | em 7  | ft 2  | it 7  | mk 3  | ns 5  | po 3  | so 19 | tm 20 |      |
| bp 3  | da 7  | eo 10 | hc 13 | lc 9  | na 4  | nx 19 | ps 2  | sp 5  | tr 10 |      |
| br 4  | de 7  | ev 17 | hw 13 | lg 10 | ne 3  | os 5  | rd 18 | ss 2  | uf 10 |      |
| c2 10 | di 7  | ex 18 | hy 13 | li 10 | nf 4  | pc 14 | rm 7  | sv 5  | ul 10 |      |

**Escape Sequences for Characters, Indicators, and Functions**

| <i>Section Reference</i> | <i>Escape Sequence</i> | <i>Meaning</i>                                                       |
|--------------------------|------------------------|----------------------------------------------------------------------|
| 10.1                     | \\                     | \ (to prevent or delay the interpretation of \)                      |
| 10.1                     | \e                     | Printable version of the <i>current</i> escape character.            |
| 2.1                      | \`                     | ` (acute accent); equivalent to \aa                                  |
| 2.1                      | \^                     | ^ (grave accent); equivalent to \ga                                  |
| 2.1                      | \-                     | - Minus sign in the <i>current</i> font                              |
| 7                        | \.                     | Period (dot) (see de)                                                |
| 11.1                     | \(space)               | Unpaddable space-size space character                                |
| 11.1                     | \0                     | Digit width space                                                    |
| 11.1                     | \                      | 1/6 em narrow space character (zero width in NROFF)                  |
| 11.1                     | \^                     | 1/12 em half-narrow space character (zero width in NROFF)            |
| 4.1                      | \&                     | Non-printing, zero width character                                   |
| 10.6                     | \!                     | Transparent line indicator                                           |
| 10.7                     | \*                     | Beginning of comment                                                 |
| 7.3                      | \\$N                   | Interpolate argument $1 \leq N \leq 9$                               |
| 13                       | \%                     | Default optional hyphenation character                               |
| 2.1                      | \(xx                   | Character named xx                                                   |
| 7.1                      | \=x, \=(xx)            | Interpolate string x or xx                                           |
| 9.1                      | \a                     | Non-interpreted leader character                                     |
| 12.3                     | \b'abc...'             | Bracket building function                                            |
| 4.2                      | \c                     | Interrupt text processing                                            |
| 11.1                     | \d                     | Forward (down) 1/2 em vertical motion (1/2 line in NROFF)            |
| 2.2                      | \fx,\f(xx),\fN         | Change to font named x or xx, or position N.                         |
| 11.1                     | \h'N'                  | Local horizontal motion; move right N ( <i>negative left</i> )       |
| 11.3                     | \kx                    | Mark horizontal <i>input</i> place in register x                     |
| 12.4                     | \l'Nc'                 | Horizontal line drawing function (optionally with c)                 |
| 12.4                     | \L'Nc'                 | Vertical line drawing function (optionally with c)                   |
| 8                        | \nx,\n(xx)             | Interpolate number register x or xx                                  |
| 12.1                     | \o'abc...'             | Overstrike characters a, b, c, ...                                   |
| 4.1                      | \p                     | Break and spread output line                                         |
| 11.1                     | \r                     | Reverse 1 em vertical motion (reverse line in NROFF)                 |
| 2.3                      | \sN,\s±N               | Point-size change function                                           |
| 9.1                      | \t                     | Non-interpreted horizontal tab                                       |
| 11.1                     | \u                     | Reverse (up) 1/2 em vertical motion (1/2 line in NROFF)              |
| 11.1                     | \v'N'                  | Local vertical motion; move down N ( <i>negative up</i> )            |
| 11.2                     | \w'string'             | Interpolate width of <i>string</i>                                   |
| 5.2                      | \x'N'                  | Extra line-space function ( <i>negative before, positive after</i> ) |
| 12.2                     | \zc                    | Print c with zero width (without spacing)                            |
| 16                       | \{                     | Begin conditional input                                              |
| 16                       | \}                     | End conditional input                                                |
| 10.7                     | \(newline)             | Concealed (ignored) newline                                          |
| -                        | \X                     | X, any character <i>not</i> listed above                             |

The escape sequences \\, \., \', \\$, \\*, \a, \n, \t, and \ (newline) are interpreted in *copy mode* (§7.2).

**Predefined General Number Registers**

| <i>Section Reference</i> | <i>Register Name</i> | <i>Description</i>                                                     |
|--------------------------|----------------------|------------------------------------------------------------------------|
| 3                        | %                    | Current page number.                                                   |
| 11.2                     | ct                   | Character type (set by <i>width</i> function).                         |
| 7.4                      | dl                   | Width (maximum) of last completed diversion.                           |
| 7.4                      | dn                   | Height (vertical size) of last completed diversion.                    |
| -                        | dw                   | Current day of the week (1-7).                                         |
| -                        | dy                   | Current day of the month (1-31).                                       |
| 11.3                     | hp                   | Current horizontal place on <i>input</i> line.                         |
| 15                       | ln                   | Output line number.                                                    |
| -                        | mo                   | Current month (1-12).                                                  |
| 4.1                      | nl                   | Vertical position of last printed text base-line.                      |
| 11.2                     | sb                   | Depth of string below base line (generated by <i>width</i> function).  |
| 11.2                     | st                   | Height of string above base line (generated by <i>width</i> function). |
| -                        | yr                   | Last two digits of current year.                                       |

**Predefined Read-Only Number Registers**

| <i>Section Reference</i> | <i>Register Name</i> | <i>Description</i>                                                                 |
|--------------------------|----------------------|------------------------------------------------------------------------------------|
| 7.3                      | .S                   | Number of arguments available at the current macro level.                          |
| -                        | .A                   | Set to 1 in TROFF, if <i>-a</i> option used; always 1 in NROFF.                    |
| 11.1                     | .H                   | Available horizontal resolution in basic units.                                    |
| -                        | .T                   | Set to 1 in NROFF, if <i>-T</i> option used; always 0 in TROFF.                    |
| 11.1                     | .V                   | Available vertical resolution in basic units.                                      |
| 5.2                      | .a                   | Post-line extra line-space most recently utilized using <i>\x'N'</i> .             |
| -                        | .c                   | Number of <i>lines</i> read from current input file.                               |
| 7.4                      | .d                   | Current vertical place in current diversion; equal to <i>nl</i> , if no diversion. |
| 2.2                      | .f                   | Current font as physical quadrant (1-4).                                           |
| 4                        | .h                   | Text base-line high-water mark on current page or diversion.                       |
| 6                        | .i                   | Current indent.                                                                    |
| 6                        | .l                   | Current line length.                                                               |
| 4                        | .n                   | Length of text portion on previous output line.                                    |
| 3                        | .o                   | Current page offset.                                                               |
| 3                        | .p                   | Current page length.                                                               |
| 2.3                      | .s                   | Current point size.                                                                |
| 7.5                      | .t                   | Distance to the next trap.                                                         |
| 4.1                      | .u                   | Equal to 1 in fill mode and 0 in nofill mode.                                      |
| 5.1                      | .v                   | Current vertical line spacing.                                                     |
| 11.2                     | .w                   | Width of previous character.                                                       |
| -                        | .x                   | Reserved version-dependent register.                                               |
| -                        | .y                   | Reserved version-dependent register.                                               |
| 7.4                      | .z                   | Name of current diversion.                                                         |

## REFERENCE MANUAL

### 1. General Explanation

*1.1. Form of input.* Input consists of *text lines*, which are destined to be printed, interspersed with *control lines*, which set parameters or otherwise control subsequent processing. Control lines begin with a *control character*—normally . (period) or ' (acute accent)—followed by a one or two character name that specifies a basic *request* or the substitution of a user-defined *macro* in place of the control line. The control character ' suppresses the *break* function—the forced output of a partially filled line—caused by certain requests. The control character may be separated from the request/macro name by white space (spaces and/or tabs) for esthetic reasons. Names must be followed by either space or newline. Control lines with unrecognized names are ignored.

Various special functions may be introduced anywhere in the input by means of an *escape* character, normally \. For example, the function \nR causes the interpolation of the contents of the *number register* R in place of the function; here R is either a single character name as in \nx, or left-parenthesis-introduced, two-character name as in \n(xx).

*1.2. Formatter and device resolution.* TROFF internally uses 432 units/inch, corresponding to the Graphic Systems phototypesetter which has a horizontal resolution of 1/432 inch and a vertical resolution of 1/144 inch. NROFF internally uses 240 units/inch, corresponding to the least common multiple of the horizontal and vertical resolutions of various typewriter-like output devices. TROFF rounds horizontal/vertical numerical parameter input to the actual horizontal/vertical resolution of the Graphic Systems typesetter. NROFF similarly rounds numerical input to the actual resolution of the output device indicated by the -T option (default Model 37 Teletype).

*1.3. Numerical parameter input.* Both NROFF and TROFF accept numerical input with the appended scale indicators shown in the following table, where S is the current type size in points, V is the current vertical line spacing in basic units, and C is a *nominal character width* in basic units.

| Scale Indicator | Meaning             | Number of basic units |               |
|-----------------|---------------------|-----------------------|---------------|
|                 |                     | TROFF                 | NROFF         |
| i               | Inch                | 432                   | 240           |
| c               | Centimeter          | 432×50/127            | 240×50/127    |
| P               | Pica = 1/6 inch     | 72                    | 240/6         |
| m               | Em = S points       | 6×S                   | C             |
| n               | En = Em/2           | 3×S                   | C, same as Em |
| p               | Point = 1/72 inch   | 6                     | 240/72        |
| u               | Basic unit          | 1                     | 1             |
| v               | Vertical line space | V                     | V             |
| none            | Default, see below  |                       |               |

In NROFF, both the em and the en are taken to be equal to the C, which is output-device dependent; common values are 1/10 and 1/12 inch. Actual character widths in NROFF need not be all the same and constructed characters such as -> (→) are often extra wide. The default scaling is ems for the horizontally-oriented requests and functions ll, in, tl, ta, lt, po, mc, \h, and \l; Vs for the vertically-oriented requests and functions pl, wh, ch, dt, sp, sv, ne, rt, \v, \x, and \L; p for the vs request; and u for the requests nr, if, and ie. All other requests ignore any scale indicators. When a number register containing an already appropriately scaled number is interpolated to provide numerical input, the unit scale indicator u may need to be appended to prevent an additional inappropriate default scaling.



The number,  $N$ , may be specified in decimal-fraction form but the parameter finally stored is rounded to an integer number of basic units.

The *absolute position* indicator | may be prepended to a number  $N$  to generate the distance to the vertical or horizontal place  $N$ . For vertically-oriented requests and functions, | $N$  becomes the distance in basic units from the current vertical place on the page or in a *diversion* (§7.4) to the the vertical place  $N$ . For all other requests and functions, | $N$  becomes the distance from the current horizontal place on the *input* line to the horizontal place  $N$ . For example,

`.sp |3.2c`

will space *in the required direction* to 3.2 centimeters from the top of the page.

**1.4. Numerical expressions.** Wherever numerical input is expected an expression involving parentheses, the arithmetic operators +, -, /, \*, % (mod), and the logical operators <, >, <=, >=, = (or ==), & (and), : (or) may be used. Except where controlled by parentheses, evaluation of expressions is left-to-right; there is no operator precedence. In the case of certain requests, an initial + or - is stripped and interpreted as an increment or decrement indicator respectively. In the presence of default scaling, the desired scale indicator must be attached to *every* number in an expression for which the desired and default scaling differ. For example, if the number register x contains 2 and the current point size is 10, then

`.ll (4.25i+\nxP+3)/2u`

will set the line length to 1/2 the sum of 4.25 inches + 2 picas + 30 points.

