

NAME

Intro – introduction to user-level library functions

DESCRIPTION

Section 3 describes user-level library routines. In this release, most user-library routines are listed in alphabetical order regardless of their subsection headings. (This eliminates having to search through several subsections of the manual.) However, due to their special-purpose nature, the routines from the following libraries are broken out into the indicated subsections:

- The Lightweight Processes Library, in subsection 3L.
- The RPC Services Library, in subsection 3R.
- The System V Compatibility Library, in subsection 3V. This library contains System V versions of functions that are not yet merged into the standard Sun libraries. To use them functions, compile programs with `/usr/5bin/cc`, instead of `/usr/bin/cc`.

The main C library, `/usr/lib/libc.a`, contains many of the functions described in this section, along with entry points for the system calls described in Section 2. This library also includes the Internet networking routines listed under the 3N subsection heading, and routines provided for compatibility with other UNIX operating systems, listed under 3C. Functions associated with the “standard I/O library” are listed under 3S.

User-level routines for access to data structures within the kernel and other processes are listed under 3K. To use these functions, compile programs with the `-lkvm` option for the C compiler, `cc(1V)`.

Math library functions are listed under 3M. To use them, compile programs with the the `-lm cc(1V)` option.

Various specialized libraries, the routines they contain, and the compiler options needed to link with them, are listed under 3X.

FILES

<code>/usr/lib/libc.a</code>	C Library (2, 3, 3N and 3C)
<code>/usr/lib/lib*.a</code>	other “standard” C libraries
<code>/usr/lib/lib*.a</code>	special-purpose C libraries
<code>/usr/5bin/cc</code>	

SEE ALSO

`cc(1V)`, `ld(1)`, `nm(1)`, `intro(2)`

LIST OF LIBRARY FUNCTIONS

Name	Appears on Page	Description
–	<code>bstring(3)</code>	bit and byte string operations
–	<code>byteorder(3N)</code>	convert values between host and network byte order
–	<code>ctime(3)</code>	convert date and time
–	<code>ctype(3)</code>	character classification and conversion macros and functions
–	<code>curses(3X)</code>	cursor addressing and screen display library
–	<code>dbm(3X)</code>	data base subroutines
–	<code>directory(3)</code>	directory operations
–	<code>ethers(3N)</code>	Ethernet address mapping operations
–	<code>inet(3N)</code>	Internet address manipulation
–	<code>intro(3L)</code>	introduction to the lightweight process library (LWP)
–	<code>intro(3M)</code>	introduction to mathematical library functions
–	<code>intro(3R)</code>	introduction to RPC service library and protocols
–	<code>intro(3V)</code>	introduction to System V functions
–	<code>mp(3X)</code>	multiple precision integer arithmetic
–	<code>ndbm(3)</code>	data base subroutines
–	<code>plot(3X)</code>	graphics interface

–	rpc(3N)	library routines for remote procedure calls
–	string(3)	string operations
–	termcap(3X)	terminal independent operation routines
–	values(3)	machine-dependent values
–	xdr(3N)	library routines for external data representation
_crypt()	crypt(3)	password and data encryption
a64l()	a64l(3)	convert between long integer and base-64 ASCII string
abort()	abort(3)	generate a fault
abs()	abs(3)	integer absolute value
addexportent()	exportent(3)	get exported file system information
addexportent()	exportent(3)	get exported file system information
addmntent()	getmntent(3)	get file system descriptor file entry
addmntent()	getmntent(3)	get file system descriptor file entry
alloca()	malloc(3)	memory allocator
alloca()	malloc(3)	memory allocator
alphasort()	scandir(3)	scan a directory
alphasort()	scandir(3)	scan a directory
arc()	plot(3X)	graphics interface
asctime()	ctime(3)	convert date and time
assert()	assert(3)	program verification
atof()	strtod(3)	convert string to double-precision number
atoi()	strtol(3)	convert string to integer
atol()	strtol(3)	convert string to integer
audit()	getacinfo(3)	get audit control file information
audit_args()	audit_args(3)	produce text audit message
audit_text()	audit_args(3)	produce text audit message
auth_destroy()	rpc(3N)	RPC services routines
authdes_create()	rpc(3N)	RPC services routines
authdes_getcred()	rpc(3N)	RPC services routines
authnon_create()	rpc(3N)	RPC services routines
authunix_create()	rpc(3N)	RPC services routines
authunix_create_default()	rpc(3N)	RPC services routines
bcmp()	bstring(3)	bit and byte string operations
bcopy()	bstring(3)	bit and byte string operations
bindresvport()	bindresvport(3N)	bind a socket to a privileged IP port
bsearch()	bsearch(3)	binary search a sorted table
bzero()	bstring(3)	bit and byte string operations
calloc()	malloc(3)	memory allocator
callrpc()	rpc(3N)	RPC services routines
cbc_crypt()	des_crypt(3)	fast DES encryption
cfree()	malloc(3)	memory allocator
circle()	plot(3X)	graphics interface
clearerr()	ferror(3S)	stream status inquiries
clnt_broadcast()	rpc(3N)	RPC services routines
clnt_call()	rpc(3N)	RPC services routines
clnt_destroy()	rpc(3N)	RPC services routines
clnt_freeres()	rpc(3N)	RPC services routines
clnt_geterr()	rpc(3N)	RPC services routines
clnt_pcreateerror()	rpc(3N)	RPC services routines
clnt_perrno()	rpc(3N)	RPC services routines
clnt_perror()	rpc(3N)	RPC services routines
clnt_sperrno()	rpc(3N)	RPC services routines
clnt_sperror()	rpc(3N)	RPC services routines

clntraw_create()	rpc(3N)	RPC services routines
clnttcp_create()	rpc(3N)	RPC services routines
clntudp_create()	rpc(3N)	RPC services routines
clock()	clock(3C)	report CPU time used
closedir()	directory(3)	directory operations
closelog()	syslog(3)	control system log
closepl()	plot(3X)	graphics interface
cont()	plot(3X)	graphics interface
control()	getacinfo(3)	get audit control file information
crypt()	crypt(3)	password and data encryption
ctermid()	ctermid(3S)	generate filename for terminal
cuserid()	cuserid(3S)	get character login name of the user
dbm_clearerr()	ndbm(3)	data base subroutines
dbm_close()	ndbm(3)	data base subroutines
dbm_delete()	ndbm(3)	data base subroutines
dbm_error()	ndbm(3)	data base subroutines
dbm_fetch()	ndbm(3)	data base subroutines
dbm_firstkey()	ndbm(3)	data base subroutines
dbm_nextkey()	ndbm(3)	data base subroutines
dbm_open()	ndbm(3)	data base subroutines
dbm_store()	ndbm(3)	data base subroutines
dbm_init()	dbm(3X)	data base subroutines
decimal_to_double()	decimal_to_floating(3)	convert decimal record to floating-point value
decimal_to_extended()	decimal_to_floating(3)	convert decimal record to floating-point value
decimal_to_single()	decimal_to_floating(3)	convert decimal record to floating-point value
delete()	dbm(3X)	data base subroutines
des_crypt()	des_crypt(3)	fast DES encryption
des_setparity()	des_crypt(3)	fast DES encryption
dn_comp()	resolver(3)	resolver routines
dn_expand()	resolver(3)	resolver routines
double_to_decimal()	floating_to_decimal(3)	convert floating-point value to decimal record
drand48()	drand48(3)	generate uniformly distributed pseudo-random numbers
dysize()	ctime(3)	convert date and time
ecb_crypt()	des_crypt(3)	fast DES encryption
econvert()	econvert(3)	output conversion
ecvt()	econvert(3)	output conversion
edata()	end(3)	last locations in program
encrypt()	crypt(3)	password and data encryption
end()	end(3)	last locations in program
endac()	getacinfo(3)	get audit control file information
endexportent()	exportent(3)	get exported file system information
endfsent()	getfsent(3)	get file system descriptor file entry
endgraent()	getgraent(3)	get group adjunct file entry
endgrent()	getgrent(3)	get group file entry
endhostent()	gethostent(3N)	get network host entry
endmntent()	getmntent(3)	get file system descriptor file entry
endnetent()	getnetent(3N)	get network entry
endnetgrent()	getnetgrent(3N)	get network group entry
endprotoent()	getprotoent(3N)	get protocol entry
endpwaent()	getpwaent(3)	get password adjunct file entry
endpwent()	getpwent(3)	get password file entry
endservent()	getservent(3N)	get service entry
endttyent()	getttyent(3)	get ttytab file entry

endusershell()	getusershell(3)	get legal user shells
erand48()	drand48(3)	generate uniformly distributed pseudo-random numbers
erase()	plot(3X)	graphics interface
errno()	perror(3)	system error messages
etext()	end(3)	last locations in program
ether_aton()	ethers(3N)	Ethernet address mapping operations
ether_hostton()	ethers(3N)	Ethernet address mapping operations
ether_line()	ethers(3N)	Ethernet address mapping operations
ether_ntoa()	ethers(3N)	Ethernet address mapping operations
ether_ntohost()	ethers(3N)	Ethernet address mapping operations
execl()	execl(3)	execute a file
execle()	execle(3)	execute a file
execlp()	execlp(3)	execute a file
execv()	execv(3)	execute a file
execvp()	execvp(3)	execute a file
exit()	exit(3)	terminate a process after performing cleanup
exportent()	exportent(3)	get exported file system information
extended_to_decimal()	floating_to_decimal(3)	convert floating-point value to decimal record
fclose()	fclose(3S)	close or flush a stream
fconvert()	econvert(3)	output conversion
fcvt()	econvert(3)	output conversion
fdate()	fdate(3)	return date and time in an ASCII string
fdopen()	fopen(3S)	open a stream
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
ffs()	bstring(3)	bit and byte string operations
fgetc()	getc(3S)	get character or integer from stream
fgetgraent()	getgraent(3)	get group adjunct file entry
fgetgrent()	getgrent(3)	get group file entry
fgetpwaent()	getpwaent(3)	get password adjunct file entry
fgetpwent()	getpwent(3)	get password file entry
fgets()	gets(3S)	get a string from a stream
file()	getacinfo(3)	get audit control file information
file_to_decimal()	string_to_decimal(3)	parse characters into decimal record
fileno()	ferror(3S)	stream status inquiries
firstkey()	dbm(3X)	data base subroutines
floatingpoint()	floatingpoint(3)	IEEE floating point definitions
fopen()	fopen(3S)	open a stream
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
free()	malloc(3)	memory allocator
freopen()	fopen(3S)	open a stream
fscanf()	scanf(3S)	formatted input conversion
fseek()	fseek(3S)	reposition a stream
ftell()	fseek(3S)	reposition a stream
ftime()	time(3C)	get date and time
ftok()	ftok(3)	standard interprocess communication package
ftw()	ftw(3)	walk a file tree
func_to_decimal()	string_to_decimal(3)	parse characters into decimal record

fwrite()	fread(3S)	buffered binary input/output
gcd()	mp(3X)	multiple precision integer arithmetic
gconvert()	econvert(3)	output conversion
gcvf()	econvert(3)	output conversion
get()	getacinfo(3)	get audit control file information
get_myaddress()	rpc(3N)	RPC services routines
getacdir()	getacinfo(3)	get audit control file information
getacflg()	getacinfo(3)	get audit control file information
getacmin()	getacinfo(3)	get audit control file information
getauditflagsbin()	getauditflags(3)	convert audit flag specifications
getauditflagschar()	getauditflags(3)	convert audit flag specifications
getc()	getc(3S)	get character or integer from stream
getchar()	getc(3S)	get character or integer from stream
getcwd()	getcwd(3)	get pathname of current working directory
getenv()	getenv(3)	return value for environment name
getexportent()	exportent(3)	get exported file system information
getexportopt()	exportent(3)	get exported file system information
getfauditflags()	getfaudflgs(3)	generates the process audit state
getfsent()	getfsent(3)	get file system descriptor file entry
getfsfile()	getfsent(3)	get file system descriptor file entry
getfsspec()	getfsent(3)	get file system descriptor file entry
getfstype()	getfsent(3)	get file system descriptor file entry
getgraent()	getgraent(3)	get group adjunct file entry
getgranam()	getgraent(3)	get group adjunct file entry
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
getlogin()	getlogin(3)	get login name
getmntent()	getmntent(3)	get file system descriptor file entry
getnetbyaddr()	getnetent(3N)	get network entry
getnetbyname()	getnetent(3N)	get network entry
getnetent()	getnetent(3N)	get network entry
getnetgrent()	getnetgrent(3N)	get network group entry
getnetname()	rpc(3N)	RPC services routines
getopt()	getopt(3)	get option letter from argument vector
getpass()	getpass(3)	read a password
getprotobyname()	getprotoent(3N)	get protocol entry
getprotobynumber()	getprotoent(3N)	get protocol entry
getprotoent()	getprotoent(3N)	get protocol entry
getpw()	getpw(3)	get name from uid
getpwaent()	getpwaent(3)	get password adjunct file entry
getpwanam()	getpwaent(3)	get password adjunct file entry
getpwent()	getpwent(3)	get password file entry
getpwnam()	getpwent(3)	get password file entry
getpwuid()	getpwent(3)	get password file entry
getrpcbyname()	getrpcent(3N)	get RPC entry
getrpcbynumber()	getrpcent(3N)	get RPC entry
getrpcent()	getrpcent(3N)	get RPC 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
getttyent()	getttyent(3)	get ttytab file entry
getttynam()	getttyent(3)	get ttytab file entry
getusershell()	getusershell(3)	get legal user shells
getw()	getc(3S)	get character or integer from stream
getwd()	getwd(3)	get current working directory pathname
gmtime()	ctime(3)	convert date and time
grpauth()	pwdauth(3)	password authentication routines
gsignal()	ssignal(3)	software signals
gtty()	stty(3C)	set and get terminal state
hasmntopt()	getmntent(3)	get file system descriptor file entry
hcreate()	hsearch(3)	manage hash search tables
hdestroy()	hsearch(3)	manage hash search tables
host2netname()	rpc(3N)	RPC services routines
hsearch()	hsearch(3)	manage hash search tables
htonl()	byteorder(3N)	convert values between host and network byte order
htons()	byteorder(3N)	convert values between host and network byte order
index()	string(3)	string operations
inet_addr()	inet(3N)	Internet address manipulation
inet_lnaof()	inet(3N)	Internet address manipulation
inet_makeaddr()	inet(3N)	Internet address manipulation
inet_netof()	inet(3N)	Internet address manipulation
inet_network()	inet(3N)	Internet address manipulation
inet_ntoa()	inet(3N)	Internet address manipulation
information()	getacinfo(3)	get audit control file information
initgroups()	initgroups(3)	initialize group access list
initstate()	random(3)	routines for changing random number generators
innetgr()	getnetgrent(3N)	get network group entry
insque()	insque(3)	insert/remove element from a queue
isalnum()	ctype(3)	character classification and conversion macros and functions
isalpha()	ctype(3)	character classification and conversion macros and functions
isascii()	ctype(3)	character classification and conversion macros and functions
isatty()	ttyname(3)	find name of a terminal
iscntrl()	ctype(3)	character classification and conversion macros and functions
isdigit()	ctype(3)	character classification and conversion macros and functions
isgraph()	ctype(3)	character classification and conversion macros and functions
islower()	ctype(3)	character classification and conversion macros and functions
isprint()	ctype(3)	character classification and conversion macros and functions
ispunct()	ctype(3)	character classification and conversion macros and functions
issecure()	issecure(3)	indicates whether system is running secure
isspace()	ctype(3)	character classification and conversion macros and functions
isupper()	ctype(3)	character classification and conversion macros and functions
isxdigit()	ctype(3)	character classification and conversion macros and functions
itom()	mp(3X)	multiple precision integer arithmetic
jrand48()	drand48(3)	generate uniformly distributed pseudo-random numbers
key_decryptsession()	rpc(3N)	RPC services routines
key_encryptsession()	rpc(3N)	RPC services routines
key_gendes()	rpc(3N)	RPC services routines
key_setsecret()	rpc(3N)	RPC services routines
kvm_close()	kvm_open(3K)	specify a kernel to examine
kvm_getcmd()	kvm_getu(3K)	get the u-area or invocation arguments for a process
kvm_getproc()	kvm_nextproc(3K)	read system process structures

kvm_getu()	kvm_getu(3K)	get the u-area or invocation arguments for a process
kvm_nextproc()	kvm_nextproc(3K)	read system process structures
kvm_nlist()	kvm_nlist(3K)	obtain kernel symbol table information
kvm_open()	kvm_open(3K)	specify a kernel to examine
kvm_read()	kvm_read(3K)	copy data to or from a kernel image or running system
kvm_setproc()	kvm_nextproc(3K)	read system process structures
kvm_write()	kvm_read(3K)	copy data to or from a kernel image or running system
l64a()	a64l(3)	convert between long integer and base-64 ASCII string
label()	plot(3X)	graphics interface
lcong48()	drand48(3)	generate uniformly distributed pseudo-random numbers
ldaclose()	ldclose(3X)	close a COFF file
ldahread()	ldahread(3X)	read the archive header of a member of a COFF archive file
ldaopen()	ldopen(3X)	open a COFF file for reading
ldclose()	ldclose(3X)	close a COFF file
ldfcn()	ldfcn(3)	common object file access routines
ldfhread()	ldfhread(3X)	read the file header of a COFF file
ldgetname()	ldgetname(3X)	retrieve symbol name for COFF file symbol table entry
ldlinit()	ldlread(3X)	manipulate line number entries of a COFF file function
ldlitem()	ldlread(3X)	manipulate line number entries of a COFF file function
ldlread()	ldlread(3X)	manipulate line number entries of a COFF file function
ldlseek()	ldlseek(3X)	seek to line number entries of a section of a COFF file
ldnseek()	ldlseek(3X)	seek to line number entries of a section of a COFF file
ldnrseek()	ldrseek(3X)	seek to relocation entries of a section of a COFF file
ldnshread()	ldshread(3X)	read an indexed/named section header of a COFF file
ldnsseek()	ldsseek(3X)	seek to an indexed/named section of a COFF file
ldohseek()	ldohseek(3X)	seek to the optional file header of a COFF file
ldopen()	ldopen(3X)	open a COFF file for reading
ldrseek()	ldrseek(3X)	seek to relocation entries of a section of a COFF file
ldshread()	ldshread(3X)	read an indexed/named section header of a COFF file
ldsseek()	ldsseek(3X)	seek to an indexed/named section of a COFF file
ldtbindex()	ldtbindex(3X)	compute the index of a symbol table entry of a COFF file
ldtbread()	ldtbread(3X)	read an indexed symbol table entry of a COFF file
ldtbseek()	ldtbseek(3X)	seek to the symbol table of a COFF file
lfind()	lsearch(3)	linear search and update
line()	plot(3X)	graphics interface
linemod()	plot(3X)	graphics interface
localtime()	ctime(3)	convert date and time
lockf()	lockf(3)	advisory record locking on files
longjmp()	setjmp(3)	non-local goto
lrand48()	drand48(3)	generate uniformly distributed pseudo-random numbers
lsearch()	lsearch(3)	linear search and update
madd()	mp(3X)	multiple precision integer arithmetic
malloc()	malloc(3)	memory allocator
malloc_debug()	malloc(3)	memory allocator
malloc_verify()	malloc(3)	memory allocator
mdiv()	mp(3X)	multiple precision integer arithmetic
memalign()	malloc(3)	memory allocator
memccpy()	memory(3)	memory operations
memchr()	memory(3)	memory operations
memcmp()	memory(3)	memory operations
memcpy()	memory(3)	memory operations
memory()	memory(3)	memory operations
memset()	memory(3)	memory operations

mfree()	mp(3X)	multiple precision integer arithmetic
min()	mp(3X)	multiple precision integer arithmetic
mkstemp()	mktemp(3)	make a unique file name
mktemp()	mktemp(3)	make a unique file name
moncontrol()	monitor(3)	prepare execution profile
monitor()	monitor(3)	prepare execution profile
monstartup()	monitor(3)	prepare execution profile
mout()	mp(3X)	multiple precision integer arithmetic
move()	plot(3X)	graphics interface
rand48()	drand48(3)	generate uniformly distributed pseudo-random numbers
msub()	mp(3X)	multiple precision integer arithmetic
mtox()	mp(3X)	multiple precision integer arithmetic
mult()	mp(3X)	multiple precision integer arithmetic
netname2host()	rpc(3N)	RPC services routines
netname2user()	rpc(3N)	RPC services routines
nextkey()	dbm(3X)	data base subroutines
nice()	nice(3C)	change priority of a process
nlist()	nlist(3)	get entries from name list
rand48()	drand48(3)	generate uniformly distributed pseudo-random numbers
ntohl()	byteorder(3N)	convert values between host and network byte order
ntohs()	byteorder(3N)	convert values between host and network byte order
on_exit()	on_exit(3)	name termination handler
opendir()	directory(3)	directory operations
openlog()	syslog(3)	control system log
openpl()	plot(3X)	graphics interface
optarg()	getopt(3)	get option letter from argument vector
optind()	getopt(3)	get option letter from argument vector
pause()	pause(3C)	stop until signal
pclose()	popen(3S)	open or close a pipe (for I/O) from or to a process
perror()	perror(3)	system error messages
pmap_getmaps()	rpc(3N)	RPC services routines
pmap_getport()	rpc(3N)	RPC services routines
pmap_rmtcall()	rpc(3N)	RPC services routines
pmap_set()	rpc(3N)	RPC services routines
pmap_unset()	rpc(3N)	RPC services routines
point()	plot(3X)	graphics interface
popen()	popen(3S)	open or close a pipe (for I/O) from or to a process
pow()	mp(3X)	multiple precision integer arithmetic
printf()	printf(3S)	formatted output conversion
prof()	prof(3)	profile within a function
psignal()	psignal(3)	system signal messages
putc()	putc(3S)	put character or word on a stream
putchar()	putc(3S)	put character or word on a stream
putenv()	putenv(3)	change or add value to environment
putpwent()	putpwent(3)	write password file entry
puts()	puts(3S)	put a string on a stream
putw()	putc(3S)	put character or word on a stream
pwdauth()	pwdauth(3)	password authentication routines
qsort()	qsort(3)	quicker sort
rand()	rand(3C)	simple random number generator
random()	random(3)	routines for changing random number generators
rcmd()	rcmd(3N)	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
regex()	regex(3)	regular expression handler
regexp()	regexp(3)	regular expression compile and match routines
registerrpc()	rpc(3N)	RPC services routines
remexportent()	exportent(3)	get exported file system information
remque()	insque(3)	insert/remove element from a queue
res_init()	resolver(3)	resolver routines
res_mkquery()	resolver(3)	resolver routines
res_send()	resolver(3)	resolver routines
resolver()	resolver(3)	resolver routines
rewind()	fseek(3S)	reposition a stream
rewinddir()	directory(3)	directory operations
rexec()	rexec(3N)	return stream to a remote command
rindex()	string(3)	string operations
rpc_createrr()	rpc(3N)	RPC services routines
rpow()	mp(3X)	multiple precision integer arithmetic
rresvport()	rcmd(3N)	routines for returning a stream to a remote command
rtime()	rtime(3N)	get remote time
ruserok()	rcmd(3N)	routines for returning a stream to a remote command
scandir()	scandir(3)	scan a directory
scanf()	scanf(3S)	formatted input conversion
seconvert()	econvert(3)	output conversion
seed48()	drand48(3)	generate uniformly distributed pseudo-random numbers
seekdir()	directory(3)	directory operations
setac()	getacinfo(3)	get audit control file information
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
setexportent()	exportent(3)	get exported file system information
setfsent()	getfsent(3)	get file system descriptor file entry
setgid()	setuid(3)	set user and group ID
setgraent()	getgraent(3)	get group adjunct file entry
setgrent()	getgrent(3)	get group file entry
sethostent()	gethostent(3N)	get network host entry
setjmp()	setjmp(3)	non-local goto
setkey()	crypt(3)	password and data encryption
setlinebuf()	setbuf(3S)	assign buffering to a stream
setlogmask()	syslog(3)	control system log
setmntent()	getmntent(3)	get file system descriptor file entry
setnetent()	getnetent(3N)	get network entry
setnetgrent()	getnetgrent(3N)	get network group entry
setprotoent()	getprotoent(3N)	get protocol entry
setpwaent()	getpwaent(3)	get password adjunct file entry
setpwent()	getpwent(3)	get password file entry
setpwfile()	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)	routines for changing random number generators
settyent()	getttyent(3)	get ttytab file entry

setuid()	setuid(3)	set user and group ID
setusershell()	getusershell(3)	get legal user shells
setvbuf()	setbuf(3S)	assign buffering to a stream
sfconvert()	econvert(3)	output conversion
sgconvert()	econvert(3)	output conversion
sigfpe()	sigfpe(3)	signal handling for specific SIGFPE codes
siginterrupt()	siginterrupt(3)	allow signals to interrupt system calls
signal()	signal(3)	simplified software signal facilities
single_to_decimal()	floating_to_decimal(3)	convert floating-point value to decimal record
sleep()	sleep(3)	suspend execution for interval
space()	plot(3X)	graphics interface
sprintf()	printf(3S)	formatted output conversion
srand()	rand(3C)	simple random number generator
srand48()	drand48(3)	generate uniformly distributed pseudo-random numbers
srandom()	random(3)	routines for changing random number generators
sscanf()	scanf(3S)	formatted input conversion
ssignal()	ssignal(3)	software signals
store()	dbm(3X)	data base subroutines
strcat()	string(3)	string operations
strchr()	string(3)	string operations
strcmp()	string(3)	string operations
strcpy()	string(3)	string operations
strcspn()	string(3)	string operations
strdup()	string(3)	string operations
string_to_decimal()(3)	string_to_decimal(3)	parse characters into decimal record
strlen()	string(3)	string operations
strncat()	string(3)	string operations
strncmp()	string(3)	string operations
strncpy()	string(3)	string operations
strpbrk()	string(3)	string operations
strrchr()	string(3)	string operations
strspn()	string(3)	string operations
strtod()	strtod(3)	convert string to double-precision number
strtok()	string(3)	string operations
strtol()	strtol(3)	convert string to integer
stty()	stty(3C)	set and get terminal state
svc_destroy()	rpc(3N)	RPC services routines
svc_fds()	rpc(3N)	RPC services routines
svc_freeargs()	rpc(3N)	RPC services routines
svc_getargs()	rpc(3N)	RPC services routines
svc_getcaller()	rpc(3N)	RPC services routines
svc_getreq()	rpc(3N)	RPC services routines
svc_register()	rpc(3N)	RPC services routines
svc_run()	rpc(3N)	RPC services routines
svc_sendreply()	rpc(3N)	RPC services routines
svc_unregister()	rpc(3N)	RPC services routines
svcerr_auth()	rpc(3N)	RPC services routines
svcerr_decode()	rpc(3N)	RPC services routines
svcerr_noproc()	rpc(3N)	RPC services routines
svcerr_noprogram()	rpc(3N)	RPC services routines
svcerr_progvers()	rpc(3N)	RPC services routines
svcerr_systemerr()	rpc(3N)	RPC services routines
svcerr_weakauth()	rpc(3N)	RPC services routines

svcfld_create()	rpc(3N)	RPC services routines
svcrow_create()	rpc(3N)	RPC services routines
svctcp_create()	rpc(3N)	RPC services routines
svcudp_create()	rpc(3N)	RPC services routines
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
tdelete()	tsearch(3)	manage binary search trees
telldir()	directory(3)	directory operations
tempnam()	tmpnam(3S)	create a name for a temporary file
tfind()	tsearch(3)	manage binary search trees
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
timegm()	ctime(3)	convert date and time
timelocal()	ctime(3)	convert date and time
times()	times(3C)	get process times
timezone()	timezone(3C)	get time zone name given offset from GMT
tmpfile()	tmpfile(3S)	create a temporary file
tmpnam()	tmpnam(3S)	create a name for a temporary file
toascii()	ctype(3)	character classification and conversion macros and functions
tolower()	ctype(3)	character classification and conversion macros and functions
toupper()	ctype(3)	character classification and conversion macros and functions
tputs()	termcap(3X)	terminal independent operation routines
tsearch()	tsearch(3)	manage binary search trees
ttyname()	ttyname(3)	find name of a terminal
ttyslot()	ttyslot(3)	find the slot in the utmp file of the current process
twalk()	tsearch(3)	manage binary search trees
tzset()	ctime(3)	convert date and time
tzsetwall()	ctime(3)	convert date and time
ualarm()	ualarm(3)	schedule signal after interval in microseconds
ulimit()	ulimit(3C)	get and set user limits
ungetc()	ungetc(3S)	push character back into input stream
user2netname()	rpc(3N)	RPC services routines
usleep()	usleep(3)	suspend execution for interval in microseconds
utime()	utime(3C)	set file times
valloc()	malloc(3)	memory allocator
varargs()	varargs(3)	handle variable argument list
fprintf()	fprintf(3S)	print formatted output of a varargs argument list
vlimit()	vlimit(3C)	control maximum system resource consumption
vprintf()	vprintf(3S)	print formatted output of a varargs argument list
vsprintf()	vprintf(3S)	print formatted output of a varargs argument list
vtimes()	vtimes(3C)	get information about resource utilization
xdr_accepted_reply()	xdr(3N)	XDR functions
xdr_array()	xdr(3N)	XDR functions
xdr_authunix_parms()	xdr(3N)	XDR functions
xdr_bool()	xdr(3N)	XDR functions

xdr_bytes()	xdr(3N)	XDR functions
xdr_callhdr()	xdr(3N)	XDR functions
xdr_callmsg()	xdr(3N)	XDR functions
xdr_char()	xdr(3N)	XDR functions
xdr_destroy()	xdr(3N)	XDR functions
xdr_double()	xdr(3N)	XDR functions
xdr_enum()	xdr(3N)	XDR functions
xdr_float()	xdr(3N)	XDR functions
xdr_getpos()	xdr(3N)	XDR functions
xdr_inline()	xdr(3N)	XDR functions
xtom()	mp(3X)	multiple precision integer arithmetic
yp_all()	ypclnt(3N)	Yellow Pages client interface
yp_bind()	ypclnt(3N)	Yellow Pages client interface
yp_first()	ypclnt(3N)	Yellow Pages client interface
yp_get_default_domain()	ypclnt(3N)	Yellow Pages client interface
yp_master()	ypclnt(3N)	Yellow Pages client interface
yp_match()	ypclnt(3N)	Yellow Pages client interface
yp_next()	ypclnt(3N)	Yellow Pages client interface
yp_order()	ypclnt(3N)	Yellow Pages client interface
yp_unbind()	ypclnt(3N)	Yellow Pages client interface
yp_update()	ypupdate(3N)	update YP information
ypclnt()	ypclnt(3N)	Yellow Pages client interface
yperr_string()	ypclnt(3N)	Yellow Pages client interface
ypprot_err()	ypclnt(3N)	Yellow Pages client interface

NAME

a64l, l64a – convert between long integer and base-64 ASCII string

SYNOPSIS

```
long a64l(s)  
char *s;  
char *l64a(l)  
long l;
```

DESCRIPTION

These functions are used to maintain numbers stored in *base-64* ASCII characters. This is a notation by which long integers can be represented by up to six characters; each character represents a “digit” in a radix-64 notation.

The characters used to represent “digits” are ‘.’ for 0, ‘/’ for 1, 0 through 9 for 2–11, A through Z for 12–37, and a through z for 38–63.

a64l() takes a pointer to a NULL-terminated base-64 representation and returns a corresponding long value. If the string pointed to by *s* contains more than six characters, **a64l()** will use the first six.

l64a() takes a long argument and returns a pointer to the corresponding base-64 representation. If the argument is 0, **l64a()** returns a pointer to a NULL string.

BUGS

The value returned by **l64a()** is a pointer into a static buffer, the contents of which are overwritten by each call.

NAME

abort – generate a fault

SYNOPSIS

abort()

DESCRIPTION

abort() first closes all open files if possible, then sends an IOT signal to the process. This signal usually results in termination with a core dump, which may be used for debugging.

It is possible for **abort()** to return control if SIGIOT is caught or ignored, in which case the value returned is that of the **kill(2V)** system call.

SEE ALSO

adb(1), **exit(2)**, **kill(2V)**, **signal(3)**

DIAGNOSTICS

If SIGIOT is neither caught nor ignored, and the current directory is writable, a core dump is produced and the message '**abort – core dumped**' is written by the shell.

NAME

abs – integer absolute value

SYNOPSIS

```
abs(i)  
int i;
```

DESCRIPTION

abs() returns the absolute value of its integer operand.

SEE ALSO

rint(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

alarm – schedule signal after specified time

SYNOPSIS

alarm(seconds)
unsigned seconds;

DESCRIPTION

alarm() sends signal **SIGALRM**, see **sigvec(2)**, 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(3)**, **sleep(3)**, **ualarm(3)**, **usleep(3)**

NAME

assert – program verification

SYNOPSIS

#include <assert.h>

assert(expression)

DESCRIPTION

assert() is a macro that indicates *expression* is expected to be true at this point in the program. It exits (see **exit(2)**) with a diagnostic comment on the standard output when *expression* is false (0). Compiling with the **cc(1V)** option **-DNDEBUG** effectively deletes **assert()** from the program.

SEE ALSO

cc(1V) **exit(2)**

DIAGNOSTICS

Assertion failed: file *f* line *n*

f is the source file and *n* the source line number of the **assert()** statement.

NAME

`audit_args`, `audit_text` – produce text audit message

SYNOPSIS

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

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

```
audit_args(event, argc, argv)
```

```
int event;
```

```
int argc;
```

```
char **argv;
```

```
audit_text(event, error, retval, argc, argv)
```

```
int event;
```

```
int error;
```

```
int retval;
```

```
int argc;
```

```
char **argv;
```

DESCRIPTION

These functions provide text interfaces to the `audit(2)` system call. In both calls, the *event* parameter identifies the event class of the action, and *argc* is the number of strings found in the vector *argv*. The *error* parameter is used to determine the failure or success of the audited operation. A negative value is always audited. A zero value is audited as a successful event. A positive value is audited as an event failure. The *retval* parameter is the return value or exit code that the invoking program will have.

`audit_args()` is equivalent to `audit_text()` with *error* and *retval* parameters of `-1`.

SEE ALSO

`audit(2)`

NAME

bindresvport – bind a socket to a privileged IP port

SYNOPSIS

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

int bindresvport(sd, sin)
int sd;
struct sockaddr_in *sin;
```

DESCRIPTION

bindresvport() is used to bind a socket descriptor to a privileged IP port, that is, a port number in the range 0-1023. The routine returns 0 if it is successful, otherwise -1 is returned and **errno** set to reflect the cause of the error. This routine differs with **rresvport** (see **rcmd(3N)**) in that this works for any IP socket, whereas **rresvport()** only works for TCP.

Only root can bind to a privileged port; this call will fail for any other users.

SEE ALSO

rcmd(3N)

NAME

bsearch – binary search a sorted table

SYNOPSIS

```
#include <search.h>
```

```
char *bsearch ((char *) key, (char *) base, nel, sizeof (*key), compar)  
unsigned nel;  
int (*compar)();
```

DESCRIPTION

bsearch() is a binary search routine generalized from Knuth (6.2.1) Algorithm B. It returns a pointer into a table indicating where a datum may be found. The table must be previously sorted in increasing order according to a provided comparison function. *key* points to a datum instance to be sought in the table. *base* points to the element at the base of the table. *nel* is the number of elements in the table. *compar* is the name of the comparison function, which is called with two arguments that point to the elements being compared. The function must return an integer less than, equal to, or greater than zero as accordingly the first argument is to be considered less than, equal to, or greater than the second.

EXAMPLE

The example below searches a table containing pointers to nodes consisting of a string and its length. The table is ordered alphabetically on the string in the node pointed to by each entry.

This code fragment reads in strings and either finds the corresponding node, in which case it prints out the string and its length, or it prints an error message.

```

#include <stdio.h>
#include <search.h>
#define TABSIZE      1000
struct node {          /* these are stored in the table */
    char *string;
    int length;
};
struct node table[TABSIZE]; /* table to be searched */
.
.
.
{
    struct node *node_ptr, node;
    int node_compare(); /* routine to compare 2 nodes */
    char str_space[20]; /* space to read string into */
    .
    .
    .
    node.string = str_space;
    while (scanf("%s", node.string) != EOF) {
        node_ptr = (struct node *)bsearch((char *)&node,
            (char *)table, TABSIZE,
            sizeof(struct node), node_compare);
        if (node_ptr != NULL) {
            (void)printf("string = %20s, length = %d\n",
                node_ptr->string, node_ptr->length);
        } else {
            (void)printf("not found: %s\n", node.string);
        }
    }
}
/*
    This routine compares two nodes based on an
    alphabetical ordering of the string field.
*/
int
node_compare(node1, node2)
struct node *node1, *node2;
{
    return strcmp(node1->string, node2->string);
}

```

NOTES

The pointers to the key and the element at the base of the table should be of type pointer-to-element, and cast to type pointer-to-character.

The comparison function need not compare every byte, so arbitrary data may be contained in the elements in addition to the values being compared.

Although declared as type pointer-to-character, the value returned should be cast into type pointer-to-element.

SEE ALSO

hsearch(3), lsearch(3), qsort(3), tsearch(3)

DIAGNOSTICS

A NULL pointer is returned if the key cannot be found in the table.

NAME

bstring, bcopy, bcmp, bzero, ffs – bit and byte string operations

SYNOPSIS

bcopy(b1, b2, length)

char *b1, *b2;

int length;

int bcmp(b1, b2, length)

char *b1, *b2;

int length;

bzero(b, length)

char *b;

int length;

int 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*. Overlapping strings are handled correctly.

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. **bcmp()** of length zero bytes always returns zero.

bzero places *length* 0 bytes in the string *b*.

ffs finds the first bit set in the argument passed it and returns the index of that bit. Bits are numbered starting at 1 from the right. A return value of zero indicates that the value passed is zero.

CAVEAT

The **bcmp()** and **bcopy()** routines take parameters backwards from **strcmp()** and **strcpy**.

SEE ALSO

string(3)

NAME

byteorder, 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 Sun-2, Sun-3 and Sun-4 systems, these routines are defined as NULL macros in the include file <netinet/in.h>. On Sun386i systems, these routines are functional since its host byte order is different from network byte order.

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)`

NAME

clock – report CPU time used

SYNOPSIS

long clock ()

DESCRIPTION

clock() returns the amount of CPU time (in microseconds) used since the first call to **clock**. The time reported is the sum of the user and system times of the calling process and its terminated child processes for which it has executed **wait(2)** or **system(3)**.

The resolution of the clock is 16.667 milliseconds.

SEE ALSO

wait(2), **system(3)**, **times(3C)**, **times(3V)**

BUGS

The value returned by **clock()** is defined in microseconds for compatibility with systems that have CPU clocks with much higher resolution. Because of this, the value returned will wrap around after accumulating only 2147 seconds of CPU time (about 36 minutes).

NAME

`crypt`, `_crypt`, `setkey`, `encrypt` – password and data encryption

SYNOPSIS

`char *crypt(key, salt)`

`char *key, *salt;`

`char * _crypt(key, salt)`

`char *key, *salt;`

`setkey(key)`

`char *key;`

`encrypt(block, edflag)`

`char *block;`

DESCRIPTION

`crypt()` is the password encryption routine, 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./]. Unless it starts with '##' or '#\$', 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.

If the *salt* string starts with '##', `pwdauth(3)` is called. If `pwdauth` returns TRUE, the salt is returned from `crypt`. Otherwise, NULL is returned. If the *salt* string starts with '#\$', `grpauth` (see `pwauth(3)`) is called. If `grpauth` returns TRUE, the salt is returned from `crypt`. Otherwise, NULL is returned. If there is a valid reason not to have this authentication happen, calling `_crypt` avoids authentication.

The `setkey` and `encrypt` entries provide (rather primitive) access to the DES algorithm. The argument of `setkey` is a character array of length 64 containing only the characters with numerical value 0 and 1. If this string is divided into groups of 8, the low-order bit in each group is ignored; this gives a 56-bit key which is set into the machine. This is the key that will be used with the above mentioned algorithm to encrypt or decrypt the string *block* with the function `encrypt`.

The argument to the `encrypt` entry is a character array of length 64 containing only the characters with numerical value 0 and 1. The argument array is modified in place to a similar array representing the bits of the argument after having been subjected to the DES algorithm using the key set by `setkey`. If *edflag* is zero, the argument is encrypted; if non-zero, it is decrypted.

SEE ALSO

`login(1)`, `passwd(1)`, `getpass(3)`, `pwdauth(3)`, `passwd(5)`

BUGS

The return value points to static data whose content is overwritten by each call.

NAME

ctermid – generate filename for terminal

SYNOPSIS

```
#include <stdio.h>
char *ctermid (s)
char *s;
```

DESCRIPTION

ctermid() generates the pathname of the controlling terminal for the current process, and stores it in a string.

If *s* is a NULL pointer, the string is stored in an internal static area, the contents of which are overwritten at the next call to **ctermid**, and the address of which is returned. Otherwise, *s* is assumed to point to a character array of at least **L_ctermid** elements; the path name is placed in this array and the value of *s* is returned. The constant **L_ctermid** is defined in the **<stdio.h>** header file.

NOTES

The difference between **ctermid()** and **ttyname(3)** is that **ttyname()** must be handed a file descriptor and returns the actual name of the terminal associated with that file descriptor, while **ctermid()** returns a string (**/dev/tty**) that will refer to the terminal if used as a file name. Thus **ttyname()** is useful only if the process already has at least one file open to a terminal. **ctermid()** is useful largely for making code portable to (non-UNIX) systems where the current terminal is referred to by a name other than **/dev/tty**.

SEE ALSO

ttyname(3)

NAME

`ctime`, `localtime`, `gmtime`, `asctime`, `dysize`, `timelocal`, `timegm`, `tzset`, `tzsetwall` – convert date and time

SYNOPSIS

```
#include <time.h>

struct tm *localtime(clock)
long *clock;

struct tm *gmtime(clock)
long *clock;

char *asctime(tm)
struct tm *tm;

char *ctime(clock)
long *clock;

int dysize(y)
int y;

time_t timelocal(tm)
struct tm *tm;

time_t timegm(tm)
struct tm *tm;

void tzset()

void tzsetwall()
```

DESCRIPTION

`localtime()` and `gmtime()` return pointers to structures containing the time, broken down into various components of that time represented in a particular time zone. `localtime()` breaks down a time specified by the `clock()` argument, correcting for the time zone and any time zone adjustments (such as Daylight Savings Time). Before doing so, `localtime()` calls `tzset` (if `tzset` has not been called in the current process). `gmtime()` breaks down a time specified by the `clock()` argument into GMT, which is the time the system uses.

`asctime` converts a time value contained in a “tm” structure to a 26-character string of the form:

```
Sun Sep 16 01:03:52 1973\n\0
```

Each field has a constant width. `asctime` returns a pointer to the string.

`ctime()` converts a long integer, pointed to by `clock`, to a 26-character string of the form produced by `asctime`. It first breaks down `clock()` to a `struct tm` by calling `localtime()`, and then calls `asctime` to convert that `struct tm` to a string.

`dysize` returns the number of days in the argument year, either 365 or 366.

`timelocal()` and `timegm()` convert the time specified by the `tm` argument to a time value that represents that time expressed as the number of seconds since Jan. 1, 1970, 00:00, Greenwich Mean Time. `timelocal()` converts a `struct tm` that represents local time, correcting for the time zone and any time zone adjustments (such as Daylight Savings Time). Before doing so, `timelocal()` calls `tzset` (if `tzset` has not been called in the current process). `timegm()` converts a `struct tm` that represents GMT.

`tzset` uses the value of the environment variable `TZ` to set time conversion information used by `localtime()`. If `TZ` is absent from the environment, the best available approximation to local wall clock time is used by `localtime()`. If `TZ` appears in the environment but its value is a NULL string, Greenwich Mean Time is used; if `TZ` appears and begins with a slash, it is used as the absolute pathname of the `tzfile-format` (see `tzfile(5)`) file from which to read the time conversion information; if `TZ` appears and begins with a character other than a slash, it is used as a pathname relative to a system time conversion information directory.

tzsetwall sets things up so that `localtime()` returns the best available approximation of local wall clock time.

Declarations of all the functions and externals, and the “tm” structure, are in the `<time.h>` header file. The structure (of type) `struct tm` includes the following fields:

```

    int tm_sec;      /* seconds (0 - 59) */
    int tm_min;     /* minutes (0 - 59) */
    int tm_hour;    /* hours (0 - 23) */
    int tm_mday;    /* day of month (1 - 31) */
    int tm_mon;     /* month of year (0 - 11) */
    int tm_year;    /* year - 1900 */
    int tm_wday;    /* day of week (Sunday = 0) */
    int tm_yday;    /* day of year (0 - 365) */
    int tm_isdst;   /* 1 if DST in effect */
    char *tm_zone; /* abbreviation of timezone name */
    long tm_gmtoff; /* offset from GMT in seconds */

```

`tm_isdst` is non-zero if Daylight Savings Time is in effect. `tm_zone` points to a string that is the name used for the local time zone at the time being converted. `tm_gmtoff` is the offset (in seconds) of the time represented from GMT, with positive values indicating East of Greenwich.

FILES

```

/usr/share/lib/zoneinfo standard time conversion information directory
/usr/share/lib/zoneinfo/localtime
                           local time zone file

```

SEE ALSO

`gettimeofday(2)`, `ctime(3V)`, `getenv(3)`, `time(3C)`, `environ(5V)`, `tzfile(5)`

BUGS

The return values point to static data, whose contents are overwritten by each call. The `tm_zone` field of a returned `struct tm` points to a static array of characters, which will also be overwritten at the next call (and by calls to *tzset* or *tzsetwall*).

NAME

`ctype`, `isalpha`, `isupper`, `islower`, `isdigit`, `isxdigit`, `isalnum`, `isspace`, `ispunct`, `isprint`, `iscntrl`, `isascii`, `isgraph`, `toupper`, `tolower`, `toascii` – character classification and conversion macros and functions

SYNOPSIS