**1.5. Notation.** Numerical parameters are indicated in this manual in two ways.  $\pm N$  means that the argument may take the forms  $N$ ,  $+N$ , or  $-N$  and that the corresponding effect is to set the affected parameter to  $N$ , to increment it by  $N$ , or to decrement it by  $N$  respectively. Plain  $N$  means that an initial algebraic sign is *not* an increment indicator, but merely the sign of  $N$ . Generally, unreasonable numerical input is either ignored or truncated to a reasonable value. For example, most requests expect to set parameters to non-negative values; exceptions are `sp`, `wh`, `ch`, `nr`, and `if`. The requests `ps`, `ft`, `po`, `vs`, `ls`, `ll`, `in`, and `lt` restore the *previous* parameter value in the *absence* of an argument.

Single character arguments are indicated by single lower case letters and one/two character arguments are indicated by a pair of lower case letters. Character string arguments are indicated by multi-character mnemonics.

## 2. Font and Character Size Control

**2.1. Character set.** The TROFF character set consists of the Graphics Systems Commercial II character set plus a Special Mathematical Font character set—each having 102 characters. These character sets are shown in the attached Table I. All ASCII characters are included, with some on the Special Font. With three exceptions, the ASCII characters are input as themselves, and non-ASCII characters are input in the form `\(xx` where `xx` is a two-character name given in the attached Table II. The three ASCII exceptions are mapped as follows:

| ASCII Input |              | Printed by TROFF |             |
|-------------|--------------|------------------|-------------|
| Character   | Name         | Character        | Name        |
| '           | acute accent | '                | close quote |
| `           | grave accent | '                | open quote  |
| -           | minus        | -                | hyphen      |

The characters ' , ` , and - may be input by \', \`, and \- respectively or by their names (Table II). The ASCII characters @, #, ", ' , ` , < , > , \ , { , } , ~ , ^ , and \_ exist only on the Special Font and are printed as a 1-em space if that Font is not mounted.

NROFF understands the entire TROFF character set, but can in general print only ASCII characters, additional characters as may be available on the output device, such characters as may be able to be constructed by overstriking or other combination, and those that can reasonably be mapped into other printable characters. The exact behavior is determined by a driving table prepared for each device. The

characters ` , ` , and \_ print as themselves.

2.2. *Fonts.* The default mounted fonts are Times Roman (R), Times Italic (I), Times Bold (B), and the Special Mathematical Font (S) on physical typesetter positions 1, 2, 3, and 4 respectively. These fonts are used in this document. The *current* font, initially Roman, may be changed (among the mounted fonts) by use of the *ft* request, or by imbedding at any desired point either  $\backslash fx$ ,  $\backslash f(xx)$ , or  $\backslash fN$  where *x* and *xx* are the name of a mounted font and *N* is a numerical font position. It is *not* necessary to change to the Special font; characters on that font are automatically handled. A request for a named but not-mounted font is *ignored*. TROFF can be informed that any particular font is mounted by use of the *fp* request. The list of known fonts is installation dependent. In the subsequent discussion of font-related requests, *F* represents either a one/two-character font name or the numerical font position, 1-4. The current font is available (as numerical position) in the read-only number register .*f*.

NROFF understands font control and normally underlines Italic characters (see §10.5).

2.3. *Character size.* Character point sizes available on the Graphic Systems typesetter are 6, 7, 8, 9, 10, 11, 12, 14, 16, 18, 20, 22, 24, 28, and 36. This is a range of 1/12 inch to 1/2 inch. The *ps* request is used to change or restore the point size. Alternatively the point size may be changed between any two characters by imbedding a  $\backslash sN$  at the desired point to set the size to *N*, or a  $\backslash s\pm N$  ( $1 \leq N \leq 9$ ) to increment/decrement the size by *N*;  $\backslash s0$  restores the *previous* size. Requested point size values that are between two valid sizes yield the larger of the two. The current size is available in the .*s* register. NROFF ignores type size control.

| <i>Request Form</i>    | <i>Initial Value</i> | <i>If No Argument</i> | <i>Notes*</i> | <i>Explanation</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------------------------|----------------------|-----------------------|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| . <i>ps</i> $\pm N$    | 10 point             | previous              | E             | Point size set to $\pm N$ . Alternatively imbed $\backslash sN$ or $\backslash s\pm N$ . Any positive size value may be requested; if invalid, the next larger valid size will result, with a maximum of 36. A paired sequence $+N, -N$ will work because the previous requested value is also remembered. Ignored in NROFF.                                                                                                                                                                                                                                                                                                                |
| . <i>ss</i> <i>N</i>   | 12/36 em             | ignored               | E             | Space-character size is set to <i>N</i> /36 ems. This size is the minimum word spacing in adjusted text. Ignored in NROFF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| . <i>cs</i> <i>FNM</i> | off                  | -                     | P             | Constant character space (width) mode is set on for font <i>F</i> (if mounted); the width of every character will be taken to be <i>N</i> /36 ems. If <i>M</i> is absent, the em is that of the character's point size; if <i>M</i> is given, the em is <i>M</i> -points. All affected characters are centered in this space, including those with an actual width larger than this space. Special Font characters occurring while the current font is <i>F</i> are also so treated. If <i>N</i> is absent, the mode is turned off. The mode must be still or again in effect when the characters are physically printed. Ignored in NROFF. |
| . <i>bd</i> <i>FN</i>  | off                  | -                     | P             | The characters in font <i>F</i> will be artificially emboldened by printing each one twice, separated by <i>N</i> -1 basic units. A reasonable value for <i>N</i> is 3 when the character size is in the vicinity of 10 points. If <i>N</i> is missing the embolden mode is turned off. The column heads above were printed with . <i>bd</i> I 3. The mode must be still or again in effect when the characters are physically printed. Ignored in NROFF.                                                                                                                                                                                   |

\*Notes are explained at the end of the Summary and Index above.

|                        |         |          |   |                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------|---------|----------|---|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>.bd</b> <i>S FN</i> | off     | -        | P | The characters in the Special Font will be emboldened whenever the current font is <i>F</i> . This manual was printed with <b>.bdSB3</b> . The mode must be still or again in effect when the characters are physically printed.                                                                                                                                                                  |
| <b>.ft</b> <i>F</i>    | Roman   | previous | E | Font changed to <i>F</i> . Alternatively, imbed <code>\fF</code> . The font name <i>P</i> is reserved to mean the previous font.                                                                                                                                                                                                                                                                  |
| <b>.fp</b> <i>N F</i>  | R,I,B,S | ignored  | - | Font position. This is a statement that a font named <i>F</i> is mounted on position <i>N</i> (1-4). It is a fatal error if <i>F</i> is not known. The phototypesetter has four fonts physically mounted. Each font consists of a film strip which can be mounted on a numbered quadrant of a wheel. The default mounting sequence assumed by TROFF is R, I, B, and S on positions 1, 2, 3 and 4. |

### 3. Page control

Top and bottom margins are *not* automatically provided; it is conventional to define two *macros* and to set *traps* for them at vertical positions 0 (top) and  $-N$  ( $N$  from the bottom). See §7 and Tutorial Examples §T2. A pseudo-page transition onto the *first* page occurs either when the first *break* occurs or when the first *non-diverted* text processing occurs. Arrangements for a trap to occur at the top of the first page must be completed before this transition. In the following, references to the *current diversion* (§7.4) mean that the mechanism being described works during both ordinary and diverted output (the former considered as the top diversion level).

The useable page width on the Graphic Systems phototypesetter is about 7.54 inches, beginning about 1/27 inch from the left edge of the 8 inch wide, continuous roll paper. The physical limitations on NROFF output are output-device dependent.

| <i>Request Form</i> | <i>Initial Value</i> | <i>If No Argument</i> | <i>Notes</i> | <i>Explanation</i>                                                                                                                                                                                                                                                                                                                    |
|---------------------|----------------------|-----------------------|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>.pl</b> $\pm N$  | 11 in                | 11 in                 | v            | Page length set to $\pm N$ . The internal limitation is about 75 inches in TROFF and about 136 inches in NROFF. The current page length is available in the <b>.p</b> register.                                                                                                                                                       |
| <b>.bp</b> $\pm N$  | $N=1$                | -                     | B*,v         | Begin page. The current page is ejected and a new page is begun. If $\pm N$ is given, the new page number will be $\pm N$ . Also see request <b>ns</b> .                                                                                                                                                                              |
| <b>.pn</b> $\pm N$  | $N=1$                | ignored               | -            | Page number. The next page (when it occurs) will have the page number $\pm N$ . A <b>pn</b> must occur before the initial pseudo-page transition to effect the page number of the first page. The current page number is in the <b>%</b> register.                                                                                    |
| <b>.po</b> $\pm N$  | 0; 26/27 in†         | previous              | v            | Page offset. The current <i>left margin</i> is set to $\pm N$ . The TROFF initial value provides about 1 inch of paper margin including the physical typesetter margin of 1/27 inch. In TROFF the maximum (line-length) + (page-offset) is about 7.54 inches. See §6. The current page offset is available in the <b>.o</b> register. |
| <b>.ne</b> <i>N</i> | -                    | $N=1$ <i>V</i>        | D,v          | Need <i>N</i> vertical space. If the distance, <i>D</i> , to the next trap position (see §7.5) is less than <i>N</i> , a forward vertical space of size <i>D</i> occurs, which will spring the trap. If there are no remaining traps on the page, <i>D</i> is the                                                                     |

\*The use of " " as control character (instead of ".") suppresses the break function.

†Values separated by ";" are for NROFF and TROFF respectively.

distance to the bottom of the page. If  $D < V$ , another line could still be output and spring the trap. In a diversion,  $D$  is the distance to the *diversion trap*, if any, or is very large.

|             |      |          |     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------|------|----------|-----|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| .mk $R$     | none | internal | D   | Mark the <i>current</i> vertical place in an internal register (both associated with the current diversion level), or in register $R$ , if given. See <code>rt</code> request.                                                                                                                                                                                                                                                                                                               |
| .rt $\pm N$ | none | internal | D,v | Return <i>upward only</i> to a marked vertical place in the current diversion. If $\pm N$ (w.r.t. current place) is given, the place is $\pm N$ from the top of the page or diversion or, if $N$ is absent, to a place marked by a previous <code>mk</code> . Note that the <code>sp</code> request (§5.3) may be used in all cases instead of <code>rt</code> by spacing to the absolute place stored in an explicit register; e. g. using the sequence <code>.mk R ... .sp   \nRu</code> . |

#### 4. Text Filling, Adjusting, and Centering

**4.1. Filling and adjusting.** Normally, words are collected from input text lines and assembled into a output text line until some word doesn't fit. An attempt is then made to hyphenate the word in effort to assemble a part of it into the output line. The spaces between the words on the output line are then increased to spread out the line to the current *line length* minus any current *indent*. A *word* is any string of characters delimited by the *space* character or the beginning/end of the input line. Any adjacent pair of words that must be kept together (neither split across output lines nor spread apart in the adjustment process) can be tied together by separating them with the *unpaddable space* character "\ " (backslash-space). The adjusted word spacings are uniform in TROFF and the minimum interword spacing can be controlled with the `ss` request (§2). In NROFF, they are normally nonuniform because of quantization to character-size spaces; however, the command line option `-e` causes uniform spacing with full output device resolution. Filling, adjustment, and hyphenation (§13) can all be prevented or controlled. The *text length* on the last line output is available in the `.n` register, and text base-line position on the page for this line is in the `nl` register. The text base-line high-water mark (lowest place) on the current page is in the `.h` register.

An input text line ending with `.`, `?`, or `!` is taken to be the end of a *sentence*, and an additional space character is automatically provided during filling. Multiple inter-word space characters found in the input are retained, except for trailing spaces; initial spaces also cause a *break*.

When filling is in effect, a `\p` may be imbedded or attached to a word to cause a *break* at the *end* of the word and have the resulting output line *spread out* to fill the current line length.

A text input line that happens to begin with a control character can be made to not look like a control line by prefacing it with the non-printing, zero-width filler character `\&`. Still another way is to specify output translation of some convenient character into the control character using `tr` (§10.5).