```
#include <ctype.h>
```

```
isalpha(c)
```

```
...
```

CHARACTER CLASSIFICATION MACROS

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(c)` is true and on the single non-ASCII value EOF (see `stdio(3S)`).

<code>isalpha(c)</code>	<code>c</code> is a letter
<code>isupper(c)</code>	<code>c</code> is an upper case letter
<code>islower(c)</code>	<code>c</code> is a lower case letter
<code>isdigit(c)</code>	<code>c</code> is a digit [0-9].
<code>isxdigit(c)</code>	<code>c</code> is a hexadecimal digit [0-9], [A-F], or [a-f].
<code>isalnum(c)</code>	<code>c</code> is an alphanumeric character, that is, <code>c</code> is a letter or a digit
<code>isspace(c)</code>	<code>c</code> is a space, tab, carriage return, newline, vertical tab, or formfeed
<code>ispunct(c)</code>	<code>c</code> is a punctuation character (neither control nor alphanumeric)
<code>isprint(c)</code>	<code>c</code> is a printing character, code 040(8) (space) through 0176 (tilde)
<code>iscntrl(c)</code>	<code>c</code> is a delete character (0177) or ordinary control character (less than 040).
<code>isascii(c)</code>	<code>c</code> is an ASCII character, code less than 0200
<code>isgraph(c)</code>	<code>c</code> is a visible graphic character, code 041 (exclamation mark) through 0176 (tilde).

CHARACTER CONVERSION MACROS

These macros perform simple conversions on single characters.

<code>toupper(c)</code>	converts <code>c</code> to its upper-case equivalent. Note: this <i>only</i> works where <code>c</code> is known to be a lower-case character to start with (presumably checked using <code>islower</code>).
<code>tolower(c)</code>	converts <code>c</code> to its lower-case equivalent. Note: this <i>only</i> works where <code>c</code> is known to be an upper-case character to start with (presumably checked using <code>isupper</code>).
<code>toascii(c)</code>	masks <code>c</code> with the correct value so that <code>c</code> is guaranteed to be an ASCII character in the range 0 through 0x7f.

DIAGNOSTICS

If the argument to any of these macros is not in the domain of the function, the result is undefined.

SEE ALSO

`ctype(3V)`, `stdio(3S)`, `ascii(7)`

NAME

curses – cursor addressing and screen display library

SYNOPSIS

cc [*flags*] *files* -lcurses -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

`tic(8V)`, `ioctl(2)`, `curses(3V)`, `getenv(3)`, `tty(4)`, `terminfo(5V)`, `term(5V)`, `termcap(5)`

Programming Utilities and Libraries

Curses Functions

<code>addch(<i>ch</i>)</code>	add a character to <i>stdscr</i>
<code>addstr(<i>str</i>)</code>	add a string to <i>stdscr</i>
<code>box(<i>win,vert,hor</i>)</code>	draw a box around a window
<code>cbreak()</code>	set cbreak mode
<code>clear()</code>	clear <i>stdscr</i>
<code>clearok(<i>scr,boolf</i>)</code>	set clear flag for <i>scr</i>
<code>clrtobot()</code>	clear to bottom on <i>stdscr</i>
<code>clrtoeol()</code>	clear to end of line on <i>stdscr</i>
<code>delch()</code>	delete a character
<code>deleteln()</code>	delete a line
<code>delwin(<i>win</i>)</code>	delete <i>win</i>
<code>echo()</code>	set echo mode
<code>endwin()</code>	end window modes
<code>erase()</code>	erase <i>stdscr</i>
<code>flusok(<i>win,boolf</i>)</code>	set flush-on-refresh flag for <i>win</i>
<code>getch()</code>	get a char through <i>stdscr</i>
<code>getcap(<i>name</i>)</code>	get terminal capability <i>name</i>
<code>getstr(<i>str</i>)</code>	get a string through <i>stdscr</i>
<code>gettmode()</code>	get tty modes
<code>getyx(<i>win,y,x</i>)</code>	get (<i>y,x</i>) co-ordinates
<code>inch()</code>	get char at current (<i>y,x</i>) co-ordinates
<code>initscr()</code>	initialize screens
<code>insch(<i>c</i>)</code>	insert a char
<code>insertln()</code>	insert a line
<code>leaveok(<i>win,boolf</i>)</code>	set leave flag for <i>win</i>
<code>longname(<i>termbuf,name</i>)</code>	get long name from <i>termbuf</i>
<code>move(<i>y,x</i>)</code>	move to (<i>y,x</i>) on <i>stdscr</i>
<code>mvcur(<i>lasty,lastx,newy,newx</i>)</code>	actually move cursor
<code>newwin(<i>lines,cols,begin_y,begin_x</i>)</code>	create a new window
<code>nl()</code>	set NEWLINE mapping
<code>nocbreak()</code>	unset cbreak mode
<code>noecho()</code>	unset echo mode
<code>nonl()</code>	unset NEWLINE mapping
<code>noraw()</code>	unset raw mode
<code>overlay(<i>win1,win2</i>)</code>	overlay <i>win1</i> on <i>win2</i>
<code>overwrite(<i>win1,win2</i>)</code>	overwrite <i>win1</i> on top of <i>win2</i>
<code>printw(<i>fmt,arg1,arg2,...</i>)</code>	printf on <i>stdscr</i>
<code>raw()</code>	set raw mode

refresh()	make current screen look like <i>stdscr</i>
resetty()	reset tty flags to stored value
savetty()	stored current tty flags
scanw(fmt,arg1,arg2,...)	scanf through <i>stdscr</i>
scroll(win)	scroll <i>win</i> one line
scrollok(win,boolf)	set scroll flag
setterm(name)	set term variables for name
standend()	end standout mode
standout()	start standout mode
subwin(win,lines,cols,begin_y,begin_x)	create a subwindow
touchline(win,y,sx,sy)	mark line <i>y</i> <i>sx</i> through <i>sy</i> as changed
touchoverlap(win1,win2)	mark overlap of <i>win1</i> on <i>win2</i> as changed
touchwin(win)	“change” all of <i>win</i>
unctrl(ch)	printable version of <i>ch</i>
waddch(win,ch)	add char to <i>win</i>
waddstr(win,str)	add string to <i>win</i>
wclear(win)	clear <i>win</i>
wclrto bot(win)	clear to bottom of <i>win</i>
wclrtoeol(win)	clear to end of line on <i>win</i>
wdelch(win,c)	delete char from <i>win</i>
wdeleteln(win)	delete line from <i>win</i>
werase(win)	erase <i>win</i>
wgetch(win)	get a char through <i>win</i>
wgetstr(win,str)	get a string through <i>win</i>
winch(win)	get char at current (<i>y,x</i>) in <i>win</i>
winsch(win,c)	insert character into <i>win</i>
winsertln(win)	insert line into <i>win</i>
wmove(win,y,x)	set current (<i>y,x</i>) co-ordinates on <i>win</i>
wprintw(win,fmt,arg1,arg2,...)	printf on <i>win</i>
wrefresh(win)	make screen look like <i>win</i>
wscanw(win,fmt,arg1,arg2,...)	scanf through <i>win</i>
wstandend(win)	end standout mode on <i>win</i>
wstandout(win)	start standout mode on <i>win</i>

NAME

cuserid – get character login name of the user

SYNOPSIS

```
#include <stdio.h>
```

```
char *cuserid (s)
```

```
char *s;
```

DESCRIPTION

cuserid() generates a character-string representation of the login name that the owner of the current process is logged in under. If *s* is a NULL pointer, this representation is generated in an internal static area, the address of which is returned. Otherwise, *s* is assumed to point to an array of at least **l_cuserid** characters; the representation is left in this array. The constant **l_cuserid** is defined in the `<stdio.h>` header file.

SEE ALSO

getlogin(3), **getpwent(3)**

DIAGNOSTICS

If the login name cannot be found, **cuserid()** returns a NULL pointer; if *s* is not a NULL pointer, a NULL character ('\0') will be placed at **s[0]**.

NAME

dbm, dbminit, dbmclose, fetch, store, delete, firstkey, nextkey – data base subroutines

SYNOPSIS

```
#include <dbm.h>

typedef struct {
    char *dptr;
    int dsize;
} datum;

dbminit(file)
char *file;

dbmclose()

datum fetch(key)
datum key;

store(key, content)
datum key, content;

delete(key)
datum key;

datum firstkey()

datum nextkey(key)
datum key;
```

DESCRIPTION

Note: the `dbm()` library has been superceded by `ndbm(3)`, and is now implemented using `ndbm`.

These functions maintain key/content pairs in a data base. The functions will handle very large (a billion blocks) databases and will access a keyed item in one or two file system accesses. The functions are obtained with the loader option `-ldb`.

keys and *contents* are described by the `datum` typedef. A `datum` specifies a string of *dsize* bytes pointed to by *dptr*. Arbitrary binary data, as well as normal ASCII strings, are allowed. The data base is stored in two files. One file is a directory containing a bit map and has `.dir` as its suffix. The second file contains all data and has `.pag` as its suffix.

Before a database can be accessed, it must be opened by `dbminit`. At the time of this call, the files `file.dir` and `file.pag` must exist. (An empty database is created by creating zero-length `.dir` and `.pag` files.)

A database may be closed by calling `dbmclose`. You must close a database before opening a new one.

Once open, the data stored under a key is accessed by `fetch()` and data is placed under a key by `store`. A key (and its associated contents) is deleted by `delete`. A linear pass through all keys in a database may be made, in an (apparently) random order, by use of `firstkey()` and `nextkey`. `firstkey()` will return the first key in the database. With any key `nextkey()` will return the next key in the database. This code will traverse the data base:

```
for (key = firstkey(); key.dptr != NULL; key = nextkey(key))
```

SEE ALSO

`ndbm(3)`

DIAGNOSTICS

All functions that return an `int` indicate errors with negative values. A zero return indicates no error. Routines that return a `datum` indicate errors with a `NULL (0) dptr`.

BUGS

The `.pag` file will contain holes so that its apparent size is about four times its actual content. Older versions of the UNIX operating system may create real file blocks for these holes when touched. These files

cannot be copied by normal means (`cp(1)`, `cat(1V)`, `tp(5)`, `tar(1)`, `ar(1)`) without filling in the holes.

dptr pointers returned by these subroutines point into static storage that is changed by subsequent calls.

The sum of the sizes of a key/content pair must not exceed the internal block size (currently 1024 bytes). Moreover all key/content pairs that hash together must fit on a single block. `store()` will return an error in the event that a disk block fills with inseparable data.

`delete()` does not physically reclaim file space, although it does make it available for reuse.

The order of keys presented by `firstkey()` and `nextkey()` depends on a hashing function, not on anything interesting.

There are no interlocks and no reliable cache flushing; thus concurrent updating and reading is risky.

NAME

`decimal_to_single`, `decimal_to_double`, `decimal_to_extended` – convert decimal record to floating-point value

SYNOPSIS

```
#include <floatingpoint.h>

void decimal_to_single(px, pm, pd, ps)
single *px ;
decimal_mode *pm;
decimal_record *pd;
fp_exception_field_type *ps;

void decimal_to_double(px, pm, pd, ps)
double *px ;
decimal_mode *pm;
decimal_record *pd;
fp_exception_field_type *ps;

void decimal_to_extended(px, pm, pd, ps)
extended *px ;
decimal_mode *pm;
decimal_record *pd;
fp_exception_field_type *ps;
```

DESCRIPTION

The `decimal_to_floating()` functions convert the decimal record at `*pd` into a floating-point value at `*px`, observing the modes specified in `*pm` and setting exceptions in `*ps`. If there are no IEEE exceptions, `*ps` will be zero.

`pd->sign` and `pd->fpclass` are always taken into account. `pd->exponent` and `pd->ds` are used when `pd->fpclass` is `fp_normal` or `fp_subnormal`. In these cases `pd->ds` must contain one or more ascii digits followed by a NULL. `*px` is set to a correctly rounded approximation to

$$(pd->sign)*(pd->ds)*10^{(pd->exponent)}$$

Thus if `pd->exponent == -2` and `pd->ds == "1234"`, `*px` will get 12.34 rounded to storage precision. `pd->ds` cannot have more than `DECIMAL_STRING_LENGTH-1` significant digits because one character is used to terminate the string with a NULL. If `pd->more != 0` on input then additional nonzero digits follow those in `pd->ds`; `fp_inexact` is set accordingly on output in `*ps`.

`*px` is correctly rounded according to the IEEE rounding modes in `pm->rd`. `*ps` is set to contain `fp_inexact`, `fp_underflow`, or `fp_overflow` if any of these arise.

`pd->ndigits`, `pm->df`, and `pm->ndigits` are not used.

`strtod(3)`, `scanf(3)`, `fscanf(3)`, and `sscanf(3)` all use `decimal_to_double`.

SEE ALSO

`scanf(3S)`, `scanf(3V)`, `strtod(3)`

NAME

`des_crypt`, `ecb_crypt`, `cbc_crypt`, `des_setparity` – fast DES encryption

SYNOPSIS

```
#include <des_crypt.h>

int ecb_crypt(key, data, datalen, mode)
char *key;
char *data;
unsigned datalen;
unsigned mode;

int cbc_crypt(key, data, datalen, mode, ivec)
char *key;
char *data;
unsigned datalen;
unsigned mode;
char *ivec;

void des_setparity(key)
char *key;
```

DESCRIPTION

`ecb_crypt()` and `cbc_crypt()` implement the NBS DES (Data Encryption Standard). These routines are faster and more general purpose than `crypt(3)`. They also are able to utilize DES hardware if it is available. `ecb_crypt()` encrypts in ECB (Electronic Code Book) mode, which encrypts blocks of data independently. `cbc_crypt()` encrypts in CBC (Cipher Block Chaining) mode, which chains together successive blocks. CBC mode protects against insertions, deletions and substitutions of blocks. Also, regularities in the clear text will not appear in the cipher text.

Here is how to use these routines. The first parameter, *key*, is the 8-byte encryption key with parity. To set the key's parity, which for DES is in the low bit of each byte, use *des_setparity*. The second parameter, *data*, contains the data to be encrypted or decrypted. The third parameter, *datalen*, is the length in bytes of *data*, which must be a multiple of 8. The fourth parameter, *mode*, is formed by OR'ing together some things. For the encryption direction 'or' in either `DES_ENCRYPT` or `DES_DECRYPT`. For software versus hardware encryption, 'or' in either `DES_HW` or `DES_SW`. If `DES_HW` is specified, and there is no hardware, then the encryption is performed in software and the routine returns `DESERR_NOHWDEVICE`. For `cbc_crypt`, the parameter *ivec* is the the 8-byte initialization vector for the chaining. It is updated to the next initialization vector upon return.

SEE ALSO

`des(1)`, `crypt(3)`

DIAGNOSTICS

<code>DESERR_NONE</code>	No error.
<code>DESERR_NOHWDEVICE</code>	Encryption succeeded, but done in software instead of the requested hardware.
<code>DESERR_HWERR</code>	An error occurred in the hardware or driver.
<code>DESERR_BADPARAM</code>	Bad parameter to routine.

Given a result status *stat*, the macro `DES_FAILED(stat)` is false only for the first two statuses.

RESTRICTIONS

These routines are not available for export outside the U.S.

NAME

directory, opendir, readdir, telldir, seekdir, rewinddir, closedir – directory operations

SYNOPSIS

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

DIR *opendir(filename)
char *filename;

struct dirent
*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 is not a directory, 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.

EXAMPLES

Sample code which searches a **directory** for entry “name” is:

```
dirp = opendir(".");
for (dp = readdir(dirp); dp != NULL; dp = readdir(dirp))
    if (!strcmp(dp->d_name, name)) {
        closedir (dirp);
        return FOUND;
    }
closedir (dirp);
return NOT_FOUND;
```

NOTES

The `directory()` library routines now use a new include file, `<dirent.h>`. This replaces the file, `<sys/dir.h>`, used in previous releases. Furthermore, with the use of this new file, the `readdir()` routine returns directory entries whose structure is named `struct dirent` rather than `struct direct` as before. The file, `<sys/dir.h>`, is retained in the current SunOS release for purposes of backwards source code compatibility; programs which use the `directory()` library and the file, `<sys/dir.h>`, will continue to compile and run without source code modifications. However, existing programs should convert to the use of the new include file, `<dirent.h>`, as `<sys/dir.h>` will be removed in a future major release.

SEE ALSO

`close(2)`, `lseek(2)`, `open(2V)`, `read(2V)`, `getwd(3)`, `malloc(3)`, `dir(5)`

NAME

drand48, **erand48**, **lrand48**, **nrand48**, **mrnd48**, **jrand48**, **srnd48**, **seed48**, **lcong48** – generate uniformly distributed pseudo-random numbers

SYNOPSIS

```

double drand48()
double erand48(xsubi)
unsigned short xsubi[3];
long lrand48()
long nrand48(xsubi)
unsigned short xsubi[3];
long mrnd48()
long jrand48(xsubi)
unsigned short xsubi[3];
void srnd48(seedval)
long seedval;
unsigned short *seed48(seed16v)
unsigned short seed16v[3];
void lcong48(param)
unsigned short param[7];

```

DESCRIPTION

This family of functions generates pseudo-random numbers using the well-known linear congruential algorithm and 48-bit integer arithmetic.

Functions **drand48()** and **drand48()** return non-negative double-precision floating-point values uniformly distributed over the interval (0.0, 1.0).

Functions **drand48()** and **drand48()** return non-negative long integers uniformly distributed over the interval (0, 2^{31}).

Functions **drand48()** and **drand48()** return signed long integers uniformly distributed over the interval (-2^{31} , 2^{31}).

Functions **drand48**, **seed48**, and **lcong48()** are initialization entry points, one of which should be invoked before either **drand48**, **drand48**, or **drand48()** is called. (Although it is not recommended practice, constant default initializer values will be supplied automatically if **drand48**, **drand48**, or **drand48()** is called without a prior call to an initialization entry point.) Functions **drand48**, **drand48**, and **drand48()** do not require an initialization entry point to be called first.

All the routines work by generating a sequence of 48-bit integer values, X_i , according to the linear congruential formula

$$X_{n+1} = (aX_n + c)_{\text{mod } m} \quad n \geq 0.$$

The parameter $m = 2^{48}$; hence 48-bit integer arithmetic is performed. Unless **lcong48()** has been invoked, the multiplier value a and the addend value c are given by

$$a = 5DEECE66D_{16} = 273673163155_8$$

$$c = B_{16} = 13_8.$$

The value returned by any of the functions **drand48**, **drand48**, **drand48**, **drand48**, **drand48**, or **drand48()** is computed by first generating the next 48-bit X_i in the sequence. Then the appropriate number of bits, according to the type of data item to be returned, are copied from the high-order (leftmost) bits of X_i and transformed into the returned value.

The functions **drand48**, **drand48**, and **drand48()** store the last 48-bit X_i generated in an internal buffer; that is why they must be initialized prior to being invoked. The functions **drand48**, **drand48**, and **drand48()** require the calling program to provide storage for the successive X_i values in the array specified as an argument when the functions are invoked. That is why these routines do not have to be initialized; the calling program merely has to place the desired initial value of X_i into the array and pass it as an argument. By using different arguments, functions **drand48**, **drand48**, and **drand48()** allow separate modules of a large program to generate several *independent* streams of pseudo-random numbers, that is, the sequence of numbers in each stream will *not* depend upon how many times the routines have been called to generate numbers for the other streams.

The initializer function **drand48()** sets the high-order 32 bits of X_i to the 32 bits contained in its argument. The low-order 16 bits of X_i are set to the arbitrary value $330E_{16}$.

The initializer function **seed48()** sets the value of X_i to the 48-bit value specified in the argument array. In addition, the previous value of X_i is copied into a 48-bit internal buffer, used only by **seed48**, and a pointer to this buffer is the value returned by **seed48**. This returned pointer, which can just be ignored if not needed, is useful if a program is to be restarted from a given point at some future time — use the pointer to get at and store the last X_i value, and then use this value to reinitialize via **seed48()** when the program is restarted.

The initialization function **lcong48()** allows the user to specify the initial X_i , the multiplier value a , and the addend value c . Argument array elements *param*[0-2] specify X_i , *param*[3-5] specify the multiplier a , and *param*[6] specifies the 16-bit addend c . After **lcong48()** has been called, a subsequent call to either **drand48()** or **seed48()** will restore the “standard” multiplier and addend values, a and c , specified on the previous page.

SEE ALSO

rand(3C)

NAME

`econvert`, `fconvert`, `gconvert`, `seconvert`, `sfconvert`, `sgconvert`, `ecvt`, `fcvt`, `gcvt` – output conversion

SYNOPSIS

```
#include <floatingpoint.h>
```

```
char *econvert(value, ndigit, decpt, sign, buf)
double value;
int ndigit, *decpt, *sign;
char *buf;
```

```
char *fconvert(value, ndigit, decpt, sign, buf)
double value;
int ndigit, *decpt, *sign;
char *buf;
```

```
char *gconvert(value, ndigit, trailing, buf)
double value;
int ndigit;
int trailing;
char *buf;
```

```
char *seconvert(value, ndigit, decpt, sign, buf)
single *value;
int ndigit, *decpt, *sign;
char *buf;
```

```
char *sfconvert(value, ndigit, decpt, sign, buf)
single *value;
int ndigit, *decpt, *sign;
char *buf;
```

```
char *sgconvert(value, ndigit, trailing, buf)
single *value;
int ndigit;
int trailing;
char *buf;
```

```
char *ecvt(value, ndigit, decpt, sign)
double value;
int ndigit, *decpt, *sign;
```

```
char *fcvt(value, ndigit, decpt, sign)
double value;
int ndigit, *decpt, *sign;
```

```
char *gcvt(value, ndigit, buf)
double value;
int ndigit;
char *buf;
```

DESCRIPTION

`econvert()` converts the *value* to a NULL-terminated string of *ndigit* ASCII digits in *buf* and returns a pointer to *buf*. *buf* should contain at least *ndigit+1* characters. The position of the decimal point relative to the beginning of the string is stored indirectly through *decpt*. Thus *buf* == "314" and **decpt* == 1 corresponds to the numerical value 3.14, while *buf* == "314" and **decpt* == -1 corresponds to the numerical value .0314. If the sign of the result is negative, the word pointed to by *sign* is nonzero; otherwise it is zero. The least significant digit is rounded.

fconvert is identical to **econvert**, except that the correct digit has been rounded for Fortran F-format output with *ndigit* digits to the right of the decimal point. *ndigit* can be negative to indicate rounding to the left of the decimal point. The return value is a pointer to *buf*. *buf* should contain at least $310 + \max(0, ndigit)$ characters to accommodate any double-precision *value*.

gconvert() converts the *value* to a NULL-terminated ASCII string in *buf* and returns a pointer to *buf*. It produces *ndigit* significant digits in fixed-decimal format, like Fortran F, if possible, and otherwise in floating-decimal format, like Fortran E; in either case *buf* is ready for printing, with sign and exponent. The result corresponds to that obtained by

```
(void) sprintf(buf, "%gw.n", value);
```

If *trailing* = 0, trailing zeros and a trailing point are suppressed. If *trailing* != 0, trailing zeros and a trailing point are retained.

seconvert, **sfconvert**, and **sgconvert()** are single-precision versions of these functions, and are more efficient than the corresponding double-precision versions. A pointer rather than the value itself is passed to avoid C's usual conversion of single-precision arguments to double.

ecvt() and **fcvt()** are obsolete versions of **econvert()** and **fconvert()** that create a string in a static data area, overwritten by each call, and return values that point to that static data. These functions are therefore not reentrant.

gcvt() is an obsolete version of **gconvert()** that always suppresses trailing zeros and point.

IEEE Infinities and NaNs are treated similarly by these functions. "NaN" is returned for NaN, and "Inf" or "Infinity" for Infinity. The longer form is produced when *ndigit* >= 8.

SEE ALSO

printf(3S)

NAME

end, etext, edata – last locations in program

SYNOPSIS

```
extern end;  
extern etext;  
extern edata;
```

DESCRIPTION

These names refer neither to routines nor to locations with interesting contents. The address of *etext* is the first address above the program text, *edata* above the initialized data region, and *end()* above the uninitialized data region.

When execution begins, the program break (the first location beyond the data) coincides with *end*, but it is reset by the routines *brk(2)*, *malloc(3)*, standard input/output (*stdio(3S)* and *stdio(3V)*), the profile (*-p*) option of *cc(1V)*, and so on. Thus, the current value of the program break should be determined by *sbrk(0)* (see *brk(2)*).

SEE ALSO

cc(1V), *brk(2)*, *malloc(3)*, *stdio(3S)*, *stdio(3V)*

NAME

`ethers`, `ether_ntoa`, `ether_aton`, `ether_ntohost`, `ether_hostton`, `ether_line` – Ethernet address mapping operations

SYNOPSIS

```
#include <sys/types.h>
#include <sys/socket.h>
#include <net/if.h>
#include <netinet/in.h>
#include <netinet/if_ether.h>

char *
ether_ntoa(e)
    struct ether_addr *e;

struct ether_addr *
ether_aton(s)
    char *s;

ether_ntohost(hostname, e)
    char *hostname;
    struct ether_addr *e;

ether_hostton(hostname, e)
    char *hostname;
    struct ether_addr *e;

ether_line(l, e, hostname)
    char *l;
    struct ether_addr *e;
    char *hostname;
```

DESCRIPTION

These routines are useful for mapping 48 bit Ethernet numbers to their ASCII representations or their corresponding host names, and vice versa.

The function `ether_ntoa()` converts a 48 bit Ethernet number pointed to by `e` to its standard ASCII representation; it returns a pointer to the ASCII string. The representation is of the form: `xx:xx:xx:xx:xx:xx` where `x` is a hexadecimal number between 0 and ff. The function `ether_aton()` converts an ASCII string in the standard representation back to a 48 bit Ethernet number; the function returns NULL if the string cannot be scanned successfully.

The function `ether_ntohost()` maps an Ethernet number (pointed to by `e`) to its associated hostname. The string pointed to by `hostname` must be long enough to hold the hostname and a NULL character. The function returns zero upon success and non-zero upon failure. Inversely, the function `ether_hostton()` maps a hostname string to its corresponding Ethernet number; the function modifies the Ethernet number pointed to by `e`. The function also returns zero upon success and non-zero upon failure.

The function `ether_line()` scans a line (pointed to by `l`) and sets the hostname and the Ethernet number (pointed to by `e`). The string pointed to by `hostname` must be long enough to hold the hostname and a NULL character. The function returns zero upon success and non-zero upon failure. The format of the scanned line is described by `ethers(5)`.

FILES

`/etc/ethers` (or the Yellow Pages maps `ethers.byaddr` and `ethers.byname`)

SEE ALSO

`ethers(5)`

NAME

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

SYNOPSIS

```

execl(name, arg0, arg1, ..., argn, (char *)0)
char *name, *arg0, *arg1, ..., *argn;

execv(name, argv)
char *name, *argv[ ];

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

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

execvp(name, argv)
char *name, *argv[ ];

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 *filename* 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 `(char *)0` argument must end the argument list. The cast to type `char *` insures portability.

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.

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(5V)` 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

`/usr/bin/sh` shell, invoked if command file found by `execlp()` or `execvp()`

SEE ALSO

`csh(1)`, `sh(1)`, `execve(2)`, `fork(2)`, `a.out(5)`, `environ(5V)`

Programming Utilities and Libraries

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.

NAME

exit – terminate a process after performing cleanup

SYNOPSIS

```
exit(status)  
int status;
```

DESCRIPTION

exit() terminates a process by calling **exit(2)** after calling any termination handlers named by calls to **on_exit**. Normally, this is just the Standard I/O library function **_cleanup**. **exit()** never returns.

SEE ALSO

exit(2), **intro(3S)**, **on_exit(3)**

NAME

exportent, getexportent, setexportent, addexportent, remexportent, endexportent, getexportopt – get exported file system information

SYNOPSIS

```
#include <stdio.h>
#include <exportent.h>
FILE *setexportent()
struct exportent *getexportent(file)
    FILE *file;
int addexportent(file, dirname, options)
    FILE *file;
    char *dirname;
    char *options;
int remexportent(file, dirname)
    FILE *file;
    char *dirname;
char *getexportopt(xent, opt)
    struct exportent *xent;
    char *opt;
void endexportent(file)
    FILE *file;
```

DESCRIPTION

These routines access the exported filesystem information in */etc/xtab*.

setexportent() opens the export information file and returns a file pointer to use with *getexportent*, *addexportent*, *remexportent*, and *endexportent*. *getexportent()* reads the next line from *filep* and returns a pointer to an object with the following structure containing the broken-out fields of a line in the file, */etc/xtab*. The fields have meanings described in *exports(5)*.

```
#define ACCESS_OPT "access" /* machines that can mount fs */
#define ROOT_OPT "root" /* machines with root access of fs */
#define RO_OPT "ro" /* export read-only */
#define ANON_OPT "anon" /* uid for anonymous requests */
#define SECURE_OPT "secure" /* require secure NFS for access */
#define WINDOW_OPT "window" /* expiration window for credential */
struct exportent {
    char *xent_dirname; /* directory (or file) to export */
    char *xent_options; /* options, as above */
};
```

addexportent() adds the *exportent()* to the end of the open file *filep*. It returns 0 if successful and -1 on failure. *remexportent()* removes the indicated entry from the list. It also returns 0 on success and -1 on failure. *getexportopt()* scans the *xent_options* field of the *exportent()* structure for a substring that matches *opt*. It returns the string value of *opt*, or NULL if the option is not found.

endexportent() closes the file.

FILES

/etc/exports
/etc/xtab

SEE ALSO

exports(5), *exportfs(8)*

DIAGNOSTICS

NULL pointer (0) returned on EOF or error.

BUGS

The returned `exportent()` structure points to static information that is overwritten in each call.

NAME

fclose, fflush – close or flush a stream

SYNOPSIS

#include <stdio.h>

fclose(stream)

FILE *stream;

fflush(stream)

FILE *stream;

DESCRIPTION

fclose() writes out any buffered data for the named stream, and closes the named stream. Buffers allocated by the standard input/output system are freed.

fclose() is performed automatically for all open files upon calling **exit(3)**.

fflush() writes out any buffered data for the named output stream. The named stream remains open.

SEE ALSO

close(2), exit(3), fopen(3S), setbuf(3S)

DIAGNOSTICS

These functions return 0 for success, and EOF if any error (such as trying to write to a file that has not been opened for writing) was detected.

NAME

fdate – return date and time in an ASCII string

SYNOPSIS

subroutine fdate(string)

character*24 string

character*24 function fdate()

DESCRIPTION

fdate() returns the current date and time as a 24 character string in the format described under **ctime(3)**. Neither **NEWLINE** nor **NULL** will be included.

fdate() can be called either as a function or as a subroutine. If called as a function, the calling routine must define its type and length. For example:

character*24 fdate

write(*,*)fdate()

SEE ALSO

ctime(3), time(3F)

NAME

ferror, feof, clearerr, fileno – stream status inquiries

SYNOPSIS

#include <stdio.h>

ferror(stream)

FILE *stream;

feof(stream)

FILE *stream;

clearerr(stream)

FILE *stream;

fileno(stream)

FILE *stream;

DESCRIPTION

ferror() returns non-zero when an error has occurred reading from or writing to the named stream, otherwise zero. Unless cleared by **clearerr**, the error indication lasts until the stream is closed.

feof() returns non-zero when EOF has previously been detected reading the named input stream, otherwise zero. Unless cleared by **clearerr**, the EOF indication lasts until the stream is closed.

clearerr() resets the error indication and EOF indication to zero on the named stream.

fileno() returns the integer file descriptor associated with the stream; see **open(2V)**.

NOTE

All these functions are implemented as macros; they cannot be redeclared.

SEE ALSO

open(2V), fopen(3S)

NAME

`single_to_decimal`, `double_to_decimal`, `extended_to_decimal` – convert floating-point value to decimal record

SYNOPSIS

```
#include <floatingpoint.h>

void single_to_decimal(px, pm, pd, ps)
single *px ;
decimal_mode *pm;
decimal_record *pd;
fp_exception_field_type *ps;

void double_to_decimal(px, pm, pd, ps)
double *px ;
decimal_mode *pm;
decimal_record *pd;
fp_exception_field_type *ps;

void extended_to_decimal(px, pm, pd, ps)
extended *px ;
decimal_mode *pm;
decimal_record *pd;
fp_exception_field_type *ps;
```

DESCRIPTION

The `floating_to_decimal()` functions convert the floating-point value at `*px` into a decimal record at `*pd`, observing the modes specified in `*pm` and setting exceptions in `*ps`. If there are no IEEE exceptions, `*ps` will be zero.

If `*px` is zero, infinity, or NaN, then only `pd->sign` and `pd->fpclass` are set. Otherwise `pd->exponent` and `pd->ds` are also set so that

$$(pd->sign)*(pd->ds)*10^{(pd->exponent)}$$

is a correctly rounded approximation to `*px`. `pd->ds` has at least one and no more than `DECIMAL_STRING_LENGTH-1` significant digits because one character is used to terminate the string with a NULL.

`pd->ds` is correctly rounded according to the IEEE rounding modes in `pm->rd`. `*ps` has `fp_inexact` set if the result was inexact, and has `fp_overflow` set if the string result does not fit in `pd->ds` because of the limitation `DECIMAL_STRING_LENGTH`.

If `pm->df == floating_form`, then `pd->ds` always contains `pm->ndigits` significant digits. Thus if `*px == 12.34` and `pm->ndigits == 8`, then `pd->ds` will contain 12340000 and `pd->exponent` will contain -6.

If `pm->df == fixed_form` and `pm->ndigits >= 0`, then `pd->ds` always contains `pm->ndigits` after the point and as many digits as necessary before the point. Since the latter is not known in advance, the total number of digits required is returned in `pd->ndigits`; if that number `>= DECIMAL_STRING_LENGTH`, then `ds` is undefined. `pd->exponent` always gets `-pm->ndigits`. Thus if `*px == 12.34` and `pm->ndigits == 1`, then `pd->ds` gets 123, `pd->exponent` gets -1, and `pd->ndigits` gets 3.

If `pm->df == fixed_form` and `pm->ndigits < 0`, then `pm->ds` always contains `-pm->ndigits` trailing zeros; in other words, rounding occurs `-pm->ndigits` to the left of the decimal point, but the digits rounded away are retained as zeros. The total number of digits required is in `pd->ndigits`. `pd->exponent` always gets 0. Thus if `*px == 12.34` and `pm->ndigits == -1`, then `pd->ds` gets 10, `pd->exponent` gets 0, and `pd->ndigits` gets 2.

`pd->more` is not used.

econvert(3), fconvert, gconvert, printf(3S), and sprintf, all use double_to_decimal.

SEE ALSO

econvert(3), printf(3S)

NAME

floatingpoint – IEEE floating point definitions

SYNOPSIS

```
#include <sys/ieeefp.h>
#include <floatingpoint.h>
```

DESCRIPTION

This file defines constants, types, variables, and functions used to implement standard floating point according to ANSI/IEEE Std 754-1985. The variables and functions are implemented in `libc.a`. The included file `<sys/ieeefp.h>` defines certain types of interest to the kernel.

IEEE Rounding Modes:

fp_direction_type The type of the IEEE rounding direction mode. Note: the order of enumeration varies according to hardware.

fp_direction The IEEE rounding direction mode currently in force. This is a global variable that is intended to reflect the hardware state, so it should only be written indirectly through a function like `ieeefp_set_direction(set, direction, ...)` that also sets the hardware state.

fp_precision_type The type of the IEEE rounding precision mode, which only applies on systems that support extended precision such as Sun-3 systems with 68881's.

fp_precision The IEEE rounding precision mode currently in force. This is a global variable that is intended to reflect the hardware state on systems with extended precision, so it should only be written indirectly through a function like `ieeefp_set_precision(set, "precision", ...)`.

SIGFPE handling:

sigfpe_code_type The type of a SIGFPE code.

sigfpe_handler_type The type of a user-definable SIGFPE exception handler called to handle a particular SIGFPE code.

SIGFPE_DEFAULT A macro indicating the default SIGFPE exception handling, namely to perform the exception handling specified by calls to `ieeefp_handler(3M)`, if any, and otherwise to dump core using `abort(3)`.

SIGFPE_IGNORE A macro indicating an alternate SIGFPE exception handling, namely to ignore and continue execution.

SIGFPE_ABORT A macro indicating an alternate SIGFPE exception handling, namely to abort with a core dump.

IEEE Exception Handling:

N_IEEE_EXCEPTION The number of distinct IEEE floating-point exceptions.

fp_exception_type The type of the `N_IEEE_EXCEPTION` exceptions. Each exception is given a bit number.

fp_exception_field_type The type intended to hold at least `N_IEEE_EXCEPTION` bits corresponding to the IEEE exceptions numbered by `fp_exception_type`. Thus `fp_inexact` corresponds to the least significant bit and `fp_invalid` to the fifth least significant bit. Note: some operations may set more than one exception.

fp_accrued_exceptions

The IEEE exceptions between the time this global variable was last cleared, and the last time a function like `ieeefp_get_exception(set, "exception", ...)` was called to update the variable by obtaining the hardware state.

ieee_handlers An array of user-specifiable signal handlers for use by the standard SIGFPE handler for IEEE arithmetic-related SIGFPE codes. Since IEEE trapping modes correspond to hardware modes, elements of this array should only be modified with a function like **ieee_handler(3M)** that performs the appropriate hardware mode update. If no **sigfpe_handler** has been declared for a particular IEEE-related SIGFPE code, then the related **ieee_handlers** will be invoked.

IEEE Formats and Classification:

single;extended Definitions of IEEE formats.

fp_class_type An enumeration of the various classes of IEEE values and symbols.

IEEE Base Conversion:

The functions described under **floating_to_decimal(3)** and **decimal_to_floating(3)** not only satisfy the IEEE Standard, but also the stricter requirements of correct rounding for all arguments.

DECIMAL_STRING_LENGTH

The length of a **decimal_string**.

decimal_string The digit buffer in a **decimal_record**.

decimal_record The canonical form for representing an unpacked decimal floating-point number.

decimal_form The type used to specify fixed or floating binary to decimal conversion.

decimal_mode A struct that contains specifications for conversion between binary and decimal.

decimal_string_form An enumeration of possible valid character strings representing floating-point numbers, infinities, or NaNs.

FILES

/usr/include/sys/ieeefp.h
/usr/include/floatingpoint.h
/usr/lib/libc.a

SEE ALSO

abort(3), **decimal_to_floating(3)**, **econvert(3)**, **floating_to_decimal(3)**, **ieee_flags(3M)**, **ieee_handler(3M)**, **sigfpe(3)**, **string_to_decimal(3)**, **strtod(3)**

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. If the open succeeds, **fopen()** returns a pointer to be used to identify the stream in subsequent operations.

filename points to a character string that contains the name of the file to be opened.

type is a character string having one of the following values:

r	open for reading
w	truncate or create for writing
a	append: open for writing at end of file, or create for writing
r+	open for update (reading and writing)
w+	truncate or create for update
a+	append; open or create for update at EOF

freopen() opens the file named by *filename* and associates the stream pointed to by *stream* with it. The *type* argument is used just as in **fopen**. The original stream is closed, regardless of whether the open ultimately succeeds. If the open succeeds, **freopen()** returns the original value of *stream*.

freopen() is typically used to attach the preopened streams associated with **stdin**, **stdout**, and **stderr** to other files.

fdopen() associates a stream with the file descriptor *fildes*. File descriptors are obtained from calls like **open**, **dup**, **creat**, or **pipe(2)**, which open files but do not return streams. Streams are necessary input for many of the Section 3S library routines. The *type* of the stream must agree with the mode of the open file.

When a file is opened for update, both input and output may be done on the resulting stream. However, output may not be directly followed by input without an intervening **fseek()** or **rewind**, and input may not be directly followed by output without an intervening **fseek**, **rewind**, or an input operation which encounters end-of-file.

SEE ALSO

open(2V), **pipe(2)**, **fclose(3S)**, **fopen(3V)**, **fseek(3S)**

DIAGNOSTICS

fopen, **freopen**, and **fdopen()** return a NULL pointer on failure.

BUGS

In order to support the same number of open files that the system does, **fopen()** must allocate additional memory for data structures using **calloc()** after 64 files have been opened. This confuses some programs which use their own memory allocators.

NAME

fread, fwrite – buffered binary input/output

SYNOPSIS

#include <stdio.h>

fread(ptr, size, nitems, stream)

FILE *stream;

fwrite(ptr, size, nitems, stream)

FILE *stream;

DESCRIPTION

fread() reads, into a block pointed to by *ptr*, *nitems* of data from the named input stream, where an item of data is a sequence of bytes (not necessarily terminated by a NULL byte) of length *size*. It returns the number of items actually read. **fread()** stops appending bytes if an EOF or error condition is encountered while reading stream, or if *nitems* items have been read. **fread()** leaves the file pointer in stream, if defined, pointing to the byte following the last byte read if there is one. **fread()** does not change the contents of stream.

fwrite() appends at most *nitems* of data from the block pointed to by *ptr* to the named output stream. It returns the number of items actually written. **fwrite()** stops appending when it has appended *nitems* items of data or if an error condition is encountered on stream. **fwrite()** does not change the contents of the block pointed to by *ptr*.

The argument *size* is typically `sizeof(*ptr)` where the pseudo-function `sizeof` specifies the length of an item pointed to by *ptr*. If *ptr* points to a data type other than *char* it should be cast into a pointer to *char*.

If *size* or *nitems* is non-positive, no characters are read or written and 0 is returned by both **fread()** and **fwrite()**.

SEE ALSO

read(2V), write(2V), fopen(3S), getc(3S), gets(3S), putc(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)

FILE *stream;

DESCRIPTION

fseek() sets the position of the next input or output operation on the stream. The new position is at the signed distance *offset* bytes from the beginning, the current position, or the end of the file, according as *ptrname* has the value 0, 1, or 2.

rewind(stream) is equivalent to **fseek(stream, 0L, 0)**, except that no value is returned.

fseek() and **rewind()** undo any effects of **ungetc(3S)**.

After **fseek()** or **rewind()**, the next operation on a file opened for update may be either input or output.

ftell() returns the offset of the current byte relative to the beginning of the file associated with the named stream.

SEE ALSO

lseek(2), fopen(3S), popen(3S), ungetc(3S)

DIAGNOSTICS

fseek() returns **-1** for improper seeks, otherwise zero. An improper seek can be, for example, an **fseek()** done on a file that has not been opened using **fopen**; in particular, **fseek()** may not be used on a terminal, or on a file opened using **popen(3S)**.

WARNING

Although on the UNIX system an offset returned by **ftell()** is measured in bytes, and it is permissible to seek to positions relative to that offset, portability to a (non-UNIX) system requires that an offset be used by **fseek()** directly. Arithmetic may not meaningfully be performed on such an offset, which is not necessarily measured in bytes.

NAME

ftok – standard interprocess communication package

SYNOPSIS

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

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

```
key_t ftok(path, id)
```

```
char *path;
```

```
char id;
```

DESCRIPTION

All interprocess communication facilities require the user to supply a key to be used by the `msgget(2)`, `semget(2)`, and `shmget(2)` system calls to obtain interprocess communication identifiers. One suggested method for forming a key is to use the `ftok()` subroutine described below. Another way to compose keys is to include the project ID in the most significant byte and to use the remaining portion as a sequence number. There are many other ways to form keys, but it is necessary for each system to define standards for forming them. If some standard is not adhered to, it will be possible for unrelated processes to unintentionally interfere with each other's operation. Therefore, it is strongly suggested that the most significant byte of a key in some sense refer to a project so that keys do not conflict across a given system.

`ftok()` returns a key based on *path* and ID that is usable in subsequent `msgget`, `semget`, and `shmget()` system calls. *path* must be the path name of an existing file that is accessible to the process. ID is a character which uniquely identifies a project. Note: `ftok()` will return the same key for linked files when called with the same ID and that it will return different keys when called with the same file name but different IDs.

SEE ALSO

`intro(2)`, `msgget(2)`, `semget(2)`, `shmget(2)`

DIAGNOSTICS

`ftok()` returns `(key_t) -1` if *path* does not exist or if it is not accessible to the process.

WARNING

If the file whose *path* is passed to `ftok()` is removed when keys still refer to the file, future calls to `ftok()` with the same *path* and ID will return an error. If the same file is recreated, then `ftok()` is likely to return a different key than it did the original time it was called.

NAME

ftw – walk a file tree

SYNOPSIS

```
#include <ftw.h>

int ftw(path, fn, depth)
char *path;
int (*fn)();
int depth;
```

DESCRIPTION

ftw() recursively descends the directory hierarchy rooted in *path*. For each object in the hierarchy, **ftw()** calls *fn*, passing it a pointer to a NULL-terminated character string containing the name of the object, a pointer to a **stat()** structure (see **stat(2)**) containing information about the object, and an integer. Possible values of the integer, defined in the **<ftw.h>** header file, are **FTW_F** for a file, **FTW_D** for a directory, **FTW_DNR** for a directory that cannot be read, and **FTW_NS** for an object for which **stat()** could not successfully be executed. If the integer is **FTW_DNR**, descendants of that directory will not be processed. If the integer is **FTW_NS**, the **stat()** structure will contain garbage. An example of an object that would cause **FTW_NS** to be passed to *fn* would be a file in a directory with read but without execute (search) permission.

ftw() visits a directory before visiting any of its descendants.

The tree traversal continues until the tree is exhausted, an invocation of *fn* returns a nonzero value, or some error is detected within **ftw()** (such as an I/O error). If the tree is exhausted, **ftw()** returns zero. If *fn* returns a nonzero value, **ftw()** stops its tree traversal and returns whatever value was returned by *fn*. If **ftw()** detects an error, it returns **-1**, and sets the error type in **errno**.

ftw() uses one file descriptor for each level in the tree. The *depth* argument limits the number of file descriptors so used. If *depth* is zero or negative, the effect is the same as if it were 1. *depth* must not be greater than the number of file descriptors currently available for use. **ftw()** will run more quickly if *depth* is at least as large as the number of levels in the tree.

SEE ALSO

stat(2), **malloc(3)**

BUGS

Because **ftw()** is recursive, it is possible for it to terminate with a memory fault when applied to very deep file structures.

It could be made to run faster and use less storage on deep structures at the cost of considerable complexity.

ftw() uses **malloc(3)** to allocate dynamic storage during its operation. If **ftw()** is forcibly terminated, such as by **longjmp()** being executed by *fn* or an interrupt routine, **ftw()** will not have a chance to free that storage, so it will remain permanently allocated. A safe way to handle interrupts is to store the fact that an interrupt has occurred, and arrange to have *fn* return a nonzero value at its next invocation.

NAME

getacinfo, getacdir, getacflg, getacmin, setac, endac – get audit control file information

SYNOPSIS

```
int getacdir(dir, len)  
char *dir;  
int len;  
  
int getacmin(min_val)  
int *min_val;  
  
int getacflg(auditstring, len)  
char *auditstring;  
int len;  
  
void setac()  
  
void endac()
```

DESCRIPTION

When first called, **getacdir()** provides information about the first audit directory in the **audit_control** file; thereafter, it returns the next directory in the file. Successive calls list all the directories listed in **audit_control(5)**. The parameter *len* specifies the length of the buffer *dir*. On return, *dir* points to the directory entry.

getacmin() reads the minimum value from the **audit_control** file and returns the value in *min_val*. The minimum value specifies how full the file system to which the audit files are being written can get before the script **audit_warn** is invoked.

getacflg() reads the system audit value from the **audit_control** file and returns the value in *auditstring*. The parameter *len* specifies the length of the buffer *auditstring*.

Calling **setac** rewinds the **audit_control** file to allow repeated searches.

Calling **endac** closes the **audit_control** file when processing is complete.

RETURN VALUE

Upon successful completion of all **getac...** functions, a zero is returned. **getacmin()** and **getacflg()** return a 1 on EOF. Only **getacdir()** returns a 2 if the directory search had to start from the beginning because another **getac...** function was called between calls to **getacdir**.

ERRORS

Upon unsuccessful completion, a value of -2 is returned and **errno** is set to indicate the error. -3 is returned if the directory entry format in the **audit_control** file is incorrect. If the input buffer is too short to accommodate the record, **getacdir()** and **getacflg()** return an error code of -3. Only **getacdir()** returns a -1 on end of file.

SEE ALSO

audit_control(5)

NAME

getauditflagsbin, getauditflagschar – convert audit flag specifications

SYNOPSIS

```
#include <sys/label.h>
#include <sys/audit.h>
#include <sys/auevents.h>

int getauditflagsbin(auditstring, masks)
char *auditstring;
audit_state_t *masks;

int getauditflagschar(auditstring, masks, verbose)
char *auditstring;
audit_state_t *masks;
int verbose;
```

DESCRIPTION

`getauditflagsbin()` converts the character representation of audit values pointed to by *auditstring* into `audit_state_t` fields pointed to by *masks*. These fields indicate which events are to be audited when they succeed and which are to be audited when they fail. The character string syntax is described in `audit_control(5)`.

`getauditflagschar()` converts the `audit_state_t` fields pointed to by *masks* into a string pointed to by *auditstring*. If *verbose* is zero, the short (2-character) flag names are used. If *verbose* is non-zero, the long flag names are used. *auditstring* should be large enough to contain the ASCII representation of the events.

auditstring contains a series of event names, each one identifying a single audit class, separated by commas. The `audit_state_t` fields pointed to by *masks* correspond to binary values defined in *audit.h*.

DIAGNOSTICS

-1 is returned on error and 0 on success.

SEE ALSO

`audit.log(5)`, `audit_control(5)`

BUGS

This is not a very extensible interface.

NAME

getc, getchar, fgetc, getw – get character or integer 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 (that is, byte) from the named input stream, as an integer. It also moves the file pointer, if defined, ahead one character in stream. **getchar()** is defined as **getc(stdin)**. **getc** and **getchar** are macros.

fgetc() behaves like **getc**, but is a function rather than a macro. **fgetc()** runs more slowly than **getc**, but it takes less space per invocation and its name can be passed as an argument to a function.

getw() returns the next C **int** (*word*) from the named input stream. **getw()** increments the associated file pointer, if defined, to point to the next word. The size of a word is the size of an integer and varies from machine to machine. **getw()** assumes no special alignment in the file.

SEE ALSO

ferror(3S), fopen(3S), fread(3S), gets(3S), putc(3S), scanf(3S), ungetc(3S)

DIAGNOSTICS

These functions return the integer constant EOF at EOF or upon an error. The EOF condition is remembered, even on a terminal, and all subsequent attempts to read will return EOF until the condition is cleared with **clearerr** (see **ferror(3S)**) Because EOF is a valid integer, **ferror(3S)** should be used to detect **getw()** errors.

WARNING

If the integer value returned by **getc**, **getchar**, or **fgetc** is stored into a character variable and then compared against the integer constant EOF, the comparison may never succeed, because sign-extension of a character on widening to integer is machine-dependent.

BUGS

Because it is implemented as a macro, **getc()** treats a stream argument with side effects incorrectly. In particular, **getc(*f++)** does not work sensibly. **fgetc()** should be used instead.

Because of possible differences in word length and byte ordering, files written using **putw()** are machine-dependent, and may not be readable using **getw()** on a different processor.

NAME

`getcwd` – get pathname of current working directory

SYNOPSIS

```
char *getcwd (buf, size)
char *buf;
int size;
```

DESCRIPTION

`getcwd()` returns a pointer to the current directory pathname. The value of *size* must be at least two greater than the length of the pathname to be returned.

If *buf* is a NULL pointer, `getcwd()` will obtain *size* bytes of space using `malloc(3)`. In this case, the pointer returned by `getcwd()` may be used as the argument in a subsequent call to `free`.

The function is implemented by using `popen(3S)` to pipe the output of the `pwd(1)` command into the specified string space.

EXAMPLE

```
char *cwd, *getcwd();
.
.
.
if ((cwd = getcwd((char *)NULL, 64)) == NULL) {
    perror ("pwd");
    exit (1);
}
printf ("%s\n", cwd);
```

SEE ALSO

`malloc(3)`, `popen(3S)`, `pwd(1)`

DIAGNOSTICS

Returns NULL with `errno` set if *size* is not large enough, or if an error occurs in a lower-level function.

BUGS

Since this function uses `popen()` to create a pipe to the `pwd` command, it is slower than `getwd()` and gives poorer error diagnostics. `getcwd()` is provided only for compatibility with other UNIX operating systems.

NAME

getenv – return value for environment name

SYNOPSIS

```
char *getenv(name)  
char *name;
```

DESCRIPTION

getenv() searches the environment list (see **environ(5V)**) for a string of the form *name=value*, and returns a pointer to the string *value* if such a string is present, otherwise NULL pointer.

SEE ALSO

environ(5V), **execve(2)**, **putenv(3)**

NAME

`getfauditflags` – generates the process audit state

SYNOPSIS

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

void getfauditflags(usremasks, usrdmasks, lastmasks)
audit_state_t *usremasks;
audit_state_t *usrdmasks;
audit_state_t *lastmasks;
```

DESCRIPTION

`getfauditflags` generates the process audit state from the user audit value as input to `getfauditflags` and the system audit value as specified in the `audit_control` file. `getfauditflags` obtains the system audit value by calling `getacflg`. The user audit value, pointed to by `usremasks` and `usrdmasks` is passed into `getfauditflags`.

`usremasks` points to `audit_state_t` fields which contains two values. The first value defines which events are *always* to be audited when they they succeed. The second value defines defines which events are *always* to be audited when they they fail.

`usrdmasks` also points to `audit_state_t` fields which contains two values. The first value defines which events are *never* to be audited when they they succeed. The second value defines defines which events are *never* to be audited when they they fail.

The structures pointed to by `usremasks` and `usrdmasks` may be obtained from the `passwd.adjunct` file by calling `getpwaent()` which returns a pointer to a strucure containing all `passwd.adjunct` fields for a user.

`lastmasks` points to the return `audit_state_t` structure. This structure contains two values which define the process audit state. The first value defines which events are to be audited when they succeed and the second value defines which events are to be audited when they fail.

Both `usremasks` and `usrdmasks` override the values in the system audit values.

SEE ALSO

`getauditflags(3)`, `getacinfo(3)`, `audit.log(5)`, `audit_control(5)`

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

These routines are included for compatibility with 4.2 BSD; they have been superseded by the *getmntent(3)* library routines.

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.

getfsent() 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

The return value points to static information which is overwritten in each call.

NAME

getgraent, getgranam, setgraent, endgraent, fgetgraent – get group adjunct file entry

SYNOPSIS

```
#include <grpadj.h>

struct group_adjunct *getgraent()

struct group_adjunct *getgranam(name)
char *name;

struct group_adjunct *fgetgraent(f)
FILE *f;

void setgraent()

void endgraent()
```

DESCRIPTION

getgraent() and **getgranam()** each return pointers to an object with the following structure containing the broken-out fields of a line in the group adjunct file. Each line contains a **group_adjunct** structure, defined in the **<grpadj.h>** header file.

```
struct group_adjunct {
    char *gra_name;    /* the name of the group */
    char *gra_passwd; /* the encrypted group password */
};
```

When first called, **getgraent()** returns a pointer to a **group_adjunct** structure corresponding to the first line in the file. Thereafter, it returns a pointer to the next **group_adjunct** structure in the file. So successive calls may be used to traverse the entire file.

For locating a particular group, **getgranam()** searches through the file until it finds group *filename*, then returns a pointer to that structure.

A call to **getgraent()** rewinds the group adjunct file to allow repeated searches. A call to **endgraent()** closes the group adjunct file when processing is complete.

SEE ALSO

getlogin(3), **getgrent(3)**, **getpwaent(3)**, **getpwent(3)**, **ypserv(8)**

DIAGNOSTICS

A NULL pointer is returned on end-of-file or error.

BUGS

All information is contained in a static area, so it must be copied if it is to be saved.

NAME

getgrent, getgrgid, getgrnam, setgrent, endgrent, fgetgrent – 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()

struct group *fgetgrent(f)
FILE *f;
```

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. Each line contains a “group” structure, defined in the <grp.h> header file.

```
struct group {
    char    *gr_name;
    char    *gr_passwd;
    int     gr_gid;
    char    **gr_mem;
};
```

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	A NULL-terminated array of pointers to the individual member names.

getgrent() when first called returns a pointer to the first group structure in the file; thereafter, it returns a pointer to the next group structure in the file; so, successive calls may be used to search the entire file. getgrgid() searches from the beginning of the file until a numerical group ID matching gid is found and returns a pointer to the particular structure in which it was found. getgrnam() searches from the beginning of the file until a group name matching name is found and returns a pointer to the particular structure in which it was found. If an end-of-file or an error is encountered on reading, these functions return a NULL pointer.

A call to getgrent() has the effect of rewinding the group file to allow repeated searches. endgrent() may be called to close the group file when processing is complete.

fgetgrent() returns a pointer to the next group structure in the stream f, which must refer to an open file in the same format as the group file /etc/group.

DIAGNOSTICS

A NULL pointer is returned on EOF or error.

FILES

/etc/group

SEE ALSO

getlogin(3), getpwent(3), group(5), ypserv(8)

BUGS

All information is contained in a static area, so it must be copied if it is to be saved.

Unlike the corresponding routines for passwords (see `getpwent(3)`), which always search the entire file, these routines start searching from the current file location.

WARNING

The above routines use `<stdio.h>`, which increases the size of programs, not otherwise using standard I/O, more than might be expected.

NAME

gethostent, gethostbyaddr, gethostbyname, sethostent, endhostent – get network host entry

SYNOPSIS

```
#include <sys/types.h>
#include <sys/socket.h>
#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 end-of-file is encountered. Host addresses are supplied in network order.

FILES

/etc/hosts

SEE ALSO

hosts(5), yperv(8)

DIAGNOSTICS

A NULL pointer is returned on end-of-file or error.

BUGS

All information is contained in a static area so it must be copied if it is to be saved. Only the Internet address format is currently understood.

NAME

getlogin – get login name

SYNOPSIS

char *getlogin()

DESCRIPTION

getlogin() returns a pointer to the login name as found in */etc/utmp*. It may be used in conjunction with **getpwnam** to locate the correct password file entry when the same user ID is shared by several login names.

If **getlogin()** is called within a process that is not attached to a terminal, or if there is no entry in */etc/utmp* for the process's terminal, it returns a NULL pointer. The correct procedure for determining the login name is to call *cuserid*, or to call **getlogin()** and, if it fails, to call **getpwuid(getuid())**.

FILES

/etc/utmp

SEE ALSO

cuserid(3S), **getpwent(3)**, **utmp(5)**

DIAGNOSTICS

Returns a NULL pointer if the name is not found.

BUGS

The return values point to static data whose content is overwritten by each call.

getlogin() does not work for processes running under a *pty* (for example, emacs shell buffers, or shell tools) unless the program “fakes” the login name in the */etc/utmp* file.

NAME

getmntent, setmntent, addmntent, endmntent, hasmntopt – get file system descriptor file entry

SYNOPSIS

```
#include <stdio.h>
#include <mntent.h>

FILE *setmntent(FILE, type)
char *filep;
char *type;

struct mntent *getmntent(FILE)
FILE *filep;

int addmntent(FILE, mnt)
FILE *filep;
struct mntent *mnt;

char *hasmntopt(mnt, opt)
struct mntent *mnt;
char *opt;

int endmntent(FILE)
FILE *filep;
```

DESCRIPTION

These routines replace the `getfsent()` routines for accessing the file system description file `/etc/fstab`. They are also used to access the mounted file system description file `/etc/mntab`.

`setmntent()` opens a file system description file and returns a file pointer which can then be used with `getmntent`, `addmntent`, or `endmntent`. The *type* argument is the same as in `fopen(3)`. `getmntent()` reads the next line from *filep* and returns a pointer to an object with the following structure containing the broken-out fields of a line in the filesystem description file, `<mntent.h>`. The fields have meanings described in `fstab(5)`.

```
struct mntent {
    char *mnt_fsname; /* file system name */
    char *mnt_dir; /* file system path prefix */
    char *mnt_type; /* 4.2, nfs, swap, or xx */
    char *mnt_opts; /* ro, quota, etc. */
    int mnt_freq; /* dump frequency, in days */
    int mnt_passno; /* pass number on parallel fsck */
};
```

`addmntent()` adds the `mntent` structure *mnt* to the end of the open file *filep*. Note: *filep* has to be opened for writing if this is to work. `hasmntopt()` scans the `mnt_opts` field of the `mntent` structure *mnt* for a substring that matches *opt*. It returns the address of the substring if a match is found, 0 otherwise. `endmntent()` closes the file.

FILES

`/etc/fstab`
`/etc/mntab`

SEE ALSO

`fopen(3S)`, `getfsent(3)`, `fstab(5)`

DIAGNOSTICS

NULL pointer (0) returned on EOF or error.

BUGS

The returned `mntent` structure points to static information that is overwritten in each call.

NAME

getnetent, getnetbyaddr, getnetbyname, setnetent, endnetent – get network entry

SYNOPSIS

```
#include <netdb.h>

struct netent *getnetent()

struct netent *getnetbyname(name)
char *name;

struct netent *getnetbyaddr(net, type)
long net;
int type;

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.

getnetent() 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 and type is found, or until end-of-file is encountered. Network numbers are supplied in host order.

FILES

/etc/networks

SEE ALSO

networks(5), yperv(8)

DIAGNOSTICS

A NULL pointer is returned on end-of-file 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.

NAME

getnetgrent, setnetgrent, endnetgrent, innetgr – get network group entry

SYNOPSIS

```
innetgr(netgroup, machine, user, domain)
char *netgroup, *machine, *user, *domain;

setnetgrent(netgroup)
char *netgroup

endnetgrent()

getnetgrent(machinep, userp, domainp)
char **machinep, **userp, **domainp;
```

DESCRIPTION

innetgr returns 1 or 0, depending on whether *netgroup* contains the machine, user, domain triple as a member. Any of the three strings *machine*, *user*, or *domain* can be NULL, in which case it signifies a wild card.

getnetgrent() returns the next member of a network group. After the call, *machinep* will contain a pointer to a string containing the name of the machine part of the network group member, and similarly for *userp* and *domainp*. If any of *machinep*, *userp* or *domainp* is returned as a NULL pointer, it signifies a wild card. **getnetgrent**() will use **malloc(3)** to allocate space for the name. This space is released when a **endnetgrent**() call is made. **getnetgrent**() returns 1 if it succeeded in obtaining another member of the network group, 0 if it has reached the end of the group.

getnetgrent() establishes the network group from which **getnetgrent**() will obtain members, and also restarts calls to **getnetgrent**() from the beginning of the list. If the previous **setnetgrent**() call was to a different network group, a **endnetgrent**() call is implied. **endnetgrent**() frees the space allocated during the **getnetgrent**() calls.

FILES

/etc/netgroup

NAME

`getopt`, `optarg`, `optind` – get option letter from argument vector

SYNOPSIS

```
int getopt(argc, argv, optstring)
int argc;
char **argv;
char *optstring;

extern char *optarg;
extern int optind, opterr;
```

DESCRIPTION

`getopt()` returns the next option letter in *argv* that matches a letter in *optstring*. *optstring* must contain the option letters the command using `getopt()` will recognize; if a letter is followed by a colon, the option is expected to have an argument, or group of arguments, which must be separated from it by white space.

optarg is set to point to the start of the option argument on return from `getopt`.

`getopt()` places in `optind()` the *argv* index of the next argument to be processed. `optind()` is external and is initialized to 1 before the first call to `getopt`.

When all options have been processed (that is, up to the first non-option argument), `getopt()` returns `-1`. The special option “—” may be used to delimit the end of the options; when it is encountered, `-1` will be returned, and “—” will be skipped.

DIAGNOSTICS

`getopt()` prints an error message on the standard error and returns a question mark (?) when it encounters an option letter not included in *optstring* or no option-argument after an option that expects one. This error message may be disabled by setting `opterr` to 0.

EXAMPLE

The following code fragment shows how one might process the arguments for a command that can take the mutually exclusive options `a` and `b`, and the option `o`, which requires an option argument:

```
main(argc, argv)
int argc;
char **argv;
{
    int c;
    extern char *optarg;
    extern int optind;
    .
    .
    .
    while ((c = getopt(argc, argv, "abo:")) != -1)
        switch (c) {
            case 'a':
                if (bflg)
                    errflg++;
                else
                    aflg++;
                break;
            case 'b':
                if (aflg)
                    errflg++;
                else
                    bproc ();
                break;
```



```
        case 'o':
            ofile = optarg;
            break;
        case '?':
            errflg++;
    }
    if (errflg) {
        (void)fprintf(stderr, "usage: ... ");
        exit (2);
    }
    for (; optind < argc; optind++) {
        if (access(argv[optind], 4)) {
            .
            .
            .
        }
    }
```

SEE ALSO**getopts(1)****WARNING**

Changing the value of the variable `optind`, or calling `getopt()` with different values of `argv`, may lead to unexpected results.

NAME

getpass – read a password

SYNOPSIS

```
char *getpass(prompt)
char *prompt;
```

DESCRIPTION

getpass() reads up to a NEWLINE or EOF from the file */dev/tty*, 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. An interrupt will terminate input and send an interrupt signal to the calling program before returning.

FILES

/dev/tty

SEE ALSO

crypt(3), getpass(3V)

WARNING

The above routine uses *<stdio.h>*, which increases the size of programs not otherwise using standard I/O, more than might be expected.

BUGS

The return value points to static data whose content is overwritten by each call.

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 */
    int     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.

getprotoent() 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 end-of-file is encountered.

FILES

/etc/protocols

SEE ALSO

protocols(5), ypserv(8)

DIAGNOSTICS

A NULL pointer is returned on end-of-file or error.

BUGS

All information is contained in a static area so it must be copied if it is to be saved. Only the Internet protocols are currently understood.

NAME

getpw – get name from uid

SYNOPSIS

getpw(uid, buf)
char *buf;

DESCRIPTION

Getpw is made obsolete by **getpwent(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

getpwaent, getpwanam, setpwaent, endpwaent, fgetpwaent – get password adjunct file entry

SYNOPSIS

```
#include <sys/types.h>
#include <sys/label.h>
#include <sys/audit.h>
#include <pwdadj.h>

struct passwd_adjunct *getpwaent()
struct passwd_adjunct *getpwanam(name)
char *name;

struct passwd_adjunct *fgetpwaent(f)
FILE *f;

void setpwaent()
void endpwaent()
```

DESCRIPTION

Both `getpwaent()` and `getpwanam()` return a pointer to an object with the following structure containing the broken-out fields of a line in the password adjunct file. Each line in the file contains a `passwd_adjunct` structure, declared in the `<pwdadj.h>` header file:

```
struct passwd_adjunct {
char      *pwa_name;
char      *pwa_passwd;
blabel_t  pwa_minimum;
blabel_t  pwa_maximum;
blabel_t  pwa_def;
audit_state_t pwa_au_always;
audit_state_t pwa_au_never;
int       pwa_version;
};
```

When first called, `getpwaent()` returns a pointer to a `passwd_adjunct` structure describing data from the first line in the file. Thereafter, it returns a pointer to a `passwd_adjunct` structure describing data from the next line in the file. So successive calls can be used to search the entire file.

`getpwanam()` searches from the beginning of the file until it finds a login name matching *name*, then returns a pointer to the particular structure in which it was found.

Calling `getpwaent()` rewinds the password adjunct file to allow repeated searches. Calling `endpwaent()` closes the password adjunct file when processing is complete.

DIAGNOSTICS

A NULL pointer is returned on end-of-file or error.

SEE ALSO

`getpwent(3)`, `getgrent(3)`, `passwd.adjunct(5)`, `ypserv(8)`

BUGS

All information is contained in a static area, so it must be copied if it is to be saved.

NAME

getpwent, getpwuid, getpwnam, setpwent, endpwent, setpwfile, fgetpwent – 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()

setpwfile(name)
char *name;

struct passwd *fgetpwent(f)
FILE *f;
```

DESCRIPTION

`getpwent`, `getpwuid()` and `getpwnam()` each return a pointer to an object with the following structure containing the broken-out fields of a line in the password file. Each line in the file contains a “passwd” structure, declared in the `<pwd.h>` header 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();
```

This structure is declared in `<pwd.h>` so it is not necessary to redeclare it.

The fields `pw_quota` and `pw_comment` are unused; the others have meanings described in `passwd(5)`. When first called, `getpwent()` returns a pointer to the first `passwd` structure in the file; thereafter, it returns a pointer to the next `passwd` structure in the file; so successive calls can be used to search the entire file. `getpwuid()` searches from the beginning of the file until a numerical user ID matching `uid` is found and returns a pointer to the particular structure in which it was found. `getpwnam()` searches from the beginning of the file until a login name matching `name` is found, and returns a pointer to the particular structure in which it was found. If an end-of-file or an error is encountered on reading, these functions return a NULL pointer.

A call to `getpwent()` has the effect of rewinding the password file to allow repeated searches. `endpwent()` may be called to close the password file when processing is complete.

`setpwfile()` changes the default password file to `name` thus allowing alternate password files to be used. Note: it does *not* close the previous file. If this is desired, `endpwent()` should be called prior to it.

`fgetpwent()` returns a pointer to the next `passwd` structure in the stream `f`, which matches the format of the password file `/etc/passwd`.

FILES

/etc/passwd

SEE ALSO

getgrent(3), getlogin(3), getpwent(3V), passwd(5), ypserv(8)

DIAGNOSTICS

A NULL pointer is returned on end-of-file or error.

BUGS

All information is contained in a static area, so it must be copied if it is to be saved.

NAME

getrpcent, getrpcbyname, getrpcbynumber – get RPC entry

SYNOPSIS

```
#include <netdb.h>

struct rpcent *getrpcent()
struct rpcent *getrpcbyname(name)
char *name;

struct rpcent *getrpcbynumber(number)
int number;

setrpcent (stayopen)
int stayopen

endrpcent ()
```

DESCRIPTION

getrpcent, getrpcbyname, and getrpcbynumber() each return a pointer to an object with the following structure containing the broken-out fields of a line in the rpc program number data base, */etc/rpc*.

```
struct rpcent {
    char    *r_name;        /* name of server for this rpc program */
    char    **r_aliases;    /* alias list */
    long    r_number;       /* rpc program number */
};
```

The members of this structure are:

r_name	The name of the server for this rpc program.
r_aliases	A zero terminated list of alternate names for the rpc program.
r_number	The rpc program number for this service.

getrpcent() reads the next line of the file, opening the file if necessary.

getrpcent() opens and rewinds the file. If the *stayopen* flag is non-zero, the net data base will not be closed after each call to getrpcent() (either directly, or indirectly through one of the other “getrpc” calls).

endrpcent closes the file.

getrpcbyname() and getrpcbynumber() sequentially search from the beginning of the file until a matching rpc program name or program number is found, or until end-of-file is encountered.

FILES

/etc/rpc

SEE ALSO

rpc(5), rpcinfo(8C), yperv(8)

DIAGNOSTICS

A NULL pointer 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

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 characters from the standard input stream, **stdin**, into the array pointed to by *s*, until a NEWLINE character is read or an EOF condition is encountered. The NEWLINE character is discarded and the string is terminated with a NULL character. **gets()** returns its argument.

fgets() reads characters from the stream into the array pointed to by *s*, until *n*–1 characters are read, a NEWLINE character is read and transferred to *s*, or an EOF condition is encountered. The string is then terminated with a NULL character. **fgets()** returns its first argument.

SEE ALSO

puts(3S), getc(3S), scanf(3S), fread(3S), ferror(3S)

DIAGNOSTICS

If EOF is encountered and no characters have been read, no characters are transferred to *s* and a NULL pointer is returned. If a read error occurs, such as trying to use these functions on a file that has not been opened for reading, a NULL pointer is returned. Otherwise *s* is returned.

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 */
    int     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.

getservent() 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 end-of-file 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), *ypserv*(8)

DIAGNOSTICS

A NULL pointer is returned on end-of-file 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

getttyent, getttynam, setttyent, endtttyent – get ttytab file entry

SYNOPSIS

```
#include <ttyent.h>

struct ttyent *getttyent()
struct ttyent *getttynam(name)
char *name;

setttyent()

endtttyent()
```

DESCRIPTION

getttyent() and getttynam() each return a pointer to an object with the following structure containing the broken-out fields of a line from the tty description file.

```
struct ttyent {
    char *ty_name;      /* terminal device name */
    char *ty_getty;    /* command to execute, usually getty */
    char *ty_type;     /* terminal type for termcap (3X) */
    int ty_status;     /* status flags (see below for defines) */
    char *ty_window;   /* command to start up window manager */
    char *ty_comment;  /* usually the location of the terminal */
};
#define TTY_ON        0x1  /* enable logins (startup getty) */
#define TTY_SECURE    0x2  /* allow root to login */
```

ty_name is the name of the character-special file in the directory /dev. For various reasons, it must reside in the directory /dev.

ty_getty is the command (usually getty(8)) which is invoked by init to initialize tty line characteristics. In fact, any arbitrary command can be used; a typical use is to initiate a terminal emulator in a window system.

ty_type is the name of the default terminal type connected to this tty line. This is typically a name from the termcap(5) data base. The environment variable TERM is initialized with this name by getty(8) or login(1).

ty_status is a mask of bit fields which indicate various actions to be allowed on this tty line. The following is a description of each flag.

TTY_ON

Enables logins (that is, init(8) will start the specified “getty” command on this entry).

TTY_SECURE

Allows root to login on this terminal. Note: TTY_ON must be included for this to be useful.

ty_window is the command to execute for a window system associated with the line. The window system will be started before the command specified in the ty_getty entry is executed. If none is specified, this will be NULL.

ty_comment is the trailing comment field, if any; a leading delimiter and white space will be removed.

getttyent() reads the next line from the ttytab file, opening the file if necessary; getttynam() rewinds the file; endtttyent() closes it.

gettynam() searches from the beginning of the file until a matching *name* is found (or until EOF is encountered).

FILES

/etc/ttytab

SEE ALSO

login(1), ttyslot(3), ttyslot(3V), gettytab(5), ttytab(5), termcap(5), getty(8), init(8)

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

getusershell, setusershell, endusershell – get legal user shells

SYNOPSIS

char *getusershell()

setusershell()

endusershell()

DESCRIPTION

getusershell() returns a pointer to a legal user shell as defined by the system manager in the file `/etc/shells`. If `/etc/shells` does not exist, the two standard system shells `/usr/bin/sh` and `/usr/bin/csh` are returned.

getusershell() reads the next line (opening the file if necessary); **setusershell()** rewinds the file; **endusershell()** closes it.

FILES

`/etc/shells`

`/usr/bin/csh`

DIAGNOSTICS

The routine **getusershell()** returns a NULL pointer (0) 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

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

```
char *getwd(pathname)  
char pathname[MAXPATHLEN];
```

DESCRIPTION

`getwd()` copies the absolute pathname of the current working directory to *pathname* and returns a pointer to the result.

DIAGNOSTICS

`getwd()` returns zero and places a message in *pathname* if an error occurs.

FILES

`/tmp/.getwd` It exists for the sole purpose of the `getwd()` library routine; no other software should depend on its existence or contents.

NAME

hsearch, hcreate, hdestroy – manage hash search tables

SYNOPSIS

```
#include <search.h>

ENTRY *hsearch (item, action)
ENTRY item;
ACTION action;

int hcreate (nel)
unsigned nel;

void hdestroy ( )
```

DESCRIPTION

hsearch() is a hash-table search routine generalized from Knuth (6.4) Algorithm D. It returns a pointer into a hash table indicating the location at which an entry can be found. *item* is a structure of type **ENTRY** (defined in the `<search.h>` header file) containing two pointers: *item.key* points to the comparison key, and *item.data* points to any other data to be associated with that key. (Pointers to types other than character should be cast to pointer-to-character.) *action* is a member of an enumeration type **ACTION** indicating the disposition of the entry if it cannot be found in the table. **ENTER** indicates that the item should be inserted in the table at an appropriate point. **FIND** indicates that no entry should be made. Unsuccessful resolution is indicated by the return of a **NULL** pointer.

hcreate() allocates sufficient space for the table, and must be called before **hsearch()** is used. *nel* is an estimate of the maximum number of entries that the table will contain. This number may be adjusted upward by the algorithm in order to obtain certain mathematically favorable circumstances.

hdestroy() destroys the search table, and may be followed by another call to **hcreate**.

NOTES

hsearch() uses open addressing with a *multiplicative* hash function.

EXAMPLE

The following example will read in strings followed by two numbers and store them in a hash table, discarding duplicates. It will then read in strings and find the matching entry in the hash table and print it out.

```
#include <stdio.h>
#include <search.h>
struct info {          /* this is the info stored in the table */
    int age, room;    /* other than the key. */
};
#define
NUM_EMPL  5000  /* # of elements in search table */
main( )
{
    /* space to store strings */
    char string_space[NUM_EMPL*20];
    /* space to store employee info */
    struct info info_space[NUM_EMPL];
    /* next avail space in string_space */
    char *str_ptr = string_space;
    /* next avail space in info_space */
    struct info *info_ptr = info_space;
    ENTRY item, *found_item, *hsearch( );
    /* name to look for in table */
    char name_to_find[30];
    int i = 0;
    /* create table */
```

```

        (void) hcreate(NUM_EMPL);
        while (scanf("%s%d%d", str_ptr, &info_ptr->age,
                    &info_ptr->room) !=
EOF && i++ <
NUM_EMPL) {
        /* put info in structure, and structure in item */
        item.key = str_ptr;
        item.data = (char *)info_ptr;
        str_ptr += strlen(str_ptr) + 1;
        info_ptr++;
        /* put item into table */
        (void) hsearch(item,
ENTER);
    }
    /* access table */
    item.key = name_to_find;
    while (scanf("%s", item.key) != EOF) {
        if ((found_item = hsearch(item,
FIND)) != NULL) {
            /* if item is in the table */
            (void)printf("found %s, age = %d, room = %d\n",
                found_item->key,
                ((struct info *)found_item->data)->age,
                ((struct info *)found_item->data)->room);
        } else {
            (void)printf("no such employee %s\n",
                name_to_find)
        }
    }
}

```

SEE ALSO

bsearch(3), lsearch(3), malloc(3), string(3), tsearch(3)

DIAGNOSTICS

hsearch() returns a NULL pointer if either the action is FIND and the item could not be found or the action is ENTER and the table is full.

hcreate() returns zero if it cannot allocate sufficient space for the table.

WARNING

hsearch() and **hcreate()** use **malloc(3)** to allocate space.

BUGS

Only one hash search table may be active at any given time.

NAME

inet inet_addr, inet_network, inet_makeaddr, inet_lnaof, inet_netof, inet_ntoa – Internet address manipulation

SYNOPSIS

```
#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>

unsigned long
inet_addr(cp)
char *cp;

inet_network(cp)
char *cp;

struct in_addr
inet_makeaddr(net, lna)
int net, lna;

inet_lnaof(in)
struct in_addr in;

inet_netof(in)
struct in_addr in;

char *
inet_ntoa(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_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.

The routine `inet_ntoa()` returns a pointer to a string in the base 256 notation “d.d.d.d” described below.

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: when an Internet address is viewed as a 32-bit integer quantity on Sun386i systems, the bytes referred to above appear as d.c.b.a. That is, Sun386i 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 (that is, 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 return value from `inet_ntoa()` points to static information which is overwritten in each call.

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** call (see **getgroups(2)**), 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

getgroups(2), **getgrent(3)**

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**.

NAME

insque, remque – insert/remove element from a queue

SYNOPSIS

```
struct qelem {
    struct qelem *q_forw;
    struct qelem *q_back;
    char    q_data[];
};

insque(elem, pred)
struct qelem *elem, *pred;

remque(elem)
struct qelem *elem;
```

DESCRIPTION

insque() and **remque()** manipulate queues built from doubly linked lists. Each element in the queue must be in the form of "struct qelem". **insque()** inserts *elem* in a queue immediately after *pred*; **remque()** removes an entry *elem* from a queue.

NAME

issecure – indicates whether system is running secure

SYNOPSIS

int issecure()

DESCRIPTION

This function tells whether the system has been configured to run in secure mode. It returns 0 if the system is not running secure, and non-zero if the system is running secure.

NAME

`kvm_getu`, `kvm_getcmd` – get the u-area or invocation arguments for a process

SYNOPSIS

```
#include <kvm.h>
#include <sys/param.h>
#include <sys/user.h>
#include <sys/proc.h>

struct user *kvm_getu(kd, proc)
kvm_t *kd;
struct proc *proc;

int kvm_getcmd(kd, proc, u, arg, env)
kvm_t *kd;
struct proc *proc;
struct user *u;
char ***arg;
char ***env;
```

DESCRIPTION

`kvm_getu()` reads the u-area of the process specified by *proc* to an area of static storage associated with *kd* and returns a pointer to it. Subsequent calls to `kvm_getu()` will overwrite the static u-area.

kd is a pointer to a kernel identifier returned by `kvm_open(3K)`. *proc* is a pointer to a copy (in the current process' address space) of a *proc* structure (obtained, for instance, by a prior `kvm_nextproc(3K)` call).

`kvm_getcmd()` constructs a list of string pointers that represent the command arguments and environment that were used to initiate the process specified by *proc*.

kd is a pointer to a kernel identifier returned by `kvm_open(3K)`. *u* is a pointer to a copy (in the current process' address space) of a *user* structure (obtained, for instance, by a prior `kvm_getu()` call). If *arg* is not NULL, then the command line arguments are formed into a NULL-terminated array of string pointers. The address of the first such pointer is returned in *arg*. If *env* is not NULL, then the environment is formed into a NULL-terminated array of string pointers. The address of the first of these is returned in *env*.

The pointers returned in *arg* and *env* refer to data allocated by `malloc(3)` and should be freed (by a call to `free` (see `malloc(3)`) when no longer needed. Both the string pointers and the strings themselves are deallocated when freed.

Since the environment and command line arguments may have been modified by the user process, there is no guarantee that it will be possible to reconstruct the original command at all. Thus, `kvm_getcmd()` will make the best attempt possible, returning `-1` if the user process data is unrecognizable.

RETURN VALUE

`kvm_getu()` returns a NULL pointer if an error occurred.

`kvm_getcmd()` returns a value of 0 on successful completion. Otherwise `-1` is returned.

SEE ALSO

`execve(2)`, `kvm_nextproc(3K)`, `kvm_open(3K)`, `kvm_read(3K)`, `malloc(3)`

NOTES

If `kvm_getcmd()` returns `-1`, the caller still has the option of using the command line fragment that is stored in the u-area.

NAME

`kvm_getproc`, `kvm_nextproc`, `kvm_setproc` – read system process structures

SYNOPSIS

```
#include <kvm.h>
#include <sys/param.h>
#include <sys/time.h>
#include <sys/proc.h>

struct proc *kvm_getproc(kd, pid)
kvm_t *kd;
int pid;

struct proc *kvm_nextproc(kd)
kvm_t *kd;

int kvm_setproc(kd)
kvm_t *kd;
```

DESCRIPTION

`kvm_nextproc()` may be used to sequentially read all of the system process structures from the kernel identified by *kd* (see `kvm_open(3K)`). Each call to `kvm_nextproc()` returns a pointer to the static memory area that contains a copy of the next valid process table entry. There is no guarantee that the data will remain valid across calls to `kvm_nextproc`, `kvm_setproc`, or `kvm_getproc`. Therefore, if the process structure must be saved, it should be copied to non-volatile storage.

For performance reasons, many implementations will cache a set of system process structures. Since the system state is liable to change between calls to `kvm_nextproc`, and since the cache may contain obsolete information, there is no guarantee that *every* process structure returned refers to an active process, nor is it certain that *all* processes will be reported.

`kvm_setproc()` rewinds the process list, enabling `kvm_nextproc()` to rescan from the beginning of the system process table. `kvm_setproc()` will always flush the process structure cache, allowing an application to re-scan the process table of a running system.

`kvm_getproc()` locates the *proc* structure of the process specified by *pid* and returns a pointer to it. `kvm_getproc()` does not interact with the process table pointer manipulated by `kvm_nextproc`, however, the restrictions regarding the validity of the data still apply.

RETURN VALUE

`kvm_getproc()` and `kvm_nextproc()` return a NULL pointer if an error has occurred.

`kvm_setproc()` returns a value of 0 on successful completion. Otherwise -1 is returned.

SEE ALSO

`kvm_getu(3K)`, `kvm_open(3K)`, `kvm_read(3K)`

NAME

`kvm_nlist` – get entries from kernel symbol table

SYNOPSIS

```
#include <kvm.h>
#include <nlist.h>

int kvm_nlist(kd, nl)
kvm_t *kd;
struct nlist *nl;
```

DESCRIPTION

`kvm_nlist()` examines the symbol table from the kernel image identified by *kd* (see `kvm_open(3K)`) and selectively extracts a list of values and puts them in the array of `nlist()` structures pointed to by *nl*. The name list pointed to by *nl* consists of an array of structures containing names, types and values. The *n_name* field of each such structure is taken to be a pointer to a character string representing a symbol name. The list is terminated by an entry with a NULL pointer (or a pointer to a NULL string) in the *n_name* field. For each entry in *nl*, if the named symbol is present in the kernel symbol table, its value and type are placed in the *n_value* and *n_type* fields. If a symbol cannot be located, the corresponding *n_type* field of *nl* is set to zero.

RETURN VALUE

Upon normal completion, `kvm_nlist()` returns the number of symbols that were not located in the symbol table. If an error occurs, `nlist()` returns -1 and sets all of the *n_type* fields in members of the array pointed to by *nl* to zero.

SEE ALSO

`kvm_open(3K)`, `kvm_read(3K)`, `nlist(3)`, `a.out(5)`

NAME

`kvm_open`, `kvm_close` – specify a kernel to examine

SYNOPSIS

```
#include <kvm.h>
#include <fcntl.h>

kvm_t *kvm_open(namelist, corefile, swapfile, flag, errstr)
char *namelist, *corefile, *swapfile;
int flag;
char *errstr;

int kvm_close(kd)
kvm_t *kd;
```

DESCRIPTION

`kvm_open()` initializes a set of file descriptors to be used in subsequent calls to Kernel VM routines. It returns a pointer to a kernel identifier that must be used as the *kd* argument in subsequent Kernel VM function calls.

The *namelist* argument specifies an unstripped executable file whose symbol table will be used to locate various offsets in *corefile*. If *namelist* is NULL, the symbol table of the currently running kernel is used to determine offsets in the core image. In this case, it is up to the implementation to select an appropriate way to resolve symbolic references (for instance, using */vmunix* as a default *namelist* file).

corefile specifies a file that contains an image of physical memory, for instance, a kernel crash dump file (see `savecore(8)`) or the special device */dev/mem*. If *corefile* is NULL, the currently running kernel is accessed (using */dev/mem* and */dev/kmem*).

swapfile specifies a file that represents the swap device. If both *corefile* and *swapfile* are NULL, the swap device of the “currently running kernel” is accessed. Otherwise, if *swapfile* is NULL, `kvm_open()` may succeed but subsequent `kvm_getu(3K)` function calls may fail if the desired information is swapped out.

flag is used to specify read or write access for *corefile* and may have one of the following values:

<code>O_RDONLY</code>	open for reading
<code>O_RDWR</code>	open for reading and writing

errstr is used to control error reporting. If it is a NULL pointer, no error messages will be printed. If it is non-NULL, it is assumed to be the address of a string that will be used to prefix error messages generated by `kvm_open`. Errors are printed to `stderr`. A useful value to supply for *errstr* would be `argv[0]`. This has the effect of printing the process name in front of any error messages.

`kvm_close()` closes all file descriptors that were associated with *kd*. These files are also closed on `exit(2)` and `execve(2)`. `kvm_close()` also resets the `proc` pointer associated with `kvm_nextproc(3K)` and flushes any cached kernel data.

RETURN VALUE

Upon successful completion, `kvm_open()` returns a non-NULL value that is suitable for use with subsequent Kernel VM function calls. If an error occurs, no files are opened and 0 is returned.

Upon successful completion, `kvm_close()` closes all file descriptors associated with *kd* and returns zero. Otherwise `-1` is returned.

FILES

```
/vmunix
/dev/kmem
/dev/mem
/dev/drum
```

SEE ALSO

execve(2), exit(2), kvm_getu(3K), kvm_nextproc(3K), kvm_nlist(3K), kvm_read(3K), savecore(8)

NAME

`kvm_read`, `kvm_write` – copy data to or from a kernel image or running system

SYNOPSIS

```
#include <kvm.h>

int kvm_read(kd, addr, buf, nbytes)
kvm_t *kd;
unsigned long addr;
char *buf;
unsigned nbytes;

int kvm_write(kd, addr, buf, nbytes)
kvm_t *kd;
unsigned long addr;
char *buf;
unsigned nbytes;
```

DESCRIPTION

`kvm_read()` transfers data from the kernel image specified by *kd* (see `kvm_open(3K)`) to the address space of the process. *nbytes* bytes of data are copied from the kernel virtual address given by *addr* to the buffer pointed to by *buf*.

`kvm_write()` is like `kvm_read`, except that the direction of data transfer is reversed. In order to use this function, the `kvm_open(3K)` call that returned *kd* must have specified write access.

RETURN VALUE

Upon normal completion, the number of bytes successfully transferred is returned. Otherwise `-1` is returned.

SEE ALSO

`kvm_getu(3K)`, `kvm_nlist(3K)`, `kvm_open(3K)`

NAME

ldahread – read the archive header of a member of a COFF archive file

SYNOPSIS

```
#include <stdio.h>
#include <ar.h>
#include <filehdr.h>
#include <ldfcn.h>

int ldahread (ldptr, arhead)
LDFILE *ldptr;
ARCHDR *arhead;
```

AVAILABILITY

Sun386i systems only.

DESCRIPTION

If **TYPE(ldptr)** is the archive file magic number, **ldahread** reads the archive header of the COFF file currently associated with *ldptr* into the area of memory beginning at *arhead*.

ldahread returns **SUCCESS** or **FAILURE**. **ldahread** will fail if **TYPE(ldptr)** does not represent an archive file, or if it cannot read the archive header.

The program must be loaded with the object file access routine library **libld.a**.

SEE ALSO

ldclose(3X), **ldopen(3X)**, **intro(5)**, **ldfcn(5)**

NAME

`ldclose`, `ldaclose` – close a COFF file

SYNOPSIS

```
#include <stdio.h>
#include <filehdr.h>
#include <ldfcn.h>

int ldclose (ldptr)
LDFILE *ldptr;

int ldaclose (ldptr)
LDFILE *ldptr;
```

AVAILABILITY

Sun386i systems only.

DESCRIPTION

`ldopen(3X)` and `ldclose()` are designed to provide uniform access to both simple COFF object files and COFF object files that are members of archive files. Thus an archive of COFF files can be processed as if it were a series of simple COFF files.

If `TYPE(ldptr)` does not represent an archive file, `ldclose()` will close the file and free the memory allocated to the `LDFILE` structure associated with `ldptr`. If `TYPE(ldptr)` is the magic number of an archive file, and if there are any more files in the archive, `ldclose()` will reinitialize `OFFSET(ldptr)` to the file address of the next archive member and return `FAILURE`. The `LDFILE` structure is prepared for a subsequent `ldopen(3X)`. In all other cases, `ldclose()` returns `SUCCESS`.

`ldaclose()` closes the file and frees the memory allocated to the `LDFILE` structure associated with `ldptr` regardless of the value of `TYPE(ldptr)`. `ldaclose()` always returns `SUCCESS`. The function is often used in conjunction with `ldaopen`.

The program must be loaded with the object file access routine library `libld.a`.

`intro(5)` describes `INCDIR` and `LIBDIR`.

SEE ALSO

`fclose(3S)`, `ldfcn(3)`, `ldopen(3X)`, `intro(5)`

NAME

ldfcn – common object file access routines

SYNOPSIS

```
#include <stdio.h>
#include <filehdr.h>
#include <ldfcn.h>
```

AVAILABILITY

Sun386i systems only.

DESCRIPTION

These routines are for reading COFF object files and archives containing COFF object files. Although the calling program must know the detailed structure of the parts of the object file that it processes, the routines effectively insulate the calling program from knowledge of the overall structure of the object file.

The interface between the calling program and the object file access routines is based on the defined type **LDFILE**, defined as **struct ldfile**, declared in the header file **ldfcn.h**. The primary purpose of this structure is to provide uniform access to both simple object files and to object files that are members of an archive file.

The function **ldopen(3X)** allocates and initializes the **LDFILE** structure and returns a pointer to the structure to the calling program. The fields of the **LDFILE** structure may be accessed individually through macros defined in **ldfcn.h** and contain the following information:

LDFILE	*ldptr;
TYPE(ldptr)	The file magic number used to distinguish between archive members and simple object files.
IOPTR(ldptr)	The file pointer returned by <i>fopen</i> and used by the standard input/output functions.
OFFSET(ldptr)	The file address of the beginning of the object file; the offset is non-zero if the object file is a member of an archive file.
HEADER(ldptr)	The file header structure of the object file.

The object file access functions themselves may be divided into four categories:

(1) Functions that open or close an object file

```
ldopen(3X) and ldaopen(see ldopen(3X))
    open a common object file
ldclose(3X) and ldaclose[see ldclose(3X)]
    close a common object file
```

(2) Functions that read header or symbol table information

```
ldahread(3X)
    read the archive header of a member of an archive file
ldfhread(3X)
    read the file header of a common object file
ldshread(3X) and ldnshread[see ldshread(3X)]
    read a section header of a common object file
ldtbread(3X)
    read a symbol table entry of a common object file
ldgetname(3X)
    retrieve a symbol name from a symbol table entry or from the string table
```

(3) Functions that position an object file at (seek to) the start of the section, relocation, or line number information for a particular section.

ldohseek(3X)

seek to the optional file header of a common object file

ldsseek(3X) and ldnsseek[see ldsseek(3X)]

seek to a section of a common object file

ldrseek(3X) and ldnrseek[see ldrseek(3X)]

seek to the relocation information for a section of a common object file

ldlseek(3X) and ldnlseek[see ldlseek(3X)]

seek to the line number information for a section of a common object file

ldtbseek(3X)

seek to the symbol table of a common object file

- (4) The uncton **ldtbindex(3X)**, which returns the index of a particular common object file symbol table entry.

These functions are described in detail on their respective manual pages.

All the functions except **ldopen(3X)**, **ldgetname(3X)**, **ldtbindex(3X)** return either **SUCCESS** or **FAILURE**, both constants defined in **ldfcn.h**. **ldopen(3X)** and **ldaopen[see ldopen(3X)]** both return pointers to an **LDFILE** structure.

Additional access to an object file is provided through a set of macros defined in **ldfcn.h**. These macros parallel the standard input/output file reading and manipulating functions, translating a reference of the **LDFILE** structure into a reference to its file descriptor field.

The following macros are provided:

```

GETC(ldptr)
FGETC(ldptr)
GETW(ldptr)
UNGETC(c, ldptr)
FGETS(s, n, ldptr)
FREAD((char *) ptr, sizeof (*ptr), nitems, ldptr)
FSEEK(ldptr, offset, ptrname)
FTELL(ldptr)
REWIND(ldptr)
FEOF(ldptr)
FERROR(ldptr)
FILENO(ldptr)
SETBUF(ldptr, buf)
STROFFSET(ldptr)

```

The **STROFFSET** macro calculates the address of the string table. See the manual entries for the corresponding standard input/output library functions for details on the use of the rest of the macros.

The program must be loaded with the object file access routine library **libld.a**.

SEE ALSO

fseek(3S), **ldahread(3X)**, **ldclose(3X)**, **ldgetname(3X)**, **ldhread(3X)**, **ldhread(3X)**, **ldlread(3X)**, **ldlseek(3X)**, **ldohseek(3X)**, **ldopen(3X)**, **ldrseek(3X)**, **ldlseek(3X)**, **ldshread(3X)**, **ldtbindex(3X)**, **ldtbread(3X)**, **ldtbseek(3X)**, **stdio(3S)**, **intro(5)**

WARNING

The macro **FSEEK** defined in the header file **ldfcn.h** translates into a call to the standard input/output function **fseek(3S)**. **FSEEK** should not be used to seek from the end of an archive file since the end of an archive file may not be the same as the end of one of its object file members.

NAME

ldfhead – read the file header of a COFF file

SYNOPSIS

```
#include <stdio.h>
#include <filehdr.h>
#include <ldfcn.h>
```

```
int ldfhead (ldptr, filehead)
LDFILE *ldptr;
FILHDR *filehead;
```

AVAILABILITY

Sun386i systems only.

DESCRIPTION

ldfhead() reads the file header of the COFF file currently associated with *ldptr* into the area of memory beginning at *filehead*.

ldfhead() returns SUCCESS or FAILURE. **ldfhead()** will fail if it cannot read the file header.

In most cases the use of **ldfhead()** can be avoided by using the macro **HEADER(*ldptr*)** defined in **ldfcn.h** (see **ldfcn(3)**). The information in any field, *fieldname*, of the file header may be accessed using **HEADER(*ldptr*).fieldname**.

The program must be loaded with the object file access routine library **libld.a**.

SEE ALSO

ldclose(3X), **ldfcn(3)**, **ldopen(3X)**

NAME

ldgetname – retrieve symbol name for COFF file symbol table entry

SYNOPSIS

```
#include <stdio.h>
#include <filehdr.h>
#include <syms.h>
#include <ldfcn.h>

char *ldgetname (ldptr, symbol)
LDFILE *ldptr;
SYMENT *symbol;
```

AVAILABILITY

Sun386i systems only.

DESCRIPTION

ldgetname() returns a pointer to the name associated with **symbol** as a string. The string is contained in a static buffer local to **ldgetname()** that is overwritten by each call to **ldgetname()**, and therefore must be copied by the caller if the name is to be saved.

ldgetname() can be used to retrieve names from object files without any backward compatibility problems. **ldgetname()** will return NULL (defined in **stdio.h**) for an object file if the name cannot be retrieved. This situation can occur:

- if the “string table” cannot be found,
- if not enough memory can be allocated for the string table,
- if the string table appears not to be a string table (for example, if an auxiliary entry is handed to **ldgetname()** that looks like a reference to a name in a nonexistent string table), or
- if the name’s offset into the string table is past the end of the string table.

Typically, **ldgetname()** will be called immediately after a successful call to **ldtbread()** to retrieve the name associated with the symbol table entry filled by **ldtbread()**.

The program must be loaded with the object file access routine library **libld.a**.

SEE ALSO

ldclose(3X), **ldfcn(3)**, **ldopen(3X)**, **ldtbread(3X)**, **ldtbseek(3X)**

NAME

ldread, ldlimit, lditem – manipulate line number entries of a COFF file function

SYNOPSIS

```
#include <stdio.h>
#include <filehdr.h>
#include <linenum.h>
#include <ldfcn.h>

int ldread(ldptr, fcnindx, linenum, linent)
LDFILE *ldptr;
long fcnindx;
unsigned short linenum;
LINENO *linent;

int ldlimit(ldptr, fcnindx)
LDFILE *ldptr;
long fcnindx;

int lditem(ldptr, linenum, linent)
LDFILE *ldptr;
unsigned short linenum;
LINENO *linent;
```

AVAILABILITY

Sun386i systems only.

DESCRIPTION

ldread() searches the line number entries of the COFF file currently associated with *ldptr*. **ldread()** begins its search with the line number entry for the beginning of a function and confines its search to the line numbers associated with a single function. The function is identified by *fcnindx*, the index of its entry in the object file symbol table. **ldread()** reads the entry with the smallest line number equal to or greater than *linenum* into the memory beginning at *linent*.

ldlimit() and **lditem()** together perform exactly the same function as **ldread()**. After an initial call to **ldread()** or **ldlimit()**, **lditem()** may be used to retrieve a series of line number entries associated with a single function. **ldlimit()** simply locates the line number entries for the function identified by *fcnindx*. **lditem()** finds and reads the entry with the smallest line number equal to or greater than *linenum* into the memory beginning at *linent*.

ldread(), **ldlimit()**, and **lditem()** each return either **SUCCESS** or **FAILURE**. **ldread()** will fail if there are no line number entries in the object file, if *fcnindx* does not index a function entry in the symbol table, or if it finds no line number equal to or greater than *linenum*. **ldlimit()** will fail if there are no line number entries in the object file or if *fcnindx* does not index a function entry in the symbol table. **lditem()** will fail if it finds no line number equal to or greater than *linenum*.

The programs must be loaded with the object file access routine library **libld.a**.

SEE ALSO

ldclose(3X), ldfcn(3), ldopen(3X), ldtbodyindex(3X)

NAME

ldlseek, ldnlseek – seek to line number entries of a section of a COFF file

SYNOPSIS

```
#include <stdio.h>
#include <filehdr.h>
#include <ldfcn.h>

int ldlseek (ldptr, sectindx)
LDFILE *ldptr;
unsigned short sectindx;

int ldnlseek (ldptr, sectname)
LDFILE *ldptr;
char *sectname;
```

AVAILABILITY

Sun386i systems only.

DESCRIPTION

ldlseek() seeks to the line number entries of the section specified by *sectindx* of the COFF file currently associated with *ldptr*.

ldnlseek() seeks to the line number entries of the section specified by *sectname*.

ldlseek() and **ldnlseek()** return **SUCCESS** or **FAILURE**. **ldlseek()** will fail if *sectindx* is greater than the number of sections in the object file; **ldnlseek()** will fail if there is no section name corresponding with **sectname*. Either function will fail if the specified section has no line number entries or if it cannot seek to the specified line number entries.

Note that the first section has an index of one.

The program must be loaded with the object file access routine library **libld.a**.

SEE ALSO

ldclose(3X), **ldfcn(3)**, **ldopen(3X)**, **ldhread(3X)**

NAME

ldohseek – seek to the optional file header of a COFF file

SYNOPSIS

```
#include <stdio.h>
#include <filehdr.h>
#include <ldfcn.h>

int ldohseek (ldptr)
LDFILE *ldptr;
```

AVAILABILITY

Sun386i systems only.

DESCRIPTION

ldohseek() seeks to the optional file header of the COFF file currently associated with *ldptr*.

ldohseek() returns **SUCCESS** or **FAILURE**. **ldohseek()** will fail if the object file has no optional header or if it cannot seek to the optional header.

The program must be loaded with the object file access routine library **libld.a**.

SEE ALSO

ldclose(3X), **ldfcn(3)**, **ldopen(3X)**, **ldhread(3X)**

NAME

`ldopen`, `ldaopen` – open a COFF file for reading

SYNOPSIS

```
#include <stdio.h>
#include <filehdr.h>
#include <ldfcn.h>

LDFILE *ldopen (filename, ldptr)
char *filename;
LDFILE *ldptr;

LDFILE *ldaopen (filename, oldptr)
char *filename;
LDFILE *oldptr;
```

AVAILABILITY

Sun386i systems only.

DESCRIPTION

`ldopen()` and `ldclose(3X)` are designed to provide uniform access to both simple object files and object files that are members of archive files. Thus an archive of COFF files can be processed as if it were a series of simple COFF files.

If `ldptr` has the value `NULL`, then `ldopen()` will open *filename* and allocate and initialize the `LDFILE` structure, and return a pointer to the structure to the calling program.

If `ldptr` is valid and if `TYPE(ldptr)` is the archive magic number, `ldopen()` will reinitialize the `LDFILE` structure for the next archive member of *filename*.

`ldopen()` and `ldclose(3X)` are designed to work in concert. `ldclose` will return `FAILURE` only when `TYPE(ldptr)` is the archive magic number and there is another file in the archive to be processed. Only then should `ldopen()` be called with the current value of `ldptr`. In all other cases, in particular whenever a new *filename* is opened, `ldopen()` should be called with a `NULL` `ldptr` argument.

The following is a prototype for the use of `ldopen()` and `ldclose(3X)`.