**4.2. Interrupted text.** The copying of a input line in *nofill* (non-fill) mode can be *interrupted* by terminating the partial line with a `\c`. The *next* encountered input text line will be considered to be a continuation of the same line of input text. Similarly, a word within *filled* text may be interrupted by terminating the word (and line) with `\c`; the next encountered text will be taken as a continuation of the interrupted word. If the intervening control lines cause a *break*, any partial line will be forced out along with any partial word.

| <i>Request Form</i> | <i>Initial Value</i> | <i>If No Argument</i> | <i>Notes</i> | <i>Explanation</i>                                                                                                                                                                                               |
|---------------------|----------------------|-----------------------|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| .br                 | -                    | -                     | B            | Break. The filling of the line currently being collected is stopped and the line is output without adjustment. Text lines beginning with space characters and empty text lines (blank lines) also cause a break. |

|              |          |        |     |                                                                                                                                                                                                                      |
|--------------|----------|--------|-----|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>.fi</b>   | fill on  | -      | B,E | Fill subsequent output lines. The register <b>.u</b> is 1 in fill mode and 0 in nofill mode.                                                                                                                         |
| <b>.nf</b>   | fill on  | -      | B,E | Nofill. Subsequent output lines are <i>neither filled nor</i> adjusted. Input text lines are copied directly to output lines <i>without regard</i> for the current line length.                                      |
| <b>.ad c</b> | adj,both | adjust | E   | Line adjustment is begun. If fill mode is not on, adjustment will be deferred until fill mode is back on. If the type indicator <i>c</i> is present, the adjustment type is changed as shown in the following table. |

| Indicator | Adjust Type              |
|-----------|--------------------------|
| l         | adjust left margin only  |
| r         | adjust right margin only |
| c         | center                   |
| b or n    | adjust both margins      |
| absent    | unchanged                |

|              |        |     |     |                                                                                                                                                                                                                                                         |
|--------------|--------|-----|-----|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>.na</b>   | adjust | -   | E   | Noadjust. Adjustment is turned off; the right margin will be ragged. The adjustment type for <b>ad</b> is not changed. Output line filling still occurs if fill mode is on.                                                                             |
| <b>.ce N</b> | off    | N=1 | B,E | Center the next <i>N</i> input text lines within the current (line-length minus indent). If <i>N</i> =0, any residual count is cleared. A break occurs after each of the <i>N</i> input lines. If the input line is too long, it will be left adjusted. |

## 5. Vertical Spacing

**5.1. Base-line spacing.** The vertical spacing (*V*) between the base-lines of successive output lines can be set using the **vs** request with a resolution of 1/144 inch = 1/2 point in TROFF, and to the output device resolution in NROFF. *V* must be large enough to accommodate the character sizes on the affected output lines. For the common type sizes (9-12 points), usual typesetting practice is to set *V* to 2 points greater than the point size; TROFF default is 10-point type on a 12-point spacing (as in this document). The current *V* is available in the **.v** register. Multiple-*V* line separation (e.g. double spacing) may be requested with **ls**.

**5.2. Extra line-space.** If a word contains a vertically tall construct requiring the output line containing it to have extra vertical space before and/or after it, the *extra-line-space* function **\x'N'** can be imbedded in or attached to that word. In this and other functions having a pair of delimiters around their parameter (here **'**), the delimiter choice is arbitrary, except that it can't look like the continuation of a number expression for *N*. If *N* is negative, the output line containing the word will be preceded by *N* extra vertical space; if *N* is positive, the output line containing the word will be followed by *N* extra vertical space. If successive requests for extra space apply to the same line, the maximum values are used. The most recently utilized post-line extra line-space is available in the **.a** register.

**5.3. Blocks of vertical space.** A block of vertical space is ordinarily requested using **sp**, which honors the *no-space* mode and which does not space *past* a trap. A contiguous block of vertical space may be reserved using **sv**.

| <b>Request Form</b> | <b>Initial Value</b> | <b>If No Argument</b> | <b>Notes</b> | <b>Explanation</b>                                                                                                                                                                               |
|---------------------|----------------------|-----------------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>.vs N</b>        | 1/6in;12pts          | previous              | E,p          | Set vertical base-line spacing size <i>V</i> . Transient <i>extra</i> vertical space available with <b>\x'N'</b> (see above).                                                                    |
| <b>.ls N</b>        | N=1                  | previous              | E            | Line spacing set to $\pm N$ . <i>N</i> -1 <i>V</i> s ( <i>blank lines</i> ) are appended to each output text line. Appended blank lines are omitted, if the text or previous appended blank line |

|                  |       |                      |     |                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------|-------|----------------------|-----|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                  |       |                      |     | reached a trap position.                                                                                                                                                                                                                                                                                                                                                                                            |
| .sp <i>N</i>     | -     | <i>N</i> =1 <i>V</i> | B,v | Space vertically in <i>either</i> direction. If <i>N</i> is negative, the motion is <i>backward</i> (upward) and is limited to the distance to the top of the page. Forward (downward) motion is truncated to the distance to the nearest trap. If the no-space mode is on, no spacing occurs (see <i>ns</i> , and <i>rs</i> below).                                                                                |
| .sv <i>N</i>     | -     | <i>N</i> =1 <i>V</i> | v   | Save a contiguous vertical block of size <i>N</i> . If the distance to the next trap is greater than <i>N</i> , <i>N</i> vertical space is output. No-space mode has <i>no</i> effect. If this distance is less than <i>N</i> , no vertical space is immediately output, but <i>N</i> is remembered for later output (see <i>os</i> ). Subsequent <i>sv</i> requests will overwrite any still remembered <i>N</i> . |
| .os              | -     | -                    | -   | Output saved vertical space. No-space mode has <i>no</i> effect. Used to finally output a block of vertical space requested by an earlier <i>sv</i> request.                                                                                                                                                                                                                                                        |
| .ns              | space | -                    | D   | No-space mode turned on. When on, the no-space mode inhibits <i>sp</i> requests and <i>bp</i> requests <i>without</i> a next page number. The no-space mode is turned off when a line of output occurs, or with <i>rs</i> .                                                                                                                                                                                         |
| .rs              | space | -                    | D   | Restore spacing. The no-space mode is turned off.                                                                                                                                                                                                                                                                                                                                                                   |
| Blank text line. | -     | -                    | B   | Causes a break and output of a blank line exactly like <i>sp</i> 1.                                                                                                                                                                                                                                                                                                                                                 |

## 6. Line Length and Indenting

The maximum line length for fill mode may be set with *ll*. The indent may be set with *in*; an indent applicable to *only* the *next* output line may be set with *ti*. The line length includes indent space but *not* page offset space. The line-length minus the indent is the basis for centering with *ce*. The effect of *ll*, *in*, or *ti* is delayed, if a partially collected line exists, until after that line is output. In fill mode the length of text on an output line is less than or equal to the line length minus the indent. The current line length and indent are available in registers *.l* and *.i* respectively. The length of *three-part titles* produced by *tl* (see §14) is *independently* set by *lt*.

| <i>Request Form</i> | <i>Initial Value</i> | <i>If No Argument</i> | <i>Notes</i> | <i>Explanation</i>                                                                                                                                                                                               |
|---------------------|----------------------|-----------------------|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| .ll ± <i>N</i>      | 6.5 in               | previous              | E,m          | Line length is set to ± <i>N</i> . In TROFF the maximum (line-length) + (page-offset) is about 7.54 inches.                                                                                                      |
| .in ± <i>N</i>      | <i>N</i> =0          | previous              | B,E,m        | Indent is set to ± <i>N</i> . The indent is prepended to each output line.                                                                                                                                       |
| .ti ± <i>N</i>      | -                    | ignored               | B,E,m        | Temporary indent. The <i>next</i> output text line will be indented a distance ± <i>N</i> with respect to the current indent. The resulting total indent may not be negative. The current indent is not changed. |

## 7. Macros, Strings, Diversion, and Position Traps

**7.1. Macros and strings.** A *macro* is a named set of arbitrary *lines* that may be invoked by name or with a *trap*. A *string* is a named string of *characters*, *not* including a newline character, that may be interpolated by name at any point. Request, macro, and string names share the *same* name list. Macro and string names may be one or two characters long and may usurp previously defined request, macro, or string names. Any of these entities may be renamed with *rn* or removed with *rm*. Macros are created by *de* and *di*, and appended to by *am* and *da*; *di* and *da* cause normal output to be stored in a macro. Strings are created by *ds* and appended to by *as*. A macro is invoked in the same way as a request; a

control line beginning `.xx` will interpolate the contents of macro `xx`. The remainder of the line may contain up to nine *arguments*. The strings `x` and `xx` are interpolated at any desired point with `\*x` and `\*(xx` respectively. String references and macro invocations may be nested.

**7.2. Copy mode input interpretation.** During the definition and extension of strings and macros (not by diversion) the input is read in *copy mode*. The input is copied without interpretation *except* that:

- The contents of number registers indicated by `\n` are interpolated.
- Strings indicated by `\*` are interpolated.
- Arguments indicated by `\$` are interpolated.
- Concealed newlines indicated by `\(newline)` are eliminated.
- Comments indicated by `\*` are eliminated.
- `\t` and `\a` are interpreted as ASCII horizontal tab and SOH respectively (§9).
- `\\` is interpreted as `\`.
- `\.` is interpreted as `"."`.

These interpretations can be suppressed by prepending a `\`. For example, since `\\` maps into a `\`, `\\n` will copy as `\n` which will be interpreted as a number register indicator when the macro or string is reread.

**7.3. Arguments.** When a macro is invoked by name, the remainder of the line is taken to contain up to nine arguments. The argument separator is the space character, and arguments may be surrounded by double-quotes to permit imbedded space characters. Pairs of double-quotes may be imbedded in double-quoted arguments to represent a single double-quote. If the desired arguments won't fit on a line, a concealed newline may be used to continue on the next line.

When a macro is invoked the *input level* is *pushed down* and any arguments available at the previous level become unavailable until the macro is completely read and the previous level is restored. A macro's own arguments can be interpolated at *any* point within the macro with `\$N`, which interpolates the *N*th argument ( $1 \leq N \leq 9$ ). If an invoked argument doesn't exist, a null string results. For example, the macro `xx` may be defined by

```
.de xx \"begin definition
Today is \\$1 the \\$2.
.. \"end definition
```

and called by

```
.xx Monday 14th
```

to produce the text

```
Today is Monday the 14th.
```

Note that the `\$` was concealed in the definition with a prepended `\`. The number of currently available arguments is in the `.$` register.

No arguments are available at the top (non-macro) level in this implementation. Because string referencing is implemented as an input-level push down, no arguments are available from *within* a string. No arguments are available within a trap-invoked macro.

Arguments are copied in *copy mode* onto a stack where they are available for reference. The mechanism does not allow an argument to contain a direct reference to a *long* string (interpolated at copy time) and it is advisable to conceal string references (with an extra `\`) to delay interpolation until argument reference time.

**7.4. Diversions.** Processed output may be diverted into a macro for purposes such as footnote processing (see Tutorial §T5) or determining the horizontal and vertical size of some text for conditional changing of pages or columns. A single diversion trap may be set at a specified vertical position. The number registers `dn` and `dl` respectively contain the vertical and horizontal size of the most recently ended diversion. Processed text that is diverted into a macro retains the vertical size of each of its lines when reread in *nofill* mode regardless of the current *V*. Constant-spaced (`cs`) or emboldened (`bd`) text that is diverted can be reread correctly only if these modes are again or still in effect at reread time. One way

to do this is to imbed in the diversion the appropriate *cs* or *bd* requests with the *transparent* mechanism described in §10.6.

Diversions may be nested and certain parameters and registers are associated with the current diversion level (the top non-diversion level may be thought of as the 0th diversion level). These are the diversion trap and associated macro, no-space mode, the internally-saved marked place (see *mk* and *rt*), the current vertical place (*.d* register), the current high-water text base-line (*.h* register), and the current diversion name (*.z* register).

**7.5. Traps.** Three types of trap mechanisms are available—page traps, a diversion trap, and an input-line-count trap. Macro-invocation traps may be planted using *wh* at any page position including the top. This trap position may be changed using *ch*. Trap positions at or below the bottom of the page have no effect unless or until moved to within the page or rendered effective by an increase in page length. Two traps may be planted at the *same* position only by first planting them at different positions and then moving one of the traps; the first planted trap will conceal the second unless and until the first one is moved (see Tutorial Examples §T5). If the first one is moved back, it again conceals the second trap. The macro associated with a page trap is automatically invoked when a line of text is output whose vertical size *reaches* or *sweeps past* the trap position. Reaching the bottom of a page springs the top-of-page trap, if any, provided there is a next page. The distance to the next trap position is available in the *.t* register; if there are no traps between the current position and the bottom of the page, the distance returned is the distance to the page bottom.

A macro-invocation trap effective in the current diversion may be planted using *dt*. The *.t* register works in a diversion; if there is no subsequent trap a *large* distance is returned. For a description of input-line-count traps, see it below.

| <i>Request Form</i>  | <i>Initial Value</i> | <i>If No Argument</i> | <i>Notes</i> | <i>Explanation</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------|----------------------|-----------------------|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>.de xx yy</i>     | -                    | <i>.yy=..</i>         | -            | Define or redefine the macro <i>xx</i> . The contents of the macro begin on the next input line. Input lines are copied in <i>copy mode</i> until the definition is terminated by a line beginning with <i>.yy</i> , whereupon the macro <i>yy</i> is called. In the absence of <i>yy</i> , the definition is terminated by a line beginning with <i>..</i> . A macro may contain <i>de</i> requests provided the terminating macros differ or the contained definition terminator is concealed. <i>..</i> can be concealed as <i>\\..</i> which will copy as <i>\..</i> and be reread as <i>..</i> . |
| <i>.am xx yy</i>     | -                    | <i>.yy=..</i>         | -            | Append to macro (append version of <i>de</i> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <i>.ds xx string</i> | -                    | ignored               | -            | Define a string <i>xx</i> containing <i>string</i> . Any initial double-quote in <i>string</i> is stripped off to permit initial blanks.                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <i>.as xx string</i> | -                    | ignored               | -            | Append <i>string</i> to string <i>xx</i> (append version of <i>ds</i> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <i>.rm xx</i>        | -                    | ignored               | -            | Remove request, macro, or string. The name <i>xx</i> is removed from the name list and any related storage space is freed. Subsequent references will have no effect.                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <i>.rn xx yy</i>     | -                    | ignored               | -            | Rename request, macro, or string <i>xx</i> to <i>yy</i> . If <i>yy</i> exists, it is first removed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <i>.di xx</i>        | -                    | end                   | D            | Divert output to macro <i>xx</i> . Normal text processing occurs during diversion except that page offsetting is not done. The diversion ends when the request <i>di</i> or <i>da</i> is encountered without an argument; extraneous requests of this type should not appear when nested diversions are being used.                                                                                                                                                                                                                                                                                   |



|                        |      |      |     |                                                                                                                                                                                                                                                                                                                                                              |
|------------------------|------|------|-----|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>.da</b> <i>xx</i>   | -    | end  | D   | Divert, appending to <i>xx</i> (append version of <b>di</b> ).                                                                                                                                                                                                                                                                                               |
| <b>.wh</b> <i>N xx</i> | -    | -    | v   | Install a trap to invoke <i>xx</i> at page position <i>N</i> ; a <i>negative N</i> will be interpreted with respect to the page <i>bottom</i> . Any macro previously planted at <i>N</i> is replaced by <i>xx</i> . A zero <i>N</i> refers to the <i>top</i> of a page. In the absence of <i>xx</i> , the first found trap at <i>N</i> , if any, is removed. |
| <b>.ch</b> <i>xx N</i> | -    | -    | v   | Change the trap position for macro <i>xx</i> to be <i>N</i> . In the absence of <i>N</i> , the trap, if any, is removed.                                                                                                                                                                                                                                     |
| <b>.dt</b> <i>N xx</i> | -    | off  | D,v | Install a diversion trap at position <i>N</i> in the <i>current</i> diversion to invoke macro <i>xx</i> . Another <b>dt</b> will redefine the diversion trap. If no arguments are given, the diversion trap is removed.                                                                                                                                      |
| <b>.it</b> <i>N xx</i> | -    | off  | E   | Set an input-line-count trap to invoke the macro <i>xx</i> after <i>N</i> lines of <i>text</i> input have been read (control or request lines don't count). The text may be in-line text or text interpolated by inline or trap-invoked macros.                                                                                                              |
| <b>.em</b> <i>xx</i>   | none | none | -   | The macro <i>xx</i> will be invoked when all input has ended. The effect is the same as if the contents of <i>xx</i> had been at the end of the last file processed.                                                                                                                                                                                         |

### 8. Number Registers

A variety of parameters are available to the user as predefined, named *number registers* (see Summary and Index, page 7). In addition, the user may define his own named registers. Register names are one or two characters long and *do not* conflict with request, macro, or string names. Except for certain predefined read-only registers, a number register can be read, written, automatically incremented or decremented, and interpolated into the input in a variety of formats. One common use of user-defined registers is to automatically number sections, paragraphs, lines, etc. A number register may be used any time numerical input is expected or desired and may be used in numerical *expressions* (§1.4).

Number registers are created and modified using **nr**, which specifies the name, numerical value, and the auto-increment size. Registers are also modified, if accessed with an auto-incrementing sequence. If the registers *x* and *xx* both contain *N* and have the auto-increment size *M*, the following access sequences have the effect shown:

| Sequence                 | Effect on Register                | Value Interpolated |
|--------------------------|-----------------------------------|--------------------|
| <b>\n</b> <i>x</i>       | none                              | <i>N</i>           |
| <b>\n</b> ( <i>xx</i> )  | none                              | <i>N</i>           |
| <b>\n</b> + <i>x</i>     | <i>x</i> incremented by <i>M</i>  | <i>N+M</i>         |
| <b>\n</b> - <i>x</i>     | <i>x</i> decremented by <i>M</i>  | <i>N-M</i>         |
| <b>\n</b> +( <i>xx</i> ) | <i>xx</i> incremented by <i>M</i> | <i>N+M</i>         |
| <b>\n</b> -( <i>xx</i> ) | <i>xx</i> decremented by <i>M</i> | <i>N-M</i>         |

When interpolated, a number register is converted to decimal (default), decimal with leading zeros, lower-case Roman, upper-case Roman, lower-case sequential alphabetic, or upper-case sequential alphabetic according to the format specified by **af**.

| <b>Request Form</b>     | <b>Initial Value</b> | <b>If No Argument</b> | <b>Notes</b> | <b>Explanation</b>                                                                                                                                              |
|-------------------------|----------------------|-----------------------|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>.nr</b> <i>R ±NM</i> | -                    | -                     | u            | The number register <i>R</i> is assigned the value $\pm N$ with respect to the previous value, if any. The increment for auto-incrementing is set to <i>M</i> . |

`.af R c` arabic - - Assign format *c* to register *R*. The available formats are:

| Format     | Numbering Sequence                   |
|------------|--------------------------------------|
| <b>1</b>   | 0,1,2,3,4,5,...                      |
| <b>001</b> | 000,001,002,003,004,005,...          |
| <b>i</b>   | 0,i,ii,iii,iv,v,...                  |
| <b>I</b>   | 0,I,II,III,IV,V,...                  |
| <b>a</b>   | 0,a,b,c,....,z,aa,ab,....,zz,aaa,... |
| <b>A</b>   | 0,A,B,C,....,Z,AA,AB,....,ZZ,AAA,... |

An arabic format having *N* digits specifies a field width of *N* digits (example 2 above). The read-only registers and the *width* function (§11.2) are always arabic.

`.rr R` - - ignored - - Remove register *R*. If many registers are being created dynamically, it may become necessary to remove no longer used registers to recapture internal storage space for newer registers.

## 9. Tabs, Leaders, and Fields

**9.1. Tabs and leaders.** The ASCII horizontal tab character and the ASCII SOH (hereafter known as the *leader* character) can both be used to generate either horizontal motion or a string of repeated characters. The length of the generated entity is governed by internal *tab stops* specifiable with `ta`. The default difference is that tabs generate motion and leaders generate a string of periods; `tc` and `lc` offer the choice of repeated character or motion. There are three types of internal tab stops—*left* adjusting, *right* adjusting, and *centering*. In the following table: *D* is the distance from the current position on the *input* line (where a tab or leader was found) to the next tab stop; *next-string* consists of the input characters following the tab (or leader) up to the next tab (or leader) or end of line; and *W* is the width of *next-string*.

| Tab type | Length of motion or repeated characters | Location of <i>next-string</i>    |
|----------|-----------------------------------------|-----------------------------------|
| Left     | <i>D</i>                                | Following <i>D</i>                |
| Right    | <i>D - W</i>                            | Right adjusted within <i>D</i>    |
| Centered | <i>D - W/2</i>                          | Centered on right end of <i>D</i> |

The length of generated motion is allowed to be negative, but that of a repeated character string cannot be. Repeated character strings contain an integer number of characters, and any residual distance is prepended as motion. Tabs or leaders found after the last tab stop are ignored, but may be used as *next-string* terminators.

Tabs and leaders are not interpreted in *copy mode*. `\t` and `\a` always generate a non-interpreted tab and leader respectively, and are equivalent to actual tabs and leaders in *copy mode*.

**9.2. Fields.** A *field* is contained between a *pair* of *field delimiter* characters, and consists of sub-strings separated by *padding* indicator characters. The field length is the distance on the *input* line from the position where the field begins to the next tab stop. The difference between the total length of all the sub-strings and the field length is incorporated as horizontal padding space that is divided among the indicated padding places. The incorporated padding is allowed to be negative. For example, if the field delimiter is `#` and the padding indicator is `^`, `#^xxx^right#` specifies a right-adjusted string with the string `xxx` centered in the remaining space.

| <i>Request Form</i> | <i>Initial Value</i> | <i>If No Argument</i> | <i>Notes</i> | <i>Explanation</i>                                                                                                                                                                                                                                                                                    |
|---------------------|----------------------|-----------------------|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| .ta <i>Nt</i> ...   | 0.8; 0.5in           | none                  | E,m          | Set tab stops and types. <i>t</i> =R, right adjusting; <i>t</i> =C, centering; <i>t</i> absent, left adjusting. TROFF tab stops are preset every 0.5in.; NROFF every 0.8in. The stop values are separated by spaces, and a value preceded by + is treated as an increment to the previous stop value. |
| .tc <i>c</i>        | none                 | none                  | E            | The tab repetition character becomes <i>c</i> , or is removed specifying motion.                                                                                                                                                                                                                      |
| .lc <i>c</i>        | .                    | none                  | E            | The leader repetition character becomes <i>c</i> , or is removed specifying motion.                                                                                                                                                                                                                   |
| .fc <i>a b</i>      | off                  | off                   | -            | The field delimiter is set to <i>a</i> ; the padding indicator is set to the <i>space</i> character or to <i>b</i> , if given. In the absence of arguments the field mechanism is turned off.                                                                                                         |

## 10. Input and Output Conventions and Character Translations

*10.1. Input character translations.* Ways of inputting the graphic character set were discussed in §2.1. The ASCII control characters horizontal tab (§9.1), SOH (§9.1), and backspace (§10.3) are discussed elsewhere. The newline delimits input lines. In addition, STX, ETX, ENQ, ACK, and BEL are accepted, and may be used as delimiters or translated into a graphic with tr (§10.5). *All others are ignored.*

The *escape* character \ introduces *escape sequences*—causes the following character to mean another character, or to indicate some function. A complete list of such sequences is given in the Summary and Index on page 6. \ should not be confused with the ASCII control character ESC of the same name. The escape character \ can be input with the sequence \\. The escape character can be changed with ec, and all that has been said about the default \ becomes true for the new escape character. \e can be used to print whatever the current escape character is. If necessary or convenient, the escape mechanism may be turned off with eo, and restored with ec.

| <i>Request Form</i> | <i>Initial Value</i> | <i>If No Argument</i> | <i>Notes</i> | <i>Explanation</i>                                    |
|---------------------|----------------------|-----------------------|--------------|-------------------------------------------------------|
| .ec <i>c</i>        | \                    | \                     | -            | Set escape character to \, or to <i>c</i> , if given. |
| .eo                 | on                   | -                     | -            | Turn escape mechanism off.                            |

*10.2. Ligatures.* Five ligatures are available in the current TROFF character set — fi, fl, ff, ffi, and ffi. They may be input (even in NROFF) by \fi, \fl, \ff, \ffi, and \ffi respectively. The ligature mode is normally on in TROFF, and *automatically* invokes ligatures during input.

| <i>Request Form</i> | <i>Initial Value</i> | <i>If No Argument</i> | <i>Notes</i> | <i>Explanation</i>                                                                                                                                                                                                                                                                                     |
|---------------------|----------------------|-----------------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| .lg <i>N</i>        | off; on              | on                    | -            | Ligature mode is turned on if <i>N</i> is absent or non-zero, and turned off if <i>N</i> =0. If <i>N</i> =2, only the two-character ligatures are automatically invoked. Ligature mode is inhibited for request, macro, string, register, or file names, and in <i>copy mode</i> . No effect in NROFF. |

*10.3. Backspacing, underlining, overstriking, etc.* Unless in *copy mode*, the ASCII backspace character is replaced by a backward horizontal motion having the width of the space character. Underlining as a form of line-drawing is discussed in §12.4. A generalized overstriking function is described in §12.1.