```
/* for each filename to be processed */
ldptr = NULL;
do
{
    if ( (ldptr = ldopen(filename, ldptr)) != NULL )
    {
        /* check magic number */
        /* process the file */
    }
} while (ldclose(ldptr) == FAILURE );
```

If the value of `oldptr` is not `NULL`, `ldaopen()` will open *filename* anew and allocate and initialize a new `LDFILE` structure, copying the `TYPE`, `OFFSET`, and `HEADER` fields from `oldptr`. `ldaopen()` returns a pointer to the new `LDFILE` structure. This new pointer is independent of the old pointer, `oldptr`. The two pointers may be used concurrently to read separate parts of the object file. For example, one pointer may be used to step sequentially through the relocation information, while the other is used to read indexed symbol table entries.

Both `ldopen()` and `ldaopen()` open *filename* for reading. Both functions return `NULL` if *filename* cannot be opened, or if memory for the `LDFILE` structure cannot be allocated. A successful open does not insure that the given file is a COFF file or an archived object file.

The program must be loaded with the object file access routine library **libld.a**.

SEE ALSO

fopen(3S), ldclose(3X), ldfcn(3)

NAME

ldrseek, ldnrseek – seek to relocation entries of a section of a COFF file

SYNOPSIS

```
#include <stdio.h>
#include <filehdr.h>
#include <ldfcn.h>

int ldrseek (ldptr, sectindx)
LDFILE *ldptr;
unsigned short sectindx;

int ldnrseek (ldptr, sectname)
LDFILE *ldptr;
char *sectname;
```

AVAILABILITY

Sun386i systems only.

DESCRIPTION

ldrseek() seeks to the relocation entries of the section specified by *sectindx* of the COFF file currently associated with *ldptr*.

ldnrseek() seeks to the relocation entries of the section specified by *sectname*.

ldrseek() and ldnrseek() return SUCCESS or FAILURE. ldrseek() will fail if *sectindx* is greater than the number of sections in the object file; ldnrseek() will fail if there is no section name corresponding with *sectname*. Either function will fail if the specified section has no relocation entries or if it cannot seek to the specified relocation entries.

Note: the first section has an index of one.

The program must be loaded with the object file access routine library libld.a.

SEE ALSO

ldclose(3X), ldfcn(3), ldopen(3X), ldshread(3X)

NAME

ldshread, ldnsread – read an indexed/named section header of a COFF file

SYNOPSIS

```
#include <stdio.h>
#include <filehdr.h>
#include <scnhdr.h>
#include <ldfcn.h>

int ldshread (ldptr, sectindx, secthead)
LDFILE *ldptr;
unsigned short sectindx;
SCNHDR *secthead;

int ldnsread (ldptr, sectname, secthead)
LDFILE *ldptr;
char *sectname;
SCNHDR *secthead;
```

AVAILABILITY

Sun386i systems only.

DESCRIPTION

ldshread() reads the section header specified by *sectindx* of the COFF file currently associated with *ldptr* into the area of memory beginning at *secthead*.

ldnsread() reads the section header specified by *sectname* into the area of memory beginning at *secthead*.

ldshread() and ldnsread() return SUCCESS or FAILURE. ldshread() will fail if *sectindx* is greater than the number of sections in the object file; ldnsread() will fail if there is no section name corresponding with *sectname*. Either function will fail if it cannot read the specified section header.

Note: the first section header has an index of *one*.

The program must be loaded with the object file access routine library *libld.a*.

SEE ALSO

ldclose(3X), ldfcn(3), ldopen(3X)

NAME

ldsseek, ldnsseek – seek to an indexed/named section of a COFF file

SYNOPSIS

```
#include <stdio.h>
#include <filehdr.h>
#include <ldfcn.h>

int ldsseek (ldptr, sectindx)
LDFILE *ldptr;
unsigned short sectindx;

int ldnsseek (ldptr, sectname)
LDFILE *ldptr;
char *sectname;
```

AVAILABILITY

Sun386i systems only.

DESCRIPTION

ldsseek() seeks to the section specified by *sectindx* of the COFF file currently associated with *ldptr*.

ldnsseek() seeks to the section specified by *sectname*.

ldsseek() and **ldnsseek()** return **SUCCESS** or **FAILURE**. **ldsseek()** will fail if *sectindx* is greater than the number of sections in the object file; **ldnsseek()** will fail if there is no section name corresponding with *sectname*. Either function will fail if there is no section data for the specified section or if it cannot seek to the specified section.

Note: the first section has an index of *one*.

The program must be loaded with the object file access routine library **libld.a**.

SEE ALSO

ldclose(3X), **ldfcn(3)**, **ldopen(3X)**, **ldhread(3X)**

NAME

ldtbindex – compute the index of a symbol table entry of a COFF file

SYNOPSIS

```
#include <stdio.h>
#include <filehdr.h>
#include <syms.h>
#include <ldfcn.h>

long ldtbindex (ldptr)
LDFILE *ldptr;
```

AVAILABILITY

Sun386i systems only.

DESCRIPTION

ldtbindex() returns the (**long**) index of the symbol table entry at the current position of the COFF file associated with *ldptr*.

The index returned by **ldtbindex()** may be used in subsequent calls to **ldtbread(3X)**. However, since **ldtbindex()** returns the index of the symbol table entry that begins at the current position of the object file, if **ldtbindex()** is called immediately after a particular symbol table entry has been read, it will return the index of the next entry.

ldtbindex() will fail if there are no symbols in the object file, or if the object file is not positioned at the beginning of a symbol table entry.

Note that the first symbol in the symbol table has an index of *zero*.

The program must be loaded with the object file access routine library **libld.a**.

SEE ALSO

ldclose(3X), **ldfcn(3)**, **ldopen(3X)**, **ldtbread(3X)**, **ldtbseek(3X)**

NAME

ldtbread – read an indexed symbol table entry of a COFF file

SYNOPSIS

```
#include <stdio.h>
#include <filehdr.h>
#include <syms.h>
#include <ldfcn.h>
```

```
int ldtbread (ldptr, symindex, symbol)
LDFILE *ldptr;
long symindex;
SYMENT *symbol;
```

AVAILABILITY

Sun386i systems only.

DESCRIPTION

ldtbread() reads the symbol table entry specified by *symindex* of the COFF file currently associated with *ldptr* into the area of memory beginning at *symbol*.

ldtbread() returns SUCCESS or FAILURE. **ldtbread()** will fail if *symindex* is greater than or equal to the number of symbols in the object file, or if it cannot read the specified symbol table entry.

Note: the first symbol in the symbol table has an index of *zero*.

The program must be loaded with the object file access routine library **libld.a**.

SEE ALSO

ldclose(3X), **ldfcn(3)**, **ldopen(3X)**, **ldtbseek(3X)**, **ldgetname(3X)**

NAME

ldtbseek – seek to the symbol table of a COFF file

SYNOPSIS

```
#include <stdio.h>
#include <filehdr.h>
#include <ldfcn.h>

int ldtbseek (ldptr)
LDFILE *ldptr;
```

AVAILABILITY

Sun386i systems only.

DESCRIPTION

ldtbseek() seeks to the symbol table of the COFF file currently associated with *ldptr*.

ldtbseek() returns **SUCCESS** or **FAILURE**. **ldtbseek()** will fail if the symbol table has been stripped from the object file, or if it cannot seek to the symbol table.

The program must be loaded with the object file access routine library **libld.a**.

SEE ALSO

ldclose(3X), **ldfcn(3)**, **ldopen(3X)**, **ldtbread(3X)**

NAME

lockf – advisory record locking on files

SYNOPSIS

```
#include <unistd.h>

#define F_ULOCK    0    /* Unlock a previously locked section */
#define F_LOCK    1    /* Lock a section for exclusive use */
#define F_TLOCK    2    /* Test and lock a section (non-blocking) */
#define F_TEST    3    /* Test section for other process' locks */

int lockf(fd, cmd, size)
int fd, cmd;
long size;
```

DESCRIPTION

lockf() may be used to test, apply, or remove an *advisory* record lock on the file associated with the open descriptor *fd*. (See *fcntl(2V)* for more information about advisory record locking.)

A lock is obtained by specifying a *cmd* parameter of *F_LOCK* or *F_TLOCK*. To unlock an existing lock, the *F_ULOCK* *cmd* is used. *F_TEST* is used to detect if a lock by another process is present on the specified segment.

F_LOCK and *F_TLOCK* requests differ only by the action taken if the lock may not be immediately granted. *F_TLOCK* returns a *-1* by the function and sets *errno* to *EAGAIN* if the section is already locked by another process. *F_LOCK* will cause the process to sleep until the lock may be granted or a signal is caught.

size is the number of contiguous bytes to be locked or unlocked. The lock starts at the current file offset in the file and extends forward for a positive *size* or backward for a negative *size* (preceeding but not including the current offset). A segment need not be allocated to the file in order to be locked; however, a segment may not extend to a negative offset relative to the beginning of the file. If *size* is zero, the lock will extend from the current offset through the EOF. If such a lock starts at offset 0, then the entire file will be locked (regardless of future file extensions).

NOTES

The descriptor *fd* must have been opened with *O_WRONLY* or *O_RDWR* permission in order to establish locks with this function call.

All locks associated with a file for a given process are removed when the file is closed or the process terminates. Locks are not inherited by the child process in a *fork(2)* system call.

RETURN VALUE

Zero is returned on success, *-1* on error, with an error code stored in *errno*.

ERRORS

lockf() will fail if one or more of the following are true:

EBADF	<i>fd</i> is not a valid open descriptor.
EBADF	<i>cmd</i> is <i>F_LOCK</i> or <i>F_TLOCK</i> and the process does not have write permission on the file.
EAGAIN	<i>cmd</i> is <i>F_TLOCK</i> or <i>F_TEST</i> and the section is already locked by another process.
EINTR	<i>cmd</i> is <i>F_LOCK</i> and a signal interrupted the process while it was waiting for the lock to be granted.
ENOLCK	<i>cmd</i> is <i>F_LOCK</i> , <i>F_TLOCK</i> , or <i>F_ULOCK</i> and there are no more file lock entries available.

SEE ALSO

fcntl(2V), flock(2), fork(2), lockd(8C)

BUGS

File locks obtained through the **lockf()** mechanism do not interact in any way with those acquired using **flock(2)**. They do, however, work correctly with the locks claimed by **fcntl(2V)**.

NAME

`lsearch`, `lfind` – linear search and update

SYNOPSIS

```
#include <stdio.h>
#include <search.h>

char *lsearch ((char *)key, (char *)base, nelp, sizeof(*key), compar)
unsigned *nelp;
    int (*compar)( );

char *lfind ((char *)key, (char *)base, nelp, sizeof(*key), compar)
unsigned *nelp;
int
    (*compar)( );
```

DESCRIPTION

`lsearch()` is a linear search routine generalized from Knuth (6.1) Algorithm S. It returns a pointer into a table indicating where a datum may be found. If the datum does not occur, it is added at the end of the table. `key()` points to the datum to be sought in the table. `base` points to the first element in the table. `nelp` points to an integer containing the current number of elements in the table. The integer is incremented if the datum is added to the table. `compar` is the name of the comparison function which the user must supply (`strcmp`, for example). It is called with two arguments that point to the elements being compared. The function must return zero if the elements are equal and non-zero otherwise.

`lfind()` is the same as `lsearch()` except that if the datum is not found, it is not added to the table. Instead, a NULL pointer is returned.

NOTES

The pointers to the key and the element at the base of the table should be of type pointer-to-element, and cast to type pointer-to-character.

The comparison function need not compare every byte, so arbitrary data may be contained in the elements in addition to the values being compared.

Although declared as type pointer-to-character, the value returned should be cast into type pointer-to-element.

EXAMPLE

This fragment will read in \leq TABSIZE strings of length \leq ELSIZE and store them in a table, eliminating duplicates.

```
#include <stdio.h>
#include <search.h>
#define
TABSIZE 50
#define
ELSIZE 120
    char line[ELSIZE], tab[TABSIZE][ELSIZE], *lsearch( );
    unsigned nel = 0;
    int strcmp( );
    ...
    while (fgets(line,
ELSIZE, stdin) != NULL &&
        nel < TABSIZE)
        (void) lsearch(line, (char *)tab, &nel, ELSIZE, strcmp);
    ...
```

SEE ALSO

bsearch(3), hsearch(3), tsearch(3)

DIAGNOSTICS

If the searched for datum is found, both **lsearch()** and **lfind()** return a pointer to it. Otherwise, **lfind()** returns NULL and **lsearch()** returns a pointer to the newly added element.

BUGS

Undefined results can occur if there is not enough room in the table to add a new item.

NAME

`malloc`, `free`, `realloc`, `calloc`, `cfree`, `memalign`, `valloc`, `alloca`, `malloc_debug`, `malloc_verify` – 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;

cfree(ptr)
char *ptr;

char *memalign(alignment, size)
unsigned alignment;
unsigned size;

char *valloc(size)
unsigned size;

#include <alloca.h>
char *alloca(size)
int size;

```

DESCRIPTION

These routines provide a general-purpose memory allocation package. They maintain a table of free blocks for efficient allocation and coalescing of free storage. When there is no suitable space already free, the allocation routines call `sbrk()` (see `brk(2)`) to get more memory from the system.

Each of the allocation routines returns a pointer to space suitably aligned for storage of any type of object. Each returns a NULL pointer if the request cannot be completed (see `DIAGNOSTICS`).

`malloc()` returns a pointer to a block of at least *size* bytes, which is appropriately aligned. A NULL (0) pointer is returned if *size* is 0 or if *size* bytes of memory cannot be allocated.

`free()` releases a previously allocated block. Its argument is a pointer to a block previously allocated by `malloc`, `calloc`, `realloc`, `malloc`, or `memalign`.

`realloc()` changes the size of the block referenced 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. For backwards compatibility, `realloc()` accepts a pointer to a block freed since the most recent call to `malloc`, `calloc`, `realloc`, `valloc`, or `memalign`. Note: using `realloc()` with a block freed *before* the most recent call to `malloc`, `calloc`, `realloc`, `valloc`, or `memalign` is an error.

`calloc()` uses `malloc()` to allocate space for an array of *nelem* elements of size *elsize*, initializes the space to zeros, and returns a pointer to the initialized block. The block can be freed with `free()` or `cfree`.

`memalign()` allocates *size* bytes on a specified alignment boundary, and returns a pointer to the allocated block. The value of the returned address is guaranteed to be an even multiple of *alignment*. Note: the value of *alignment* must be a power of two, and must be greater than or equal to the size of a word.

`valloc(size)` is equivalent to `memalign(getpagesize(), size)`.

`alloca()` allocates *size* bytes of space in the stack frame of the caller, and returns a pointer to the allocated block. This temporary space is automatically freed when the caller returns.

ERRORS

malloc, **calloc**, **realloc**, **valloc**, **memalign**, **cfree**, and **free()** will each fail if one or more of the following are true:

- | | |
|---------------|---|
| EINVAL | The requested allocation size is zero or an invalid argument was specified. The value of <i>ptr</i> passed to free , cfree , or realloc() must be a pointer to a block previously allocated by malloc , calloc , realloc , valloc , or memalign . |
| EINVAL | The allocation heap is found to have been corrupted. More detailed information may be obtained by enabling range checks using malloc_debug . |
| ENOMEM | <i>size</i> bytes of memory could not be allocated. |

DIAGNOSTICS

More detailed diagnostics can be made available to programs using **malloc**, **calloc**, **realloc**, **valloc**, **memalign**, **cfree**, and **free**, by including a special relocatable object file at link time (see **FILES**). This file also provides routines for control of error handling and diagnosis, as defined below. Note: these routines are *not* defined in the standard library.

```
int malloc_debug(level)
```

```
int level;
```

```
int malloc_verify()
```

malloc_debug() sets the level of error diagnosis and reporting during subsequent calls to **malloc**, **calloc**, **realloc**, **valloc**, **memalign**, **cfree**, and **free**. The value of *level* is interpreted as follows:

- | | |
|----------------|--|
| Level 0 | malloc , calloc , realloc , valloc , memalign , cfree , and free behave the same as in the standard library. |
| Level 1 | The routines abort with a message to the standard error if errors are detected in arguments or in the heap. If a bad block is encountered, its address and size are included in the message. |
| Level 2 | Same as level 1, except that the entire heap is examined on every call to the above routines. |

malloc_debug() returns the previous error diagnostic level. The default level is 1.

malloc_verify() attempts to determine if the heap has been corrupted. It scans all blocks in the heap (both free and allocated) looking for strange addresses or absurd sizes, and also checks for inconsistencies in the free space table. **malloc_verify()** returns 1 if all checks pass without error, and otherwise returns 0. The checks can take a significant amount of time, so it should not be used indiscriminately.

FILES

/usr/lib/debug/malloc.o diagnostic versions of **malloc()** routines.

BUGS

alloca() is both machine- and compiler-dependent; its use is discouraged.

Since **realloc()** accepts a pointer to a block freed since the last call to **malloc**, **calloc**, **realloc**, **valloc**, or **memalign**, a degradation of performance results. The semantics of **free()** should be changed so that the contents of a previously freed block are undefined.

SEE ALSO**brk(2)**

Stephenson, C.J., *Fast Fits*, in *Proceedings of the ACM 9th Symposium on Operating Systems*, SIGOPS *Operating Systems Review*, vol. 17, no. 5, October 1983.
Core Wars, in *Scientific American*, May 1984.

NAME

memory, memccpy, memchr, memcmp, memcpy, memset – memory operations

SYNOPSIS

```
#include <memory.h>

char *memccpy (s1, s2, c, n)
char *s1, *s2;
int c, n;

char *memchr (s, c, n)
char *s;
int c, n;

int memcmp (s1, s2, n)
char *s1, *s2;
int n;

char *memcpy (s1, s2, n)
char *s1, *s2;
int n;

char *memset (s, c, n)
char *s;
int c, n;
```

DESCRIPTION

These functions operate as efficiently as possible on memory areas (arrays of characters bounded by a count, not terminated by a NULL character). They do not check for the overflow of any receiving memory area.

memccpy() copies characters from memory area *s2* into *s1*, stopping after the first occurrence of character *c* has been copied, or after *n* characters have been copied, whichever comes first. It returns a pointer to the character after the copy of *c* in *s1*, or a NULL pointer if *c* was not found in the first *n* characters of *s2*.

memchr() returns a pointer to the first occurrence of character *c* in the first *n* characters of memory area *s*, or a NULL pointer if *c* does not occur.

memcmp() compares its arguments, looking at the first *n* characters only, and returns an integer less than, equal to, or greater than 0, according as *s1* is lexicographically less than, equal to, or greater than *s2*.

memcpy() copies *n* characters from memory area *s2* to *s1*. It returns *s1*.

memset() sets the first *n* characters in memory area *s* to the value of character *c*. It returns *s*.

NOTE

For user convenience, all these functions are declared in the **<memory.h>** header file.

BUGS

memcmp() uses native character comparison, which is signed on some machines and unsigned on other machines. Thus the sign of the value returned when one of the characters has its high-order bit set is implementation-dependent.

Character movement is performed differently in different implementations. Thus overlapping moves may yield surprises.

NAME

mktemp, mkstemp – make a unique file name

SYNOPSIS

```
char *mktemp(template)
char *template;

mkstemp(template)
char *template;
```

DESCRIPTION

mktemp() creates a unique file name, typically in a temporary filesystem, by replacing *template* with a unique file name, and always returns the address of *template*. The string in *template* should contain a file name with six trailing Xs; **mktemp()** replaces the Xs with a letter and the current process ID. The letter will be chosen so that the resulting name does not duplicate an existing file. **mkstemp()** makes the same replacement to the template but returns a file descriptor for the template file open for reading and writing. **mkstemp()** avoids the race between testing whether the file exists and opening it for use.

Notes:

- **mktemp()** and **mkstemp()** actually *change* the template string which you pass; this means that you cannot use the same template string more than once — you need a fresh template for every unique file you want to open.
- When **mktemp()** or **mkstemp()** are creating a new unique filename they check for the prior existence of a file with that name. This means that if you are creating more than one unique filename, it is bad practice to use the same root template for multiple invocations of **mktemp()** or **mkstemp()**.

SEE ALSO

getpid(2), open(2V), tmpfile(3S), tmpnam(3S)

DIAGNOSTICS

mkstemp() returns an open file descriptor upon success. It returns **-1** if no suitable file could be created.

BUGS

It is possible to run out of letters.

NAME

monitor, monstartup, moncontrol – prepare execution profile

SYNOPSIS

```
#include <a.out.h>

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 `profil(1)` monitor, and includes an initial call with default parameters to its start-up routine `monstartup`. In this case, `monitor()` need not be called explicitly, except to gain fine control over `profil(2)` buffer allocation. An executable program created by 'cc -pg' automatically includes calls for the `gprofil(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` (see `brk(2)`) and passes it to `monitor()` (as described below) to record a histogram of program-counter values, and calls to certain functions. Only calls to functions compiled with 'cc -p' are recorded.

On Sun-2, Sun-3, and Sun-4 systems, an entire program can be profiled with:

```
extern etext();
...
monstartup(N_TXTOFF(0), etext);
```

On Sun386i systems, the equivalent code sequence is:

```
extern etext();
extern _start();
...
monstartup(_start, etext);
```

`etext` lies just above all the program text, see `end(3)`.

To stop execution monitoring and post results to the file `mon.out`, use:

```
monitor(0);
```

`profil(1)` can then be used to examine the results.

`moncontrol()` is used to selectively control profiling within a program. This works with both `profil(1)` and `gprofil(1)`. Profiling begins when the program starts. To stop the collection of profiling statistics, use:

```
moncontrol(0)
```

To resume the collection of statistics, use:

```
moncontrol(1)
```

This allows you to measure the cost of particular functions. Note: an output file is be produced upon program exit, regardless 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 `cc -p`.

To profile the entire program on Sun-2, Sun-3, and Sun-4 systems using the low-level interface to `profil(2)`, it is sufficient to use

```
extern etext();
...
monitor(N_TXTOFF(0), etext, buf, bufsize, nfunc);
```

On Sun386i systems, the equivalent calls are:

```
extern etext();
extern _start();
...
monitor(_start, etext, buf, bufsize, nfunc);
```

FILES

`mon.out`

SEE ALSO

`cc(1V)`, `prof(1)`, `gprof(1)`, `brk(2)`, `profil(2)`, `end(3)`

NAME

`mp`, `itom`, `madd`, `msub`, `mult`, `mdiv`, `min`, `mout`, `pow`, `gcd`, `rpow`, `xtom`, `mtox`, `mfree` – multiple precision integer arithmetic

SYNOPSIS

```
#include <mp.h>

madd(a, b, c)
MINT *a, *b, *c;

msub(a, b, c)
MINT *a, *b, *c;

mult(a, b, c)
MINT *a, *b, *c;

mdiv(a, b, q, r)
MINT *a, *b, *q, *r;

min(a)
MINT *a;

mout(a)
MINT *a;

pow(a, b, c, d)
MINT *a, *b, *c, *d;

gcd(a, b, c)
MINT *a, *b, *c;

rpow(a, n, b)
MINT *a, *b;
short n;

msqrt(a, b, r)
MINT *a, *b, *r;

sdiv(a, n, q, r)
MINT *a, *q;
short n, *r;

MINT *itom(n)
short n;

MINT *xtom(s)
char *s;

char *mtox(a)
MINT *a;

void mfree(a)
MINT *a;
```

DESCRIPTION

These routines perform arithmetic on integers of arbitrary length. The integers are stored using the defined type `MINT`. Pointers to a `MINT` should be initialized using the function `itom`, which sets the initial value to `n`. Alternatively, `xtom` may be used to initialize a `MINT` from a string of hexadecimal digits. `mfree()` may be used to release the storage allocated by these routines.

`madd`, `msub()` and `mult()` assign to their third arguments the sum, difference, and product, respectively, of their first two arguments. `mdiv()` assigns the quotient and remainder, respectively, to its third and fourth arguments. `sdiv` is like `mdiv()` except that the divisor is an ordinary integer. `msqrt`

produces the square root and remainder of its first argument. `mpow` calculates a raised to the power b , while `pow()` calculates this reduced modulo m . `min()` and `mout()` do decimal input and output. `mtox()` provides the inverse of `xtom`.

Use the `-lmp` loader option to obtain access to these functions.

DIAGNOSTICS

Illegal operations and running out of memory produce messages and core images.

FILES

`/usr/lib/libmp.a`

NAME

`ndbm`, `dbm_open`, `dbm_close`, `dbm_fetch`, `dbm_store`, `dbm_delete`, `dbm_firstkey`, `dbm_nextkey`, `dbm_error`, `dbm_clearerr` – data base subroutines

SYNOPSIS

```
#include <ndbm.h>

typedef struct {
    char *dptr;
    int dsize;
} datum;

DBM *dbm_open(file, flags, mode)
char *file;
int flags, mode;

void dbm_close (db)
DBM *db;

datum dbm_fetch(db, key)
DBM *db;
datum key;

int dbm_store(db, key, content, flags)
DBM *db;
datum key, content;
int flags;

int dbm_delete(db, key)
DBM *db;
datum key;

datum dbm_firstkey(db)
DBM *db;

datum dbm_nextkey(db)
DBM *db;

int dbm_error(db)
DBM *db;

int dbm_clearerr(db)
DBM *db;
```

DESCRIPTION

These functions maintain key/content pairs in a data base. The functions will handle very large (a billion blocks) databases and will access a keyed item in one or two file system accesses. This package replaces the earlier `dbm(3X)` library, which managed only a single database.

keys and *contents* are described by the `datum` typedef. A `datum` specifies a string of *dsize* bytes pointed to by *dptr*. Arbitrary binary data, as well as normal ASCII strings, are allowed. The data base is stored in two files. One file is a directory containing a bit map and has `.dir` as its suffix. The second file contains all data and has `.pag` as its suffix.

Before a database can be accessed, it must be opened by `dbm_open`. This will open and/or create the files `file.dir` and `file.pag` depending on the `flags` parameter (see `open(2V)`).

A database is closed by calling `dbm_close`.

Once open, the data stored under a key is accessed by `dbm_fetch()` and data is placed under a key by `dbm_store`. The `flags` field can be either `DBM_INSERT` or `DBM_REPLACE`. `DBM_INSERT` will only insert new entries into the database and will not change an existing entry with the same key. `DBM_REPLACE` will replace an existing entry if it has the same key. A key (and its associated contents) is

contents) is deleted by `dbm_delete`. A linear pass through all keys in a database may be made, in an (apparently) random order, by use of `dbm_firstkey()` and `dbm_nextkey`. `dbm_firstkey()` will return the first key in the database. `dbm_nextkey()` will return the next key in the database. This code will traverse the data base:

```
for (key = dbm_firstkey(db); key.dptr != NULL; key = dbm_nextkey(db))
```

`dbm_error()` returns non-zero when an error has occurred reading or writing the database. `dbm_clearerr()` resets the error condition on the named database.

SEE ALSO

`open(2V)`, `dbm(3X)`

DIAGNOSTICS

All functions that return an `int` indicate errors with negative values. A zero return indicates no error. Routines that return a `datum` indicate errors with a `NULL (0) dptr`. If `dbm_store` called with a `flags` value of `DBM_INSERT` finds an existing entry with the same key it returns 1.

BUGS

The `.pag` file will contain holes so that its apparent size is about four times its actual content. Older versions of the UNIX operating system may create real file blocks for these holes when touched. These files cannot be copied by normal means (`cp(1)`, `cat(1V)`, `tar(1)`, `ar(1)`) without filling in the holes.

`dptr` pointers returned by these subroutines point into static storage that is changed by subsequent calls.

The sum of the sizes of a key/content pair must not exceed the internal block size (currently 4096 bytes). Moreover all key/content pairs that hash together must fit on a single block. `dbm_store()` will return an error in the event that a disk block fills with inseparable data.

`dbm_delete()` does not physically reclaim file space, although it does make it available for reuse.

The order of keys presented by `dbm_firstkey()` and `dbm_nextkey()` depends on a hashing function, not on anything interesting.

There are no interlocks and no reliable cache flushing; thus concurrent updating and reading is risky.

NAME

nice – change priority of a process

SYNOPSIS

int nice(incr)

DESCRIPTION

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 undue impact on system performance.

Negative increments are illegal, except when specified by the super-user. The priority is limited to the range -20 (most urgent) to 20 (least). Requests for values above or below these limits result in the scheduling priority being set to the corresponding limit.

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).

RETURN VALUE

Upon successful completion, `nice()` returns 0. Otherwise, a value of -1 is returned and `errno` is set to indicate the error.

ERRORS

The priority is not changed if:

EACCES The value of *incr* specified was negative, and the effective user ID is not super-user.

SEE ALSO

nice(1), fork(2), getpriority(2), renice(8)

NAME

`nlist` – get entries from symbol table

SYNOPSIS

```
#include <nlist.h>

int nlist(filename, nl)
char *filename;
struct nlist *nl;
```

DESCRIPTION

`nlist()` examines the symbol table from the executable image whose name is pointed to by *filename*, and selectively extracts a list of values and puts them in the array of `nlist()` structures pointed to by *nl*. The name list pointed to by `nl()` consists of an array of structures containing names, types and values. The *n_name* field of each such structure is taken to be a pointer to a character string representing a symbol name. The list is terminated by an entry with a NULL pointer (or a pointer to a NULL string) in the *n_name* field. For each entry in *nl*, if the named symbol is present in the executable image's symbol table, its value and type are placed in the *n_value* and *n_type* fields. If a symbol cannot be located, the corresponding *n_type* field of `nl()` is set to zero.

RETURN VALUE

Upon normal completion, `nlist()` returns the number of symbols that were not located in the symbol table. If an error occurs, `nlist()` returns `-1` and sets all of the *n_type* fields in members of the array pointed to by `nl()` to zero.

SEE ALSO

`a.out(5)`
`coff(5)`

DIAGNOSTICS

On Sun-2, Sun-3, and Sun-4 systems, type entries are set to 0 if the file cannot be read or if it does not contain a valid name list.

On Sun386i systems, the type entries may be zero even when the name list succeeded, but the value entries will be zero only when the file cannot be read or does not contain a valid name list. Therefore, on Sun386i systems, the value entry can be used to determine whether the command succeeded.

NAME

on_exit – name termination handler

SYNOPSIS

```
int on_exit(proc, arg)  
void (*proc)();  
caddr_t arg;
```

DESCRIPTION

on_exit() names a routine to be called after a program calls **exit(3)** or returns normally, and before its process terminates. The routine named is called as

```
(*proc)(status, arg);
```

where *status* is the argument with which **exit()** was called, or zero if *main* returns. Typically, *arg* is the address of an argument vector to *(*proc)*, but may be an integer value. Several calls may be made to **on_exit**, specifying several termination handlers. The order in which they are called is the reverse of that in which they were given to **on_exit**.

SEE ALSO

gprof(1), tcov(1), exit(3)

DIAGNOSTICS

on_exit() returns zero normally, or nonzero if the procedure name could not be stored.

BUGS

Currently there is a limit of 20 termination handlers, including any invoked implicitly (for example, by **gprof(1)** or **tcov(1)** processing). Calls to **on_exit()** beyond this number will fail.

NOTES

This call is specific to the SunOS operating system and should not be used if portability is a concern. Standard I/O exit processing is always done last.

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(2V) or an interval timer, see getitimer(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(2V), getitimer(2), select(2), sigpause(2)

NAME

perror, sys_errlist, sys_nerr, errno – system error messages

SYNOPSIS

```
perror(s)  
char *s;  
int sys_nerr;  
char *sys_errlist[ ];  
int errno;
```

DESCRIPTION

perror() produces a short error message on the standard error describing the last error encountered during a call to a system or library function. If *s* is not a NULL pointer and does not point to a null string, the string it points to is printed, followed by a colon, followed by a space, followed by the message and a NEWLINE. If *s* is a NULL pointer or points to a null string, just the message is printed, followed by a NEWLINE. To be of most use, the argument string should include the name of the program that incurred the error. The error number is taken from the external variable **errno** (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

plot, openpl, erase, label, line, circle, arc, move, cont, point, linemod, space, closepl – graphics interface

SYNOPSIS

```

openpl()
erase()
label(s)
char s[ ];
line(x1, y1, x2, y2)
circle(x, y, r)
arc(x, y, x0, y0, x1, y1)
move(x, y)
cont(x, y)
point(x, y)
linemod(s)
char s[ ];
space(x0, y0, x1, y1)
closepl()

```

DESCRIPTION

These subroutines generate graphic output in a relatively device-independent manner. See plot(5) for a description of their effect. openpl() must be used before any of the others to open the device for writing. closepl() flushes the output.

String arguments to label() and linemod() are NULL-terminated, and do not contain NEWLINE characters.

Various flavors of these functions exist for different output devices. They are obtained by the following ld(1) options:

```

-lplot      device-independent graphics stream on standard output for plot(1G) filters
-l300       GSI 300 terminal
-l300s      GSI 300S terminal
-l450       GSI 450 terminal
-l4014      Tektronix 4014 terminal
-lplotaed   AED 512 color graphics terminal
-lplotbg    BBN bitgraph graphics terminal
-lplotdumb  Dumb terminals without cursor addressing or line printers
-lplotgigi  DEC Gigi terminals
-lplot2648  Hewlett Packard 2648 graphics terminal
-lplot7221  Hewlett Packard 7221 graphics terminal
-lplotimagen Imagen laser printer (default 240 dots-per-inch resolution).

```

FILES

/usr/lib/libplot.a
/usr/lib/lib300.a
/usr/lib/lib300s.a
/usr/lib/lib450.a
/usr/lib/lib4014.a
/usr/lib/libplotaed.a
/usr/lib/libplotbg.a
/usr/lib/libplotdumb.a
/usr/lib/libplotgigi.a
/usr/lib/libplot2648.a
/usr/lib/libplot7221.a
/usr/lib/libplotimagen.a

SEE ALSO

graph(1G), ld(1), plot(1G), plot(5)

NAME

popen, pclose – open or close a pipe (for I/O) from or to a process

SYNOPSIS

```
#include <stdio.h>
```

```
FILE *popen(command, type)
```

```
char *command, *type;
```

```
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 such that one can write to the standard input of the command, if the I/O mode is **w**, by writing to the file stream; and one can read from the standard output of the command, if the I/O mode is **r**, by reading from the file stream.

A stream opened by **popen()** should be closed by **pclose**, which waits for the associated process to terminate and returns the exit status of the command.

Because open files are shared, a type **r** command may be used as an input filter, reading its standard input (which is also the standard input of the process doing the **popen**) and providing filtered input on the stream, and a type **w** command may be used as an output filter, reading a stream of output written to the stream process doing the **popen()** and further filtering it and writing it to its standard output (which is also the standard input of the process doing the **popen**).

popen() always calls **sh(1)**, never **csh(1)**.

SEE ALSO

csh(1), **sh(1)**, **pipe(2)**, **wait(2)**, **fclose(3S)**, **fopen(3S)**, **system(3)**

DIAGNOSTICS

popen() returns a NULL pointer if the pipe or process cannot be created, or if it cannot allocate as much memory as it needs.

pclose() returns **-1** if stream is not associated with a 'popened' command.

BUGS

If the original and 'popened' processes concurrently read or write a common file, neither should use buffered I/O, because the buffering gets all mixed up. Similar problems with an output filter may be forestalled by careful buffer flushing, for instance, with **fflush**; see **fclose(3S)**.

NAME

printf, fprintf, sprintf – formatted output conversion

SYNOPSIS

```
#include <stdio.h>
int printf(format [ , arg ] ... )
char *format;

int fprintf(stream, format [ , arg ] ... )
FILE *stream;
char *format;

char *sprintf(s, format [ , arg ] ... )
char *s, *format;

#include <varargs.h>
int _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”, followed by the NULL character (\0), in consecutive bytes starting at *s; it is the user’s responsibility to ensure that enough storage is available. printf() and fprintf() return the number of characters transmitted. The return value of sprintf() is not normally used, but cast to type void instead. printf() and fprintf() return an EOF if an output error was encountered.

Each of these functions converts, formats, and prints its *args* under control of the *format*. The *format* is a character string 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 zero or more *args*. The results are undefined if there are insufficient *args* for the format. If the format is exhausted while *args* remain, the excess *args* are simply ignored.

Each conversion specification is introduced by the character %. After the %, the following appear in sequence:

Zero or more *flags*, which modify the meaning of the conversion specification.

An optional decimal digit string specifying a minimum *field width*. If the converted value has fewer characters than the field width, it will be padded on the left (or right, if the left-adjustment flag ‘-’, described below, has been given) to the field width. The padding is with blanks unless the field width digit string starts with a zero, in which case the padding is with zeros.

A *precision* that gives the minimum number of digits to appear for the d, i, o, u, x, or X conversions, the number of digits to appear after the decimal point for the e, E, and f conversions, the maximum number of significant digits for the g and G conversion, or the maximum number of characters to be printed from a string in s conversion. The precision takes the form of a period (.) followed by a decimal digit string; a NULL digit string is treated as zero. Padding specified by the precision overrides the padding specified by the field width.

An optional l (ell) specifying that a following d, i, o, u, x, or X conversion character applies to a long integer *arg*. An l before any other conversion character is ignored.

A character that indicates the type of conversion to be applied.

A field width or precision or both may be indicated by an asterisk (*) instead of a digit string. In this case, an integer *arg* supplies the field width or precision. The *arg* that is actually converted is not fetched until the conversion letter is seen, so the *args* specifying field width or precision must appear

before the *arg* (if any) to be converted. A negative field width argument is taken as a '-' flag followed by a positive field width. If the precision argument is negative, it will be changed to zero.

The flag characters and their meanings are:

- The result of the conversion will be left-justified within the field.
- + The result of a signed conversion will always begin with a sign (+ or -).
- blank If the first character of a signed conversion is not a sign, a blank will be prefixed to the result. This implies that if the blank and + flags both appear, the blank flag will be ignored.
- # This flag specifies that the value is to be converted to an "alternate form." For c, d, i, s, and u conversions, the flag has no effect. For o conversion, it increases the precision to force the first digit of the result to be a zero. For x or X conversion, a non-zero result will have 0x or 0X prefixed to it. For e, E, f, g, and G conversions, the result will always contain a decimal point, even if no digits follow the point (normally, a decimal point appears in the result of these conversions only if a digit follows it). For g and G conversions, trailing zeroes will *not* be removed from the result (which they normally are).

The conversion characters and their meanings are:

- d,i,o,u,x,X The integer *arg* is converted to signed decimal (d or i), unsigned octal (o), unsigned decimal (u), or unsigned hexadecimal notation (x and X), respectively; the letters abcdef are used for x conversion and the letters ABCDEF for X conversion. The precision specifies the minimum number of digits to appear; if the value being converted can be represented in fewer digits, it will be expanded with leading zeroes. (For compatibility with older versions, padding with leading zeroes may alternatively be specified by prepending a zero to the field width. This does not imply an octal value for the field width.) The default precision is 1. The result of converting a zero value with a precision of zero is a NULL string.
- f The float or double *arg* is converted to decimal notation in the style "[−]ddd.ddd" where the number of digits after the decimal point is equal to the precision specification. If the precision is missing, 6 digits are given; if the precision is explicitly 0, no digits and no decimal point are printed.
- e,E The float or double *arg* is converted in the style "[−]d.ddde±ddd," where there is one digit before the decimal point and the number of digits after it is equal to the precision; when the precision is missing, 6 digits are produced; if the precision is zero, no decimal point appears. The E format code will produce a number with E instead of e introducing the exponent. The exponent always contains at least two digits.
- g,G The float or double *arg* is printed in style f or e (or in style E in the case of a G format code), with the precision specifying the number of significant digits. The style used depends on the value converted: style e or E will be used only if the exponent resulting from the conversion is less than −4 or greater than the precision. Trailing zeroes are removed from the result; a decimal point appears only if it is followed by a digit.

The e, E, f, g, and G formats print IEEE indeterminate values (infinity or not-a-number) as "Infinity" or "NaN" respectively.

- c The character *arg* is printed.
- s The *arg* is taken to be a string (character pointer) and characters from the string are printed until a NULL character (\0) is encountered or until the number of characters indicated by the precision specification is reached. If the precision is missing, it is taken to be infinite, so all characters up to the first NULL character are printed. A NULL value for *arg* will yield undefined results.
- % Print a %; no argument is converted.

In no case does a non-existent or small field width cause truncation of a field; if the result of a conversion is wider than the field width, the field is simply expanded to contain the conversion result. Padding takes place only if the specified field width exceeds the actual width. Characters generated

by `printf()` and `fprintf()` are printed as if `putc(3S)` had been called.

EXAMPLES

To print a date and time in the form "Sunday, July 3, 10:02," where *weekday* and *month* are pointers to NULL-terminated strings:

```
printf("%s, %s %i, %d:%.2d", weekday, month, day, hour, min);
```

To print π to 5 decimal places:

```
printf("pi = %.5f", 4 * atan(1. 0));
```

NOTE

These routines call `_doprnt`, which is an implementation-dependent routine. Each uses the variable-length argument facilities of `varargs(3)`. Although it is possible to use `_doprnt` to take a list of arguments and pass them on to a routine like `printf`, not all implementations have such a routine. We strongly recommend that you use the routines described in `vprintf(3S)` instead.

SEE ALSO

`econvert(3)`, `printf(3V)`, `putc(3S)`, `scanf(3S)`, `varargs(3)`, `vprintf(3S)`

BUGS

Very wide fields (>128 characters) fail.

NAME

prof – profile within a function

SYNOPSIS

```
#define MARK
#include <prof.h>

void MARK (name)
```

DESCRIPTION

MARK will introduce a mark called *name* that will be treated the same as a function entry point. Execution of the mark will add to a counter for that mark, and program-counter time spent will be accounted to the immediately preceding mark or to the function if there are no preceding marks within the active function.

name may be any combination of up to six letters, numbers or underscores. Each *name* in a single compilation must be unique, but may be the same as any ordinary program symbol.

For marks to be effective, the symbol MARK must be defined before the header file <prof.h> is included. This may be defined by a preprocessor directive as in the synopsis, or by a command line argument, such as:

```
cc -p -DMARK foo.c
```

If MARK is not defined, the MARK (name) statements may be left in the source files containing them and will be ignored.

EXAMPLE

In this example, marks can be used to determine how much time is spent in each loop. Unless this example is compiled with MARK defined on the command line, the marks are ignored.

```
#include <prof.h>
func( )
{
    int i, j;
    .
    .
    .
    MARK (loop1);
    for (i = 0; i < 2000; i++) {
        ...
    }
    MARK (loop2);
    for (j = 0; j < 2000; j++) {
        ...
    }
}
```

SEE ALSO

prof(1), profil(2), monitor(3)

NAME

psignal, sys_siglist – system signal messages

SYNOPSIS

```
psignal(sig, s)  
unsigned sig;  
char *s;  
char *sys_siglist[];
```

DESCRIPTION

psignal() produces a short message on the standard error file describing the indicated signal. First the argument string *s* is printed, then a colon, then the name of the signal and a NEWLINE. Most usefully, the argument string is the name of the program which incurred the signal. The signal number should be from among those found in `<signal.h>`.

To simplify variant formatting of signal names, the vector of message strings **sys_siglist()** is provided; the signal number can be used as an index in this table to get the signal name without the newline. The define `NSIG` defined in `<signal.h>` is the number of messages provided for in the table; it should be checked because new signals may be added to the system before they are added to the table.

SEE ALSO

perror(3), signal(3)

NAME

putc, putchar, fputc, putw – put character or word on a stream

SYNOPSIS

```
#include <stdio.h>

int putc(c, stream)
char c;
FILE *stream;

int putchar(c)
char c;

int fputc(c, stream)
char c;
FILE *stream;

int putw(w, stream)
int w;
FILE *stream;
```

DESCRIPTION

putc() writes the character *c* onto the standard I/O output stream *stream* (at the position where the file pointer, if defined, is pointing). It returns the character written.

putchar(c) is defined as **putc(c, stdout)**. **putc()** and **putchar()** are macros.

fputc() behaves like **putc**, but is a function rather than a macro. **fputc()** runs more slowly than **putc**, but it takes less space per invocation and its name can be passed as an argument to a function.

putw() writes the C int (word) *w* to the standard I/O output stream *stream* (at the position of the file pointer, if defined). The size of a word is the size of an integer and varies from machine to machine. **putw()** neither assumes nor causes special alignment in the file.

Output streams are by default buffered if the output refers to a file and line-buffered if the output refers to a terminal. When an output stream is unbuffered, information is queued for writing on the destination file or terminal as soon as written; when it is buffered, many characters are saved up and written as a block. When it is line-buffered, each line of output is queued for writing on the destination terminal as soon as the line is completed (that is, as soon as a NEWLINE character is written or terminal input is requested). **setbuf(3S)**, **setbuffer**, or **setvbuf** may be used to change the stream's buffering strategy.

SEE ALSO

fclose(3S), **ferror(3S)**, **fopen(3S)**, **fread(3S)**, **getc(3S)**, **printf(3S)**, **puts(3S)**, **setbuf(3S)**,

DIAGNOSTICS

On success, **putc**, **fputc**, and **putchar** return the value that was written. On error, those functions return the constant EOF. **putw()** returns **ferror(stream)**, so that it returns 0 on success and 1 on failure.

BUGS

Because it is implemented as a macro, **putc()** treats a *stream* argument with side effects improperly. In particular, **putc(c, *f++)**; does not work sensibly. **fputc()** should be used instead.

Errors can occur long after the call to **putc**.

Because of possible differences in word length and byte ordering, files written using **putw()** are machine-dependent, and may not be read using **getw()** on a different processor.

NAME

putenv – change or add value to environment

SYNOPSIS

```
int putenv(string)
char *string;
```

DESCRIPTION

`string()` points to a string of the form '*name = value*'. `putenv()` makes the value of the environment variable *name* equal to *value* by altering an existing variable or creating a new one. In either case, the string pointed to by `string()` becomes part of the environment, so altering the string will change the environment. The space used by `string()` is no longer used once a new string-defining *name* is passed to `putenv`.

SEE ALSO

`execve(2)`, `getenv(3)`, `malloc(3)`, `environ(5V)`.

DIAGNOSTICS

`putenv()` returns non-zero if it was unable to obtain enough space using `malloc(3)` for an expanded environment, otherwise zero.

WARNINGS

`putenv()` manipulates the environment pointed to by *environ*, and can be used in conjunction with `getenv`. However, *envp* (the third argument to *main*) is not changed.

This routine uses `malloc(3)` to enlarge the environment.

After `putenv()` is called, environmental variables are not in alphabetical order.

A potential error is to call `putenv()` with an automatic variable as the argument, then exit the calling function while `string()` is still part of the environment.

NAME

putpwent – write password file entry

SYNOPSIS

```
#include <pwd.h>

int putpwent(p, f)
struct passwd *p;
FILE *f;
```

DESCRIPTION

putpwent() is the inverse of getpwent(3). Given a pointer to a passwd structure created by getpwent() (or getpwuid() or getpwnam), putpwent() writes a line on the stream *f*, which matches the format of lines in the password file */etc/passwd*.

FILES

/etc/passwd

SEE ALSO

getpwent(3)

DIAGNOSTICS

putpwent() returns non-zero if an error was detected during its operation, otherwise zero.

WARNING

The above routine uses *<stdio.h>*, which increases the size of programs, not otherwise using standard I/O, more than might be expected.

BUGS

This routine is of limited utility, since most password files are maintained as Yellow Pages files, and cannot be updated with this routine.

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() writes the NULL-terminated string pointed to by *s*, followed by a NEWLINE character, to the standard output stream **stdout**.

fputs() writes the NULL-terminated string pointed to by *s* to the named output stream.

Neither function writes the terminal SM NULL character.

DIAGNOSTICS

Both routines return EOF on error. This will happen if the routines try to write on a file that has not been opened for writing.

NOTES

puts() appends a NEWLINE while **fputs()** does not.

SEE ALSO

ferror(3S), fopen(3S), fread(3S), printf(3S),putc(3S)

NAME

pwdauth, grpauth – password authentication routines

SYNOPSIS

```
int pwdauth(user, password)
```

```
char *user;
```

```
char *password;
```

```
int grpauth(group, password)
```

```
char *group;
```

```
char *password;
```

DESCRIPTION

pwdauth() and **grpauth()** determine whether the given guess at a *password* is valid for the given *user* or *group*. If the *password* is valid, the functions return 0.

A *password* is valid if the password when encrypted matches the encrypted password in the appropriate file. For **pwdauth**, if the **password.adjunct** file exists, the encrypted password will be in either the local or the Yellow Pages version of that file. Otherwise, either the local or YP **passwd** file will be used. For **grpauth**, the **group.adjunct** file (if it exists) or the **group** file (otherwise) will be checked on the local machine and then using the YP. In all cases, the local files will be checked before the YP files. Also, if the adjunct files exist, the main file will never be used for authentication even if they include encrypted passwords.

Both **pwdauth()** and **grpauth()** interface to the authentication daemon, **rpc.pwdauthd**, to do the checking of the adjunct files. This daemon must be running on any system that provides password authentication.

FILES

/etc/passwd

/etc/group

SEE ALSO

getgraent(3), **getgrent(3)**, **getpwaent(3)**, **getpwent(3)**, **pwdauthd(8C)**

NAME

qsort – quicker sort

SYNOPSIS

```
qsort(base, nel, width, compar)
char *base;
int (*compar)();
```

DESCRIPTION

qsort() is an implementation of the quicker-sort algorithm. It sorts a table of data in place.

base points to the element at the base of the table. *nel* is the number of elements in the table. *width* is the size, in bytes, of each element in the table. *compar* is the name of the comparison function, which is called with two arguments that point to the elements being compared. As the function must return an integer less than, equal to, or greater than zero, so must the first argument to be considered be less than, equal to, or greater than the second.

NOTES

The pointer to the base of the table should be of type pointer-to-element, and cast to type pointer-to-character.

The comparison function need not compare every byte, so arbitrary data may be contained in the elements in addition to the values being compared.

The order in the output of two items which compare as equal is unpredictable.

SEE ALSO

sort(1V), **bsearch(3)**, **lsearch(3)**, **string(3)**

EXAMPLE

The following program sorts a simple array:

```
static int intcompare(i,j)
    int *i, *j;
{
    return(*i - *j);
}

main()
{
    int a[10];
    int i;

    a[0] = 9;
    a[1] = 8;
    a[2] = 7;
    a[3] = 6;
    a[4] = 5;
    a[5] = 4;
    a[6] = 3;
    a[7] = 2;
    a[8] = 1;
    a[9] = 0;

    qsort(a,10,sizeof(int),intcompare)
    for (i=0; i<10; i++) printf(" %d,a[i]);
    printf "0);
}
```

NAME

rand, srand – simple random number generator

SYNOPSIS

srand(seed)

int seed;

rand()

DESCRIPTION

rand() uses a multiplicative congruential random number generator with period 2^{32} to return successive pseudo-random numbers in the range from 0 to $2^{31}-1$.

srand() can be called at any time to reset the random-number generator to a random starting point. The generator is initially seeded with a value of 1.

NOTE

The spectral properties of **rand()** leave a great deal to be desired. **drand48(3)** and **random(3)** provide much better, though more elaborate, random-number generators.

SEE ALSO

drand48(3), random(3), rand(3V)

BUGS

The low bits of the numbers generated are not very random; use the middle bits. In particular the lowest bit alternates between 0 and 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 $16 \times (2^{31}-1)$.

random/srandom have (almost) the same calling sequence and initialization properties as **rand/srand**. The difference is that **rand(3C)** 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(3C)**, 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 previous state array; its 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.

SEE ALSO

rand(3C)

EXAMPLE

```

/* Initialize and array and pass it in to initState. */
static long state1[32] = {
    3,
    0x9a319039, 0x32d9c024, 0x9b663182, 0x5da1f342,
    0x7449e56b, 0xbcb1dbb0, 0xab5c5918, 0x946554fd,
    0x8c2e680f, 0xeb3d799f, 0xb11ee0b7, 0x2d436b86,
    0xda672e2a, 0x1588ca88, 0xe369735d, 0x904f35f7,
    0xd7158fd6, 0x6fa6f051, 0x616e6b96, 0xac94efdc,
    0xde3b81e0, 0xdf0a6fb5, 0xf103bc02, 0x48f340fb,
    0x36413f93, 0xc622c298, 0xf5a42ab8, 0x8a88d77b,
    0xf5ad9d0e, 0x8999220b, 0x27fb47b9
};

main()
{
    unsigned seed;
    int n;

    seed = 1;
    n = 128;
    initState(seed, state1, n);

    setstate(state1);
    printf("%d0,random());
}

```

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.

BUGS

About 2/3 the speed of `rand(3C)`.

NAME

rcmd, rresvport, ruserok – routines for returning a stream to a remote command

SYNOPSIS

```
int rcmd(ahost, inport, locuser, remuser, cmd, fd2p)
char **ahost;
int inport;
char *locuser, *remuser, *cmd;
int *fd2p

int rresvport(port)
int *port;

ruserok(rhost, super-user, ruser, luser)
char *rhost;
int super-user;
char *ruser, *luser;
```

DESCRIPTION

rcmd() is a routine used by the super-user to execute a command on a remote machine using an authentication scheme based on reserved port numbers. **rresvport()** is a routine which returns a descriptor to a socket with an address in the privileged port space. **ruserok()** is a routine used by servers to authenticate clients requesting service with **rcmd**. All three functions are present in the same file and are used by the **rshd(8C)** server (among others).

rcmd() looks up the host **ahost* using **gethostbyname** (see **gethostent(3N)**), returning **-1** if the host does not exist. Otherwise **ahost* is set to the standard name of the host and a connection is established to a server residing at the well-known Internet port *inport*.

If the connection succeeds, a socket in the Internet domain of type **SOCK_STREAM** is returned to the caller, and given to the remote command as its standard input (file descriptor 0) and standard output (file descriptor 1). If *fd2p* is non-zero, then an auxiliary channel to a control process will be set up, and a descriptor for it will be placed in **fd2p*. The control process will return diagnostic output from the command (file descriptor 2) on this channel, and will also accept bytes on this channel as signal numbers, to be forwarded to the process group of the command. If *fd2p* is 0, then the standard error (file descriptor 2) of the remote command will be made the same as its standard output and no provision is made for sending arbitrary signals to the remote process, although you may be able to get its attention by using out-of-band data.

The protocol is described in detail in **rshd(8C)**.

The **rresvport()** routine is used to obtain a socket with a privileged address bound to it. This socket is suitable for use by **rcmd()** and several other routines. Privileged Internet ports are those in the range 0 to 1023. Only the super-user is allowed to bind an address of this sort to a socket.

ruserok() takes a remote host's name, as returned by a **gethostbyaddr** (see **gethostent(3N)**) routine, two user names and a flag indicating whether the local user's name is that of the super-user. It then checks the files **/etc/hosts.equiv** and, possibly, **.rhosts** in the local user's home directory to see if the request for service is allowed. A 0 is returned if the machine name is listed in the **/etc/hosts.equiv** file, or the host and remote user name are found in the **.rhosts** file; otherwise **ruserok()** returns **-1**. If the super-user flag is 1, the checking of the **/etc/hosts.equiv** file is bypassed.

FILES

```
/etc/hosts.equiv
.rhosts
```

SEE ALSO

rlogin(1C), **rsh(1C)**, **intro(2)**, **gethostent(3N)**, **rexec(3N)**, **rexecd(8C)**, **rlogind(8C)**, **rshd(8C)**

DIAGNOSTICS

rcmd() returns a valid socket descriptor on success. It returns -1 on error and prints a diagnostic message on the standard error.

rresvport() returns a valid, bound socket descriptor on success. It returns -1 on error with the global value **errno** set according to the reason for failure. The error code **EAGAIN** is overloaded to mean "All network ports in use."

NAME

`realpath` – returns the real file name.

SYNOPSIS

```
char *realpath(file_name, resolved_name)
char *file_name, resolved_name;
```

DESCRIPTION

`realpath()` resolves all links and all references to "." and ".." in *file_name* and stores it in *resolved_name*.

It can handle both relative and absolute path names.

RETURN VALUE

If there is no error, it returns a pointer to the *resolved_name*. Otherwise it returns a NULL pointer and places the name of the offending file in *resolved_name*.

SEE ALSO

`getwd(3)`

WARNINGS

It operates on null-terminated strings.

It does not check for overflow of the receiving string.

NAME

regex, 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 a NULL pointer 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 NULL characters. The regular expressions recognized are described in the manual entry for **ed(1)**, given the above difference.

SEE ALSO

ed(1), ex(1), grep(1V)

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

regex – regular expression compile and match routines

SYNOPSIS

```
#define INT <declarations>
#define GETC() <getc code>
#define PEEKC() <peekc code>
#define UNGETC(c) <ungetc code>
#define RETURN(pointer) <return code>
#define ERROR(val) <error code>

#include <regex.h>

char *compile(instring, expbuf, endbuf, eof)
char *instring, *expbuf, *endbuf;
int eof;

int step(string, expbuf)
char *string, *expbuf;

extern char *loc1, *loc2, *locs;
extern int circf, sed, nbra;
```

DESCRIPTION

This page describes general-purpose regular expression matching routines.

The interface to this file is unpleasantly complex. Programs that include this file must have the following five macros declared before the '#include <regex.h>' statement. These macros are used by the *compile* routine.

GETC() Return the value of the next character in the regular expression pattern. Successive calls to GETC() should return successive characters of the regular expression.

PEEKC() Return the next character in the regular expression. Successive calls to PEEKC() should return the same character, which should also be the next character returned by GETC().

UNGETC(c) Returns the argument *c* by the next call to GETC() or PEEKC(). No more than one character of pushback is ever needed and this character is guaranteed to be the last character read by GETC(). The value of the macro UNGETC(*c*) is always ignored.

RETURN(pointer) This macro is used on normal exit of the *compile* routine. The value of the argument *pointer* is a pointer to the character after the last character of the compiled regular expression. This is useful to programs that have memory allocation to manage.

ERRORS

ERROR(val) This is the abnormal return from the *compile()* routine. The argument *val* is an error number (see table below for meanings). This call should never return.

ERROR	MEANING
11	Range endpoint too large.
16	Bad number.
25	“\ digit” out of range.
36	Illegal or missing delimiter.
41	No remembered search string.
42	\(\) imbalance.
43	Too many \(.

44	More than 2 numbers given in \{ \}.
45	} expected after \.
46	First number exceeds second in \{ \}.
49	[] imbalance.
50	Regular expression too long.

The syntax of the `compile()` routine is as follows:

`compile(instring, expbuf, endbuf, eof)`

The first parameter *instring* is never used explicitly by the `compile()` routine but is useful for programs that pass down different pointers to input characters. It is sometimes used in the `INIT()` declaration (see below). Programs that call functions to input characters or have characters in an external array can pass down a value of `((char *) 0)` for this parameter.

The next parameter *expbuf* is a character pointer. It points to the place where the compiled regular expression will be placed.

The parameter *endbuf* is one more than the highest address where the compiled regular expression may be placed. If the compiled expression cannot fit in `(endbuf-expbuf)` bytes, a call to `ERROR(50)` is made.

The parameter *eof* is the character that marks the end of the regular expression. For example, in an editor like `ed(1)`, this character would usually a `'/'`.

Each program that includes this file must have a `#define` statement for `INIT()`. This definition will be placed right after the declaration for the function `compile()` and `{` (opening curly brace). It is used for dependent declarations and initializations. Most often it is used to set a register variable to point the beginning of the regular expression so that this register variable can be used in the declarations for `GETC()`, `PEEKC()`, and `UNGETC()`. Otherwise it can be used to declare external variables that might be used by `GETC()`, `PEEKC()`, and `UNGETC()`. See the example below of the declarations taken from `grep(1V)`.

There are other functions in this file that perform actual regular expression matching, one of which is the function `step()`. The call to `step()` is as follows:

`step(string, expbuf)`

The first parameter to `step()` is a pointer to a string of characters to be checked for a match. This string should be `NULL`-terminated.

The second parameter *expbuf* is the compiled regular expression that was obtained by a call of the function `compile`.

The function `step()` returns non-zero if the given string matches the regular expression, and zero if the expressions do not match. If there is a match, two external character pointers are set as a side effect to the call to `step()`. The variable set in `step()` is *loc1*. This is a pointer to the first character that matched the regular expression. The variable *loc2*, which is set by the function `advance()`, points to the character after the last character that matches the regular expression. Thus if the regular expression matches the entire line, *loc1* will point to the first character of *string* and *loc2* will point to the `NULL` at the end of *string*.

`step()` uses the external variable `circf` which is set by `compile()` if the regular expression begins with `^`. If this is set then `step()` will try to match the regular expression to the beginning of the string only. If more than one regular expression is to be compiled before the first is executed the value of `circf` should be saved for each compiled expression and `circf` should be set to that saved value before each call to `step()`.

The function `advance()` is called from `step()` with the same arguments as `step()`. The purpose of `step()` is to step through the *string* argument and call `advance()` until `advance()` returns non-zero indicating a match or until the end of *string* is reached. If one wants to constrain *string* to the beginning of the line in all cases, `step()` need not be called; simply call `advance()`.

When `advance()` encounters a `*` or `\{ \}` sequence in the regular expression, it will advance its pointer to the string to be matched as far as possible and will recursively call itself trying to match the rest of the string to the rest of the regular expression. As long as there is no match, `advance()` will back up along the string until it finds a match or reaches the point in the string that initially matched the `*` or `\{ \}`. It is sometimes desirable to stop this backing up before the initial point in the string is reached. If the external character pointer `loc` is equal to the point in the string at sometime during the backing up process, `advance()` will break out of the loop that backs up and will return zero. This could be used by an editor like `ed(1)` or `sed(1)` for substitutions done globally (not just the first occurrence, but the whole line) so, for example, expressions like `s/y*//g` do not loop forever.

The additional external variables `sed` and `nbra` are used for special purposes.

EXAMPLES

The following is an example of how the regular expression macros and calls could look in a command like `grep(1V)`:

```
#define INIT    register char *sp = instring;
#define GETC() (*sp++)
#define PEEKC()  (*sp)
#define UNGETC(c)  (—sp)
#define RETURN(c)  return;
#define ERROR(c)  regerr()

#include <regex.h>
...
                (void) compile(*argv, expbuf, &expbuf[ESIZE], ^0');
...
                if (step(linebuf, expbuf)
                    succeed ());
```

FILES

`/usr/include/regex.h`

SEE ALSO

`ed(1)`, `grep(1V)`, `sed(1V)`

BUGS

The handling of `circf` is difficult.

NAME

resolver, res_mkquery, res_send, res_init, dn_comp, dn_expand – resolver routines

SYNOPSIS

```
#include <sys/types.h>
#include <netinet/in.h>
#include <arpa/nameser.h>
#include <arpa/resolv.h>

res_mkquery(op, dname, class, type, data, datalen, newrr, buf, buflen)
int op;
char *dname;
int class, type;
char *data;
int datalen;
struct rrec *newrr;
char *buf;
int buflen;

res_send(msg, msglen, answer, anslen)
char *msg;
int msglen;
char *answer;
int anslen;

res_init()

dn_comp(exp_dn, comp_dn, length, dnptrs, lastdnptr)
char *exp_dn, *comp_dn;
int length;
char **dnptrs, **lastdnptr;

dn_expand(msg, msglen, comp_dn, exp_dn, length)
char *msg, *comp_dn, exp_dn;
int msglen, length;
```

DESCRIPTION

These routines are used for making, sending and interpreting packets to Internet domain name servers. Global information that is used by the resolver routines is kept in the variable *_res*. Most of the values have reasonable defaults and can be ignored. Options stored in *_res.options* are defined in *resolv.h* and are as follows. Options are a simple bit mask and are OR'ed in to enable.

RES_INIT	True if the initial name server address and default domain name are initialized (that is, <i>res_init</i> has been called).
RES_DEBUG	Print debugging messages.
RES_AAONLY	Accept authoritative answers only. <i>res_send</i> will continue until it finds an authoritative answer or finds an error. Currently this is not implemented.
RES_USEVC	Use TCP connections for queries instead of UDP.
RES_STAYOPEN	Used with RES_USEVC to keep the TCP connection open between queries. This is useful only in programs that regularly do many queries. UDP should be the normal mode used.
RES_IGNTC	Unused currently (ignore truncation errors, that is, do not retry with TCP).
RES_RECURSE	Set the recursion desired bit in queries. This is the default. (<i>res_send</i> does not do iterative queries and expects the name server to handle recursion.)
RES_DEFNAMES	Append the default domain name to single label queries. This is the default.

res_init reads the initialization file to get the default domain name and the Internet address of the initial hosts running the name server. If this line does not exist, the host running the resolver is tried. *res_mkquery* makes a standard query message and places it in *buf*. *res_mkquery* will return the size of the query or -1 if the query is larger than *buflen*. *op* is usually QUERY but can be any of the query types defined in *nameser.h*. *dname* is the domain name. If *dname* consists of a single label and the RES_DEFNAMES flag is enabled (the default), *dname* will be appended with the current domain name. The current domain name is defined in a system file and can be overridden by the environment variable LOCALDOMAIN. *newrr* is currently unused but is intended for making update messages.

res_send sends a query to name servers and returns an answer. It will call *res_init* if RES_INIT is not set, send the query to the local name server, and handle timeouts and retries. The length of the message is returned or -1 if there were errors.

dn_expand Expands the compressed domain name *comp_dn* to a full domain name. Expanded names are converted to upper case. *msg* is a pointer to the beginning of the message, *exp_dn* is a pointer to a buffer of size *length* for the result. The size of compressed name is returned or -1 if there was an error.

dn_comp Compresses the domain name *exp_dn* and stores it in *comp_dn*. The size of the compressed name is returned or -1 if there were errors. *length* is the size of the array pointed to by *comp_dn*. *dnptrs* is a list of pointers to previously compressed names in the current message. The first pointer points to the beginning of the message and the list ends with NULL. *lastdnptr* is a pointer to the end of the array pointed to *dnptrs*. A side effect is to update the list of pointers for labels inserted into the message by *dn_comp* as the name is compressed. If *dnptr* is NULL, we do not try to compress names. If *lastdnptr* is NULL, we do not update the list.

FILES

/etc/resolve.conf see *resolve.conf(5)*

SEE ALSO

resolve.conf(5), *named(8)*

NAME

rexec – return stream to a remote command

SYNOPSIS

```
rem = rexec(ahost, inport, user, passwd, cmd, fd2p);
char **ahost;
u_short inport;
char *user, *passwd, *cmd;
int *fd2p;
```

DESCRIPTION

rexec() looks up the host *ahost* using `gethostbyname` (see `gethostent(3N)`), returning `-1` if the host does not exist. Otherwise *ahost* is set to the standard name of the host. If a username and password are both specified, then these are used to authenticate to the foreign host; otherwise the environment and then the user's `.netrc` file in his home directory are searched for appropriate information. If all this fails, the user is prompted for the information.

The port *inport* specifies which well-known DARPA Internet port to use for the connection; it will normally be the value returned from the call `getservbyname("exec", "tcp")` (see `getservent(3N)`). The protocol for connection is described in detail in `rexecd(8C)`.

If the call succeeds, a socket of type `SOCK_STREAM` is returned to the caller, and given to the remote command as its standard input and standard output. If *fd2p* is non-zero, then a auxiliary channel to a control process will be setup, and a descriptor for it will be placed in *fd2p*. The control process will return diagnostic output from the command (unit 2) on this channel, and will also accept bytes on this channel as signal numbers, to be forwarded to the process group of the command. If *fd2p* is 0, then the standard error (unit 2 of the remote command) will be made the same as its standard output and no provision is made for sending arbitrary signals to the remote process, although you may be able to get its attention by using out-of-band data.

SEE ALSO

`gethostent(3N)`, `getservent(3N)`, `rcmd(3N)`, `rexecd(8C)`

BUGS

There is no way to specify options to the `socket()` call that `rexec()` makes.

NAME

rpc – library routines for remote procedure calls

SYNOPSIS AND DESCRIPTION

These routines allow C programs to make procedure calls on other machines across the network. First, the client calls a procedure to send a data packet to the server. Upon receipt of the packet, the server calls a dispatch routine to perform the requested service, and then sends back a reply. Finally, the procedure call returns to the client.

```
#include <rpc/rpc.h>
```

```
void
```

```
auth_destroy(auth)
```

```
AUTH *auth;
```

A macro that destroys the authentication information associated with *auth*. Destruction usually involves deallocation of private data structures. The use of *auth* is undefined after calling `auth_destroy()`.

```
AUTH *
```

```
authnone_create()
```

Create and returns an RPC authentication handle that passes nonusable authentication information with each remote procedure call. This is the default authentication used by RPC.

```
AUTH *
```

```
authdes_create(name, window, syncaddr, ckey)
```

```
char *name;
```

```
unsigned window;
```

```
struct sockaddr *addr;
```

```
des_block *ckey;
```

`authdes_create()` is the first of two routines which interface to the RPC secure authentication system, known as DES authentication. The second is `authdes_getucred()`, below. Note: the keyserver daemon `keyserv(8C)` must be running for the DES authentication system to work.

`authdes_create()`, used on the client side, returns an authentication handle that will enable the use of the secure authentication system. The first parameter *name* is the network name, or *netname*, of the owner of the server process. This field usually represents a *hostname* derived from the utility routine `host2netname`, but could also represent a user name using `user2netname`. The second field is window on the validity of the client credential, given in seconds. A small window is more secure than a large one, but choosing too small of a window will increase the frequency of resynchronizations because of clock drift. The third parameter *syncaddr* is optional. If it is NULL, then the authentication system will assume that the local clock is always in sync with the server's clock, and will not attempt resynchronizations. If an address is supplied, however, then the system will use the address for consulting the remote time service whenever resynchronization is required. This parameter is usually the address of the RPC server itself. The final parameter *ckey* is also optional. If it is NULL, then the authentication system will generate a random DES key to be used for the encryption of credentials. If it is supplied, however, then it will be used instead.

```

authdes_getucrd(adc, uid, gid, grouplen, groups)
struct authdes_cred *adc;
short *uid;
short *gid;
short *grouplen;
int *groups;

```

authdes_getucrd(), the second of the two DES authentication routines, is used on the server side for converting a DES credential, which is operating system independent, into a credential. This routine differs from utility routine **netname2user** in that **authdes_getucrd()** pulls its information from a cache, and does not have to do a Yellow Pages lookup every time it is called to get its information.

```

AUTH *
authunix_create(host, uid, gid, len, aup_gids)
char *host;
int uid, gid, len, *aup.gids;

```

Create and return an RPC authentication handle that contains authentication information. The parameter *host* is the name of the machine on which the information was created; *uid* is the user's user ID ; *gid* is the user's current group ID ; *len* and *aup_gids* refer to a counted array of groups to which the user belongs. It is easy to impersonate a user.

```

AUTH *
authunix_create_default()

```

Calls **authunix_create()** with the appropriate parameters.

```

callrpc(host, prognum, versnum, procnum, inproc, in, outproc, out)
char *host;
u_long prognum, versnum, procnum;
char *in, *out;
xdrproc_t inproc, outproc;

```

Call the remote procedure associated with *prognum*, *versnum*, and *procnum* on the machine, *host*. The parameter *in* is the address of the procedure's argument(s), and *out* is the address of where to place the result(s); *inproc* is used to encode the procedure's parameters, and *outproc* is used to decode the procedure's results. This routine returns zero if it succeeds, or the value of **enum clnt_stat** cast to an integer if it fails. The routine **clnt_perrno()** is handy for translating failure statuses into messages.

Warning: calling remote procedures with this routine uses UDP/IP as a transport; see **clntudp_create()** for restrictions. You do not have control of timeouts or authentication using this routine.

```
enum clnt_stat
clnt_broadcast(prognum, versnum, procnum, inproc, in, outproc, out, eachresult)
u_long prognum, versnum, procnum;
char *in, *out;
xdrproc_t inproc, outproc;
resultproc_t eachresult;
```

Like `callrpc()`, except the call message is broadcast to all locally connected broadcast nets. Each time it receives a response, this routine calls `eachresult()`, whose form is:

```
eachresult(out, addr)
char *out;
struct sockaddr_in *addr;
```

where `out` is the same as `out` passed to `clnt_broadcast()`, except that the remote procedure's output is decoded there; `addr` points to the address of the machine that sent the results. If `eachresult()` returns zero, `clnt_broadcast()` waits for more replies; otherwise it returns with appropriate status.

Warning: broadcast sockets are limited in size to the maximum transfer unit of the data link. For ethernet, this value is 1500 bytes.

```
enum clnt_stat
clnt_call(clnt, procnum, inproc, in, outproc, out, tout)
CLIENT *clnt;
u_long
procnum;
xdrproc_t inproc, outproc;
char *in, *out;
struct timeval tout;
```

A macro that calls the remote procedure `procnum` associated with the client handle, `clnt`, which is obtained with an RPC client creation routine such as `clnt_create()`. The parameter `in` is the address of the procedure's argument(s), and `out` is the address of where to place the result(s); `inproc` is used to encode the procedure's parameters, and `outproc` is used to decode the procedure's results; `tout` is the time allowed for results to come back.

```
clnt_destroy(clnt)
CLIENT *clnt;
```

A macro that destroys the client's RPC handle. Destruction usually involves deallocation of private data structures, including `clnt` itself. Use of `clnt` is undefined after calling `clnt_destroy()`. If the RPC library opened the associated socket, it will close it also. Otherwise, the socket remains open.

```
CLIENT *
clnt_create(host, prog, vers, proto)
char *host;
u_long prog, vers;
char *proto;
```

Generic client creation routine. `host` identifies the name of the remote host where the server is located. `proto` indicates which kind of transport protocol to use. The currently supported values for this field are "udp" and "tcp". Default timeouts are set, but can be modified using `clnt_control()`.

Warning: Using UDP has its shortcomings. Since UDP-based RPC messages can only hold up to 8 Kbytes of encoded data, this transport cannot be used for procedures that take large arguments or return huge results.

```

bool_t
clnt_control(cl, req, info)
CLIENT *cl;
char *info;

```

A macro used to change or retrieve various information about a client object. *req* indicates the type of operation, and *info* is a pointer to the information. For both UDP and TCP, the supported values of *req* and their argument types and what they do are:

CLSET_TIMEOUT	struct timeval	set total timeout
CLGET_TIMEOUT	struct timeval	get total timeout

Note: if you set the timeout using `clnt_control()`, the timeout parameter passed to `clnt_call()` will be ignored in all future calls.

CLGET_SERVER_ADDR	struct sockaddr	get server's address
-------------------	-----------------	----------------------

The following operations are valid for UDP only:

CLSET_RETRY_TIMEOUT	struct timeval	set the retry timeout
CLGET_RETRY_TIMEOUT	struct timeval	get the retry timeout

The retry timeout is the time that UDP RPC waits for the server to reply before retransmitting the request.

```

clnt_freeres(clnt, outproc, out)
CLIENT *clnt;
xdrproc_t outproc;
char *out;

```

A macro that frees any data allocated by the RPC/XDR system when it decoded the results of an RPC call. The parameter *out* is the address of the results, and *outproc* is the XDR routine describing the results. This routine returns one if the results were successfully freed, and zero otherwise.

```

void
clnt_geterr(clnt, errp)
CLIENT *clnt;
struct rpc_err *errp;

```

A macro that copies the error structure out of the client handle to the structure at address *errp*.

```

void
clnt_pcreateerror(s)
char *s;

```

Print a message to standard error indicating why a client RPC handle could not be created. The message is prepended with string *s* and a colon. Used when a `clnt_create()`, `clntraw_create()`, `clnttcp_create()`, or `clntudp_create()` call fails.

```

void
clnt_perrno(stat)
enum clnt_stat stat;
    Print a message to standard error corresponding to the condition indicated by stat. Used after
    callrpc().

clnt_perror(clnt, s)
CLIENT *clnt;
char *s;
    Print a message to standard error indicating why an RPC call failed; clnt is the handle used to
    do the call. The message is prepended with string s and a colon. Used after clnt_call().

char *
clnt_spcreateerror
char *s;
    Like clnt_pcreateerror(), except that it returns a string instead of printing to the standard
    error.
    Bugs: returns pointer to static data that is overwritten on each call.

char *
clnt_sperrno(stat)
enum clnt_stat stat;
    Take the same arguments as clnt_perrno(), but instead of sending a message to the standard
    error indicating why an RPC call failed, return a pointer to a string which contains the mes-
    sage. The string ends with a NEWLINE.
    clnt_sperrno() is used instead of clnt_perrno() if the program does not have a standard
    error (as a program running as a server quite likely does not), or if the programmer does not
    want the message to be output with printf, or if a message format different than that sup-
    ported by clnt_perrno() is to be used. Note: unlike clnt_sperror() and clnt_spcreaterror(),
    clnt_sperrno() does not return pointer to static data so the result will not get overwritten on
    each call.

char *
clnt_sperror(rpch, s)
CLIENT *rpch;
char *s;
    Like clnt_perror(), except that (like clnt_sperrno()) it returns a string instead of printing to
    standard error.
    Bugs: returns pointer to static data that is overwritten on each call.

CLIENT *
clntraw_create(prognum, versnum)
u_long prognum, versnum;
    This routine creates a toy RPC client for the remote program prognum, version versnum. The
    transport used to pass messages to the service is actually a buffer within the process's address
    space, so the corresponding RPC server should live in the same address space; see
    svcrw_create(). This allows simulation of RPC and acquisition of RPC overheads, such as
    round trip times, without any kernel interference. This routine returns NULL if it fails.

```


CLIENT *

```

clnttcp_create(addr, prognum, versnum, sockp, sendsz, recvsz)
struct sockaddr_in *addr;
u_long prognum, versnum;
int *sockp;
u_int sendsz, recvsz;

```

This routine creates an RPC client for the remote program *prognum*, version *versnum*; the client uses TCP/IP as a transport. The remote program is located at Internet address **addr*. If *addr->sin_port* is zero, then it is set to the actual port that the remote program is listening on (the remote *portmap* service is consulted for this information). The parameter *sockp* is a socket; if it is *RPC_ANYSOCK*, then this routine opens a new one and sets *sockp*. Since TCP-based RPC uses buffered I/O, the user may specify the size of the send and receive buffers with the parameters *sendsz* and *recvsz*; values of zero choose suitable defaults. This routine returns NULL if it fails.

CLIENT *

```

clntudp_create(addr, pronum, versnum, wait, sockp)
struct sockaddr_in *addr;
u_long prognum, versnum;
struct timeval wait;
int *sockp;

```

This routine creates an RPC client for the remote program *prognum*, version *versnum*; the client uses UDP/IP as a transport. The remote program is located at Internet address *addr*. If *addr->sin_port* is zero, then it is set to actual port that the remote program is listening on (the remote *portmap* service is consulted for this information). The parameter *sockp* is a socket; if it is *RPC_ANYSOCK*, then this routine opens a new one and sets *sockp*. The UDP transport resends the call message in intervals of wait time until a response is received or until the call times out. The total time for the call to time out is specified by *clnt_call()*.

Warning: since UDP-based RPC messages can only hold up to 8 Kbytes of encoded data, this transport cannot be used for procedures that take large arguments or return huge results.

```

host2netname(name, host, domain)
char *name;
char *host;
char *domain;

```

Convert from a domain-specific hostname to an operating-system independent netname. Return TRUE if it succeeds and FALSE if it fails. Inverse of *netname2host()*.

```

key_decryptsession(remotename, deskey)
char *remotename;
des_block *deskey;

```

key_decryptsession() is an interface to the keyserver daemon, which is associated with RPC's secure authentication system (DES authentication). User programs rarely need to call it, or its associated routines *key_encryptsession()*, *key_gendes()* and *key_setsecret()*. System commands such as *login* and the RPC library are the main clients of these four routines.

key_decryptsession() takes a server netname and a des key, and decrypts the key by using the the public key of the the server and the secret key associated with the effective uid of the calling process. It is the inverse of *key_encryptsession()*.

```
key_encryptsession(remotename, deskey)  
char *remotename;  
des_block *deskey;
```

key_encryptsession() is a keyserver interface routine. It takes a server netname and a des key, and encrypts it using the public key of the the server and the secret key associated with the effective uid of the calling process. It is the inverse of **key_decryptsession()**.

```
key_gendes(deskey)  
des_block *deskey;
```

key_gendes() is a keyserver interface routine. It is used to ask the keyserver for a secure conversation key. Choosing one at "random" is usually not good enough, because the common ways of choosing random numbers, such as using the current time, are very easy to guess.

```
key_setsecret(key)  
char *key;
```

key_setsecret() is a keyserver interface routine. It is used to set the key for the effective *uid* of the calling process.

```
void  
get_myaddress(addr)  
struct sockaddr_in *addr;
```

Stuff the machine's IP address into **addr*, without consulting the library routines that deal with */etc/hosts*. The port number is always set to **htons(PMAPPORT)**.

```
getnetname(name)  
char name[MAXNETNAMELEN];
```

getnetname() installs the unique, operating-system independent netname of the caller in the fixed-length array *name*. Returns TRUE if it succeeds and FALSE if it fails.

```
netname2host(name, host, hostlen)  
char *name;  
char *host;  
int hostlen;
```

Convert from an operating-system independent netname to a domain-specific hostname. Returns TRUE if it succeeds and FALSE if it fails. Inverse of **host2netname()**.

```
netname2user(name, uidp, gidp, gidlenp, gidlist)  
char *name;  
int *uidp;  
int *gidp;  
int *gidlenp;  
int *gidlist;
```

Convert from an operating-system independent netname to a domain-specific user ID. Returns TRUE if it succeeds and FALSE if it fails. Inverse of **user2netname()**.

```

struct pmaplist *
pmap_getmaps(addr)
struct sockaddr_in *addr;

```

A user interface to the portmap service, which returns a list of the current RPC program-to-port mappings on the host located at IP address **addr*. This routine can return NULL. The command 'rpcinfo -p' uses this routine.

```

u_short
pmap_getport(addr, prognum, versnum, protocol)
struct sockaddr_in *addr;
u_long prognum, versnum, protocol;

```

A user interface to the portmap service, which returns the port number on which waits a service that supports program number *prognum*, version *versnum*, and speaks the transport protocol associated with *protocol*. The value of *protocol* is most likely IPPROTO_UDP or IPPROTO_TCP. A return value of zero means that the mapping does not exist or that the RPC system failed to contact the remote portmap service. In the latter case, the global variable *rpc_createerr()* contains the RPC status.

```

enum clnt_stat
pmap_rmtcall(addr, prognum, versnum, procnum, inproc, in, outproc, out, tout, portp)
struct sockaddr_in *addr;
u_long prognum, versnum, procnum;
char *in, *out;
xdrproc_t inproc, outproc;
struct timeval tout;
u_long *portp;

```

A user interface to the portmap service, which instructs portmap on the host at IP address **addr* to make an RPC call on your behalf to a procedure on that host. The parameter **portp* will be modified to the program's port number if the procedure succeeds. The definitions of other parameters are discussed in *callrpc()* and *clnt_call()*. This procedure should be used for a "ping" and nothing else. See also *clnt_broadcast()*.

```

pmap_set(prognum, versnum, protocol, port)
u_long prognum, versnum, protocol;
u_short port;

```

A user interface to the portmap service, which establishes a mapping between the triple [*prognum,versnum,protocol*] and *port* on the machine's portmap service. The value of *protocol* is most likely IPPROTO_UDP or IPPROTO_TCP. This routine returns one if it succeeds, zero otherwise. Automatically done by *svc_register()*.

```

pmap_unset(prognum, versnum)
u_long prognum, versnum;

```

A user interface to the portmap service, which destroys all mapping between the triple [*prognum,versnum,**] and ports on the machine's portmap service. This routine returns one if it succeeds, zero otherwise.

```

registerrpc(prognum, versnum, procnum, procname, inproc, outproc)
u_long prognum, versnum, procnum;
char *(*procname) ();
xdrproc_t inproc, outproc;

```

Register procedure *procname* with the RPC service package. If a request arrives for program *prognum*, version *versnum*, and procedure *procnum*, *procname* is called with a pointer to its parameter(s); *progname* should return a pointer to its static result(s); *inproc* is used to decode the parameters while *outproc* is used to encode the results. This routine returns zero if the registration succeeded, -1 otherwise.

Warning: remote procedures registered in this form are accessed using the UDP/IP transport; see `svcudp_create()` for restrictions.

```

struct rpc_createerr    rpc_createerr;

```

A global variable whose value is set by any RPC client creation routine that does not succeed. Use the routine `clnt_pcreateerror()` to print the reason why.

```

svc_destroy(xpirt)
SVCXPRT *
xpirt;

```

A macro that destroys the RPC service transport handle, *xprt*. Destruction usually involves deallocation of private data structures, including *xprt* itself. Use of *xprt* is undefined after calling this routine.

```

fd_set svc_fdset;

```

A global variable reflecting the RPC service side's read file descriptor bit mask; it is suitable as a parameter to the `select` system call. This is only of interest if a service implementor does not call `svc_run()`, but rather does his own asynchronous event processing. This variable is read-only (do not pass its address to `select!`), yet it may change after calls to `svc_getreqset()` or any creation routines.

```

int svc_fds;

```

Similar to `svc_fedset()`, but limited to 32 descriptors. This interface is obsoleted by `svc_fdset()`.

```

svc_freeargs(xpirt, inproc, in)
SVCXPRT *xpirt;
xdrproc_t inproc;
char *in;

```

A macro that frees any data allocated by the RPC/XDR system when it decoded the arguments to a service procedure using `svc_getargs()`. This routine returns 1 if the results were successfully freed, and zero otherwise.

```

svc_getargs(xpirt, inproc, in)
SVCXPRT *xpirt;
xdrproc_t inproc;
char *in;

```

A macro that decodes the arguments of an RPC request associated with the RPC service transport handle, *xprt*. The parameter *in* is the address where the arguments will be placed; *inproc* is the XDR routine used to decode the arguments. This routine returns one if decoding succeeds, and zero otherwise.

```

struct sockaddr_in *
svc_getcaller(xprt)
SVCXPRT *xprt;

```

The approved way of getting the network address of the caller of a procedure associated with the RPC service transport handle, *xprt*.

```

svc_getreqset(rdfds)
fd_set *rdfds;

```

This routine is only of interest if a service implementor does not call *svc_run()*, but instead implements custom asynchronous event processing. It is called when the *select* system call has determined that an RPC request has arrived on some RPC socket(s); *rdfds* is the resultant read file descriptor bit mask. The routine returns when all sockets associated with the value of *rdfds* have been serviced.

```

svc_getreq(rdfds)
int rdfds;

```

Similar to *svc_getreqset()*, but limited to 32 descriptors. This interface is obsoleted by *svc_getreqset()*.

```

svc_register(xprt, prognum, versnum, dispatch, protocol)
SVCXPRT *xprt;
u_long prognum, versnum;
void (*dispatch) ();
u_long protocol;

```

Associates *prognum* and *versnum* with the service dispatch procedure, *dispatch*. If *protocol* is zero, the service is not registered with the portmap service. If *protocol* is non-zero, then a mapping of the triple [*prognum,versnum,protocol*] to *xprt->xp_port* is established with the local portmap service (generally *protocol* is zero, *IPPROTO_UDP* or *IPPROTO_TCP*). The procedure *dispatch* has the following form:

```

dispatch(request, xprt)
struct svc_req *request;
SVCXPRT *xprt;

```

The *svc_register()* routine returns one if it succeeds, and zero otherwise.

```

svc_run()

```

This routine never returns. It waits for RPC requests to arrive, and calls the appropriate service procedure using *svc_getreq()* when one arrives. This procedure is usually waiting for a *select()* system call to return.