NROFF automatically underlines characters in the *underline* font, specifiable with uf, normally that on font position 2 (normally Times Italic, see §2.2). In addition to ft and \fF, the underline font may be selected by ul and cu. Underlining is restricted to an output-device-dependent subset of *reasonable* characters.

| <i>Request Form</i> | <i>Initial Value</i> | <i>If No Argument</i> | <i>Notes</i> | <i>Explanation</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------|----------------------|-----------------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>.ul N</code>  | off                  | $N=1$                 | E            | Underline in NROFF (italicize in TROFF) the next $N$ input text lines. Actually, switch to <i>underline</i> font, saving the current font for later restoration; <i>other</i> font changes within the span of a <code>ul</code> will take effect, but the restoration will undo the last change. Output generated by <code>tl</code> (§14) is affected by the font change, but does <i>not</i> decrement $N$ . If $N > 1$ , there is the risk that a trap interpolated macro may provide text lines within the span; environment switching can prevent this. |
| <code>.cu N</code>  | off                  | $N=1$                 | E            | A variant of <code>ul</code> that causes <i>every</i> character to be underlined in NROFF. Identical to <code>ul</code> in TROFF.                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <code>.uf F</code>  | Italic               | Italic                | -            | Underline font set to $F$ . In NROFF, $F$ may <i>not</i> be on position 1 (initially Times Roman).                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

10.4. *Control characters.* Both the control character `.` and the *no-break* control character `'` may be changed, if desired. Such a change must be compatible with the design of any macros used in the span of the change, and particularly of any trap-invoked macros.

| <i>Request Form</i> | <i>Initial Value</i> | <i>If No Argument</i> | <i>Notes</i> | <i>Explanation</i>                                                                  |
|---------------------|----------------------|-----------------------|--------------|-------------------------------------------------------------------------------------|
| <code>.cc c</code>  | .                    | .                     | E            | The basic control character is set to $c$ , or reset to <code>"."</code> .          |
| <code>.c2 c</code>  | .                    | '                     | E            | The <i>nobreak</i> control character is set to $c$ , or reset to <code>"'"</code> . |

10.5. *Output translation.* One character can be made a stand-in for another character using `tr`. All text processing (e. g. character comparisons) takes place with the input (stand-in) character which appears to have the width of the final character. The graphic translation occurs at the moment of output (including diversion).

| <i>Request Form</i>       | <i>Initial Value</i> | <i>If No Argument</i> | <i>Notes</i> | <i>Explanation</i>                                                                                                                                                                                                                                  |
|---------------------------|----------------------|-----------------------|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>.tr abcd....</code> | none                 | -                     | O            | Translate $a$ into $b$ , $c$ into $d$ , etc. If an odd number of characters is given, the last one will be mapped into the space character. To be consistent, a particular translation must stay in effect from <i>input</i> to <i>output</i> time. |

10.6. *Transparent throughput.* An input line beginning with a `\!` is read in *copy mode* and *transparently* output (without the initial `\!`); the text processor is otherwise unaware of the line's presence. This mechanism may be used to pass control information to a post-processor or to imbed control lines in a macro created by a diversion.

10.7. *Comments and concealed newlines.* An uncomfortably long input line that must stay one line (e. g. a string definition, or nofilled text) can be split into many physical lines by ending all but the last one with the escape `\`. The sequence `\(newline)` is *always* ignored—except in a comment. Comments may be imbedded at the *end* of any line by prefacing them with `\'`. The newline at the end of a comment cannot be concealed. A line beginning with `\'` will appear as a blank line and behave like `.sp 1`; a comment can be on a line by itself by beginning the line with `\'`.

## 11. Local Horizontal and Vertical Motions, and the Width Function

11.1. *Local Motions.* The functions `\v'N` and `\h'N` can be used for *local* vertical and horizontal motion respectively. The distance  $N$  may be negative; the *positive* directions are *rightward* and *downward*. A *local* motion is one contained *within* a line. To avoid unexpected vertical dislocations, it is necessary that the *net* vertical local motion within a word in filled text and otherwise within a line balance to zero. The above and certain other escape sequences providing local motion are summarized in the following table.

| Vertical<br>Local Motion | Effect in              |             | Horizontal<br>Local Motion | Effect in                   |         |
|--------------------------|------------------------|-------------|----------------------------|-----------------------------|---------|
|                          | TROFF                  | NROFF       |                            | TROFF                       | NROFF   |
| <code>\v'N'</code>       | Move distance <i>N</i> |             | <code>\h'N'</code>         | Move distance <i>N</i>      |         |
|                          |                        |             | <code>\(space)</code>      | Unpaddable space-size space |         |
|                          |                        |             | <code>\0</code>            | Digit-size space            |         |
| <code>\u</code>          | ½ em up                | ½ line up   | <code>\ </code>            | 1/6 em space                | ignored |
| <code>\d</code>          | ½ em down              | ½ line down |                            | 1/12 em space               | ignored |
| <code>\r</code>          | 1 em up                | 1 line up   |                            |                             |         |

As an example,  $E^2$  could be generated by the sequence `E\s-2\v'-0.4m'2\v'0.4m'\s+2`; it should be noted in this example that the 0.4 em vertical motions are at the smaller size.

**11.2. Width Function.** The width function `\w'string'` generates the numerical width of *string* (in basic units). Size and font changes may be safely imbedded in *string*, and will not affect the current environment. For example, `.ti -\w'1. 'u` could be used to temporarily indent leftward a distance equal to the size of the string "1. ".

The width function also sets three number registers. The registers `st` and `sb` are set respectively to the highest and lowest extent of *string* relative to the baseline; then, for example, the total height of the string is `\n(stu-\n(sbu)`. In TROFF the number register `ct` is set to a value between 0 and 3: 0 means that all of the characters in *string* were short lower case characters without descenders (like *e*); 1 means that at least one character has a descender (like *y*); 2 means that at least one character is tall (like *H*); and 3 means that both tall characters and characters with descenders are present.

**11.3. Mark horizontal place.** The escape sequence `\kx` will cause the current horizontal position in the input line to be stored in register *x*. As an example, the construction `\kx word\h'\|nxu+2u' word` will embolden *word* by backing up to almost its beginning and overprinting it, resulting in **word**.

## 12. Overstrike, Bracket, Line-drawing, and Zero-width Functions

**12.1. Overstriking.** Automatically centered overstriking of up to nine characters is provided by the *overstrike* function `\o'string'`. The characters in *string* overprinted with centers aligned; the total width is that of the widest character. *string* should not contain local vertical motion. As examples, `\o'e''` produces  $\acute{e}$ , and `\o'\(mo)\(sl'` produces  $\text{€}$ .

**12.2. Zero-width characters.** The function `\zc` will output *c* without spacing over it, and can be used to produce left-aligned overstruck combinations. As examples, `\z\(\ci)\(\pl` will produce  $\oplus$ , and `\(br\z\(\rn)\(\ul)\(br` will produce the smallest possible constructed box  $\square$ .

**12.3. Large Brackets.** The Special Mathematical Font contains a number of bracket construction pieces (`{ } [ ] { } [ ] [ ] [ ]`) that can be combined into various bracket styles. The function `\b'string'` may be used to pile up vertically the characters in *string* (the first character on top and the last at the bottom); the characters are vertically separated by 1 em and the total pile is centered 1/2 em above the current baseline (½ line in NROFF). For example, `\b'\(\lc)\(lf'E'\|b'\(\rc)\(rf'\x'-0.5m'\x'0.5m'` produces  $\left[ E \right]$ .

**12.4. Line drawing.** The function `\l'Nc'` will draw a string of repeated *c*'s towards the right for a distance *N*. (`\l` is `\(lower case L)`). If *c* looks like a continuation of an expression for *N*, it may insulated from *N* with a `\&`. If *c* is not specified, the `_` (baseline rule) is used (underline character in NROFF). If *N* is negative, a backward horizontal motion of size *N* is made before drawing the string. Any space resulting from *N*/(size of *c*) having a remainder is put at the beginning (left end) of the string. In the case of characters that are designed to be connected such as baseline-rule `_`, underrule `_`, and root-en `¯`, the remainder space is covered by over-lapping. If *N* is less than the width of *c*, a single *c* is centered on a distance *N*. As an example, a macro to underscore a string can be written

```
.de us
\\$1\l'1'0\ul'
..
```

or one to draw a box around a string

```
.de bx
\ (br\|\\$1\\(br\l'|0\rn\l'|0\ul'
..
```

such that

```
.ul "underlined words"
```

and

```
.bx "words in a box"
```

yield underlined words and words in a box.

The function `\L' Nc'` will draw a vertical line consisting of the (optional) character *c* stacked vertically apart 1 em (1 line in NROFF), with the first two characters overlapped, if necessary, to form a continuous line. The default character is the *box rule* | (`\(br)`); the other suitable character is the *bold vertical* | (`\(bv)`). The line is begun without any initial motion relative to the current base line. A positive *N* specifies a line drawn downward and a negative *N* specifies a line drawn upward. After the line is drawn *no* compensating motions are made; the instantaneous baseline is at the *end* of the line.

The horizontal and vertical line drawing functions may be used in combination to produce large boxes. The zero-width *box-rule* and the 1/2-em wide *underrule* were *designed* to form corners when using 1-em vertical spacings. For example the macro

```
.de eb
.sp -1 \ "compensate for next automatic base-line spacing
.nf \ "avoid possibly overflowing word buffer
\h'-.5n\L'|\\nau-1\l'\n(.lu+1n\ (ul\L'-|\\nau+1\l'|0u-.5n\ (ul' \ "draw box
.fi
..
```

will draw a box around some text whose beginning vertical place was saved in number register *a* (e. g. using `.mk a`) as done for this paragraph.

### 13. Hyphenation.

The automatic hyphenation may be switched off and on. When switched on with `hy`, several variants may be set. A *hyphenation indicator* character may be imbedded in a word to specify desired hyphenation points, or may be prepended to suppress hyphenation. In addition, the user may specify a small exception word list.

Only words that consist of a central alphabetic string surrounded by (usually null) non-alphabetic strings are considered candidates for automatic hyphenation. Words that were input containing hyphens (minus), em-dashes (`\(em)`), or hyphenation indicator characters—such as *mother-in-law*—are *always* subject to splitting after those characters, whether or not automatic hyphenation is on or off.

| <i>Request Form</i>        | <i>Initial Value</i> | <i>If No Argument</i> | <i>Notes</i> | <i>Explanation</i>                                                                                                                                                                                                                                                                                                             |
|----------------------------|----------------------|-----------------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>.nh</code>           | hyphenate            | -                     | E            | Automatic hyphenation is turned off.                                                                                                                                                                                                                                                                                           |
| <code>.hyN</code>          | on, N=1              | on, N=1               | E            | Automatic hyphenation is turned on for $N \geq 1$ , or off for $N=0$ . If $N=2$ , <i>last</i> lines (ones that will cause a trap) are not hyphenated. For $N=4$ and 8, the last and first two characters respectively of a word are not split off. These values are additive; i. e. $N=14$ will invoke all three restrictions. |
| <code>.hc c</code>         | \%                   | \%                    | E            | Hyphenation indicator character is set to <i>c</i> or to the default \%. The indicator does not appear in the output.                                                                                                                                                                                                          |
| <code>.hw wordl ...</code> |                      | ignored               | -            | Specify hyphenation points in words with imbedded minus signs. Versions of a word with terminal <i>s</i> are                                                                                                                                                                                                                   |

implied; i. e. *dig-it* implies *dig-its*. This list is examined initially *and* after each suffix stripping. The space available is small—about 128 characters.

#### 14. Three Part Titles.

The titling function `tl` provides for automatic placement of three fields at the left, center, and right of a line with a title-length specifiable with `lt`. `tl` may be used anywhere, and is independent of the normal text collecting process. A common use is in header and footer macros.

| <i>Request Form</i>                    | <i>Initial Value</i> | <i>If No Argument</i> | <i>Notes</i> | <i>Explanation</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------------------|----------------------|-----------------------|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>.tl 'left' center' right'</code> |                      | -                     | -            | The strings <i>left</i> , <i>center</i> , and <i>right</i> are respectively left-adjusted, centered, and right-adjusted in the current title-length. Any of the strings may be empty, and overlapping is permitted. If the page-number character (initially <code>%</code> ) is found within any of the fields it is replaced by the current page number having the format assigned to register <code>%</code> . Any character may be used as the string delimiter. |
| <code>.pc c</code>                     | <code>%</code>       | off                   | -            | The page number character is set to <code>c</code> , or removed. The page-number register remains <code>%</code> .                                                                                                                                                                                                                                                                                                                                                  |
| <code>.lt ±N</code>                    | 6.5 in               | previous              | E,m          | Length of title set to $\pm N$ . The line-length and the title-length are <i>independent</i> . Indents do not apply to titles; page-offsets do.                                                                                                                                                                                                                                                                                                                     |

#### 15. Output Line Numbering.

Automatic sequence numbering of output lines may be requested with `nm`. When in effect, a three-digit, arabic number plus a digit-space is prepended to output text lines. The text lines are thus offset by four digit-spaces, and otherwise retain their line length; a reduction in line length may be desired to keep the right margin aligned with an earlier margin. Blank lines, other vertical spaces, and lines generated by `tl` are *not* numbered. Numbering can be temporarily suspended with `6 nn`, or with an `.nm` followed by a later `.nm +0`. In addition, a line number indent *I*, and the number-text separation *S* may be specified in digit-spaces. Further, it can be specified that only those line numbers that are multiples of some number *M* are to be printed (the others will appear as blank number fields).

| <i>Request Form</i>       | <i>Initial Value</i> | <i>If No Argument</i> | <i>Notes</i> | <i>Explanation</i>                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------|----------------------|-----------------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>.nm ±N M S I</code> |                      | off                   | E            | Line number mode. If $\pm N$ is given, line numbering is turned on, and the next output line numbered is numbered $\pm N$ . Default values are $M=1$ , $S=1$ , and $I=0$ . Parameters corresponding to missing arguments are unaffected; a non-numeric argument is considered missing. In the absence of all arguments, numbering is turned off; the next line number is preserved for possible further use in number register <code>ln</code> . |
| <code>.nn N</code>        | -                    | $N=1$                 | E            | The next <i>N</i> text output lines are not numbered.                                                                                                                                                                                                                                                                                                                                                                                            |

As an example, the paragraph portions of this section are numbered with  $M=3$ : `.nm 1 3` was placed at the beginning; `.nm` was placed at the end of the first paragraph; and `.nm +0` was placed in front of this paragraph; and `.nm` finally placed at the end. Line lengths were also changed (by `\w'0000'u`) to keep the right side aligned. Another example is `.nm +5 5 x 3` which turns on numbering with the line number of the next line to be 5 greater than the last numbered line, with  $M=5$ , with spacing *S* untouched, and with the indent *I* set to 3.

## 16. Conditional Acceptance of Input

In the following, *c* is a one-character, built-in *condition* name, *!* signifies *not*, *N* is a numerical expression, *string1* and *string2* are strings delimited by any non-blank, non-numeric character *not* in the strings, and *anything* represents what is conditionally accepted.

| Request Form                                   | Initial Value | If No Argument | Notes | Explanation                                                                                                     |
|------------------------------------------------|---------------|----------------|-------|-----------------------------------------------------------------------------------------------------------------|
| .if <i>c anything</i>                          |               | -              | -     | If condition <i>c</i> true, accept <i>anything</i> as input; in multi-line case use <code>\{anything\}</code> . |
| .if <i>!c anything</i>                         |               | -              | -     | If condition <i>c</i> false, accept <i>anything</i> .                                                           |
| .if <i>N anything</i>                          |               | -              | u     | If expression $N > 0$ , accept <i>anything</i> .                                                                |
| .if <i>!N anything</i>                         |               | -              | u     | If expression $N \leq 0$ , accept <i>anything</i> .                                                             |
| .if ' <i>string1 string2</i> ' <i>anything</i> |               | -              | -     | If <i>string1</i> identical to <i>string2</i> , accept <i>anything</i> .                                        |
| .if <i>!'string1 string2'</i> <i>anything</i>  |               | -              | -     | If <i>string1</i> not identical to <i>string2</i> , accept <i>anything</i> .                                    |
| .ie <i>c anything</i>                          |               | -              | u     | If portion of if-else; all above forms (like if).                                                               |
| .el <i>anything</i>                            |               | -              | -     | Else portion of if-else.                                                                                        |

The built-in condition names are:

| Condition Name | True If                     |
|----------------|-----------------------------|
| <b>o</b>       | Current page number is odd  |
| <b>e</b>       | Current page number is even |
| <b>t</b>       | Formatter is TROFF          |
| <b>n</b>       | Formatter is NROFF          |

If the condition *c* is *true*, or if the number *N* is greater than zero, or if the strings compare identically (including motions and character size and font), *anything* is accepted as input. If a *!* precedes the condition, number, or string comparison, the sense of the acceptance is reversed.

Any spaces between the condition and the beginning of *anything* are skipped over. The *anything* can be either a single input line (text, macro, or whatever) or a number of input lines. In the multi-line case, the first line must begin with a left delimiter `\{` and the last line must end with a right delimiter `\}`.

The request *ie* (if-else) is identical to *if* except that the acceptance state is remembered. A subsequent and matching *el* (else) request then uses the reverse sense of that state. *ie* - *el* pairs may be nested.

Some examples are:

```
.if e .tl 'Even Page %''
```

which outputs a title if the page number is even; and

```
.ie \n% > 1 \\
'sp 0.5i
.tl 'Page %''
'sp |1.2i \
.el .sp |2.5i
```

which treats page 1 differently from other pages.

## 17. Environment Switching.

A number of the parameters that control the text processing are gathered together into an *environment*, which can be switched by the user. The environment parameters are those associated with requests noting *E* in their *Notes* column; in addition, partially collected lines and words are in the environment. Everything else is global; examples are page-oriented parameters, diversion-oriented parameters,



number registers, and macro and string definitions. All environments are initialized with default parameter values.

| <i>Request Form</i> | <i>Initial Value</i> | <i>If No Argument</i> | <i>Notes</i> | <i>Explanation</i>                                                                                                                                                                                                |
|---------------------|----------------------|-----------------------|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>.ev N</code>  | $N=0$                | previous              | -            | Environment switched to environment $0 \leq N \leq 2$ . Switching is done in push-down fashion so that restoring a previous environment <i>must</i> be done with <code>.ev</code> rather than specific reference. |

### 18. Insertions from the Standard Input

The input can be temporarily switched to the system *standard input* with `rd`, which will switch back when *two* newlines in a row are found (the *extra* blank line is not used). This mechanism is intended for insertions in form-letter-like documentation. On UNIX, the *standard input* can be the user's keyboard, a *pipe*, or a *file*.

| <i>Request Form</i>     | <i>Initial Value</i> | <i>If No Argument</i>    | <i>Notes</i> | <i>Explanation</i>                                                                                                                                                                                                                                                                |
|-------------------------|----------------------|--------------------------|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>.rd prompt</code> | -                    | <code>prompt=BEL-</code> | -            | Read insertion from the standard input until two newlines in a row are found. If the standard input is the user's keyboard, <i>prompt</i> (or a BEL) is written onto the user's terminal. <code>rd</code> behaves like a macro, and arguments may be placed after <i>prompt</i> . |
| <code>.ex</code>        | -                    | -                        | -            | Exit from NROFF/TROFF. Text processing is terminated exactly as if all input had ended.                                                                                                                                                                                           |

If insertions are to be taken from the terminal keyboard *while* output is being printed on the terminal, the command line option `-q` will turn off the echoing of keyboard input and prompt only with BEL. The regular input and insertion input *cannot* simultaneously come from the standard input.

As an example, multiple copies of a form letter may be prepared by entering the insertions for all the copies in one file to be used as the standard input, and causing the file containing the letter to reinvoke itself using `nx` (§19); the process would ultimately be ended by an `ex` in the insertion file.

### 19. Input/Output File Switching

| <i>Request Form</i>       | <i>Initial Value</i> | <i>If No Argument</i> | <i>Notes</i> | <i>Explanation</i>                                                                                                                                                                                                                                                                                                  |
|---------------------------|----------------------|-----------------------|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>.so filename</code> | -                    | -                     | -            | Switch source file. The top input (file reading) level is switched to <i>filename</i> . The effect of an <code>so</code> encountered in a macro is not felt until the input level returns to the file level. When the new file ends, input is again taken from the original file. <code>so</code> 's may be nested. |
| <code>.nx filename</code> | -                    | end-of-file           | -            | Next file is <i>filename</i> . The current file is considered ended, and the input is immediately switched to <i>filename</i> .                                                                                                                                                                                     |
| <code>.pi program</code>  | -                    | -                     | -            | Pipe output to <i>program</i> (NROFF only). This request must occur <i>before</i> any printing occurs. No arguments are transmitted to <i>program</i> .                                                                                                                                                             |

### 20. Miscellaneous

| <i>Request Form</i>  | <i>Initial Value</i> | <i>If No Argument</i> | <i>Notes</i> | <i>Explanation</i>                                                                                                                                                                                                                                                                      |
|----------------------|----------------------|-----------------------|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>.mc c N</code> | -                    | off                   | E,m          | Specifies that a <i>margin</i> character <i>c</i> appear a distance <i>N</i> to the right of the right margin after each non-empty text line (except those produced by <code>t1</code> ). If the output line is too-long (as can happen in <code>nofill</code> mode) the character will |

|                   |   |         |                                                                                                                                                                                                                         |                                                                                                                                                                                                                                     |
|-------------------|---|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                   |   |         | be appended to the line. If <i>N</i> is not given, the previous <i>N</i> is used; the initial <i>N</i> is 0.2 inches in NROFF and 1 em in TROFF. The margin character used with this paragraph was a 12-point box-rule. |                                                                                                                                                                                                                                     |
| .tm <i>string</i> | - | newline | -                                                                                                                                                                                                                       | After skipping initial blanks, <i>string</i> (rest of the line) is read in <i>copy mode</i> and written on the user's terminal.                                                                                                     |
| .ig <i>yy</i>     | - | .yy=..  | -                                                                                                                                                                                                                       | Ignore input lines. <i>ig</i> behaves exactly like <i>de</i> (§7) except that the input is discarded. The input is read in <i>copy mode</i> , and any auto-incremented registers will be affected.                                  |
| .pm <i>t</i>      | - | all     | -                                                                                                                                                                                                                       | Print macros. The names and sizes of all of the defined macros and strings are printed on the user's terminal; if <i>t</i> is given, only the total of the sizes is printed. The sizes is given in <i>blocks</i> of 128 characters. |
| .fl               | - | -       | B                                                                                                                                                                                                                       | Flush output buffer. Used in interactive debugging to force output.                                                                                                                                                                 |

## 21. Output and Error Messages.

The output from *tm*, *pm*, and the prompt from *rd*, as well as various *error* messages are written onto UNIX's *standard message* output. The latter is different from the *standard output*, where NROFF formatted output goes. By default, both are written onto the user's terminal, but they can be independently redirected.