```

svc_sendreply(xprt, outproc, out)
SVCXPRT *xprt;
xdrproc_t outproc;
char *out;

```

Called by an RPC service's dispatch routine to send the results of a remote procedure call. The parameter *xprt* is the request's associated transport handle; *outproc* is the XDR routine which is used to encode the results; and *out* is the address of the results. This routine returns one if it succeeds, zero otherwise.

void
svc_unregister(prognum, versnum)
u_long prognum, versnum;

Remove all mapping of the double [*prognum,versnum*] to dispatch routines, and of the triple [*prognum,versnum,**] to port number.

void
svcerr_auth(xprt, why)
SVCXPRT *xprt;
enum auth_stat why;

Called by a service dispatch routine that refuses to perform a remote procedure call due to an authentication error.

void
svcerr_decode(xprt)
SVCXPRT *xprt;

Called by a service dispatch routine that cannot successfully decode its parameters. See also **svc_getargs()**.

void
svcerr_noproc(xprt)
SVCXPRT *xprt;

Called by a service dispatch routine that does not implement the procedure number that the caller requests.

void
svcerr_noprogram(xprt)
SVCXPRT *xprt;

Called when the desired program is not registered with the RPC package. Service implementors usually do not need this routine.

void
svcerr_progvers(xprt)
SVCXPRT *xprt;

Called when the desired version of a program is not registered with the RPC package. Service implementors usually do not need this routine.

void
svcerr_systemerr(xprt)
SVCXPRT *xprt;

Called by a service dispatch routine when it detects a system error not covered by any particular protocol. For example, if a service can no longer allocate storage, it may call this routine.

void
svcerr_weakauth(xprt)
SVCXPRT *xprt;

Called by a service dispatch routine that refuses to perform a remote procedure call due to insufficient (but correct) authentication parameters. The routine calls **svcerr_auth(xprt, AUTH_TOOWEAK)**.

SVCXPRT *
svccraw_create()

This routine creates a toy RPC service transport, to which it returns a pointer. The transport is really a buffer within the process's address space, so the corresponding RPC client should live in the same address space; see `clntraw_create()`. This routine allows simulation of RPC and acquisition of RPC overheads (such as round trip times), without any kernel interference. This routine returns NULL if it fails.

SVCXPRT *
svctcp_create(sock, send_buf_size, recv_buf_size)
int sock;
u_int send_buf_size, recv_buf_size;

This routine creates a TCP/IP-based RPC service transport, to which it returns a pointer. The transport is associated with the socket `sock`, which may be `RPC_ANYSOCK`, in which case a new socket is created. If the socket is not bound to a local TCP port, then this routine binds it to an arbitrary port. Upon completion, `xprt->xp_sock` is the transport's socket descriptor, and `xprt->xp_port` is the transport's port number. This routine returns NULL if it fails. Since TCP-based RPC uses buffered I/O, users may specify the size of buffers; values of zero choose suitable defaults.

void
svcfdd_create(fd, sendsize, recvsize)
int fd;
u_int sendsize;
u_int recvsize;

Create a service on top of any open descriptor. Typically, this descriptor is a connected socket for a stream protocol such as TCP. `sendsize` and `recvsize` indicate sizes for the send and receive buffers. If they are zero, a reasonable default is chosen.

SVCXPRT *
svccudp_create(sock)
int sock;

This routine creates a UDP/IP-based RPC service transport, to which it returns a pointer. The transport is associated with the socket `sock`, which may be `RPC_ANYSOCK`, in which case a new socket is created. If the socket is not bound to a local UDP port, then this routine binds it to an arbitrary port. Upon completion, `xprt->xp_sock` is the transport's socket descriptor, and `xprt->xp_port` is the transport's port number. This routine returns NULL if it fails.

Warning: since UDP-based RPC messages can only hold up to 8 Kbytes of encoded data, this transport cannot be used for procedures that take large arguments or return huge results.

user2netname(name, uid, domain)
char *name;
int uid;
char *domain;

Convert from a domain-specific username to an operating-system independent netname. Returns TRUE if it succeeds and FALSE if it fails. Inverse of `netname2user()`.

xdr_accepted_reply(xdrs, ar)
XDR *xdrs;
struct accepted_reply *ar;

Used for encoding RPC reply messages. This routine is useful for users who wish to generate RPC-style messages without using the RPC package.

```
xdr_authunix_parms(xdrs, aupp)  
XDR *xdrs;  
struct authunix_parms *aupp;
```

Used for describing UNIX credentials. This routine is useful for users who wish to generate these credentials without using the RPC authentication package.

```
void  
xdr_callhdr(xdrs, chdr)  
XDR *xdrs;  
struct rpc_msg *chdr;
```

Used for describing RPC call header messages. This routine is useful for users who wish to generate RPC-style messages without using the RPC package.

```
xdr_callmsg(xdrs, cmsg)  
XDR *xdrs;  
struct rpc_msg *cmsg;
```

Used for describing RPC call messages. This routine is useful for users who wish to generate RPC-style messages without using the RPC package.

```
xdr_opaque_auth(xdrs, ap)  
XDR *xdrs;  
struct opaque_auth *ap;
```

Used for describing RPC authentication information messages. This routine is useful for users who wish to generate RPC-style messages without using the RPC package.

```
xdr_pmap(xdrs, regs)  
XDR *xdrs;  
struct pmap *regs;
```

Used for describing parameters to various portmap procedures, externally. This routine is useful for users who wish to generate these parameters without using the pmap interface.

```
xdr_pmaplist(xdrs, rp)  
XDR *xdrs;  
struct pmaplist **rp;
```

Used for describing a list of port mappings, externally. This routine is useful for users who wish to generate these parameters without using the pmap interface.

```
xdr_rejected_reply(xdrs, rr)  
XDR *xdrs;  
struct rejected_reply *rr;
```

Used for describing RPC reply messages. This routine is useful for users who wish to generate RPC-style messages without using the RPC package.

```
xdr_replymsg(xdrs, rmsg)  
XDR *xdrs;  
struct rpc_msg *rmsg;
```

Used for describing RPC reply messages. This routine is useful for users who wish to generate RPC style messages without using the RPC package.


```
void  
xprt_register(xprt)  
SVCXPRT *xprt;
```

After RPC service transport handles are created, they should register themselves with the RPC service package. This routine modifies the global variable `svc_fds()`. Service implementors usually do not need this routine.

```
void  
xprt_unregister(xprt)  
SVCXPRT *xprt;
```

Before an RPC service transport handle is destroyed, it should unregister itself with the RPC service package. This routine modifies the global variable `svc_fds()`. Service implementors usually do not need this routine.

SEE ALSO

`xdr(3N)`, `keyserv(8C)`

Network Programming:

NAME

rtime – get remote time

SYNOPSIS

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

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

```
#include <netinet/in.h>
```

```
int rtime(addrp, timep, timeout)
```

```
struct sockaddr_in *addrp;
```

```
struct timeval *timep;
```

```
struct timeval *timeout;
```

DESCRIPTION

rtime() consults the Internet Time Server at the address pointed to by *addrp* and returns the remote time in the *timeval* struct pointed to by *timep*. Normally, the UDP protocol is used when consulting the Time Server. The *timeout* parameter specifies how long the routine should wait before giving up when waiting for a reply. If *timeout* is specified as NULL, however, the routine will instead use TCP and block until a reply is received from the time server.

The routine returns 0 if it is successful. Otherwise, it returns -1 and *errno* is set to reflect the cause of the error.

SEE ALSO

timed(8C)

NAME

scandir, alphasort – 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)*. The second parameter is a pointer to an array of structure pointers. The third parameter is a pointer to a routine which is called with a pointer to a directory entry and should return a non zero value if the directory entry should be included in the array. If this pointer is *NULL*, then all the directory entries will be included. The last argument is a pointer to a routine 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 will sort the array alphabetically.

scandir() returns the number of entries in the array and a pointer to the array through the parameter *namelist*.

SEE ALSO

directory(3), malloc(3), qsort(3)

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

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 results are undefined in there are insufficient *args* for the format. If the format is exhausted while *args* remain, the excess *args* are simply ignored.

The control string usually contains conversion specifications, which are used to direct interpretation of input sequences. The control string may contain:

1. White-space characters (SPACE, TAB, or NEWLINE) which, except in two cases described below, cause input to be read up to the next non-white-space character.
2. An ordinary character (not '%'), which must match the next character of the input stream.
3. Conversion specifications, consisting of the character '%', an optional assignment suppressing character '*', an optional numerical maximum field width, an optional l (ell) or h indicating the size of the receiving variable, and a conversion code.

A conversion specification directs the conversion of the next input field; the result is placed in the variable pointed to by the corresponding argument, unless assignment suppression was indicated by '*'. The suppression of assignment provides a way of describing an input field which is to be skipped. An input field is defined as a string of non-space characters; it extends to the next inappropriate character or until the field width, if specified, is exhausted. For all descriptors except '[' and 'c', white space leading an input field is ignored.

The conversion character indicates the interpretation of the input field; the corresponding pointer argument must usually be of a restricted type. For a suppressed field, no pointer argument is given. 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. |
| u | An unsigned decimal integer is expected; the corresponding argument should be an unsigned integer pointer. |
| o | An octal integer is expected; the corresponding argument should be an integer pointer. |
| x | A hexadecimal integer is expected; the corresponding argument should be an integer pointer. |
| i | An integer is expected; the corresponding argument should be an integer pointer. It will store the value of the next input item interpreted according to C conventions: a leading '0' implies octal; a leading '0x' implies hexadecimal; otherwise, decimal. |
| n | Stores in an integer argument the total number of characters (including white space) that have been scanned so far since the function call. No input is consumed. |

- e,f,g A floating point number is expected; the next field is converted accordingly and stored through the corresponding argument, which should be a pointer to a *float*. The input format for floating point numbers is as described for `string_to_decimal(3)`, with *fortran_exponent* zero.
- 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 automatically. The input field is terminated by a white space character.
- c A character is expected; the corresponding argument should be a character pointer. The normal skip over white space is suppressed in this case; to read the next non-space character, use `%1s`. If a field width is given, the corresponding argument should refer to a character array, and the indicated number of characters is read.
- [Indicates string data; the normal skip over leading white space is suppressed. The left bracket is followed by a set of characters, which we will call the *scanset*, and a right bracket; the input field is the maximal sequence of input characters consisting entirely of characters in the scanset. The circumflex (`^`), when it appears as the first character in the scanset, serves as a complement operator and redefines the scanset as the set of all characters *not* contained in the remainder of the scanset string. There are some conventions used in the construction of the scanset. A range of characters may be represented by the construct *first-last*, thus `[0123456789]` may be expressed `[0-9]`. Using this convention, *first* must be lexically less than or equal to *last*, or else the dash will stand for itself. The dash will also stand for itself whenever it is the first or the last character in the scanset. To include the right square bracket as an element of the scanset, it must appear as the first character (possibly preceded by a circumflex) of the scanset, and in this case it will not be syntactically interpreted as the closing bracket. The corresponding argument must point to a character array large enough to hold the data field and the terminating `\0`, which will be added automatically. At least one character must match for this conversion to be considered successful.

The conversion characters `d`, `u`, `o`, `x`, and `i` may be preceded by `l` or `h` to indicate that a pointer to `long` or to `short` rather than to `int` is in the argument list. Similarly, the conversion characters `e`, `f`, and `g` may be preceded by `l` to indicate that a pointer to `double` rather than to `float` is in the argument list. The `l` or `h` modifier is ignored for other conversion characters.

Avoid this common error: because `printf(3S)` does not require that the lengths of conversion descriptors and actual parameters match, coders sometimes are careless with the `scanf()` functions. But converting `%f` to `&double` or `%lf` to `&float` *does not work*; the results are quite incorrect.

`scanf()` conversion terminates at EOF, at the end of the control string, or when an input character conflicts with the control string. In the latter case, the offending character is left unread in the input stream.

`scanf()` returns the number of successfully matched and assigned input items; this number can be zero in the event of an early conflict between an input character and the control string. The constant EOF is returned upon end of input. Note: 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.

If the input ends before the first conflict or conversion, EOF is returned. If the input ends after the first conflict or conversion, the number of successfully matched items is returned.

EXAMPLES

The call:

```
int i, n; float x; char name[50];
n = scanf("%d%f%s", &i, &x, name);
```

with the input line:

```
25 54.32E-1 thompson
```

will assign to *n* the value 3, to *i* the value 25, to *x* the value 5.432, and *name* will contain thompson\0. Or:

```
int i, j; float x; char name[50];
(void) scanf("%i%2d%f%*d %[0-9]", &j, &i, &x, name);
```

with input:

```
011 56789 0123 56a72
```

will assign 9 to *j*, 56 to *i*, 789.0 to *x*, skip 0123, and place the string 56\0 in *name*. The next call to `getchar()` (see `getc(3S)`) will return a. Or:

```
int i, j, s, e; char name[50];
(void) scanf("%i %i %n%s%n", &i, &j, &s, name, &e);
```

with input:

```
0x11 0xy johnson
```

will assign 17 to *i*, 0 to *j*, 6 to *s*, will place the string xy\0 in *name*, and will assign 8 to *e*. Thus, the length of *name* is $e - s = 2$. The next call to `getchar()` (see `getc(3S)`) will return a SPACE.

SEE ALSO

`getc(3S)`, `printf(3S)`, `scanf(3V)`, `stdio(3S)`, `string_to_decimal(3)`, `strtol(3)`

DIAGNOSTICS

These functions return EOF on end of input, and a short count for missing or illegal data items.

BUGS

The success of literal matches and suppressed assignments is not directly determinable.

WARNINGS

Trailing white space (including a NEWLINE) is left unread unless matched in the control string.

NAME

setbuf, setbuffer, setlinebuf, setvbuf – 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;

int setvbuf(stream, buf, type, size)
FILE *stream;
char *buf;
int type, size;
```

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`. `fflush()` (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 unbuffered by default.

`setbuf()` can be used after a stream has been opened but before it is read or written. It causes the array pointed to by `buf` to be used instead of an automatically allocated buffer. If `buf` is the NULL pointer, input/output will be completely unbuffered. A manifest constant `BUFSIZ`, defined in the `<stdio.h>` header file, tells how big an array is needed:

```
char buf[BUFSIZ];
```

`setbuffer`, an alternate form of `setbuf`, can be used after a stream has been opened but before it is read or written. It uses the character array `buf` whose size is determined by the `size` argument instead of an automatically allocated buffer. If `buf` is the NULL pointer, input/output will be completely unbuffered.

`setvbuf()` can be used after a stream has been opened but before it is read or written. `type` determines how stream will be buffered. Legal values for `type` (defined in `<stdio.h>`) are:

```
_IOFBF    fully buffers the input/output.
_IOLBF    line buffers the output; the buffer will be flushed when a NEWLINE is written, the
           buffer is full, or input is requested.
_IONBF    completely unbuffers the input/output.
```

If `buf` is not the NULL pointer, the array it points to will be used for buffering, instead of an automatically allocated buffer. `size` specifies the size of the buffer to be used.

`setlinebuf()` is used to change the buffering on a stream from block buffered or unbuffered to line buffered. Unlike `setbuf`, `setbuffer`, and `setvbuf`, 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`.

NOTE

A common source of error is allocating buffer space as an "automatic" variable in a code block, and then failing to close the stream in the same block.

SEE ALSO

`fclose(3S)`, `fopen(3S)`, `fread(3S)`, `getc(3S)`, `malloc(3)`, `printf(3S)`, `putc(3S)`, `puts(3S)`, `setbuf(3V)`

DIAGNOSTICS

If an illegal value for *type* or *size* is provided, `setvbuf()` returns a non-zero value. Otherwise, the value returned will be zero.

NAME

setjmp, longjmp, sigsetjmp, siglongjmp – non-local goto

SYNOPSIS

```
#include <setjmp.h>

int setjmp(env)
jmp_buf env;

longjmp(env, val)
jmp_buf env;
int val;

int _setjmp(env)
jmp_buf env;

_longjmp(env, val)
jmp_buf env;
int val;

int sigsetjmp(env, savemask)
sigjmp_buf env;
int savemask;

siglongjmp(env, val)
sigjmp_buf env;
int val;
```

DESCRIPTION

setjmp() and longjmp() 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. A normal call to setjmp() returns zero. setjmp() also saves the register environment. If a longjmp() call will be made, the routine which called setjmp() should not return until after the longjmp() has returned control (see below).

longjmp() restores the environment saved by the last call of setjmp, and 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; however, if *val* were zero, execution would continue as if the call of setjmp() had returned one. This ensures that a “return” from setjmp() caused by a call to longjmp() can be distinguished from a regular return from setjmp. The calling function must not itself have returned in the interim, otherwise longjmp() will be returning control to a possibly non-existent environment. All memory-bound data have values as of the time longjmp() was called. The CPU and floating-point data registers are restored to the values they had at the time that setjmp() was called. But, because the register storage class is only a hint to the C compiler, variables declared as register variables may not necessarily be assigned to machine registers, so their values are unpredictable after a longjmp. This is especially a problem for programmers trying to write machine-independent C routines.

setjmp() and longjmp() save and restore the signal mask (see sigsetmask(2)), while _setjmp and _longjmp manipulate only the C stack and registers. If the *savemask* flag to sigsetjmp is non-zero, the signal mask is saved, and a subsequent siglongjmp using the same *env* will restore the signal mask. If the *savemask* flag is zero, the signal mask is not saved, and a subsequent siglongjmp using the same *env* will not restore the signal mask. In all other ways, _setjmp and sigsetjmp function in the same way that setjmp() does, and _longjmp and siglongjmp function in the same way that longjmp() does.

None of these functions save or restore any floating-point status or control registers, in particular the MC68881 *fpsr*, *fpcr*, or *fpiar*, the Sun-3 FPA *fpamode* or *fpastatus*, and the Sun-4 *%fsr*. See *ieee_flags(3M)* to save and restore floating-point status or control information.

EXAMPLE

The following code fragment indicates the flow of control of the `setjmp()` and `longjmp()` combination:

```

function declaration
...
    jmp_buf my_environment;
    ...
    if (setjmp(my_environment)) {
        /* register variables have unpredictable values
           code after the return from longjmp
           ...
        } else {
            /* do not modify register vars
               this is the return from setjmp
               ...
            }
}

```

SEE ALSO

`cc(1V)`, `sigsetmask(2)`, `sigvec(2)`, `ieee_flags(3M)`, `signal(3)`, `setjmp(3V)`

BUGS

`setjmp()` does not save the 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.

On Sun-2 and Sun-3 systems `setjmp()` also saves the register environment. Therefore, all data that are bound to registers are restored to the values they had at the time that `setjmp()` was called. All memory-bound data have values as of the time `longjmp()` was called. However, because the register storage class is only a hint to the C compiler, variables declared as register variables may not necessarily be assigned to machine registers, so their values are unpredictable after a `longjmp`. When using compiler options that specify automatic register allocation (see `cc(1V)`), the compiler will not attempt to assign variables to registers in routines that call `setjmp`.

`longjmp()` never causes `setjmp()` to return zero in the Sun implementation; this is also true of many other implementations, including all System V implementations, so programmers should not depend on `longjmp()` being able to cause `setjmp()` to return zero.

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() (setrgid) 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

getgid(2), getuid(2), setregid(2), setreuid(2)

DIAGNOSTICS

Zero is returned if the user (group) ID is set; -1 is returned otherwise, with the global variable errno set as for setreuid() or setregid.

NAME

`sigfpe` - signal handling for specific SIGFPE codes

SYNOPSIS

```
#include <signal.h>
#include <floatingpoint.h>
sigfpe_handler_type sigfpe(code, hdl)
sigfpe_code_type code;
sigfpe_handler_type hdl;
```

DESCRIPTION

This function allows signal handling to be specified for particular SIGFPE codes. A call to `sigfpe()` defines a new handler *hdl* for a particular SIGFPE *code* and returns the old handler as the value of the function `sigfpe`. Normally handlers are specified as pointers to functions; the special cases `SIGFPE_IGNORE`, `SIGFPE_ABORT`, and `SIGFPE_DEFAULT` allow ignoring, specifying core dump using `abort(3)`, or default handling respectively.

For these IEEE-related codes:

<code>FPE_FLTINEX_TRAP</code>	<code>fp_inexact</code> - floating inexact result
<code>FPE_FLTDIV_TRAP</code>	<code>fp_division</code> - floating division by zero
<code>FPE_FLTUND_TRAP</code>	<code>fp_underflow</code> - floating underflow
<code>FPE_FLTOVF_TRAP</code>	<code>fp_overflow</code> - floating overflow
<code>FPE_FLTBSUN_TRAP</code>	<code>fp_invalid</code> - branch or set on unordered
<code>FPE_FLTOPERR_TRAP</code>	<code>fp_invalid</code> - floating operand error
<code>FPE_FLTNAN_TRAP</code>	<code>fp_invalid</code> - floating Not-A-Number

default handling is defined to be to call the handler specified to `ieee_handler(3M)`.

For all other SIGFPE codes, default handling is to core dump using `abort(3)`.

The compilation option `-ffpa` causes fpa recomputation to replace the default abort action for code `FPE_FPA_ERROR`. Note: `SIGFPE_DEFAULT` will restore abort rather than FPA recomputation for this code.

Three steps are required to intercept an IEEE-related SIGFPE code with `sigfpe`:

- 1) Set up a handler with `sigfpe`.
- 2) Enable the relevant IEEE trapping capability in the hardware, perhaps by using assembly-language instructions.
- 3) Perform a floating-point operation that generates the intended IEEE exception.

Unlike `ieee_handler(3M)`, `sigfpe()` never changes floating-point hardware mode bits affecting IEEE trapping. No IEEE-related SIGFPE signals will be generated unless those hardware mode bits are enabled.

SIGFPE signals can be handled using `sigvec(2)`, `signal(3)`, `sigfpe(3)`, or `ieee_handler(3M)`. In a particular program, to avoid confusion, use only one of these interfaces to handle SIGFPE signals.

EXAMPLE

A user-specified signal handler might look like this:

```
void sample_handler( sig, code, scp, addr )
    int sig ;          /* sig == SIGFPE always */
    int code ;
    struct sigcontext *scp ;
    char *addr ;
    {
        /*
         * Sample user-written sigfpe code handler.
         * Prints a message and continues.
         * struct sigcontext is defined in <signal.h>.
         */
        printf(" ieee exception code %x occurred at pc %X \n",code,scp->sc_pc);
    }

```

and it might be set up like this:

```
extern void sample_handler();
main()
{
    sigfpe_handler_type hdl, old_handler1, old_handler2;
    /*
     * save current overflow and invalid handlers; set the new
     * overflow handler to sample_handler() and set the new
     * invalid handler to SIGFPE_ABORT (abort on invalid)
     */
    hdl = (sigfpe_handler_type) sample_handler;
    old_handler1 = sigfpe(FPE_FLTOVF_TRAP, hdl);
    old_handler2 = sigfpe(FPE_FLTOPERR_TRAP, SIGFPE_ABORT);
    ...
    /*
     * restore old overflow and invalid handlers
     */
    sigfpe(FPE_FLTOVF_TRAP, old_handler1);
    sigfpe(FPE_FLTOPERR_TRAP, old_handler2);
}

```

FILES

```
/usr/include/floatingpoint.h
/usr/include/signal.h

```

SEE ALSO

sigvec(2), abort(3), floatingpoint(3), ieee_handler(3M), signal(3),

DIAGNOSTICS

sigfpe() returns BADSIG if *code* is not zero or a defined SIGFPE code.

NAME

siginterrupt – allow signals to interrupt system calls

SYNOPSIS

int siginterrupt(sig, flag)

int sig, flag;

DESCRIPTION

siginterrupt() is used to change the system call restart behavior when a system call is interrupted by the specified signal. If the flag is false (0), then system calls will be restarted if they are interrupted by the specified signal and no data has been transferred yet. System call restart is the default behavior on 4.2BSD, and on SunOS in the 4.2 environment, when the **signal(3)** routine is used.

If the flag is true (1), then restarting of system calls is disabled. If a system call is interrupted by the specified signal and no data has been transferred, the system call will return -1 with **errno** set to **EINTR**. Interrupted system calls that have started transferring data will return the amount of data actually transferred. System call interrupt is the signal behavior found on older version of the UNIX operating systems, such as 4.1BSD and System V UNIX. It is the default behavior on SunOS in the System V environment when the **signal()** routine is used; therefore, this routine is useful in that environment only if a signal that a **sigvec(2)** specified should restart system calls is to be changed not to restart them.

Note: the new 4.2BSD signal handling semantics are not altered in any other way. Most notably, signal handlers always remain installed until explicitly changed by a subsequent **sigvec** call, and the signal mask operates as documented in **sigvec**, unless the **SV_RESETHAND** bit has been used to specify that the pre-4.2BSD signal behavior is to be used. Programs may switch between restartable and interruptible system call operation as often as desired in the execution of a program.

Issuing a **siginterrupt()** call during the execution of a signal handler will cause the new action to take place on the next signal to be caught.

NOTES

This library routine uses an extension of the **sigvec(2)** system call that is not available in 4.2BSD, hence it should not be used if backward compatibility is needed.

RETURN VALUE

A 0 value indicates that the call succeeded. A -1 value indicates that an invalid signal number has been supplied.

SEE ALSO

sigblock(2), **sigpause(2)**, **sigsetmask(2)**, **sigvec(2)**, **signal(3)**

NAME

signal – simplified software signal facilities

SYNOPSIS

```
#include <signal.h>

void (*signal(sig, func))()
void (*func)();
```

DESCRIPTION

signal() is a simplified interface to the more general sigvec(2) facility. Programs that use signal() in preference to sigvec() are more likely to be portable to all systems.

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 termio(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 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
SIGABRT	6*	abort (generated by abort(3) routine)
SIGEMT	7*	emulator trap
SIGFPE	8*	arithmetic exception
SIGKILL	9	kill (cannot be caught, blocked, or ignored)
SIGBUS	10*	bus error
SIGSEGV	11*	segmentation violation
SIGSYS	12*	bad argument to system call
SIGPIPE	13	write on a pipe or other socket with no one to read it
SIGALRM	14	alarm clock
SIGTERM	15	software termination signal
SIGURG	16●	urgent condition present on socket
SIGSTOP	17†	stop (cannot be caught, blocked, or ignored)
SIGTSTP	18†	stop signal generated from keyboard
SIGCONT	19●	continue after stop (cannot be blocked)
SIGCHLD	20●	child status has changed
SIGTTIN	21†	background read attempted from control terminal
SIGTTOU	22†	background write attempted to control terminal
SIGIO	23●	I/O is possible on a descriptor (see fcntl(2V))
SIGXCPU	24	cpu time limit exceeded (see getrlimit(2))
SIGXFSZ	25	file size limit exceeded (see getrlimit(2))
SIGVTALRM	26	virtual time alarm (see getitimer(2))
SIGPROF	27	profiling timer alarm (see getitimer(2))
SIGWINCH	28●	window changed (see termio(4) and win(4S))
SIGLOST	29*	resource lost (see lockd(8C))
SIGUSR1	30	user-defined signal 1
SIGUSR2	31	user-defined signal 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, terminating the call prematurely, the call is automatically restarted. In particular this can occur during a read(2V) or write(2V) 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. An execve(2) resets all caught signals to the default action; ignored signals remain ignored.

NOTES

The handler routine can be declared:

```
void handler(sig, code, scp, addr)
int sig, code;
struct sigcontext *scp;
char *addr;
```

Here *sig* is the signal number; *code* is a parameter of certain signals that provides additional detail; *scp* is a pointer to the sigcontext structure (defined in <signal.h>), used to restore the context from before the signal; and *addr* is additional address information. See sigvec(2) for more details.

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	<i>sig</i> 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), execve(2), fork(2), getitimer(2), getrlimit(2), kill(2V), ptrace(2), read(2V), sigblock(2), sigpause(2), sigsetmask(2), sigstack(2), sigvec(2), vfork(2), wait(2), write(2V), setjmp(3), termio(4)

NAME

sleep – suspend execution for interval

SYNOPSIS

sleep(seconds)
unsigned seconds;

DESCRIPTION

sleep() suspends the current process from execution for the number of seconds specified by the argument. The actual suspension time may be up to 1 second less than that requested, because scheduled wakeups occur at fixed 1-second intervals, and may be an arbitrary amount longer because of other activity in the system.

sleep() is implemented by setting an interval timer and pausing until it expires. The previous state of this timer is saved and restored. If the sleep time exceeds the time to the expiration of the previous value of the timer, the process sleeps only until the timer would have expired, and the signal which occurs with the expiration of the timer is sent one second later.

SEE ALSO

getitimer(2), sigpause(2), usleep(3)

NAME

ssignal, gsignal – software signals

SYNOPSIS

```
#include <signal.h>

int (*ssignal (sig, action))()
int sig, (*action)();

int gsignal (sig)
int sig;
```

DESCRIPTION

ssignal() and **gsignal()** implement a software facility similar to **signal(3)**.

Software signals made available to users are associated with integers in the inclusive range 1 through 15. A call to **ssignal()** associates a procedure, *action*, with the software signal *sig*; the software signal, *sig*, is raised by a call to **ssignal**. Raising a software signal causes the action established for that signal to be *taken*.

The first argument to **ssignal()** is a number identifying the type of signal for which an action is to be established. The second argument defines the action; it is either the name of a (user-defined) *action function* or one of the manifest constants **SIG_DFL** (default) or **SIG_IGN** (ignore). **ssignal()** returns the action previously established for that signal type; if no action has been established or the signal number is illegal, **ssignal()** returns **SIG_DFL**.

ssignal() raises the signal identified by its argument, *sig*:

If an action function has been established for *sig*, then that action is reset to **SIG_DFL** and the action function is entered with argument *sig*. **ssignal()** returns the value returned to it by the action function.

If the action for *sig* is **SIG_IGN**, **ssignal()** returns the value 1 and takes no other action.

If the action for *sig* is **SIG_DFL**, **ssignal()** returns the value 0 and takes no other action.

If *sig* has an illegal value or no action was ever specified for *sig*, **ssignal()** returns the value 0 and takes no other action.

SEE ALSO

signal(3)

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 I/O buffering scheme. The in-line macros `getc(3S)` and `putc(3S)` handle characters quickly. The macros `getchar` and `putchar`, and the higher level routines `fgetc`, `getw`, `gets`, `fgets`, `scanf`, `fscanf`, `fread`, `fputc`, `putw`, `puts`, `fputs`, `printf`, `fprintf`, `fwrite` all use or act as if they use `getc()` and `putc()`; they can be freely intermixed.

A file with associated buffering is called a *stream*, and is declared to be a pointer to a defined type `FILE`. `fopen(3S)` creates certain descriptive data for a stream and returns a pointer to designate the stream in all further transactions. Normally, there are three open streams with constant pointers declared in the `<stdio.h>` include file and associated with the standard open files:

```
stdin      standard input file
stdout     standard output file
stderr     standard error file
```

A constant `NULL` (0) designates a nonexistent pointer.

An integer constant `EOF` (-1) is returned upon end-of-file or error by most integer functions that deal with streams (see the individual descriptions for details).

Any module that uses this package must include the header file of pertinent macro definitions, as follows:

```
#include <stdio.h>
```

The functions and constants mentioned in sections labeled 3S of this manual are declared in that header file and need no further declaration. The constants and the following ‘functions’ are implemented as macros; redeclaration of these names is perilous: `getc`, `getchar`, `putc`, `putchar`, `feof`, `ferror`, `fileno`, and `clearerr`.

Output streams, with the exception of the standard error stream `stderr`, are by default buffered if the output refers to a file and line-buffered if the output refers to a terminal. The standard error output stream `stderr` is by default unbuffered, but use of `fopen(3S)` will cause it to become buffered or line-buffered. When an output stream is unbuffered, information is written to the destination file or terminal as soon as it is output to the stream; when it is buffered, many characters are saved up and written as a block. When it is line-buffered, each line of output is written to the destination file or terminal as soon as the line is completed (that is, as soon as a `NEWLINE` character is output or, if the output stream is `stdout` or `stderr`, as soon as input is read from `stdin`). `setbuf(3S)`, `setbuffer`, `setlinebuf`, or `setvbuf` can be used to change the stream’s buffering strategy.

SEE ALSO

`open(2V)`, `close(2)`, `lseek(2)`, `pipe(2)`, `read(2V)`, `write(2V)`, `ctermid(3S)`, `cuserid(3S)`, `fclose(3S)`, `ferror(3S)`, `fopen(3S)`, `fread(3S)`, `fseek(3S)`, `getc(3S)`, `gets(3S)`, `popen(3S)`, `printf(3S)`, `putc(3S)`, `puts(3S)`, `scanf(3S)`, `setbuf(3S)`, `system(3)`, `tmpfile(3S)`, `tmpnam(3S)`, `ungetc(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.

BUGS

The standard buffered functions do not interact well with certain other library and system functions, especially `vfork(2)`.

NOTES

The line buffering of output to terminals is almost always transparent, but may cause confusion or malfunctioning of programs which use standard I/O routines but use `read(2V)` to read from the standard input, as calls to `read()` do not cause output to line-buffered streams to be flushed.

In cases where a large amount of computation is done after printing part of a line on an output terminal, it is necessary to call `fflush` (see `fclose(3S)`) on the standard output before performing the computation so that the output will appear.

NAME

string, strcat, strncat, strdup, strcmp, strncmp, strcpy, strncpy, strlen, strchr, strrchr, strpbrk, strspn, strcspn, strtok, index, rindex – string operations

SYNOPSIS

```
#include <string.h>

char *strcat(s1, s2)
char *s1, *s2;

char *strncat(s1, s2, n)
char *s1, *s2;
int n;

char *strdup(s1)
char *s1;

int strcmp(s1, s2)
char *s1, *s2;

int strncmp(s1, s2, n)
char *s1, *s2;
int n;

char *strcpy(s1, s2)
char *s1, *s2;

char *strncpy(s1, s2, n)
char *s1, *s2;
int n;

int strlen(s)
char *s;

char *strchr(s, c)
char *s;
int c;

char *strrchr(s, c)
char *s;
int c;

char *strpbrk(s1, s2)
char *s1, *s2;

int strspn(s1, s2)
char *s1, *s2;

int strcspn(s1, s2)
char *s1, *s2;

char *strtok(s1, s2)
char *s1, *s2;

#include <strings.h>

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()** appends at most *n* characters. Each returns 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 compares at most *n* characters.

strdup() returns a pointer to a new string which is a duplicate of the string pointed to by *s1*. The space for the new string is obtained using **malloc(3)**. If the new string cannot be created, a NULL pointer is returned.

strcpy() copies string *s2* to *s1*, stopping after the NULL character has been copied. **strncpy()** copies exactly *n* characters, truncating or NULL-padding *s2*. The result will not be NULL-terminated if the length of *s2* is *n* or more. Each function returns *s1*.

strlen() returns the number of characters in *s*, not including the NULL-terminating character.

strchr() (**strrchr**) returns a pointer to the first (last) occurrence of character *c* in string *s*, or a NULL pointer if *c* does not occur in the string. The NULL character terminating a string is considered to be part of the string.

index() (**rindex**) returns a pointer to the first (last) occurrence of character *c* in string *s*, or a NULL pointer if *c* does not occur in the string. These functions are identical to **strchr()** (**strchr**) and merely have different names.

strpbrk() returns a pointer to the first occurrence in string *s1* of any character from string *s2*, or a NULL pointer if no character from *s2* exists in *s1*.

strspn() (**strcspn**) returns the length of the initial segment of string *s1* which consists entirely of characters from (not from) string *s2*.

strtok() considers the string *s1* to consist of a sequence of zero or more text tokens separated by spans of one or more characters from the separator string *s2*. The first call (with pointer *s1* specified) returns a pointer to the first character of the first token, and will have written a NULL character into *s1* immediately following the returned token. The function keeps track of its position in the string between separate calls, so that subsequent calls (which must be made with the first argument a NULL pointer) will work through the string *s1* immediately following that token. In this way subsequent calls will work through the string *s1* until no tokens remain. The separator string *s2* may be different from call to call. When no token remains in *s1*, a NULL pointer is returned.

NOTE

For user convenience, all these functions, except for **index()** and **rindex**, are declared in the optional **<string.h>** header file. All these functions, including **index()** and **rindex()** but excluding **strchr**, **strrchr**, **strpbrk**, **strspn**, **strcspn**, and **strtok**, are declared in the optional **<strings.h>** include file; these headers are set this way for backward compatibility.

SEE ALSO

malloc(3), **bstring(3)**

WARNINGS

strcmp() and **strncmp()** use native character comparison, which is signed on the Sun, but may be unsigned on other machines. Thus the sign of the value returned when one of the characters has its high-order bit set is implementation-dependent.

On the Sun processor, as well as on many other machines, you can *not* use a NULL pointer to indicate a NULL string. A NULL pointer is an error and results in an abort of the program. If you wish to indicate a NULL string, you must have a pointer that points to an explicit NULL string. On some implementations of the C language on some machines, a NULL pointer, if dereferenced, would yield a NULL string; this highly non-portable trick was used in some programs. Programmers using a NULL pointer to represent an empty string should be aware of this portability issue; even on machines where dereferencing a NULL pointer does not cause an abort of the program, it does not necessarily yield a

NULL string.

Character movement is performed differently in different implementations. Thus overlapping moves may yield surprises.

NAME

`string_to_decimal`, `file_to_decimal`, `func_to_decimal` – parse characters into decimal record

SYNOPSIS

```
#include <floatingpoint.h>
#include <stdio.h>

void string_to_decimal(pc,nmax,fortran_conventions,pd,pform,pechar)
char **pc;
int nmax;
int fortran_conventions;
decimal_record *pd;
enum decimal_string_form *pform;
char **pechar;

void file_to_decimal(pc,nmax,fortran_conventions,pd,pform,pechar,pf,pnread)
char **pc;
int nmax;
int fortran_conventions;
decimal_record *pd;
enum decimal_string_form *pform;
char **pechar;
FILE *pf;
int *pnread;

void func_to_decimal(pc,nmax,fortran_conventions,pd,pform,pechar,pget,pnread,punget)
char **pc;
int nmax;
int fortran_conventions;
decimal_record *pd;
enum decimal_string_form *pform;
char **pechar;
int (*pget)();
int *pnread;
int (*punget)();
```

DESCRIPTION

The `char_to_decimal` functions parse a numeric token from at most *nmax* characters in a string ***pc* or file **pf* or function *(*pget)()* into a decimal record **pd*, classifying the form of the string in **pform* and **pechar*. The accepted syntax is intended to be sufficiently flexible to accommodate many languages:

whitespace value

or

whitespace sign value

where *whitespace* is any number of characters defined by *isspace* in `/usr/include/ctype.h`, *sign* is either of `[+-]`, and *value* can be *number*, *nan*, or *inf*. *inf* can be `INF` (*inf_form*) or `INFINITY` (*infinity_form*) without regard to case. *nan* can be `NAN` (*nan_form*) or `NAN(nstring)` (*nanstring_form*) without regard to case; *nstring* is any string of characters not containing `'` or `NULL`; *nstring* is copied to *pd->ds* and, currently, not used subsequently. *number* consists of

significand

or

significand efield

where *significand* must contain one or more digits and may contain one point; possible forms are

```

digits           (int_form)
digits.         (intdot_form)
.digits          (dotfrac_form)
digits.digits   (intdotfrac_form)

```

efield consists of

```
echar digits
```

or

```
echar sign digits
```

where *echar* is one of [Ee], and *digits* contains one or more digits.

When *fortran_conventions* is nonzero, additional input forms are accepted according to various Fortran conventions:

- 0 no Fortran conventions
- 1 Fortran list-directed input conventions
- 2 Fortran formatted input conventions, ignore blanks (BN)
- 3 Fortran formatted input conventions, blanks are zeros (BZ)

When *fortran_conventions* is nonzero, *echar* may also be one of [Dd], and *efield* may also have the form

```
sign digits
```

When *fortran_conventions* ≥ 2 , blanks may appear in the *digits* strings for the integer, fraction, and exponent fields and may appear between *echar* and the exponent sign and after the infinity and NaN forms. If *fortran_conventions* $= 2$, the blanks are ignored. When *fortran_conventions* $= 3$, the blanks that appear in *digits* strings are interpreted as zeros, and other blanks are ignored.

The form of the accepted decimal string is placed in **perform*. If an *efield* is recognized, **pechar* is set to point to the *echar*.

On input, **pc* points to the beginning of a character string buffer of length $\geq nmax$. On output, **pc* points to a character in that buffer, one past the last accepted character. *string_to_decimal()* gets its characters from the buffer; *file_to_decimal()* gets its characters from **pf* and records them in the buffer, and places a null after the last character read. *func_to_decimal()* gets its characters from an int function (**pget*()).

The scan continues until no more characters could possibly fit the acceptable syntax or until *nmax* characters have been scanned. If the *nmax* limit is not reached then at least one extra character will usually be scanned that is not part of the accepted syntax. *file_to_decimal()* and *func_to_decimal()* set **pnread* to the number of characters read from the file; if greater than *nmax*, some characters were lost. If no characters were lost, *file_to_decimal()* and *func_to_decimal()* attempt to push back, with *ungetc(3S)* or (**punget*()), as many as possible of the excess characters read, adjusting **pnread* accordingly. If all *unget* calls are successful, then ***pc* will be NULL. No push back will be attempted if (**punget*()) is NULL.

Typical declarations for **pget*() and **punget*() are:

```

int xget()
{ ... }
int (*pget)() = xget ;
int xunget(c)
char c ;
{ ... }
int (*punget)() = xunget ;

```

If no valid number was detected, *pd->fpclass* is set to *fp_signaling*, **pc* is unchanged, and **pform* is set to *invalid_form*.

atof and *strtod(3)* use *string_to_decimal*. *scanf(3S)* uses *file_to_decimal*.

FILES

/usr/include/ctype.h

SEE ALSO

scanf(3S), *strtod(3)*, *ungetc(3S)*

NAME

`strtod`, `atof` – convert string to double-precision number

SYNOPSIS

```
double strtod(str, ptr)
```

```
char *str, **ptr;
```

```
double atof(str)
```

```
char *str;
```

DESCRIPTION

`strtod()` returns as a double-precision floating-point number the value represented by the character string pointed to by *str*. The string is scanned up to the first unrecognized character, using `string_to_decimal(3)`, with `fortran_conventions` set to 0.

If the value of *ptr* is not `(char **)NULL`, a pointer to the character terminating the scan is returned in the location pointed to by *ptr*. If no number can be formed, **ptr* is set to *str*, and for historical compatibility, 0.0 is returned, although a NaN would better match the IEEE Floating-Point Standard's intent.

`atof(str)` is equivalent to `strtod(str, (char **)NULL)`. Thus, when `atof(str)` returns 0.0 there is no way to determine whether *str* contained a valid numerical string representing 0.0 or an invalid numerical string.

SEE ALSO

`scanf(3S)`, `string_to_decimal(3)`

DIAGNOSTICS

Exponent overflow and underflow produce the results specified by the IEEE Standard. In addition, `errno` is set to `ERANGE`.

NAME

`strtol`, `atol`, `atoi` – convert string to integer

SYNOPSIS

```
long strtol(str, ptr, base)
```

```
char *str, **ptr;
```

```
int base;
```

```
long atol(str)
```

```
char *str;
```

```
int atoi(str)
```

```
char *str;
```

DESCRIPTION

`strtol()` returns as a long integer the value represented by the character string pointed to by *str*. The string is scanned up to the first character inconsistent with the base. Leading “white-space” characters (as defined by `isspace()` in `ctype(3)`) are ignored.

If the value of *ptr* is not `(char **)NULL`, a pointer to the character terminating the scan is returned in the location pointed to by *ptr*. If no integer can be formed, that location is set to *str*, and zero is returned.

If *base* is positive (and not greater than 36), it is used as the base for conversion. After an optional leading sign, leading zeros are ignored, and “0x” or “0X” is ignored if *base* is 16.

If *base* is zero, the string itself determines the base thusly: after an optional leading sign a leading zero indicates octal conversion, and a leading “0x” or “0X” hexadecimal conversion. Otherwise, decimal conversion is used.

Truncation from long to int can, of course, take place upon assignment or by an explicit cast.

`atol(str)` is equivalent to `strtol(str, (char **)NULL, 10)`.

`atoi(str)` is equivalent to `(int) strtol(str, (char **)NULL, 10)`.

SEE ALSO

`ctype(3)`, `scanf(3S)`, `strtod(3)`

BUGS

Overflow conditions are ignored.

NAME

stty, gtty – set and get terminal state

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*. `stty()` 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 `ttcompat(4M)` 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)`, `ttcompat(4M)`

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 high-ender machines (IBM 360's, MC68000's, etc) and low-endian machines (such as Sun386i).

nbytes should be even.

The *from* and *to* addresses should not overlap in portable programs.

NAME

syslog, openlog, closelog, setlogmask – control system log

SYNOPSIS

```
#include <syslog.h>

openlog(ident, logopt, facility)
char *ident;

syslog(priority, message, parameters ...)
char *message;

closelog()

setlogmask(maskpri)
```

DESCRIPTION

`syslog()` passes *message* to `syslogd(8)`, which logs it in an appropriate system log, writes it to the system console, forwards it to a list of users, or forwards it to the `syslogd` on another host over the network. The message is tagged with a priority of *priority*. The message looks like a `printf(3S)` string except that `%m` is replaced by the current error message (collected from `errno`). A trailing NEWLINE is added if needed.

Priorities are encoded as a *facility* and a *level*. The facility describes the part of the system generating the message. The level is selected from an ordered list:

<code>LOG_EMERG</code>	A panic condition. This is normally broadcast to all users.
<code>LOG_ALERT</code>	A condition that should be corrected immediately, such as a corrupted system database.
<code>LOG_CRIT</code>	Critical conditions, such as hard device errors.
<code>LOG_ERR</code>	Errors.
<code>LOG_WARNING</code>	Warning messages.
<code>LOG_NOTICE</code>	Conditions that are not error conditions, but that may require special handling.
<code>LOG_INFO</code>	Informational messages.
<code>LOG_DEBUG</code>	Messages that contain information normally of use only when debugging a program.

If special processing is needed, `openlog()` can be called to initialize the log file. The parameter *ident* is a string that is prepended to every message. *logopt* is a bit field indicating logging options. Current values for *logopt* are:

<code>LOG_PID</code>	Log the process ID with each message. This is useful for identifying specific daemon processes (for daemons that fork).
<code>LOG_CONS</code>	Write messages to the system console if they cannot be sent to <code>syslogd</code> . This option is safe to use in daemon processes that have no controlling terminal, since <code>syslog()</code> forks before opening the console.
<code>LOG_NDELAY</code>	Open the connection to <code>syslogd</code> immediately. Normally the open is delayed until the first message is logged. This is useful for programs that need to manage the order in which file descriptors are allocated.
<code>LOG_NOWAIT</code>	Do not wait for child processes that have been forked to log messages onto the console. This option should be used by processes that enable notification of child termination using <code>SIGCHLD</code> , since <code>syslog()</code> may otherwise block waiting for a child whose exit status has already been collected.

The *facility* parameter encodes a default facility to be assigned to all messages that do not have an explicit facility already encoded:

LOG_KERN	Messages generated by the kernel. These cannot be generated by any user processes.
LOG_USER	Messages generated by random user processes. This is the default facility identifier if none is specified.
LOG_MAIL	The mail system.
LOG_DAEMON	System daemons, such as ftpd(8C) , routed(8C) , etc.
LOG_AUTH	The authorization system: login(1) , su(1) , getty(8) , etc.
LOG_LPR	The line printer spooling system: lpr(1) , lpc(8) , lpd(8) , etc.
LOG_NEWS	Reserved for the USENET network news system.
LOG_UUCP	Reserved for the UUCP system; it does not currently use syslog .
LOG_CRON	The cron/at facility; crontab(1) , at(1) , cron(8) , etc.
LOG_LOCAL0-7	Reserved for local use.

closelog() can be used to close the log file.

setlogmask() sets the log priority mask to *maskpri* and returns the previous mask. Calls to **syslog()** with a priority not set in *maskpri* are rejected. The mask for an individual priority *pri* is calculated by the macro **LOG_MASK(pri)**; the mask for all priorities up to and including *toppri* is given by the macro **LOG_UPTO(toppri)**. The default allows all priorities to be logged.

EXAMPLES

This call logs a message at priority **LOG_ALERT**:

```
syslog(LOG_ALERT, "who: internal error 23");
```

The FTP daemon **ftpd** would make this call to **openlog()** to indicate that all messages it logs should have an identifying string of **ftpd**, should be treated by **syslogd** as other messages from system daemons are, should include the process ID of the process logging the message:

```
openlog("ftpd", LOG_PID, LOG_DAEMON);
```

Then it would make the following call to **setlogmask()** to indicate that messages at priorities from **LOG_EMERG** through **LOG_ERR** should be logged, but that no messages at any other priority should be logged:

```
setlogmask(LOG_UPTO(LOG_ERR));
```

Then, to log a message at priority **LOG_INFO**, it would make the following call to **syslog**:

```
syslog(LOG_INFO, "Connection from host %d", CallingHost);
```

A locally-written utility could use the following call to **syslog()** to log a message at priority **LOG_INFO** to be treated by **syslogd** as other messages to the facility **LOG_LOCAL2** are:

```
syslog(LOG_INFO|LOG_LOCAL2, "error: %m");
```

SEE ALSO

at(1), **crontab(1)**, **logger(1)**, **login(1)**, **lpr(1)**, **su(1)**, **printf(3S)**, **syslog.conf(5)**, **cron(8)**, **ftpd(8C)**, **getty(8)**, **lpc(8)**, **lpd(8)**, **routed(8C)**, **syslogd(8)**

NAME

system – issue a shell command

SYNOPSIS

system(string)
char *string;

DESCRIPTION

system() gives the *string* to **sh(1)** as input, just as if the string had been typed as a command from a terminal. The current process performs a **wait(2)** system call, and waits until the shell terminates. **system()** then returns the exit status returned by **wait**. Unless the shell was interrupted by a signal, its termination status is contained in the 8 bits higher up from the low-order 8 bits of the value returned by **wait**.

SEE ALSO

sh(1), **execve(2)**, **wait(2)**, **popen(3S)**

DIAGNOSTICS

Exit status 127 (may be displayed as "32512") indicates the shell could not be executed.

NAME

termcap, 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 *bp* buffer, with the current size of the tty (usually a window). This allows pre-SunWindows programs to run in a window of arbitrary size. *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 `/etc/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 cannot write the file `/etc/termcap`. Note: if the window size changes, the “lines” and “columns” entries in *bp* are no longer correct. See the *SunView 1 Programmer's Guide* for details regarding [how to handle] this.

`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. `tgetstr()` returns the string pointer if successful. Otherwise it returns zero.

tgoto() returns a cursor addressing string decoded from *cm* to go to column *destcol* in line *destline*. It uses the external variables UP (from the up capability) and BC (if bc is given rather than bs) if necessary to avoid placing \n, ^D or ^@ in the returned string. (Programs which call **tgoto()** should be sure to turn off the XTABS bit(s), since **tgoto()** may now output a tab. Note: programs using **termcap()** should in general turn off XTABS anyway since some terminals use ^I (CTRL-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, *outr* is a routine which is called with each character in turn. The external variable *ospeed* should contain the encoded output speed of the terminal as described in **tty(4)**. 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)**, **tty(4)**, **termcap(5)**

NAME

time, ftime – get date and time

SYNOPSIS

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

time_t timeofday = time((time_t *)0)
time_t timeofday = time(tloc)
time_t *tloc;

ftime(tp)
struct timeb *tp;
```

DESCRIPTION

time() returns the time since 00:00:00 GMT, Jan. 1, 1970, measured in seconds.