Various *error* conditions may occur during the operation of NROFF and TROFF. Certain less serious errors having only local impact do not cause processing to terminate. Two examples are *word overflow*, caused by a word that is too large to fit into the word buffer (in fill mode), and *line overflow*, caused by an output line that grew too large to fit in the line buffer; in both cases, a message is printed, the offending excess is discarded, and the affected word or line is marked at the point of truncation with a \* in NROFF and a ■ in TROFF. The philosophy is to continue processing, if possible, on the grounds that output useful for debugging may be produced. If a serious error occurs, processing terminates, and an appropriate message is printed. Examples are the inability to create, read, or write files, and the exceeding of certain internal limits that make future output unlikely to be useful.

## TUTORIAL EXAMPLES

### T1. Introduction

Although NROFF and TROFF have by design a syntax reminiscent of earlier text processors\* with the intent of easing their use, it is almost always necessary to prepare at least a small set of macro definitions to describe most documents. Such common formatting needs as page margins and footnotes are deliberately not built into NROFF and TROFF. Instead, the macro and string definition, number register, diversion, environment switching, page-position trap, and conditional input mechanisms provide the basis for user-defined implementations.

The examples to be discussed are intended to be useful and somewhat realistic, but won't necessarily cover all relevant contingencies. Explicit numerical parameters are used in the examples to make them easier to read and to illustrate typical values. In many cases, number registers would really be used to reduce the number of places where numerical information is kept, and to concentrate conditional parameter initialization like that which depends on whether TROFF or NROFF is being used.

### T2. Page Margins

As discussed in §3, *header* and *footer* macros are usually defined to describe the top and bottom page margin areas respectively. A trap is planted at page position 0 for the header, and at  $-N$  ( $N$  from the page bottom) for the footer. The simplest such definitions might be

```
.de hd \define header
`sp 1i
.. \end definition
.de fo \define footer
`bp
.. \end definition
.wh 0 hd
.wh -1i fo
```

which provide blank 1 inch top and bottom margins. The header will occur on the *first* page, only if the definition and trap exist prior to the

initial pseudo-page transition (§3). In fill mode, the output line that springs the footer trap was typically forced out because some part or whole word didn't fit on it. If anything in the footer and header that follows causes a *break*, that word or part word will be forced out. In this and other examples, requests like *bp* and *sp* that normally cause breaks are invoked using the *no-break* control character ` to avoid this. When the header/footer design contains material requiring independent text processing, the environment may be switched, avoiding most interaction with the running text.

A more realistic example would be

```
.de hd \header
.if t .tl `(rn``(rn' \troff cut mark
.if \n%>1 \{\
`sp |0.5i-1 \tl base at 0.5i
.tl ``- % -" \centered page number
.ps \restore size
.ft \restore font
.vs \} \restore vs
`sp |1.0i \space to 1.0i
.ns \turn on no-space mode
..
.de fo \footer
.ps 10 \set footer/header size
.ft R \set font
.vs 12p \set base-line spacing
.if \n%=1 \{\
`sp |\n(.pu-0.5i-1 \tl base 0.5i up
.tl ``- % -" \} \first page number
`bp
..
.wh 0 hd
.wh -1i fo
```

which sets the size, font, and base-line spacing for the header/footer material, and ultimately restores them. The material in this case is a page number at the bottom of the first page and at the top of the remaining pages. If TROFF is used, a *cut mark* is drawn in the form of *root-en's* at each margin. The *sp's* refer to absolute positions to avoid dependence on the base-line spacing. Another reason for this in the footer is that the footer is invoked by printing a line whose vertical spacing swept past the trap position by possibly as

\*For example: P. A. Crisman, Ed., *The Compatible Time-Sharing System*, MIT Press, 1965, Section AH9.01 (Description of RUNOFF program on MIT's CTSS system).

much as the base-line spacing. The *no-space* mode is turned on at the end of *hd* to render ineffective accidental occurrences of *sp* at the top of the running text.

The above method of restoring size, font, etc. presupposes that such requests (that set *previous* value) are *not* used in the running text. A better scheme is *save* and *restore* both the current and previous values as shown for *size*, in the following:

```
.de fo
.nr s1 \\n(.s \ "current size
.ps
.nr s2 \\n(.s \ "previous size
. --- \ "rest of footer
..
.de hd
. --- \ "header stuff
.ps \\n(s2 \ "restore previous size
.ps \\n(s1 \ "restore current size
..
```

Page numbers may be printed in the bottom margin by a separate macro triggered during the footer's page ejection:

```
.de bn \ "bottom number
.tl "-- % -- \ "centered page number
..
.wh -0.5i-1v bn \ "tl base 0.5i up
```

### T3. Paragraphs and Headings

The housekeeping associated with starting a new paragraph should be collected in a paragraph macro that, for example, does the desired preparagraph spacing, forces the correct font, size, base-line spacing, and indent, checks that enough space remains for *more than one* line, and requests a temporary indent.

```
.de pg \ "paragraph
.br \ "break
.ft R \ "force font,
.ps 10 \ "size,
.vs 12p \ "spacing,
.in 0 \ "and indent
.sp 0.4 \ "prespace
.ne 1+\\n(.Vu \ "want more than 1 line
.tl 0.2i \ "temp indent
..
```

The first break in *pg* will force out any previous partial lines, and must occur before the *vs*. The forcing of font, etc. is partly a defense against prior error and partly to permit things like section heading macros to set parameters only once.

The prespacing parameter is suitable for TROFF; a larger space, at least as big as the output device vertical resolution, would be more suitable in NROFF. The choice of remaining space to test for in the *ne* is the smallest amount greater than one line (the *.V* is the available vertical resolution).

A macro to automatically number section headings might look like:

```
.de sc \ "section
. --- \ "force font, etc.
.sp 0.4 \ "prespace
.ne 2.4+\\n(.Vu \ "want 2.4+ lines
.fi
\\n+S.
..
.nr S 0 1 \ "init S
```

The usage is *.sc*, followed by the section heading text, followed by *.pg*. The *ne* test value includes one line of heading, 0.4 line in the following *pg*, and one line of the paragraph text. A word consisting of the next section number and a period is produced to begin the heading line. The format of the number may be set by *af* (§8).

Another common form is the labeled, indented paragraph, where the label protrudes left into the indent space.

```
.de lp \ "labeled paragraph
.pg
.in 0.5i \ "paragraph indent
.ta 0.2i 0.5i \ "label, paragraph
.tl 0
\t\\$1\t\c \ "flow into paragraph
..
```

The intended usage is *.lp label*; *label* will begin at 0.2 inch, and cannot exceed a length of 0.3 inch without intruding into the paragraph. The label could be right adjusted against 0.4 inch by setting the tabs instead with *.ta 0.4iR 0.5i*. The last line of *lp* ends with *\c* so that it will become a part of the first line of the text that follows.

### T4. Multiple Column Output

The production of multiple column pages requires the footer macro to decide whether it was invoked by other than the last column, so that it will begin a new column rather than produce the bottom margin. The header can initialize a column register that the footer will increment and test. The following is arranged for two columns, but is easily modified for more.

```
.de hd \header
. ---
.nr cl 0 1 \init column count
.mk \mark top of text
..
.de fo \footer
.ie \\n + (cl < 2) \\
.po +3.4i \next column; 3.1+0.3
.rt \back to mark
.ns \\ \no-space mode
.ei \\ \
.po \\nMu \restore left margin
. ---
'bp \\
..
.ll 3.1i \column width
.nr M \\n(.o \save left margin
```

Typically a portion of the top of the first page contains full width text; the request for the narrower line length, as well as another .mk would be made where the two column output was to begin.

#### T5. Footnote Processing

The footnote mechanism to be described is used by imbedding the footnotes in the input text at the point of reference, demarcated by an initial .fn and a terminal .ef:

```
.fn
Footnote text and control lines...
.ef
```

In the following, footnotes are processed in a separate environment and diverted for later printing in the space immediately prior to the bottom margin. There is provision for the case where the last collected footnote doesn't completely fit in the available space.

```
.de hd \header
. ---
.nr x 0 1 \init footnote count
.nr y 0-\\nb \current footer place
.ch fo -\\nbu \reset footer trap
.if \\n(dn .fz \leftover footnote
..
.de fo \footer
.nr dn 0 \zero last diversion size
.if \\nx \\
.ev 1 \expand footnotes in ev1
.nf \retain vertical size
.FN \footnotes
.rm FN \delete it
.if "\\n(.z"fy" .di \end overflow diversion
.nr x 0 \disable fx
```

```
.ev \\ \pop environment
. ---
'bp
..
.de fx \process footnote overflow
.if \\nx .di fy \divert overflow
..
.de fn \start footnote
.da FN \divert (append) footnote
.ev 1 \in environment 1
.if \\n + x = 1 .fs \if first, include separator
.fi \fill mode
..
.de ef \end footnote
.br \finish output
.nr z \\n(.v \save spacing
.ev \pop ev
.di \end diversion
.nr y -\\n(dn \new footer position,
.if \\nx = 1 .nr y -(\n(.v-\\nz) \
\uncertainty correction
.ch fo \\nyu \y is negative
.if ((\\n(nl+1v) > (\\n(.p+\\ny) \
.ch fo \\n(nlu+1v \it didn't fit
..
.de fs \separator
\l' 1i' \1 inch rule
.br
..
.de fz \get leftover footnote
.fn
.nf \retain vertical size
.fy \where fx put it
.ef
..
.nr b 1.0i \bottom margin size
.wh 0 hd \header trap
.wh 12i fo \footer trap, temp position
.wh -\\nbu fx \fx at footer position
.ch fo -\\nbu \conceal fx with fo
```

The header hd initializes a footnote count register x, and sets both the current footer trap position register y and the footer trap itself to a nominal position specified in register b. In addition, if the register dn indicates a leftover footnote, fz is invoked to reprocess it. The footnote start macro fn begins a diversion (append) in environment 1, and increments the count x; if the count is one, the footnote separator fs is interpolated. The separator is kept in a separate macro to permit user redefinition. The footnote end macro ef restores the previous environment and ends the diversion after saving the spacing size in register z. y is then decremented by the size of the

footnote, available in `dn`; then on the first footnote, `y` is further decremented by the difference in vertical base-line spacings of the two environments, to prevent the late triggering the footer trap from causing the last line of the combined footnotes to overflow. The footer trap is then set to the lower (on the page) of `y` or the current page position (`n1`) plus one line, to allow for printing the reference line. If indicated by `x`, the footer `fo` rereads the footnotes from `FN` in `nofill` mode in environment 1, and deletes `FN`. If the footnotes were too large to fit, the macro `fx` will be trap-invoked to redirect the overflow into `fy`, and the register `dn` will later indicate to the header whether `fy` is empty. Both `fo` and `fx` are planted in the nominal footer trap position in an order that causes `fx` to be concealed unless the `fo` trap is moved. The footer then terminates the overflow diversion, if necessary, and zeros `x` to disable `fx`, because the uncertainty correction together with a not-too-late triggering of the footer can result in the footnote rereading finishing before reaching the `fx` trap.

A good exercise for the student is to combine the multiple-column and footnote mechanisms.

#### T6. The Last Page

After the last input file has ended, NROFF and TROFF invoke the *end macro* (§7), if any, and when it finishes, eject the remainder of the page. During the eject, any traps encountered are processed normally. At the *end* of this last page, processing terminates *unless* a partial line, word, or partial word remains. If it is desired that another page be started, the end-macro

```
.de en *end-macro
\c
`bp
..
.em en
```

will deposit a null partial word, and effect another last page.

## Table I

### Font Style Examples

The following fonts are printed in 12-point, with a vertical spacing of 14-point, and with non-alphanumeric characters separated by 1/4 em space. The Special Mathematical Font was specially prepared for Bell Laboratories by Graphic Systems, Inc. of Hudson, New Hampshire. The Times Roman, Italic, and Bold are among the many standard fonts available from that company.

#### Times Roman

abcdefghijklmnopqrstuvwxyz  
 ABCDEFGHIJKLMNOPQRSTUVWXYZ  
 1234567890  
 ! \$ % & ( ) ' \* + - . , / : ; = ? [ ] |  
 • □ - - - ¼ ½ ¾ fi fl ff ffi ffl ° † ' ¢ ® ©

#### *Times Italic*

*abcdefghijklmnopqrstuvwxyz*  
*ABCDEFGHIJKLMNOPQRSTUVWXYZ*  
*1234567890*  
*! \$ % & ( ) ' \* + - . , / : ; = ? [ ] |*  
 • □ - - - ¼ ½ ¾ *fi fl ff ffi ffl* ° † ' ¢ ® ©

#### **Times Bold**

**abcdefghijklmnopqrstuvwxyz**  
**ABCDEFGHIJKLMNOPQRSTUVWXYZ**  
**1234567890**  
**! \$ % & ( ) ' \* + - . , / : ; = ? [ ] |**  
 • □ - - - ¼ ½ ¾ **fi fl ff ffi ffl** ° † ' ¢ ® ©

#### Special Mathematical Font

" ' \ ^ \_ ` ~ / < > { } # @ + - = \*  
 α β γ δ ε ζ η θ ι κ λ μ ν ξ ο π ρ σ τ υ φ χ ψ ω  
 Γ Δ Θ Λ Ξ Π Σ Υ Φ Ψ Ω  
 √ ∞ ≥ ≤ ≡ ~ ≈ ≠ → ← ↑ ↓ × ÷ ± ∪ ∩ ⊂ ⊃ ⊆ ⊇ ∞ ∂  
 § ∇ ∫ ∞ ∅ ∈ † ⚡ Ⓢ | ○ ( ) { } || || |

Table II

Input Naming Conventions for ', `and -  
 and for Non-ASCII Special Characters

Non-ASCII characters and *minus* on the standard fonts.

| <i>Char</i> | <i>Input Name</i> | <i>Character Name</i> | <i>Char</i> | <i>Input Name</i> | <i>Character Name</i> |
|-------------|-------------------|-----------------------|-------------|-------------------|-----------------------|
| '           |                   | close quote           | fi          | \(fi              | fi                    |
| `           |                   | open quote            | fl          | \(fl              | fl                    |
| -           | \(em              | 3/4 Em dash           | ff          | \(ff              | ff                    |
| -           | -                 | hyphen or             | ffi         | \(Fi              | ffi                   |
| -           | \(hy              | hyphen                | ffl         | \(Fl              | ffl                   |
| -           | \-                | current font minus    | °           | \(de              | degree                |
| •           | \(bu              | bullet                | †           | \(dg              | dagger                |
| □           | \(sq              | square                | '           | \(fm              | foot mark             |
| -           | \(ru              | rule                  | ¢           | \(ct              | cent sign             |
| ¼           | \(14              | 1/4                   | ®           | \(rg              | registered            |
| ½           | \(12              | 1/2                   | ©           | \(co              | copyright             |
| ¾           | \(34              | 3/4                   |             |                   |                       |

Non-ASCII characters and ', `\_, +, -, =, and \* on the special font.

The ASCII characters @, #, ", ', `;, <, >, \, {, }, ~, ^, and \_ exist *only* on the special font and are printed as a 1-em space if that font is not mounted. The following characters exist only on the special font except for the upper case Greek letter names followed by † which are mapped into upper case English letters in whatever font is mounted on font position one (default Times Roman). The special math plus, minus, and equals are provided to insulate the appearance of equations from the choice of standard fonts.

| <i>Char</i> | <i>Input Name</i> | <i>Character Name</i>      | <i>Char</i> | <i>Input Name</i> | <i>Character Name</i> |
|-------------|-------------------|----------------------------|-------------|-------------------|-----------------------|
| +           | \(pl              | math plus                  | κ           | \(*k              | kappa                 |
| -           | \(mi              | math minus                 | λ           | \(*l              | lambda                |
| =           | \(eq              | math equals                | μ           | \(*m              | mu                    |
| *           | \(**              | math star                  | ν           | \(*n              | nu                    |
| §           | \(sc              | section                    | ξ           | \(*c              | xi                    |
| '           | \(aa              | acute accent               | ο           | \(*o              | omicron               |
| `           | \(ga              | grave accent               | π           | \(*p              | pi                    |
| -           | \(ul              | underrule                  | ρ           | \(*r              | rho                   |
| /           | \(sl              | slash (matching backslash) | σ           | \(*s              | sigma                 |
| α           | \(*a              | alpha                      | ς           | \(ts              | terminal sigma        |
| β           | \(*b              | beta                       | τ           | \(*t              | tau                   |
| γ           | \(*g              | gamma                      | υ           | \(*u              | upsilon               |
| δ           | \(*d              | delta                      | φ           | \(*f              | phi                   |
| ε           | \(*e              | epsilon                    | χ           | \(*x              | chi                   |
| ζ           | \(*z              | zeta                       | ψ           | \(*q              | psi                   |
| η           | \(*y              | eta                        | ω           | \(*w              | omega                 |
| θ           | \(*h              | theta                      | Α           | \(*A              | Alpha†                |
| ι           | \(*i              | iota                       | Β           | \(*B              | Beta†                 |



| <i>Char</i>    | <i>Input Name</i>    | <i>Character Name</i> |
|----------------|----------------------|-----------------------|
| $\Gamma$       | <code>\(*G</code>    | Gamma                 |
| $\Delta$       | <code>\(*D</code>    | Delta                 |
| $E$            | <code>\(*E</code>    | Epsilon†              |
| $Z$            | <code>\(*Z</code>    | Zeta†                 |
| $H$            | <code>\(*Y</code>    | Eta†                  |
| $\Theta$       | <code>\(*H</code>    | Theta                 |
| $I$            | <code>\(*I</code>    | Iota†                 |
| $K$            | <code>\(*K</code>    | Kappa†                |
| $\Lambda$      | <code>\(*L</code>    | Lambda                |
| $M$            | <code>\(*M</code>    | Mu†                   |
| $N$            | <code>\(*N</code>    | Nu†                   |
| $\Xi$          | <code>\(*C</code>    | Xi                    |
| $O$            | <code>\(*O</code>    | Omicron†              |
| $\Pi$          | <code>\(*P</code>    | Pi                    |
| $\rho$         | <code>\(*R</code>    | Rho†                  |
| $\Sigma$       | <code>\(*S</code>    | Sigma                 |
| $T$            | <code>\(*T</code>    | Tau†                  |
| $Y$            | <code>\(*U</code>    | Upsilon               |
| $\Phi$         | <code>\(*F</code>    | Phi                   |
| $X$            | <code>\(*X</code>    | Chi†                  |
| $\Psi$         | <code>\(*Q</code>    | Psi                   |
| $\Omega$       | <code>\(*W</code>    | Omega                 |
| $\sqrt{\quad}$ | <code>\(sr</code>    | square root           |
| $\sqrt{\quad}$ | <code>\(rn</code>    | root en extender      |
| $\succcurlyeq$ | <code>\(&gt;=</code> | $\succcurlyeq$        |
| $\preccurlyeq$ | <code>\(&lt;=</code> | $\preccurlyeq$        |
| $\equiv$       | <code>\(==</code>    | identically equal     |
| $\approx$      | <code>\(≈</code>     | approx =              |
| $\sim$         | <code>\(ap</code>    | approximates          |
| $\neq$         | <code>\(!=</code>    | not equal             |
| $\rightarrow$  | <code>\(-&gt;</code> | right arrow           |
| $\leftarrow$   | <code>\(&lt;-</code> | left arrow            |
| $\uparrow$     | <code>\(ua</code>    | up arrow              |
| $\downarrow$   | <code>\(da</code>    | down arrow            |
| $\times$       | <code>\(mu</code>    | multiply              |
| $\div$         | <code>\(di</code>    | divide                |
| $\pm$          | <code>\(+-</code>    | plus-minus            |
| $\cup$         | <code>\(cu</code>    | cup (union)           |
| $\cap$         | <code>\(ca</code>    | cap (intersection)    |
| $\subset$      | <code>\(sb</code>    | subset of             |
| $\supset$      | <code>\(sp</code>    | superset of           |
| $\subsetneq$   | <code>\(ib</code>    | improper subset       |
| $\supsetneq$   | <code>\(ip</code>    | improper superset     |
| $\infty$       | <code>\(if</code>    | infinity              |
| $\partial$     | <code>\(pd</code>    | partial derivative    |
| $\nabla$       | <code>\(gr</code>    | gradient.             |
| $\neg$         | <code>\(no</code>    | not                   |
| $\int$         | <code>\(is</code>    | integral sign         |
| $\propto$      | <code>\(pt</code>    | proportional to       |
| $\emptyset$    | <code>\(es</code>    | empty set             |
| $\in$          | <code>\(mo</code>    | member of             |

| <i>Char</i> | <i>Input Name</i> | <i>Character Name</i>                          |
|-------------|-------------------|------------------------------------------------|
|             | <code>\(br</code> | box vertical rule                              |
| ‡           | <code>\(dd</code> | double dagger                                  |
| ☞           | <code>\(rh</code> | right hand                                     |
| ☜           | <code>\(lh</code> | left hand                                      |
| Ⓚ           | <code>\(bs</code> | Bell System logo                               |
|             | <code>\(or</code> | or                                             |
| ○           | <code>\(ci</code> | circle                                         |
| {           | <code>\(lt</code> | left top of big curly bracket                  |
| {           | <code>\(lb</code> | left bottom                                    |
| }           | <code>\(rt</code> | right top                                      |
| }           | <code>\(rb</code> | right bot                                      |
| {           | <code>\(lk</code> | left center of big curly bracket               |
| }           | <code>\(rk</code> | right center of big curly bracket              |
|             | <code>\(bv</code> | bold vertical                                  |
|             | <code>\(lf</code> | left floor (left bottom of big square bracket) |
|             | <code>\(rf</code> | right floor (right bottom)                     |
|             | <code>\(lc</code> | left ceiling (left top)                        |
|             | <code>\(rc</code> | right ceiling (right top)                      |

May 15, 1977

## Summary of Changes to N/TROFF Since October 1976 Manual

### Options

- h (Nroff only) Output tabs used during horizontal spacing to speed output as well as reduce output byte count. Device tab settings assumed to be every 8 nominal character widths. The default settings of input (logical) tabs is also initialized to every 8 nominal character widths.
- z Efficiently suppresses formatted output. Only message output will occur (from "tm"s and diagnostics).

### Old Requests

- .ad c The adjustment type indicator "c" may now also be a number previously obtained from the "j" register (see below).
- .so name The contents of file "name" will be interpolated at the point the "so" is encountered. Previously, the interpolation was done upon return to the file-reading input level.

### New Request

- .ab text Prints "text" on the message output and terminates without further processing. If "text" is missing, "User Abort." is printed. Does not cause a break. The output buffer is flushed.
- .fz F N forces font "F" to be in size N. N may have the form N, +N, or -N. For example,  
.fz 3 -2  
will cause an implicit \s-2 every time font 3 is entered, and a corresponding \s+2 when it is left. Special font characters occurring during the reign of font F will have the same size modification. If special characters are to be treated differently,  
.fz S F N  
may be used to specify the size treatment of special characters during font F. For example,  
.fz 3 -3  
.fz S 3 -0  
will cause automatic reduction of font 3 by 3 points while the special characters would not be affected. Any ".fp" request specifying a font on some position must precede ".fz" requests relating to that position.

### New Predefined Number Registers.

- .k Read-only. Contains the horizontal size of the text portion (without indent) of the current partially collected output line, if any, in the current environment.
- .j Read-only. A number representing the current adjustment mode and type. Can be saved and later given to the "ad" request to restore a previous mode.
- .P Read-only. 1 if the current page is being printed, and zero otherwise.
- .L Read-only. Contains the current line-spacing parameter ("ls").
- c. General register access to the input line-number in the current input file. Contains the same value as the read-only ".c" register.