If tloc is non-NULL, 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>:

```
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(1V), gettimeofday(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`:

```
struct tms {
    time_t tms_ftime;           /* user time */
    time_t tms_stime;          /* system time */
    time_t tms_cutime;         /* user time, children */
    time_t tms_cstime;         /* system time, children */
};
```

The children's times are the sum of the children's process times and their children's times.

SEE ALSO

`time(1V)`, `getrusage(2)`, `wait(2)`, `time(3C)`

NAME

timezone – get time zone name given offset from GMT

SYNOPSIS

char *timezone(zone, dst)

DESCRIPTION

timezone() attempts to return 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 Savings Time version. If the required name does not appear in a table built into the routine, the difference from GMT is produced; for instance, 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.

Note: the offset westward from Greenwich and an indication of whether Daylight Savings Time is in effect may not be sufficient to determine the name of the time zone, as the name may differ between different locations in the same time zone. Instead of using **timezone()** to determine the name of the time zone for a given time, that time should be converted to a '**struct tm**' using **localtime** (see **ctime(3)**) and the **tm_zone** field of that structure should be used. **timezone()** is retained for compatibility with existing programs.

SEE ALSO

ctime(3)

NAME

tmpfile – create a temporary file

SYNOPSIS

#include <stdio.h>

FILE *tmpfile()

DESCRIPTION

tmpfile() creates a temporary file using a name generated by **tmpnam(3S)**, and returns a corresponding **FILE** pointer. If the file cannot be opened, an error message is printed using **perror(3)**, and a **NULL** pointer is returned. The file will automatically be deleted when the process using it terminates. The file is opened for update (“w+”).

SEE ALSO

creat(2), unlink(2), fopen(3S), mktemp(3), perror(3), tmpnam(3S)

NAME

`tmpnam`, `tempnam` – create a name for a temporary file

SYNOPSIS

```
#include <stdio.h>
```

```
char *tmpnam (s)
```

```
char *s;
```

```
char *tempnam (dir, pfx)
```

```
char *dir, *pfx;
```

DESCRIPTION

These functions generate file names that can safely be used for a temporary file.

`tmpnam()` always generates a file name using the path-prefix defined as `P_tmpdir` in the `<stdio.h>` header file. If `s` is `NULL`, `tmpnam()` leaves its result in an internal static area and returns a pointer to that area. The next call to `tmpnam()` will destroy the contents of the area. If `s` is not `NULL`, it is assumed to be the address of an array of at least `L_tmpnam` bytes, where `L_tmpnam` is a constant defined in `<stdio.h>`; `tmpnam()` places its result in that array and returns `s`.

`tempnam()` allows the user to control the choice of a directory. The argument `dir` points to the name of the directory in which the file is to be created. If `dir` is `NULL` or points to a string which is not a name for an appropriate directory, the path-prefix defined as `P_tmpdir` in the `<stdio.h>` header file is used. If that directory is not accessible, `/tmp` will be used as a last resort. This entire sequence can be up-staged by providing an environment variable `TMPDIR` in the user's environment, whose value is the name of the desired temporary-file directory.

Many applications prefer their temporary files to have certain favorite initial letter sequences in their names. Use the `pfx` argument for this. This argument may be `NULL` or point to a string of up to five characters to be used as the first few characters of the temporary-file name.

`tempnam()` uses `malloc` to get space for the constructed file name, and returns a pointer to this area. Thus, any pointer value returned from `tempnam()` may serve as an argument to `free` (see `malloc(3)`). If `tempnam()` cannot return the expected result for any reason, that is, `malloc` failed, or none of the above mentioned attempts to find an appropriate directory was successful, a `NULL` pointer will be returned.

NOTES

These functions generate a different file name each time they are called.

Files created using these functions and either `fopen` or `creat` are temporary only in the sense that they reside in a directory intended for temporary use, and their names are unique. It is the user's responsibility to use `unlink(2)` to remove the file when its use is ended.

SEE ALSO

`creat(2)`, `unlink(2)`, `fopen(3S)`, `malloc(3)`, `mktemp(3)`, `tmpfile(3S)`

BUGS

If called more than 17,576 times in a single process, these functions will start recycling previously used names.

Between the time a file name is created and the file is opened, it is possible for some other process to create a file with the same name. This can never happen if that other process is using these functions or `mktemp`, and the file names are chosen so as to render duplication by other means unlikely.

NAME

tsearch, tfind, tdelete, twalk – manage binary search trees

SYNOPSIS

```
#include <search.h>

char *tsearch ((char *) key, (char **) rootp, compar)
int (*compar)( );

char *tfind ((char *) key, (char **) rootp, compar)
int (*compar)( );

char *tdelete ((char *) key, (char **) rootp, compar)
int (*compar)( );

void twalk ((char *) root, action)
void (*action)( );
```

DESCRIPTION

tsearch, tfind, tdelete, and twalk() are routines for manipulating binary search trees. They are generalized from Knuth (6.2.2) Algorithms T and D. All comparisons are done with a user-supplied routine. This routine is called with two arguments, the pointers to the elements being compared. It returns an integer less than, equal to, or greater than 0, according to whether the first argument is to be considered less than, equal to or greater than the second argument. The comparison function need not compare every byte, so arbitrary data may be contained in the elements in addition to the values being compared.

tsearch() is used to build and access the tree. *key* is a pointer to a datum to be accessed or stored. If there is a datum in the tree equal to **key* (the value pointed to by *key*), a pointer to this found datum is returned. Otherwise, **key* is inserted, and a pointer to it returned. Only pointers are copied, so the calling routine must store the data. *rootp* points to a variable that points to the root of the tree. A NULL value for the variable pointed to by *rootp* denotes an empty tree; in this case, the variable will be set to point to the datum which will be at the root of the new tree.

Like tsearch, tfind() will search for a datum in the tree, returning a pointer to it if found. However, if it is not found, tfind() will return a NULL pointer. The arguments for tfind() are the same as for tsearch.

tdelete() deletes a node from a binary search tree. The arguments are the same as for tsearch. The variable pointed to by *rootp* will be changed if the deleted node was the root of the tree. tdelete() returns a pointer to the parent of the deleted node, or a NULL pointer if the node is not found.

twalk() traverses a binary search tree. *root* is the root of the tree to be traversed. (Any node in a tree may be used as the root for a walk below that node.) *action* is the name of a routine to be invoked at each node. This routine is, in turn, called with three arguments. The first argument is the address of the node being visited. The second argument is a value from an enumeration data type `typedef enum { preorder, postorder, endorder, leaf } VISIT;` (defined in the <search.h> header file), depending on whether this is the first, second or third time that the node has been visited (during a depth-first, left-to-right traversal of the tree), or whether the node is a leaf. The third argument is the level of the node in the tree, with the root being level zero.

The pointers to the key and the root of the tree should be of type pointer-to-element, and cast to type pointer-to-character. Similarly, although declared as type pointer-to-character, the value returned should be cast into type pointer-to-element.

EXAMPLE

The following code reads in strings and stores structures containing a pointer to each string and a count of its length. It then walks the tree, printing out the stored strings and their lengths in alphabetical order.

```

#include <search.h>
#include <stdio.h>
struct node {          /* pointers to these are stored in the tree */
    char *string;
    int length;
};
char string_space[10000]; /* space to store strings */
struct node nodes[500]; /* nodes to store */
struct node *root = NULL; /* this points to the root */
main()
{
    char *strptr = string_space;
    struct node *nodeptr = nodes;
    void print_node( ), twalk( );
    int i = 0, node_compare( );
    while (gets(strptr) != NULL && i++ < 500) {
        /* set node */
        nodeptr->string = strptr;
        nodeptr->length = strlen(strptr);
        /* put node into the tree */
        (void) tsearch((char *)nodeptr, &root,
            node_compare);
        /* adjust pointers, so we don't overwrite tree */
        strptr += nodeptr->length + 1;
        nodeptr++;
    }
    twalk(root, print_node);
}
/*
This routine compares two nodes, based on an
alphabetical ordering of the string field.
*/
int
node_compare(node1, node2)
struct node *node1, *node2;
{
    return strcmp(node1->string, node2->string);
}
/*
This routine prints out a node, the first time
twalk encounters it.
*/
void
print_node(node, order, level)
struct node **node;
VISIT order;
int level;
{
    if (order == preorder || order == leaf) {
        (void)printf("string = %20s, length = %d\n",
            ((*node)->string, (*node)->length);
    }
}

```

SEE ALSO

bsearch(3), hsearch(3), lsearch(3)

DIAGNOSTICS

A NULL pointer is returned by `tsearch()` if there is not enough space available to create a new node.

A NULL pointer is returned by `tsearch`, `tfind()` and `tdelete()` if *rootp* is NULL on entry.

If the datum is found, both `tsearch()` and `tfind()` return a pointer to it. If not, `tfind()` returns NULL, and `tsearch()` returns a pointer to the inserted item.

WARNINGS

The *root* argument to `twalk()` is one level of indirection less than the *rootp* arguments to `tsearch()` and `tdelete`.

There are two nomenclatures used to refer to the order in which tree nodes are visited. `tsearch()` uses *preorder*, *postorder* and *endorder* to respectively refer to visiting a node before any of its children, after its left child and before its right, and after both its children. The alternate nomenclature uses *preorder*, *inorder* and *postorder* to refer to the same visits, which could result in some confusion over the meaning of *postorder*.

BUGS

If the calling function alters the pointer to the root, results are unpredictable.

NAME

ttyname, isatty – find name of a terminal

SYNOPSIS

char *ttyname(filedes)

isatty(filedes)

DESCRIPTION

ttyname() returns a pointer to the NULL-terminated path name of the terminal device associated with file descriptor *filedes*.

isatty() returns 1 if *filedes* is associated with a terminal device, 0 otherwise.

FILES

/dev/*

SEE ALSO

ioctl(2), ttytab(5)

DIAGNOSTICS

ttyname() returns a NULL pointer if *filedes* does not describe a terminal device in directory */dev*.

BUGS

The return value points to static data whose content is overwritten by each call.

NAME

ttyslot – find the slot in the utmp file of the current process

SYNOPSIS

ttyslot()

DESCRIPTION

ttyslot() returns the index of the current user's entry in the **/etc/utmp** file. This is accomplished by actually scanning the file **/etc/ttys** for the name of the terminal associated with the standard input, the standard output, or the error output (0, 1 or 2).

FILES

/etc/ttys
/etc/utmp

DIAGNOSTICS

A value of 0 is returned if an error was encountered while searching for the terminal name or if none of the above file descriptors is associated with a terminal device.

NAME

`ualarm` – schedule signal after interval in microseconds

SYNOPSIS

```
unsigned ualarm(value, interval)
unsigned value;
unsigned interval;
```

DESCRIPTION

This is a simplified interface to `setitimer` (see `getitimer(2)`).

`ualarm()` sends signal `SIGALRM`, see `signal(3)`, to the invoking process in a number of microseconds given by the *value* argument. Unless caught or ignored, the signal terminates the process.

If the *interval* argument is non-zero, the `SIGALRM` signal will be sent to the process every *interval* microseconds after the timer expires (for instance, after *value* microseconds have passed).

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 microseconds.

The return value is the amount of time previously remaining in the alarm clock.

SEE ALSO

`getitimer(2)`, `sigpause(2)`, `sigvec(2)`, `alarm(3C)`, `signal(3)`, `sleep(3)`, `usleep(3)`

NAME

ulimit – get and set user limits

SYNOPSIS

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

DESCRIPTION

This function is included for System V compatibility.

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

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

RETURN VALUE

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

ERRORS

The following error codes may be set in *errno*:

EPERM A user other than the super-user attempted to increase the file size limit.

SEE ALSO

brk(2), *getdtablesize(2)*, *getrlimit(2)*, *write(2V)*

NAME

`ungetc` – push character back into input stream

SYNOPSIS

```
#include <stdio.h>
```

```
ungetc(c, stream)
```

```
FILE *stream;
```

DESCRIPTION

`ungetc()` pushes the character `c` back onto an input stream. That character will be returned by the next `getc()` call on that stream. `ungetc()` returns `c`, and leaves the file stream unchanged.

One character of pushback is guaranteed provided something has been read from the stream and the stream is actually buffered. In the case that stream is `stdin`, one character may be pushed back onto the buffer without a previous read statement.

If `c` equals EOF, `ungetc()` does nothing to the buffer and returns EOF.

An `fseek(3S)` erases all memory of pushed back characters.

SEE ALSO

`fseek(3S)`, `getc(3S)`, `setbuf(3S)`

DIAGNOSTICS

`ungetc()` returns EOF if it cannot push a character back.

NAME

usleep – suspend execution for interval in microseconds

SYNOPSIS

usleep(useconds)
unsigned useconds;

DESCRIPTION

Suspend the current process for the number of microseconds specified by the argument. The actual suspension time may be an arbitrary amount longer because of other activity in the system, or because of the time spent in processing the call.

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 a short time later.

This routine is implemented using **setitimer** (see **getitimer(2)**); it requires eight system calls each time it is invoked. A similar but less compatible function can be obtained with a single **select(2)**; it would not restart after signals, but would not interfere with other uses of **setitimer**.

SEE ALSO

getitimer(2), **sigpause(2)**, **alarm(3C)**, **sleep(3)**, **ualarm(3)**

NAME

utime – set file times

SYNOPSIS

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

```
int utime(file, timep)
```

```
char *file;
```

```
time_t *timep;
```

DESCRIPTION

utime() sets the access and modification times of the file named by *file*.

If the *timep* argument is NULL, the access and modification times are set to the current time. A process must be the owner of the file or have write permission for the file to use **utime()** in this manner.

If the *timep* argument is not NULL, it is assumed to point to an array of two **time_t** values. The access time is set to the value of the first member, and the modification time is set to the value of the second member. The times contained in that array are measured in seconds since 00:00:00 GMT Jan 1, 1970. Only the owner of the file or the super-user may use **utime()** in this manner.

In either case, 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:

ENOTDIR	A component of the path prefix of <i>file</i> is not a directory.
ENAMETOOLONG	The length of a component of <i>file</i> exceeds 255 characters, or the length of <i>file</i> exceeds 1023 characters.
ENOENT	The file referred to by <i>file</i> does not exist.
EACCESS	Search permission is denied for a component of the path prefix of <i>file</i> .
ELOOP	Too many symbolic links were encountered in translating <i>file</i> .
EPERM	The effective user ID of the process is not super-user and not the owner of the file, and <i>timep</i> is not NULL.
EACCESS	The effective user ID is not super-user and not the owner of the file, write permission is denied for the file, and <i>timep</i> is NULL.
EIO	An I/O error occurred while reading from or writing to the file system.
EROFS	The file system containing the file is mounted read-only.
EFAULT	<i>file</i> or <i>timep</i> points outside the process's allocated address space.

SEE ALSO

stat(2), **utimes(2)**

NAME

values – machine-dependent values

SYNOPSIS

```
#include <values.h>
```

DESCRIPTION

This file contains a set of manifest constants, conditionally defined for particular processor architectures.

The model assumed for integers is binary representation (one's or two's complement), where the sign is represented by the value of the high-order bit.

<code>BITS(<i>type</i>)</code>	The number of bits in a specified type (for instance, <code>int</code>).
<code>HIBITS</code>	The value of a short integer with only the high-order bit set (in most implementations, <code>0x8000</code>).
<code>HIBITL</code>	The value of a long integer with only the high-order bit set (in most implementations, <code>0x80000000</code>).
<code>HIBITI</code>	The value of a regular integer with only the high-order bit set (usually the same as <code>HIBITS</code> or <code>HIBITL</code>).
<code>MAXSHORT</code>	The maximum value of a signed short integer (in most implementations, <code>0x7FFF</code> \equiv 32767).
<code>MAXLONG</code>	The maximum value of a signed long integer (in most implementations, <code>0x7FFFFFFF</code> \equiv 2147483647).
<code>MAXINT</code>	The maximum value of a signed regular integer (usually the same as <code>MAXSHORT</code> or <code>MAXLONG</code>).
<code>MAXFLOAT</code>	
<code>LN_MAXFLOAT</code>	The maximum value of a single-precision floating-point number, and its natural logarithm.
<code>MAXDOUBLE</code>	
<code>LN_MAXDOUBLE</code>	The maximum value of a double-precision floating-point number, and its natural logarithm.
<code>MINFLOAT</code>	
<code>LN_MINFLOAT</code>	The minimum positive value of a single-precision floating-point number, and its natural logarithm.
<code>MINDOUBLE</code>	
<code>LN_MINDOUBLE</code>	The minimum positive value of a double-precision floating-point number, and its natural logarithm.
<code>FSIGNIF</code>	The number of significant bits in the mantissa of a single-precision floating-point number.
<code>DSIGNIF</code>	The number of significant bits in the mantissa of a double-precision floating-point number.

FILES

`/usr/include/values.h`

SEE ALSO

`intro(3)`, `intro(3M)`

NAME

varargs – handle variable argument list

SYNOPSIS

```
#include <varargs.h>

function(va_alist) va_dcl
    va_list pvar ;
    va_start
    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(3S)`) but do not use `varargs()` are inherently nonportable, since different machines use different argument passing conventions. Routines with variable arguments lists *must* use `varargs()` functions in order to run correctly on Sun-4 systems.

`va_alist` is used in a function header to declare a variable argument list.

`va_dcl` is a declaration for `va_alist`. No semicolon should follow `va_dcl`.

`va_list` is a type defined for the variable used to traverse the list. 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`. The parameter `type` is a type name such that the type of a pointer to an object that has the specified type can be obtained simply by appending a `*` to `type`. If `type` disagrees with the type of the actual next argument (as promoted according to the default argument promotions), the behavior is undefined.

In standard C, arguments that are `char` or `short` are converted to `int` and should be accessed as `int`, arguments that are `unsigned char` or `unsigned short` are converted to `unsigned int` and should be accessed as `unsigned int`, and arguments that are `float` are converted to `double` and should be accessed as `double`. 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.

`va_alist` must encompass the entire arguments list. This insures that a `#define` statement can be used to redefine or expand its value.

The argument list (or its remainder) can be passed to another function using a pointer to a variable of type `va_list`— in which case a call to `va_arg` in the subroutine advances the argument-list pointer with respect to the caller as well.

EXAMPLE

This example is a possible implementation of `execl(3)`.

```
#include <varargs.h>
#define MAXARGS    100

/*    execl is called by
 *    execl(file, arg1, arg2, ..., (char *)0);
 */
execl (va_alist)
va_dcl
{
    va_list ap;
    char *file;
    char *args[MAXARGS];
    int argno = 0;

    va_start (ap);
    file = va_arg(ap, char *);
    while ((args[argno++] = va_arg(ap, char *)) != (char *)0)
        ;
    va_end (ap);
    return execv(file, args);
}
```

SEE ALSO

`execl(3)`, `printf(3S)`

BUGS

It is up to the calling routine to specify how many arguments there are, since it is not possible to determine this from the stack frame. For example, `execl()` is passed a zero pointer to signal the end of the list. `printf()` can tell how many arguments are supposed to be there by the format.

The macros `va_start` and `va_end` may be arbitrarily complex; for example, `va_start` might contain an opening brace, which is closed by a matching brace in `va_end`. Thus, they should only be used where they could be placed within a single complex statement.

NAME

`vlimit` – control maximum system resource consumption

SYNOPSIS

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

```
vlimit(resource, value) int 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:

<code>LIM_NORAISE</code>	A pseudo-limit; if set non-zero then the limits may not be raised. Only the super-user may remove the <i>noraise</i> restriction.
<code>LIM_CPU</code>	the maximum number of CPU-seconds to be used by each process
<code>LIM_FSIZE</code>	the largest single file which can be created
<code>LIM_DATA</code>	the maximum growth of the data+stack region using <code>sbrk</code> (see <code>brk(2)</code>) beyond the end of the program text
<code>LIM_STACK</code>	the maximum size of the automatically-extended stack region
<code>LIM_CORE</code>	the size of the largest core dump that will be created.
<code>LIM_MAXRSS</code>	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 <code>LIM_MAXRSS</code> .

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 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

`cs(1)`, `sh(1)`, `brk(2)`

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 `cs(1)`.

NAME

vprintf, vfprintf, vsprintf – print formatted output of a varargs argument list

SYNOPSIS

```
#include <stdio.h>
#include <varargs.h>

int vprintf(format, ap)
char *format;
va_list ap;

int vfprintf(stream, format, ap)
FILE *stream;
char *format;
va_list ap;

char *vsprintf(s, format, ap)
char *s, *format;
va_list ap;
```

DESCRIPTION

vprintf, vfprintf, and vsprintf() are the same as printf(3S), fprintf, and sprintf respectively, except that instead of being called with a variable number of arguments, they are called with an argument list as defined by varargs(3).

EXAMPLE

The following demonstrates how vfprintf() could be used to write an error routine.

```
#include <stdio.h>
#include <varargs.h>
...
/* error should be called like:
 * error(function_name, format, arg1, arg2...);
 * Note: function_name and format cannot be declared
 * separately because of the definition of varargs.
 */

/*VARARGS0*/
void
error (va_alist)
    va_dcl;
{
    va_list args;
    char *fmt;

    va_start(args);
    /* print name of function causing error */
    (void) fprintf(stderr, "ERROR in %s: ", va_arg(args, char *));
    fmt = va_arg(args, char *);
    /* print out remainder of message */
    (void) vfprintf(stderr, fmt, args);
    va_end(args);
    (void) abort();
}
```

SEE ALSO

printf(3S), varargs(3)

NAME

`vsyslog` – log message with a `varargs` argument list

SYNOPSIS

```
#include <syslog.h>
#include <varargs.h>

int vsyslog(priority, message, ap)
char *message;
va_list ap;
```

DESCRIPTION

`vsyslog()` is the same as `syslog(3)` except that instead of being called with a variable number of arguments, it is called with an argument list as defined by `varargs(3)`.

EXAMPLE

The following demonstrates how `vsyslog()` could be used to write an error routine.

```
#include <syslog.h>
#include <varargs.h>
...
/* error should be called like:
 *   error(pri, function_name, format, arg1, arg2...);
 * Note that pri, function_name, and format cannot be declared
 * separately because of the definition of varargs.
 */

/*VARARGS0*/
void
error(va_alist)
    va_dcl;
{
    va_list args;
    int pri;
    char *message;

    va_start(args);
    pri = va_arg(args, int);
        /* log name of function causing error */
    (void) syslog(pri, "ERROR in %s", va_arg(args, char *));
    message = va_arg(args, char *);
        /* log remainder of message */
    (void) vsyslog(pri, fmt, args);
    va_end(args);
    (void) abort();
}
```

SEE ALSO

`syslog(3)`, `varargs(3)`

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 `<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*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

`getrusage(2)`, `wait(2)`

NAME

xdr – library routines for external data representation

SYNOPSIS AND DESCRIPTION

These routines allow C programmers to describe arbitrary data structures in a machine-independent fashion. Data for remote procedure calls are transmitted using these routines.

```
xdr_array(xdrs, arrp, sizep, maxsize, elsize, elproc)
XDR *xdrs;
char **arrp;
u_int *sizep, maxsize, elsize;
xdrproc_t elproc;
```

A filter primitive that translates between variable-length arrays and their corresponding external representations. The parameter *arrp* is the address of the pointer to the array, while *sizep* is the address of the element count of the array; this element count cannot exceed *maxsize*. The parameter *elsize* is the *sizeof* each of the array's elements, and *elproc* is an XDR filter that translates between the array elements' C form, and their external representation. This routine returns one if it succeeds, zero otherwise.

```
xdr_bool(xdrs, bp)
XDR *xdrs;
bool_t *bp;
```

A filter primitive that translates between booleans (C integers) and their external representations. When encoding data, this filter produces values of either one or zero. This routine returns one if it succeeds, zero otherwise.

```
xdr_bytes(xdrs, sp, sizep, maxsize)
XDR *xdrs;
char **sp;
u_int *sizep, maxsize;
```

A filter primitive that translates between counted byte strings and their external representations. The parameter *sp* is the address of the string pointer. The length of the string is located at address *sizep*; strings cannot be longer than *maxsize*. This routine returns one if it succeeds, zero otherwise.

```
xdr_char(xdrs, cp)
XDR *xdrs;
char *cp;
```

A filter primitive that translates between C characters and their external representations. This routine returns one if it succeeds, zero otherwise. Note: encoded characters are not packed, and occupy 4 bytes each. For arrays of characters, it is worthwhile to consider `xdr_bytes()`, `xdr_opaque()` or `xdr_string()`.

```
void
xdr_destroy(xdrs)
XDR *xdrs;
```

A macro that invokes the destroy routine associated with the XDR stream, *xdrs*. Destruction usually involves freeing private data structures associated with the stream. Using *xdrs* after invoking `xdr_destroy()` is undefined.

xdr_double(xdrs, dp)

XDR *xdrs;

double *dp;

A filter primitive that translates between C double precision numbers and their external representations. This routine returns one if it succeeds, zero otherwise.

xdr_enum(xdrs, ep)

XDR *xdrs;

enum_t *ep;

A filter primitive that translates between C enums (actually integers) and their external representations. This routine returns one if it succeeds, zero otherwise.

xdr_float(xdrs, fp)

XDR *xdrs;

float *fp;

A filter primitive that translates between C floats and their external representations. This routine returns one if it succeeds, zero otherwise.

void

xdr_free(proc, objp)

xdrproc_t proc;

char *objp;

Generic freeing routine. The first argument is the XDR routine for the object being freed. The second argument is a pointer to the object itself. Note: the pointer passed to this routine is *not* freed, but what it points to *is* freed (recursively).

u_int

xdr_getpos(xdrs)

XDR *xdrs;

A macro that invokes the get-position routine associated with the XDR stream, *xdrs*. The routine returns an unsigned integer, which indicates the position of the XDR byte stream. A desirable feature of XDR streams is that simple arithmetic works with this number, although the XDR stream instances need not guarantee this.

long *

xdr_inline(xdrs, len)

XDR *xdrs;

int len;

A macro that invokes the in-line routine associated with the XDR stream, *xdrs*. The routine returns a pointer to a contiguous piece of the stream's buffer; *len* is the byte length of the desired buffer. Note: pointer is cast to long *.

Warning: *xdr_inline()* may return NULL (0) if it cannot allocate a contiguous piece of a buffer. Therefore the behavior may vary among stream instances; it exists for the sake of efficiency.

xdr_int(xdrs, ip)

XDR *xdrs;

int *ip;

A filter primitive that translates between C integers and their external representations. This routine returns one if it succeeds, zero otherwise.

```
xdr_long(xdrs, lp)
XDR *xdrs;
long *lp;
```

A filter primitive that translates between C long integers and their external representations. This routine returns one if it succeeds, zero otherwise.

```
void
xdrmem_create(xdrs, addr, size, op)
XDR *xdrs;
char *addr;
u_int size;
enum xdr_op op;
```

This routine initializes the XDR stream object pointed to by *xdrs*. The stream's data is written to, or read from, a chunk of memory at location *addr* whose length is no more than *size* bytes long. The *op* determines the direction of the XDR stream (either XDR_ENCODE, XDR_DECODE, or XDR_FREE).

```
xdr_opaque(xdrs, cp, cnt)
XDR *xdrs;
char *cp;
u_int cnt;
```

A filter primitive that translates between fixed size opaque data and its external representation. The parameter *cp* is the address of the opaque object, and *cnt* is its size in bytes. This routine returns one if it succeeds, zero otherwise.

```
xdr_pointer(xdrs, objpp, objsize, xdrobj)
XDR *xdrs;
char **objpp;
u_int objsize;
xdrproc_t xdrobj;
```

Like *xdr_reference()* except that it serializes NULL pointers, whereas *xdr_reference()* does not. Thus, *xdr_pointer()* can represent recursive data structures, such as binary trees or linked lists.

```
void
xdrrec_create(xdrs, sendsize, recvsize, handle, readit, writeit)
XDR *xdrs;
u_int sendsize, recvsize;
char *handle;
int (*readit) (), (*writeit) ();
```

This routine initializes the XDR stream object pointed to by *xdrs*. The stream's data is written to a buffer of size *sendsize*; a value of zero indicates the system should use a suitable default. The stream's data is read from a buffer of size *recvsize*; it too can be set to a suitable default by passing a zero value. When a stream's output buffer is full, *writeit* is called. Similarly, when a stream's input buffer is empty, *readit* is called. The behavior of these two routines is similar to the system calls *read* and *write*, except that *handle* is passed to the former routines as the first parameter. Note: the XDR stream's *op* field must be set by the caller.

Warning: this XDR stream implements an intermediate record stream. Therefore there are additional bytes in the stream to provide record boundary information.

xdrrec_endofrecord(xdrs, sendnow)

XDR *xdrs;
int sendnow;

This routine can be invoked only on streams created by `xdrrec_create()`. The data in the output buffer is marked as a completed record, and the output buffer is optionally written out if `sendnow` is non-zero. This routine returns one if it succeeds, zero otherwise.

xdrrec_eof(xdrs)

XDR *xdrs;
int empty;

This routine can be invoked only on streams created by `xdrrec_create()`. After consuming the rest of the current record in the stream, this routine returns one if the stream has no more input, zero otherwise.

xdrrec_skiprecord(xdrs)

XDR *xdrs;

This routine can be invoked only on streams created by `xdrrec_create()`. It tells the XDR implementation that the rest of the current record in the stream's input buffer should be discarded. This routine returns one if it succeeds, zero otherwise.

xdr_reference(xdrs, pp, size, proc)

XDR *xdrs;
char **pp;
u_int size;
xdrproc_t proc;

A primitive that provides pointer chasing within structures. The parameter `pp` is the address of the pointer; `size` is the *sizeof* the structure that `*pp` points to; and `proc` is an XDR procedure that filters the structure between its C form and its external representation. This routine returns one if it succeeds, zero otherwise.

Warning: this routine does not understand NULL pointers. Use `xdr_pointer()` instead.

xdr_setpos(xdrs, pos)

XDR *xdrs;
u_int pos;

A macro that invokes the set position routine associated with the XDR stream `xdrs`. The parameter `pos` is a position value obtained from `xdr_getpos()`. This routine returns one if the XDR stream could be repositioned, and zero otherwise.

Warning: it is difficult to reposition some types of XDR streams, so this routine may fail with one type of stream and succeed with another.

xdr_short(xdrs, sp)

XDR *xdrs;
short *sp;

A filter primitive that translates between C short integers and their external representations. This routine returns one if it succeeds, zero otherwise.

```
void
xdrstdio_create(xdrs, file, op)
XDR *xdrs;
FILE *file;
enum xdr_op op;
```

This routine initializes the XDR stream object pointed to by *xdrs*. The XDR stream data is written to, or read from, the Standard I/O stream *file*. The parameter *op* determines the direction of the XDR stream (either `XDR_ENCODE`, `XDR_DECODE`, or `XDR_FREE`).

Warning: the destroy routine associated with such XDR streams calls `fflush()` on the *file* stream, but never `fclose()`.

```
xdr_string(xdrs, sp, maxsize)
XDR
*xdrs;
char **sp;
u_int maxsize;
```

A filter primitive that translates between C strings and their corresponding external representations. Strings cannot be longer than *maxsize*. Note: *sp* is the address of the string's pointer. This routine returns one if it succeeds, zero otherwise.

```
xdr_u_char(xdrs, ucp)
XDR *xdrs;
unsigned char *ucp;
```

A filter primitive that translates between unsigned C characters and their external representations. This routine returns one if it succeeds, zero otherwise.

```
xdr_u_int(xdrs, up)
XDR *xdrs;
unsigned *up;
```

A filter primitive that translates between C unsigned integers and their external representations. This routine returns one if it succeeds, zero otherwise.

```
xdr_u_long(xdrs, ulp)
XDR *xdrs;
unsigned long *ulp;
```

A filter primitive that translates between C unsigned long integers and their external representations. This routine returns one if it succeeds, zero otherwise.

```
xdr_u_short(xdrs, usp)
XDR *xdrs;
unsigned short *usp;
```

A filter primitive that translates between C unsigned short integers and their external representations. This routine returns one if it succeeds, zero otherwise.

```
xdr_union(xdrs, dscmp, unp, choices, dfault)
XDR *xdrs;
int *dscmp;
char *unp;
struct xdr_discrim *choices;
bool_t (*defaultarm) (); /* may equal NULL */
```

A filter primitive that translates between a discriminated C union and its corresponding external representation. It first translates the discriminant of the union located at *dscmp*. This discriminant is always an `enum_t`. Next the union located at *unp* is translated. The parameter *choices* is a pointer to an array of `xdr_discrim()` structures. Each structure contains an ordered pair of [*value,proc*]. If the union's discriminant is equal to the associated *value*, then the *proc* is called to translate the union. The end of the `xdr_discrim()` structure array is denoted by a routine of value NULL. If the discriminant is not found in the *choices* array, then the *defaultarm* procedure is called (if it is not NULL). Returns one if it succeeds, zero otherwise.

```
xdr_vector(xdrs, arrp, size, elsize, elproc)
XDR *xdrs;
char *arrp;
u_int size, elsize;
xdrproc_t elproc;
```

A filter primitive that translates between fixed-length arrays and their corresponding external representations. The parameter *arrp* is the address of the pointer to the array, while *size* is the element count of the array. The parameter *elsize* is the *sizeof* each of the array's elements, and *elproc* is an XDR filter that translates between the array elements' C form, and their external representation. This routine returns one if it succeeds, zero otherwise.

```
xdr_void()
```

This routine always returns one. It may be passed to RPC routines that require a function parameter, where nothing is to be done.

```
xdr_wrapstring(xdrs, sp)
XDR *xdrs;
char **sp;
```

A primitive that calls `xdr_string(xdrs, sp,MAXUN.UNSIGNED)`; where `MAXUN.UNSIGNED` is the maximum value of an unsigned integer. `xdr_wrapstring()` is handy because the RPC package passes a maximum of two XDR routines as parameters, and `xdr_string()`, one of the most frequently used primitives, requires three. Returns one if it succeeds, zero otherwise.

SEE ALSO

`rpc(3N)`

Network Programming

NAME

ypclnt, yp_get_default_domain, yp_bind, yp_unbind, yp_match, yp_first, yp_next, yp_all, yp_order, yp_master, yperr_string, ypprot_err – Yellow Pages client interface

SYNOPSIS AND DESCRIPTION

This package of functions provides an interface to the Yellow Pages (YP) network lookup service. The package can be loaded from the standard library, `/usr/lib/libc.a`. Refer to `ypfiles(5)` and `ypserv(8)` for an overview of the Yellow Pages, including the definitions of `map` and `domain`, and a description of the various servers, databases, and commands that comprise the YP.

All input parameters names begin with *in*. Output parameters begin with *out*. Output parameters of type `char **` should be addresses of uninitialized character pointers. Memory is allocated by the YP client package using `malloc(3)`, and may be freed if the user code has no continuing need for it. For each *outkey* and *outval*, two extra bytes of memory are allocated at the end that contain NEWLINE and NULL, respectively, but these two bytes are not reflected in *outkeylen* or *outvallen*. *indomain* and *inmap* strings must be non-NULL and NULL-terminated. String parameters which are accompanied by a count parameter may not be NULL, but may point to NULL strings, with the count parameter indicating this. Counted strings need not be NULL-terminated.

All functions in this package of type *int* return 0 if they succeed, and a failure code (`YPERR_xxxx`) otherwise. Failure codes are described under DIAGNOSTICS below.

yp_bind (indomain);
char *indomain;

To use the YP services, the client process must be “bound” to a YP server that serves the appropriate domain using `yp_bind()`. Binding need not be done explicitly by user code; this is done automatically whenever a YP lookup function is called. `yp_bind()` can be called directly for processes that make use of a backup strategy (for example, a local file) in cases when YP services are not available.

void
yp_unbind (indomain)
char *indomain;

Each binding allocates (uses up) one client process socket descriptor; each bound domain costs one socket descriptor. However, multiple requests to the same domain use that same descriptor. `yp_unbind()` is available at the client interface for processes that explicitly manage their socket descriptors while accessing multiple domains. The call to `yp_unbind()` make the domain *unbound*, and free all per-process and per-node resources used to bind it.

If an RPC failure results upon use of a binding, that domain will be unbound automatically. At that point, the `ypclnt` layer will retry forever or until the operation succeeds, provided that `ypbind` is running, and either

- a) the client process cannot bind a server for the proper domain, or
- b) RPC requests to the server fail.

If an error is not RPC-related, or if `ypbind` is not running, or if a bound `ypserv` process returns any answer (success or failure), the `ypclnt` layer will return control to the user code, either with an error code, or a success code and any results.


```
yp_get_default_domain(outdomain);
char **outdomain;
```

The YP lookup calls require a map name and a domain name, at minimum. It is assumed that the client process knows the name of the map of interest. Client processes should fetch the node's default domain by calling `yp_get_default_domain()`, and use the returned *outdomain* as the *indomain* parameter to successive YP calls.

```
yp_match(indomain, inmap, inkey, inkeylen, outval, outvallen)
char *indomain;
char *inmap;
char *inkey;
int inkeylen;
char **outval;
int *outvallen;
```

`yp_match()` returns the value associated with a passed key. This key must be exact; no pattern matching is available.

```
yp_first(indomain, inmap, outkey, outkeylen, outval, outvallen)
char *indomain;
char *inmap;
char **outkey;
int *outkeylen;
char **outval;
int *outvallen;
```

`yp_first()` returns the first key-value pair from the named map in the named domain.

```
yp_next(indomain, inmap, inkey, inkeylen, outkey, outkeylen, outval, outvallen);
char *indomain;
char *inmap;
char *inkey;
int inkeylen;
char **outkey;
int *outkeylen;
char **outval;
int *outvallen;
```

`yp_next()` returns the next key-value pair in a named map. The *inkey* parameter should be the *outkey* returned from an initial call to `yp_first()` (to get the second key-value pair) or the one returned from the *n*th call to `yp_next()` (to get the *n*th + second key-value pair).

The concept of first (and, for that matter, of next) is particular to the structure of the YP map being processing; there is no relation in retrieval order to either the lexical order within any original (non-YP) data base, or to any obvious numerical sorting order on the keys, values, or key-value pairs. The only ordering guarantee made is that if the `yp_first()` function is called on a particular map, and then the `yp_next()` function is repeatedly called on the same map at the same server until the call fails with a reason of `YPERR_NOMORE`, every entry in the data base will be seen exactly once. Further, if the same sequence of operations is performed on the same map at the same server, the entries will be seen in the same order.

Under conditions of heavy server load or server failure, it is possible for the domain to become unbound, then bound once again (perhaps to a different server) while a client is running. This can cause a break in one of the enumeration rules; specific entries may be seen twice by the client, or not at all. This approach protects the client from error messages that would otherwise be returned in the midst of the enumeration. The next paragraph describes a better solution to enumerating all entries in a map.

```
yp_all(indomain, inmap, incallback);
char *indomain;
char *inmap;
struct ypall_callback incallback;
```

`yp_all()` provides a way to transfer an entire map from server to client in a single request using TCP (rather than UDP as with other functions in this package). The entire transaction takes place as a single RPC request and response. You can use `yp_all()` just like any other YP procedure, identify the map in the normal manner, and supply the name of a function which will be called to process each key-value pair within the map. You return from the call to `yp_all()` only when the transaction is completed (successfully or unsuccessfully), or your `foreach` function decides that it does not want to see any more key-value pairs.

The third parameter to `yp_all()` is

```
struct ypall_callback *incallback {
    int (*foreach)();
    char *data;
};
```

The function `foreach` is called

```
foreach(instatus, inkey, inkeylen, inval, invallen, indata);
int instatus;
char *inkey;
int inkeylen;
char *inval;
int invallen;
char *indata;
```

The `instatus` parameter will hold one of the return status values defined in `<rpcsvc/yp_prot.h>` — either `YP_TRUE` or an error code. (See `ypprot_err()`, below, for a function which converts a YP protocol error code to a `ypclnt` layer error code.)

The key and value parameters are somewhat different than defined in the synopsis section above. First, the memory pointed to by the `inkey` and `inval` parameters is private to the `yp_all()` function, and is overwritten with the arrival of each new key-value pair. It is the responsibility of the `foreach` function to do something useful with the contents of that memory, but it does not own the memory itself. Key and value objects presented to the `foreach` function look exactly as they do in the server's map — if they were not NEWLINE-terminated or NULL-terminated in the map, they will not be here either.

The `indata` parameter is the contents of the `incallback->data` element passed to `yp_all()`. The `data` element of the callback structure may be used to share state information between the `foreach` function and the mainline code. Its use is optional, and no part of the YP client package inspects its contents — cast it to something useful, or ignore it as you see fit.

The `foreach` function is a Boolean. It should return zero to indicate that it wants to be called again for further received key-value pairs, or non-zero to stop the flow of key-value pairs. If `foreach` returns a non-zero value, it is not called again; the functional value of `yp_all()` is then 0.

```

yp_order(indomain, inmap, outorder);
char *indomain;
char *inmap;
int *outorder;

```

yp_order() returns the order number for a map.

```

yp_master(indomain, inmap, outname);
char *indomain;
char *inmap;
char **outname;

```

yp_master() returns the machine name of the master YP server for a map.

```

char *yperr_string(incode)
int incode;

```

yperr_string() returns a pointer to an error message string that is NULL-terminated but contains no period or NEWLINE.

```

ypprot_err (incode)
unsigned int incode;

```

ypprot_err() takes a YP protocol error code as input, and returns a ypclnt layer error code, which may be used in turn as an input to yperr_string().

FILES

```

/usr/include/rpcsvc/ypclnt.h
/usr/include/rpcsvc/yp_prot.h
/usr/lib/libc.a

```

SEE ALSO

malloc(3), ypupdate(3N), ypfiles(5), ypserv(8)

DIAGNOSTICS

All integer functions return 0 if the requested operation is successful, or one of the following errors if the operation fails.

```

#define YPERR_BADARGS
    1      /* args to function are bad
#define YPERR_RPC
    2      /* RPC failure - domain has
#define YPERR_DOMAIN
    3      /* can't bind to server on
#define YPERR_MAP
    4      /* no such map in server's
#define YPERR_KEY
    5      /* no such key in map
#define YPERR_YPERR
    6      /* internal yp server or client
#define YPERR_RESRC
    7      /* resource allocation failure */
#define YPERR_NOMORE
    8      /* no more records in map
#define YPERR_PMAP
    9      /* can't communicate with portmapper */
#define YPERR_YPBIND

```

```
    10    /* can't communicate with ypbind */
#define YPERR_YPSErv
    11    /* can't communicate with ypserv */
#define YPERR_NODOM
    12    /* local domain name not set
#define YPERR_BADDBfR
    13    /* yp database is bad */
#define YPERR_VERSfR
    14    /* yp version mismatch */
#define YPERR_ACCESS
    15    /* access violation */
#define YPERR_BUSY
    16    /* database busy */
```

NAME

`yp_update` – changes yp information

SYNOPSIS

```
#include <rpcsvc/ypclnt.h>
```

```
yp_update(domain, map, ypop, key, keylen, data, datalen)
    char *domain;
    char *map;
    unsigned ypop
    char *key;
    int keylen;
    char *data;
    int datalen;
```

DESCRIPTION

`yp_update()` is used to make changes to the YP database. The syntax is the same as that of `yp_match()` except for the extra parameter *ypop* which may take on one of four values. If it is `YPOP_CHANGE` then the data associated with the key will be changed to the new value. If the key is not found in the database, then `yp_update()` will return `YPERR_KEY`. If *ypop* has the value `YPOP_INSERT` then the key-value pair will be inserted into the database. The error `YPERR_KEY` is returned if the key already exists in the database. To store an item into the database without concern for whether it exists already or not, pass *ypop* as `YPOP_STORE` and no error will be returned if the key already or does not exist. To delete an entry, the value of *ypop* should be `YPOP_DELETE`.

This routine depends upon secure RPC, and will not work unless the network is running secure RPC.

SEE ALSO

System and Network Administration

NAME

intro – introduction to the lightweight process library (LWP)

DESCRIPTION

The lightweight process library (LWP) provides a mechanism to support multiple threads of control that share a single address space. Under SunOS, the address space is derived from a single *forked* (“heavy-weight”) process. Each thread has its own stack segment (specified when the thread is created) so that it can access local variables and make procedure calls independently of other threads. The collection of threads sharing an address space is called a *pod*. Under SunOS, threads share all of the resources of the heavyweight process that contains the pod, including descriptors and signal handlers.

The LWP provides a means for creating and destroying threads, message exchange between threads, manipulating condition variables and monitors, handling synchronous exceptions, mapping asynchronous events into messages, mapping synchronous events into exceptions, arranging for special per-thread context, multiplexing the clock for timeouts, and scheduling threads both preemptively and non-preemptively.

The LWP system exists as a library of routines (`/usr/lib/liblwp.a`) linked in (`-llwp`) with a client program which should `#include` the file `<lwp/lwp.h>`. `main` is transparently converted into a lightweight process as soon as it attempts to use any LWP primitives.

When an object created by a LWP primitive is destroyed, every attempt is made to clean up after it. For example, if a thread dies, all threads blocked on sends to or receives from that thread are unblocked, and all monitor locks held by the dead thread are released.

Because there is no kernel support for threads at present, system calls effectively block the entire pod. By linking in the non-blocking I/O library (`-lnbio`) ahead of the LWP library, you can alleviate this problem for those system calls that can issue a signal when a system call would be profitable to try. This library (which redefines some system calls) uses asynchronous I/O and events (for example, `SIGCHLD` and `SIGIO`) to make blocking less painful. The system calls remapped by the nbio library are: `open(2V)`, `socket(2)`, `pipe(2)`, `close(2)`, `read(2V)`, `write(2V)`, `send(2)`, `recv(2)`, `accept(2)`, `connect(2)`, `select(2)`, `wait(2)`

RETURN VALUES

LWP primitives return `-1` on errors. Upon success, a non-negative integer is returned. See `lwp_perror(3L)` for details on error handling.

FILES

`/usr/lib/liblwp.a`
`/usr/lib/libnbio.a`
`/usr/include/lwp/check.h`
`/usr/include/lwp/lwp.h`
`/usr/include/lwp/lwperror.h`
`/usr/include/lwp/lwpmachdep.h`
`/usr/include/lwp/stackdep.h`

SEE ALSO

`accept(2)`, `close(2)`, `connect(2)`, `open(2V)`, `pipe(2)`, `read(2V)`, `recv(2)`, `select(2)`, `send(2)`, `socket(2)`, `wait(2)` `write(2V)`,
Lightweight Processes in the *System Services Overview*

INDEX

The following are the primitives currently supported, grouped roughly by function.

Thread Creation

`lwp_self(tid)`
`lwp_getstate(tid, statvec)`
`lwp_setregs(tid, machstate)`
`lwp_getregs(tid, machstate)`
`lwp_ping(tid)`
`lwp_create(tid, pc, prio, flags, stack, nargs, arg1, ..., argn)`
`lwp_destroy(tid)`

```

    lwp_enumerate(vec, maxsize)
    pod_setexit(status)
    pod_getexit()
    pod_exit(status)
    SAMETHREAD(t1, t2)
Thread Scheduling
    pod_setmaxpri(maxprio)
    pod_getmaxpri()
    pod_getmaxsize()
    lwp_resched(prio)
    lwp_setpri(tid, prio)
    lwp_sleep(timeout)
    lwp_suspend(tid)
    lwp_resume(tid)
    lwp_yield(tid)
    lwp_join(tid)
Error Handling
    lwp_geterr()
    lwp_perror(s)
    lwp_errstr()
Messages
    msg_send(tid, argbuf, argsize, resbuf, ressize)
    msg_rcv(tid, argbuf, argsize, resbuf, ressize, timeout)
    MSG_RECVALL(tid, argbuf, argsize, resbuf, ressize, timeout)
    msg_reply(tid)
    msg_enumsend(vec, maxsize)
    msg_enumrcv(vec, maxsize)
Event Mapping (Agents)
    agt_create(agt, event, memory)
    agt_enumerate(vec, maxsize)
    agt_trap(event)
Thread Synchronization: Monitors
    mon_create(mid)
    mon_destroy(mid)
    mon_enter(mid)
    mon_exit(mid)
    mon_enumerate(vec, maxsize)
    mon_waiters (mid, owner, vec, maxsize)
    mon_cond_enter(mid)
    mon_break(mid)
    MONITOR(mid)
    SAMEMON(m1, m2)
Thread Synchronization: Condition Variables
    cv_create(cv, mid)
    cv_destroy(cv)
    cv_wait(cv)
    cv_notify(cv)
    cv_send(cv, tid)
    cv_broadcast(cv)
    cv_enumerate(vec, maxsize)
    cv_waiters(cv, vec, maxsize)
    SAMECV(c1, c2)

```

Exception Handling

exc_handle(pattern, func, arg)
exc_unhandle()
(*exc_bound(pattern, arg))()
exc_notify(pattern)
exc_raise(pattern)
exc_on_exit(func, arg)
exc_uniqpatt()

Special Context Handling

lwp_ctxinit(tid, cookie)
lwp_ctxremove(tid, cookie)
lwp_ctxset(save, restore, ctxsize, optimise)
lwp_ctxmemget(mem, tid, ctx)
lwp_ctxmemset(mem, tid, ctx)
lwp_fpset(tid)
lwp_libcset(tid)

Stack Management

CHECK(location, result)
lwp_setstkcache(minsize, numstks)
lwp_newstk()
lwp_datastk(data, size, addr)
lwp_stkcswwset(tid, limit)
lwp_checkstkset(tid, limit)
STKTOP(s)

BUGS

There is no language support available from C.

There is no kernel support yet. Thus system calls in different threads cannot execute in parallel.

Killing a process that uses the non-blocking I/O library may leave objects (such as its standard input) in a non-blocking state. This could cause confusion to the shell.

LIST OF LWP LIBRARY FUNCTIONS

Name	Appears on Page	Description
agt_create()	agt_create(3L)	map LWP events into messages
agt_enumerate()	agt_create(3L)	map LWP events into messages
agt_trap()	agt_create(3L)	map LWP events into messages
CHECK()	lwp_newstk(3L)	LWP stack management
cv_broadcast()	cv_create(3L)	manage LWP condition variables
cv_create()	cv_create(3L)	manage LWP condition variables
cv_destroy()	cv_create(3L)	manage LWP condition variables
cv_enumerate()	cv_create(3L)	manage LWP condition variables
cv_notify()	cv_create(3L)	manage LWP condition variables
cv_send()	cv_create(3L)	manage LWP condition variables
cv_wait()	cv_create(3L)	manage LWP condition variables
cv_waiters()	cv_create(3L)	manage LWP condition variables
exc_bound()	exc_handle(3L)	LWP exception handling
exc_handle()	exc_handle(3L)	LWP exception handling
exc_notify()	exc_handle(3L)	LWP exception handling
exc_on_exit()	exc_handle(3L)	LWP exception handling
exc_raise()	exc_handle(3L)	LWP exception handling
exc_unhandle()	exc_handle(3L)	LWP exception handling
exc_uniqpatt()	exc_handle(3L)	LWP exception handling
lwp_checkstkset()	lwp_newstk(3L)	LWP stack management
lwp_create()	lwp_create(3L)	LWP thread creation and destruction primitives
lwp_ctxinit()	lwp_ctxinit(3L)	special LWP context operations
lwp_ctxremove()	lwp_ctxinit(3L)	special LWP context operations
lwp_ctxset()	lwp_ctxinit(3L)	special LWP context operations
lwp_ctxmemget()	lwp_ctxinit(3L)	special LWP context operations
lwp_ctxmemset()	lwp_ctxinit(3L)	special LWP context operations
lwp_destroy()	lwp_create(3L)	LWP thread creation and destruction primitives
lwp_enumerate()	lwp_status(3L)	LWP status information
lwp_errstr()	lwp_perror(3L)	LWP error handling
lwp_fpset()	lwp_ctxinit(3L)	special LWP context operations
lwp_geterr()	lwp_perror(3L)	LWP error handling
lwp_getstate()	lwp_status(3L)	LWP status information
lwp_getregs()	lwp_status(3L)	LWP status information
lwp_getregs()	lwp_status(3L)	LWP status information
lwp_ping()	lwp_status(3L)	LWP status information
lwp_join()	lwp_yield(3L)	control LWP scheduling
lwp_libcset()	lwp_ctxinit(3L)	special LWP context operations
lwp_newstk()	lwp_newstk(3L)	LWP stack management
lwp_datastk()	lwp_newstk(3L)	LWP stack management
lwp_perror()	lwp_perror(3L)	LWP error handling
lwp_resched()	lwp_yield(3L)	control LWP scheduling
lwp_resume()	lwp_yield(3L)	control LWP scheduling
lwp_self()	lwp_status(3L)	LWP status information
lwp_setpri()	lwp_yield(3L)	control LWP scheduling
lwp_setstkcache()	lwp_newstk(3L)	LWP stack management
lwp_sleep()	lwp_yield(3L)	control LWP scheduling
lwp_stkcswwset()	lwp_newstk(3L)	LWP stack management
lwp_suspend()	lwp_yield(3L)	control LWP scheduling
lwp_yield()	lwp_yield(3L)	control LWP scheduling
MINSTACKSZ()	lwp_newstk(3L)	LWP stack management

mon_break()	mon_create(3L)	LWP routines to manage critical sections
mon_cond_enter()	mon_create(3L)	LWP routines to manage critical sections
mon_create()	mon_create(3L)	LWP routines to manage critical sections
mon_destroy()	mon_create(3L)	LWP routines to manage critical sections
mon_enter()	mon_create(3L)	LWP routines to manage critical sections
mon_enumerate()	mon_create(3L)	LWP routines to manage critical sections
mon_exit()	mon_create(3L)	LWP routines to manage critical sections
mon_waiters()	mon_create(3L)	LWP routines to manage critical sections
MONITOR()	mon_create(3L)	LWP routines to manage critical sections
msg_enumrcv()	msg_send(3L)	LWP send and receive messages
msg_enumsend()	msg_send(3L)	LWP send and receive messages
msg_rcv()	msg_send(3L)	LWP send and receive messages
MSG_RECVALL()	msg_send(3L)	LWP send and receive messages
msg_reply()	msg_send(3L)	LWP send and receive messages
msg_send()	msg_send(3L)	LWP send and receive messages
pod_exit()	lwp_create(3L)	LWP thread creation and destruction primitives
pod_getexit()	lwp_create(3L)	LWP thread creation and destruction primitives
pod_getmaxpri()	pod_setmaxpri(3L)	control LWP scheduling priority
pod_getmaxsize()	pod_setmaxpri(3L)	control LWP scheduling priority
pod_setexit()	lwp_create(3L)	LWP thread creation and destruction primitives
pod_setmaxpri()	pod_setmaxpri(3L)	control LWP scheduling priority
SAMECV()	cv_create(3L)	manage LWP condition variables
SAMEMON()	mon_create(3L)	LWP routines to manage critical sections
SAMETHREAD()	lwp_create(3L)	LWP thread creation and destruction primitives
STKTOP()	lwp_newstk(3L)	LWP stack management

NAME

`agt_create`, `agt_enumerate`, `agt_trap` – map LWP events into messages

SYNOPSIS

```
#include <lwp/lwp.h>

thread_t
agt_create(agt, event, memory)
thread_t *agt;
int event;
caddr_t memory;

int
agt_enumerate(vec, maxsize)
thread_t vec[ ];
int maxsize;

int
agt_trap(event)
int event;
```

DESCRIPTION

Agents are entities that act like threads sending messages when an asynchronous event occurs. `agt_create()` creates an object called an *agent* which maps the asynchronous event *event* into messages that can be received with `msg_receive`. *agt* stores the handle on this object. *event* is a UNIX signal number.

`agt_trap()` causes the event, *event*, to generate an exception (see `exc_handle(3L)`). Once initialized using `agt_create()` or `agt_trap`, an event can not be remapped to a different style of handling. If traps are enabled, an event will cause the termination of the *thread* running at the time of the trap if the trap exception is not handled. If an exception handler is in place, an exception will be raised. If an agent exists for the event, the event is mapped into a message for the agent. If neither agent nor trap mapping is enabled, the default signal action (SIG_DFL) is applied to the *pod*. Use of standard UNIX signal handling facilities will defeat the event mapping mechanism.

The message sent by the agent (in the argument buffer) will look like any other message with the sender being the agent. The receive buffer is NULL. A message is always sent by an agent to the thread which created the agent.

All messages sent by an agent contain an `eventinfo_t`. This structure indicates the thread running at the time the interrupt happened, and the particular event that occurred. Some agent messages contain more information if the particular event warrants it. In this case, a struct containing an `eventinfo_t` as its first element is passed as the argument buffer. Definitions of these structures are contained in `<lwp/lwp.h>`.

An agent appears to the owning thread just like another thread. It must therefore have some memory for holding its message, as the sender and receiver must belong to the same address space. *memory* is the space an agent will use to store its message. Typically, this is on the stack of the thread that created the agent. It must be of the correct size for the kind of event being created (most events need something to store an `eventinfo_t`. SIGCHLD events need room for a `sigchld_t`.)

You should reply to an agent (using `msg_reply` (see `msg_send(3L)`) as you would reply to a thread. Although agents do not ordinarily lose events, the next agent message will not be delivered until a reply is sent to the agent. Thus, an agent appears to the client as an ordinary thread sending messages. An agent will only lose events if the total number of unreplyed-to events in a pod exceeds AGENTMEMORY.

`lwp_destroy()` is used to destroy an agent. All agents created by a thread automatically disappear when that thread dies. `agt_enumerate()` fills in a list with the ID's of all existing agents and returns the total number of agents. This primitive uses *maxsize* to avoid exceeding the capacity of the list. If the number of agents is greater than *maxsize*, only *maxsize* agents ID's are filled in *vec*. If *maxsize* is zero, `agt_enumerate()` returns the total number of agents.

The special event **LASTRITES** is caused by the termination of a thread. An agent for **LASTRITES** will be informed about every thread that terminates, regardless of cause. The **eventinfo_code** element of this agent will contain the stack argument that the dead thread was created with. Note: by allocating adjacent space above the thread stack, this argument can be used to point to private information about a thread. The **eventinfo_victimid** element will contain the id of the dead thread.

RETURN VALUE

Upon successful completion, **agt_create()** and **agt_trap()** return 0. Otherwise, -1 is returned.

agt_enumerate() returns the total number of agents.

ERRORS

agt_trap() will fail if one or more of the following is true:

LE_INVALIDARG Event specified does not exist.

LE_INUSE Agent in use for this event.

agt_create() will fail if one or more of the following are true:

LE_INVALIDARG Attempt to create agent for non-existent event.

LE_INUSE Trap mapping in use for this event.

SEE ALSO

exc_handle(3L), **msg_send(3L)**

BUGS

Signal handlers always take the **SIG_DFL** action when no agent manages the event.

If a descriptor used by a parent of the pod (such as its standard input) is marked non-blocking by a thread, it should be reset when the pod terminates to prevent the parent from receiving **EWOULDBLOCK** errors on the descriptor. There is no way to prevent this from happening if a pod is terminated with extreme prejudice (for instance, using **SIGKILL**).

If an agent reports that a descriptor has I/O available, there may be more than one occurrence of I/O available from that descriptor. Thus, being informed that **SIGIO** has occurred on socket *s* may mean that there are several messages waiting to be received from *s*. Clients should be careful to clean out all I/O from a descriptor before going back to sleep.

All system calls should be protected with loops testing for **EINTR** (and monitors if multiple threads can try to use system calls concurrently). An **lwp_sleep()** could result in a hidden clock interrupt for example.

WARNINGS

agt_trap() should not be used for asynchronous events. If an unsuspecting thread which has no exception handler is running at the time of a trapped event, it will be terminated.

Clients should not normally handle signals themselves since the agent mechanism assumes it is the only entity handling signals.

NAME

`cv_create`, `cv_destroy`, `cv_wait`, `cv_notify`, `cv_broadcast`, `cv_send`, `cv_enumerate`, `cv_waiters`, `SAMECV` – manage LWP condition variables

SYNOPSIS

```
#include <lwp/lwp.h>

cv_t
cv_create(cv, mid)
cv_t *cv;
mon_t mid;

int
cv_destroy(cv)
cv_t cv;

int
cv_wait(cv)
cv_t cv;

int
cv_notify(cv)
cv_t cv;

int
cv_send(cv, tid)
cv_t cv;
lwp_t tid

int
cv_broadcast(cv)
cv_t cv;

int
cv_enumerate(vec, maxsize)
cv_t vec[]; /* will contain list of all conditions */
int maxsize; /* maximum size of vec */

int
cv_waiters(cv, vec, maxsize)
cv_t cv; /* condition variable being interrogated */
thread_t vec[]; /* which threads are blocked on cv */
int maxsize; /* maximum size of vec */

SAMECV(c1, c2)
```

DESCRIPTION

Condition variables are useful for synchronization within monitors. By waiting on a condition variable, the currently-held monitor (a condition variable must *always* be used within a monitor) is released atomically and the invoking thread is suspended. When monitors are nested, monitor locks other than the current one are retained by the thread. At some later point, a different thread may awaken the waiting thread by issuing a notification on the condition variable. When the notification occurs, the waiting thread will queue to reacquire the monitor it gave up. It is possible to have different condition variables operating within the same monitor to allow selectivity in waking up threads.

`cv_create()` creates a new condition variable (returned in `cv`) which is bound to the monitor specified by `mid`. It is illegal to access (using `cv_wait`, `cv_notify`, `cv_send()` or `cv_broadcast`) a condition variable from a monitor other than the one it is bound to. `cv_destroy()` removes a condition variable.

cv_wait() blocks the current thread and releases the monitor lock associated with the condition (which must also be the monitor lock most recently acquired by the thread). Other monitor locks held by the thread are not affected. The blocked thread is enqueued by its scheduling priority on the condition.

cv_notify() awakens at most one thread blocked on the condition variable and causes the awakened thread to queue for access to the monitor released at the time it waited on the condition. It can be dangerous to use **cv_notify()** if there is a possibility that the thread being awakened is one of several threads that are waiting on a condition variable and the awakened thread may not be the one intended. In this case, use of **cv_broadcast()** is recommended.

cv_broadcast() is the same as **cv_notify()** except that *all* threads blocked on the condition variable are awakened. **cv_notify()** and **cv_broadcast()** do nothing if no thread is waiting on the condition. For both **cv_notify()** and **cv_broadcast**, the currently held monitor must agree with the one bound to the condition by **cv_create**.

cv_send() is like **cv_notify()** except that the particular thread **tid** is awakened. If this thread is not currently blocked on the condition, **cv_send()** reports an error.

cv_enumerate() lists the ID of all of the condition variables. The value returned is the total number of condition variables. The vector supplied is filled in with the ID's of condition variables. **cv_waiters()** lists the ID's of the threads blocked on the condition variable **cv** and returns the number of threads blocked on **cv**. For both **cv_enumerate()** and **cv_waiters**, *maxsize* is used to avoid exceeding the capacity of the list *vec*. If the number of entries to be filled is greater than *maxsize*, only *maxsize* entries are filled in *vec*. It is legal in both of these primitives to specify a *maxsize* of 0.

SAMECV is a convenient predicate used to compare two condition variables for equality.

RETURN VALUE

cv_create, **cv_destroy**, **cv_send**, **cv_wait**, **cv_notify**, **cv_broadcast**:

A 0 return indicates success; -1 indicates an error.

cv_enumerate() returns the total number of condition variables.

cv_waiters() returns the number of threads blocked on a condition variable.

ERRORS

cv_destroy() will fail if one or more of the following is true:

LE_INUSE Attempt to destroy condition variable being waited on by a thread.

LE_NONEXIST Attempt to destroy non-existent condition variable.

cv_wait() will fail if one or more of the following is true:

LE_NOTOWNED Attempt to wait on a condition without possessing the correct monitor lock.

LE_NONEXIST Attempt to wait on non-existent condition variable.

cv_notify() will fail if one or more of the following is true:

LE_NOTOWNED Attempt to notify condition variable without possessing the correct monitor.

LE_NONEXIST Attempt to notify non-existent condition variable.

cv_send() will fail if one or more of the following is true:

LE_NOTOWNED Attempt to awaken condition variable without possessing the correct monitor lock.

LE_NONEXIST Attempt to awaken non-existent condition variable.

LE_NOWAIT The specified thread is not currently blocked on the condition.

cv_broadcast() will fail if one or more of the following is true:

LE_NOTOWNED Attempt to broadcast condition without possessing the correct monitor lock.

LE_NONEXIST Attempt to broadcast non-existent condition variable.

CV_CREATE(3L)

LIGHTWEIGHT PROCESSES LIBRARY

CV_CREATE(3L)

SEE ALSO

mon_create(3L)

NAME

`exc_handle`, `exc_unhandle`, `exc_bound`, `exc_notify`, `exc_raise`, `exc_on_exit`, `exc_uniqpatt` – LWP exception handling

SYNOPSIS

```
#include <lwp/lwp.h>

int
exc_handle(pattern, func, arg)
int pattern;
caddr_t (*func)();
caddr_t arg;

int
exc_raise(pattern)
int pattern;

int
exc_unhandle()

caddr_t
(*exc_bound(pattern, arg))()
int pattern;
caddr_t *arg;

int
exc_notify(pattern)
int pattern;

int
exc_on_exit(func, arg)
void (*func)();
caddr_t arg;

int
exc_uniqpatt()
```

DESCRIPTION

These primitives can be used to manage exceptional conditions in a thread. Basically, raising an exception is a more general form of non-local goto or *longjmp*, but the invocation is pattern-based. It is also possible to *notify* an exception handler whereby a function supplied by the exception handler is invoked and control is returned to the raiser of the exception. Finally, one can establish a handler which is always invoked upon procedure exit, regardless of whether the procedure exits using a *return* or an exception raised to a handler established prior to the invocation of the exiting procedure.

`exc_handle()` is used to establish an exception handler. `exc_handle()` returns 0 to indicate that a handler has been established. A return of -1 indicates an error in trying to establish the exception handler. If it returns something else, an exception has occurred and any procedure calls deeper than the one containing the handler have disappeared. All exception handlers established by a procedure are automatically discarded when the procedure terminates.

`exc_handle()` binds a *pattern* to the handler, where a pattern is an integer, and two patterns *match* if their values are equal. When an exception is raised with `exc_raise`, the most recent handler that has established a matching pattern will catch the exception. A special pattern (CATCHALL) is provided which matches any `exc_raise()` pattern. This is useful for handlers which know that there is no chance the resources allocated in a routine can be reclaimed by previous routines in the call chain.

The other two arguments to `exc_handle()` are a function and an argument to that function. `exc_bound()` retrieves these arguments from an `exc_handle()` call made by the specified thread. By using `exc_bound()` to retrieve and call a function bound by the exception handler, a procedure can raise a *notification exception* which allows control to return to the raiser of the exception after the exception is handled.

exc_raise() allows the caller to transfer control (do a non-local goto) to the matching **exc_handle**. This matching exception handler is destroyed after the control transfer. At this time, it behaves as if **exc_handle()** returns with the *pattern* from **exc_raise()** as the return value. Note: *func* of **exc_handle()** is not called using **exc_raise()** — it is only there for notification exceptions. Because the exception handler returns the pattern that invoked it, it is possible for a handler that matches the CATCHALL pattern to *reraise* the exact exception it caught by using **exc_raise()** on the caught pattern. It is illegal to handle or raise the pattern 0 or the pattern -1. Handlers are searched for pattern matches in the reverse execution order that they are set (i.e., the most recently established handler is searched first).

exc_unhandle() destroys the most recently established exception handler set by the current thread. It is an error to destroy an exit-handler set up by **exc_on_exit**. When a procedure exits, all handlers and exit handlers set in the procedure are automatically deallocated.

exc_notify() is a convenient way to use **exc_bound**. The function which is bound to *pattern* is retrieved. If the function is not NULL, the function is called with the associated argument and the result is returned. If the function is NULL, **exc_raise(pattern)** is returned.

exc_on_exit() specifies an exit procedure and argument to be passed to the exit procedure, which is called when the procedure which sets an exit handler using **exc_on_exit()** exits. The exit procedures (more than one may be set) will be called regardless if the setting procedure is exited using a *return* or an **exc_raise**. Because the exit procedure is called as if the handling procedure had returned, the argument passed to it should not contain addresses on the handler's stack. However, any value returned by the procedure which established the exit procedure is preserved no matter what the exit procedure returns. This primitive is used in the MONITOR macro to enforce the monitor discipline on procedures.

Some signals can be considered to be synchronous traps. They are usually the starred (*) signals in the **signal(3)** man pages. These are: SIGSYS, SIGBUS, SIGEMT, SIGFPE, SIGILL, SIGTRAP, SIGSEGV. If an event is marked as a trap using **agt_trap** (see **agt_create(3L)**) the event will generate exceptions instead of agent messages. This mapping is per-pod, not per-thread. A thread which handles the signal number of one of these as the pattern for **exc_handle()** will catch such a signal as an exception. The exception will be raised as an **exc_notify()** so either escape or notification style exceptions can be used, depending on what the matching **exc_handle()** provides. If the exception is not handled, the thread will terminate. Note: it can be dangerous to supply an exception handler to treat stack overflow since the client's stack is used in raising the exception.

exc_uniqpatt() returns an exception pattern that is not any of the pre-defined patterns (any of the synchronous exceptions or -1 or CATCHALL). Each call to **exc_uniqpatt()** results in a different pattern. If **exc_uniqpatt()** cannot guarantee uniqueness, -1 is returned instead the *first* time this happens. Subsequent calls after this error result in patterns which may be duplicates.

RETURN VALUE

exc_uniqpatt() returns -1 the *first* time it fails. Otherwise, it returns a unique pattern.

When **exc_handle()** is called, a return value of 0 indicates success and -1 indicates error. When **exc_handle()** returns because of a matching **exc_raise()** call, it returns the *pattern* raised by **exc_raise**.

Upon successful completion, **exc_raise()** transfers control to the matching **exc_handle()** and does not return. If there is an error, **exc_raise()** returns -1.

Upon successful completion, **exc_unhandle()** returns 0. It returns -1 if there is an error.

exc_bound() returns a pointer to a function or 0 if no function was bound.

Upon successful completion, **exc_notify()** returns the return value of a function or transfers control to a matching **exc_handle()** and does not return. It returns -1 if there is an error.

exc_on_exit() returns 0.

ERRORS

exc_unhandle() will fail if one or more of the following is true:

LE_NONEXIST Attempt to remove an exit handler or a non-existent handler.

exc_raise() will fail if one or more of the following is true:

LE_NONEXIST No context found to raise an exception to.

LE_INVALIDARG Attempt to raise an illegal pattern (-1 or 0).

exc_handle() will fail if one or more of the following is true:

LE_INVALIDARG Attempt to handle an illegal pattern (-1 or 0).

exc_uniqpatt() will fail if one or more of the following is true:

LE_REUSE Possible reuse of existing object. **agt_create(3L)**

BUGS

The stack may not contain useful information after an exception has been caught so post-exception debugging can be difficult. The reason for this is that a given handler may call procedures that trash the stack before reraising an exception.

The distinction between traps and interrupts can be problematical.

The environment restored on **exc_raise()** consists of the registers at the time of the **exc_handle**. As a result, modifications to register variables between the times of **exc_handle()** and **exc_raise()** will not be seen. This problem does not occur in the sun4 implementation.

WARNINGS

exc_on_exit() passes a simple type as an argument to the exit routine. If you need to pass a complex type, such as a *thread_t*, *mon_t*, or *cv_t*, pass a pointer to the object instead.

NAME

`lwp_create`, `lwp_destroy`, `SAMETHREAD`, `pod_setexit`, `pod_getexit`, `pod_exit` – LWP thread creation and destruction primitives

SYNOPSIS

```
#include <lwp/lwp.h>
#include <lwp/stackdep.h>

int
lwp_create(tid, func, prio, flags, stack, nargs, arg1, ..., argn)
thread_t *tid;
void (*func)();
int prio;
int flags;
stkalign_t *stack;
int nargs;
int arg1, ..., argn;

int
lwp_destroy(tid)
thread_t tid;

void
pod_setexit(status)
int status;

int
pod_getexit(status)
int status;

void
pod_exit(status)
int status

SAMETHREAD(t1, t2)
```

DESCRIPTION

`lwp_create()` creates a lightweight process which starts at address *func* and has stack segment *stack*. If *stack* is NULL, the thread is created in a suspended state (see below) and no stack or pc is bound to the thread. *prio* is the scheduling priority of the thread (higher priorities are favored by the scheduler). The identity of the new thread is filled in the reference parameter *tid*. *flags* describes some options on the new thread. `LWPSUSPEND` creates the thread in suspended state (see `lwp_yield(3L)`). `LWPNOLASTRITES` will disable the `LASTRITES` agent message when the thread dies. The default (0) is to create the thread in running state with `LASTRITES` reporting enabled. `LWPSEVER` indicates that a thread is only viable as long as non-`LWPSEVER` threads are alive. The pod will terminate if the only living threads are marked `LWPSEVER` and blocked on a lwp resource (for instance, waiting for a message to be sent). *nargs* is the number (0 or more) of simple-type (int) arguments supplied to the thread.

The first time a lwp primitive is used, the lwp library automatically converts the caller (i.e., `main`) into a thread with the highest available scheduling priority (see `pod_setmaxpri(3L)`). The identity of this thread can be retrieved using `lwp_self` (see `lwp_status(3L)`). This thread has the normal SunOS stack given to any *forked* process.

Scheduling is, by default, non-preemptive within a priority, and within a priority, threads enter the run queue on a FIFO basis (that is, whenever a thread becomes eligible to run, it goes to the end of the run queue of its particular priority). Thus, a thread continues to run until it voluntarily relinquishes control or an event (including thread creation) occurs to enable a higher priority thread. Some primitives may cause the current thread to block, in which case the unblocked thread with the highest priority runs next. When several threads are created with the same priority, they are queued for execution in the order of creation.

This order may not be preserved as threads yield and block within a priority. If an agent owned by a thread with a higher priority is invoked, that thread will preempt the currently running one.

There is no concept of ancestry in threads: the creator of a thread has no special relation to the thread it created. When all threads have died, the pod terminates.

lwp_destroy() is a way to explicitly terminate a thread or agent (instead of having an executing thread “fall though”, which also terminates the thread). *tid* specifies the id of the thread or agent to be terminated. If *tid* is SELF, the invoking thread is destroyed. Upon termination, the resources (messages, monitor locks, agents) owned by the thread are released, in some cases resulting in another thread being notified of the death of its peer (by having a blocking primitive become unblocked with an error indication). A thread may terminate itself explicitly, although self-destruction is automatic when it returns from the procedure specified in the **lwp_create()** primitive.

pod_setexit() sets the exit status for a pod. This value will be returned to the parent process of the pod when the pod dies (default is 0). **exit(3)** terminates the current *thread*, using the argument supplied to *exit* to set the current value of the exit status. **on_exit(3)** establishes an action that will be taken when the entire pod terminates. **pod_exit()** is available to terminate the pod immediately with the final actions established by **on_exit**. If you wish to terminate the pod immediately, **pod_exit()** or **exit(2)** should be used. **pod_getexit()** returns the current value of the pod’s exit status.

SAMETHREAD is a convenient predicate used to compare two threads for equality.

RETURN VALUE

Upon successful completion, **lwp_create**, and **lwp_destroy()** return 0. Otherwise, -1 is returned. **pod_getexit()** returns the current exit status of the pod.

ERRORS

lwp_create() will fail if one or more of the following are true:

LE_NOROOM	Unable to allocate memory for thread context.
LE_INVALIDARG	Too many arguments (> 512).
LE_ILLPRIO	Illegal priority.

lwp_destroy() will fail if one or more of the following are true:

LE_NONEXIST	Attempt to destroy a thread or agent that does not exist.
-------------	---

SEE ALSO

exit(3), **lwp_yield(3L)**, **on_exit(3)**, **pod_setmaxpri(3L)**

WARNINGS

Some special threads may be created silently by the lwp library. These include an *idle* thread that runs when no other activity is going on, and a *reaper* thread that frees stacks allocated by **lwp_newstk**. These special threads will show up in status calls. A pod will terminate if these special threads are the only ones extant.

NAME

`lwp_ctxinit`, `lwp_ctxremove`, `lwp_ctxset`, `lwp_ctxmemget`, `lwp_ctxmemset`, `lwp_fpset`, `lwp_libcset` – special LWP context operations

SYNOPSIS

```
#include <lwp/lwp.h>

int
lwp_ctxset(save, restore, ctxsize, optimise)
void (*save)(/* caddr_t ctx, thread_t old, thread_t new */);
void (*restore)(/* caddr_t ctx, thread_t old, thread_t new */);
unsigned int ctxsize;
int optimise;

int
lwp_ctxinit(tid, cookie)
thread_t tid;          /* thread with special contexts */
int cookie;           /* type of context */

int
lwp_ctxremove(tid, cookie)
thread_t tid;
int cookie;

int
lwp_ctxmemget(mem, tid, ctx)
caddr_t mem;
thread_t tid;
int ctx;

int
lwp_ctxmemset(mem, tid, ctx)
caddr_t mem;
thread_t tid;
int ctx;

int
lwp_fpset(tid)
thread_t tid;          /* thread utilizing floating point hardware */

int
lwp_libcset(tid)
thread_t tid;          /* thread utilizing errno */
```

DESCRIPTION

Normally on a context switch, only machine registers are saved/restored to provide each thread its own virtual machine. However, there are other hardware and software resources which can be multiplexed in this way. For example, floating point registers can be used by several threads in a pod. As another example, the global value `errno` in the standard C library may be used by all threads making system calls.

To accommodate the variety of contexts that a thread may need without requiring all threads to pay for unneeded switching overhead, `lwp_ctxinit()` is provided. This primitive allows a client to specify that a given thread requires certain context to be saved and restored across context switches (by default just the machine registers are switched). More than one special context may be given to a thread.

To use `lwp_ctxinit`, it is first necessary to define a special context. `lwp_ctxset()` specifies save and restore routines, as well as the size of the context that will be used to hold the switchable state. The *save* routine will automatically be invoked when an active thread is blocked and the *restore* routine will be invoked when a blocked thread is restarted. These routines will be passed a pointer to a buffer (initialized to all 0's) of size *ctxsize* which is allocated by the LWP library and used to hold the volatile state. In addition, the

identity of the thread whose special context is being saved (*old*) and the identity of the thread being restarted (*new*) are passed in to the *save* and *restore* routines. `lwp_ctxset()` returns a cookie used by subsequent `lwp_ctxinit()` calls to refer to the kind of context just defined. If the *optimise* flag is TRUE, a special context switch action will not be invoked unless the thread resuming execution differs from the last thread to use the special context and also uses the special context. If the *optimise* flag is FALSE, the *save* routine will always be invoked immediately when the thread using this context is scheduled out and the *restore* routine will be invoked immediately when a new thread using this context is scheduled in. Note that an unoptimised special context is protected from threads which do not use the special context but which do affect the context state. `lwp_ctxremove()` can be used to remove a special context installed by `lwp_ctxinit`.

Because context switching is done by the scheduler on behalf of a thread, it is an error to use an LWP primitive in an action done at context switch time. Also, the stack used by the *save* and *restore* routines belongs to the scheduler, so care should be taken not to use lots of stack space. As a result of these restrictions, only knowledgeable users should write their own special context switching routines.

`lwp_ctxmemget` and `lwp_ctxmemset` are used to retrieve and set (respectively) the memory associated with a given special context (*ctx*) and a given thread (*tid*). *mem* is the address of client memory that will hold the context information being retrieved or set. Note that the special context *save* and *restore* routines may be NULL, so pure data may be associated with a given thread using these primitives.

Several kinds of special contexts are predefined. To allow a thread to share floating point hardware with other threads, the `lwp_fpset()` primitive is available. The floating-point hardware bound at compile-time is selected automatically. To multiplex the global variable `errno`, `lwp_libcset()` is used to have `errno` become part of the context of thread *tid*.

Special contexts can be used to assist in managing stacks. See `lwp_newstk(3L)` for details.

RETURN VALUE

`lwp_ctxset()` returns a cookie to be used by subsequent `lwp_ctxinit()` calls, -1 if unable to define the context.

ERRORS

`lwp_ctxinit()` will fail if one or more of the following are true:

LE_INUSE This special context already set for this thread.

`lwp_ctxremove()` will fail if one or more of the following are true:

LE_NONEXIST The specified context is not set for this thread.

`lwp_ctxset()` will fail if one or more of the following are true:

LE_NOROOM Unable to allocate memory to define special context.

SEE ALSO

`lwp_newstk(3L)`

BUGS

The floating point contexts should be initialized implicitly for those threads that use floating point.

NAME

`lwp_checkstkset`, `lwp_stkcswset`, `CHECK`, `lwp_setstkcache`, `lwp_newstk`, `lwp_datastk`, `STKTOP` – LWP stack management

SYNOPSIS

```
#include <lwp/lwp.h>
#include <lwp/check.h>
#include <lwp/lwpmachdep.h>
#include <lwp/stackdep.h>

CHECK(location, result)

int
lwp_checkstkset(tid, limit)
thread_t tid;
caddr_t limit;

int
lwp_stkcswset(tid, limit)
thread_t tid;
caddr_t limit;

int
lwp_setstkcache(minstksz, numstks)
int minstksz;
int numstks;

stkalign_t *
lwp_newstk()

stkalign_t *
lwp_datastk(data, size, addr)
caddr_t data;
int size;
caddr_t *addr;

STKTOP(s)
```

DESCRIPTION

Stacks are problematical with lightweight processes. What is desired is that stacks for each thread are red-zone protected so that one thread's stack does not unexpectedly grow into the stack of another. In addition, stacks should be of infinite length, grown as needed. The process stack is a maximum-sized segment (see `getrlimit(2)`.) This stack is redzone protected, and you can even try to extend it beyond its initial maximum size in some cases. With SunOS 4.x, it is possible to efficiently allocate large stacks that have red zone protection, and the LWP library provides some support for this. For those systems that do not have flexible memory management, the LWP library provides assistance in dealing with the problems of maintaining multiple stacks.

The stack used by *main* is the same stack that the system allocates for a *forked* process. For allocating other thread stacks, the client is free to use any statically or dynamically allocated memory (using memory from *main*'s stack is subject to the stack resource limit for any *forked* process). In addition, the `LAS-TRITES` agent message is available to free allocated resources when a thread dies. Any stack should be at least `MINSTACKSZ` *stkalign_t*'s large because the LWP library will use the client stack to execute primitives. For very fast dynamically allocated stacks, a stack cacheing mechanism is available. `lwp_setstkcache()` allocates a cache of stacks. Each time the cache is empty, it is filled with *numstks* new stacks, each containing at least *minstksz* bytes. *minstksz* will automatically be augmented to take into account the stack needs of the LWP library. `lwp_newstk()` returns a cached stack that is suitable for use in an `lwp_create()` call. `lwp_setstkcache()` must be called (once) prior to any use of `lwp_newstk`. If running under SunOS 4.x, the stacks allocated by `lwp_newstk()` will be red-zone protected (an attempt to reference below the stack bottom will result in a `SIGSEGV` event).

Threads created with stacks from `lwp_newstk()` should not use the `NOLASTRITES` flag. If they do, cached stacks will not be returned to the cache when a thread dies.

`lwp_datastk()` also returns a red-zone protected stack like `lwp_newstk()` does. It copies any amount of data (subject to the size limitations imposed by `lwp_setstkcache`) onto the stack *above* the stack top that it returns. `data` points to information of *size* bytes to be copied. The exact location where the data is stored is returned in the reference parameter `addr`. Because `lwp_create()` only passes simple types to the newly-created thread, `lwp_datastk()` is useful to pass a more complex argument: Call `lwp_datastk()` to get an initialized stack, and pass the address of the data structure (`addr`) as an argument to the new thread.

A *reaper* thread running at the maximum pod priority is created by `lwp_setstkcache`. It's action may be delayed by other threads running at that priority, so it is suggested that the maximum pod priority not be used for client-created threads when `lwp_newstk()` is being used. Altering the maximum pod priority with `pod_setmaxpri()` will have the side effect of increasing the reaper thread priority as well.

The stack address passed to `lwp_create()` represents the top of the stack: the LWP library will not use any addresses at or above it. Thus, it is safe to store information above the stack top if there is room there.

For stacks that are not protected with hardware redzones, some protection is still possible. For any thread *tid* with stack boundary *limit* made part of a special context with `lwp_checkstkset`, the `CHECK` macro may be used. This macro, if used at the beginning of each procedure (and before local storage is initialized (it is okay to *declare* locals though)), will check that the stack limit has not been violated. If it has, the non-local *location* will be set to *result* and the procedure will return. `CHECK` is not perfect, as it is possible to call a procedure with many arguments after `CHECK` validates the stack, only to have these arguments clobber the stack before the new procedure is entered.

`lwp_stkcswwset()` checks at context-switch time the stack belonging to thread *tid* for passing stack boundary *limit*. In addition, a checksum at the bottom of the stack is validated to ensure that the stack did not temporarily grow beyond its limit. This is automated and more efficient than using `CHECK`, but by the time a context switch occurs, it's too late to do much but `abort(3)` if the stack was clobbered.

To portably use statically allocated stacks, the macros in `stackdep.h` should be used. Declare a stack *s* to be an array of `stkalign_t`'s, and pass the stack to `lwp_create()` as `STKTOP(s)`.

RETURN VALUES

`lwp_newstk` and `lwp_datastk()` return 0 on failure, else a valid new stack address.

`lwp_setstkcache()` returns the actual size of the stacks allocated in the cache.

SEE ALSO

`getrlimit(2)`, `abort(3)`

BUGS

C should provide support for heap-allocated stacks at procedure entry time. The hardware should be segment-based to eliminate the problem altogether.

WARNING

`lwp_datastk()` should not be directly used in a `lwp_create()` call since C does not guarantee the order in which arguments to a function are evaluated.

NAME

`lwp_geterr`, `lwp_perror`, `lwp_errstr` – LWP error handling

SYNOPSIS

```
#include <lwp/lwp.h>
#include <lwp/lwperror.h>

lwp_err_t lwp_geterr();

void
lwp_perror(s)
char *s;

char **lwp_errstr();
```

DESCRIPTION

When a primitive fails (returns `-1`), `lwp_geterr()` can be used to obtain the identity of the error (which is part of the context for each lwp). `lwp_perror()` can be used to print an error message on the standard error file (analogous to `perror(3)`) when a lwp primitive returns an error indication. `lwp_perror()` uses the same mechanism as `lwp_geterr()` to obtain the last error. `lwp_errstr` returns a pointer to the (NULL-terminated) list of error messages.

`lwp_libcset` (see `lwp_ctxinit(3L)`) allows `errno` from the standard C library reflect a per-thread value rather than a per-pod value.

SEE ALSO

`lwp_ctxinit(3L)`, `perror(3)`

NAME

`lwp_self`, `lwp_ping`, `lwp_enumerate`, `lwp_getstate`, `lwp_setregs`, `lwp_getregs` – LWP status information

SYNOPSIS

```
#include <lwp/lwp.h>
#include <lwp/lwpmachdep.h>

int
lwp_enumerate(vec, maxsize)
thread_t vec[ ]; /* list of id's to be filled in */
int maxsize;    /* number of elements in vec */

int
lwp_ping(tid)
thread_t tid;

int
lwp_getregs(tid, machstate)
thread_t tid;
machstate_t *machstate;

int
lwp_setregs(tid, machstate)
thread_t tid;
machstate_t *machstate;

int
lwp_getstate(tid, statvec)
thread_t tid;
statvec_t *statvec;

int
lwp_self(tid)
thread_t *tid;
```

DESCRIPTION

`lwp_self()` returns the ID of the current thread in *tid*. This is the *only* way to retrieve the identity of *main*.

`lwp_enumerate()` fills in a list with the ID's of all existing threads and returns the total number of threads. This primitive will use *maxsize* to avoid exceeding the capacity of the list. If the number of threads is greater than *maxsize*, only *maxsize* thread ID's are filled in *vec*. If *maxsize* is zero, `lwp_enumerate()` just returns the total number of threads.

`lwp_getstate()` is used to retrieve the context of a given thread. It is possible to see what object (thread, monitor, etc.) if any that thread is blocked on, and the scheduling priority of the thread.

`lwp_ping` returns 0 (no error) if the thread *tid* exists. Otherwise, -1 is returned.

`lwp_setregs` sets the machine-dependent context (i.e., registers) of a thread. The next time the thread is scheduled in, this context is installed. Consult `lwpmachdep.h` for the details. `lwp_getregs` retrieves the machine-dependent context. Note: the registers may not be meaningful unless the thread in question is blocked or suspended because the state of the registers as of the most recent context switch is returned.

RETURNS

Upon successful completion, `lwp_self` and `lwp_getstate()` return 0, -1 on error.

`lwp_enumerate()` returns the total number of threads.

`lwp_ping` returns 0 if the specified thread exists, else -1.

ERRORS

`lwp_getstate`, `lwp_ping`, and `lwp_setstate()` will fail if one or more of the following is true:

LE_NONEXIST Attempt to get the status of a non-existent thread.

NAME

`lwp_yield`, `lwp_suspend`, `lwp_resume`, `lwp_join`, `lwp_setpri`, `lwp_resched`, `lwp_sleep` – control LWP scheduling

SYNOPSIS

```
#include <lwp/lwp.h>

int
lwp_yield(tid)
thread_t tid;

int
lwp_sleep(timeout)
struct timeval *timeout;

int
lwp_resched(prio)
int prio;

int
lwp_setpri(tid, prio)
thread_t tid;
int prio;

int
lwp_suspend(tid)
thread_t tid;

int
lwp_resume(tid)
thread_t tid;

int
lwp_join(tid)
thread_t tid;
```

DESCRIPTION

`lwp_yield()` allows the currently running thread to voluntarily relinquish control to another thread *with the same scheduling priority*. **On a uniprocessor, it is never the case that a thread can execute while a higher priority thread is eligible to run.** If `tid` is `THREADNULL`, the next thread in the same priority queue of the yielding thread will run and the current thread will go to the end of the scheduling queue. Otherwise, it is the ID of the thread to run next, and the current thread will take second place in the scheduling queue.

`lwp_sleep()` blocks the thread executing this primitive for at least the time specified by `timeout`.

Scheduling of threads is, by default, preemptive (higher priorities preempt lower ones) across priorities and non-preemptive within a priority. `lwp_resched()` moves the front thread for a given priority to the end of the scheduling queue. Thus, to achieve a preemptive round-robin scheduling discipline, a high priority thread can periodically wake up and shuffle the queue of threads at a lower priority. `lwp_resched()` does not affect threads which are blocked. If the priority of the rescheduled thread is the same as that of the caller, the effect is the same as `lwp_yield`.

`lwp_setpri()` is used to alter (raise or lower) the scheduling priority of the specified thread. If `tid` is `SELF`, the priority of the invoking thread is set. Note: if the priority of the affected thread becomes greater than that of the caller and the affected thread is not blocked, the caller will not run next. `lwp_setpri()` can be used on either blocked or unblocked threads.

`lwp_join()` blocks the thread issuing the join until the thread `tid` terminates. More than one thread may join `tid`.

lwp_suspend() makes the specified thread ineligible to run. If *tid* is **SELF**, the caller is itself suspended. **lwp_resume()** undoes the effect of **lwp_suspend**. If a blocked thread is suspended, it will not run until it has been unblocked as well as explicitly made eligible to run using **lwp_resume**. By suspending a thread, one can safely examine it without worrying that its execution-time state will change.

NOTE

When scheduling preemptively, be sure to use monitors to protect shared data structures such as those used by the standard I/O library.

RETURN VALUE

lwp_yield, lwp_sleep, lwp_resched, lwp_join, lwp_suspend, lwp_resume:

A 0 return indicates success; -1 indicates an error.

lwp_setpri:

Upon successful completion, the previous priority is returned. Otherwise, -1 is returned.

ERRORS

lwp_yield() will fail if one or more of the following is true:

- LE_INVALIDARG** Attempt to yield to a blocked thread.
- LE_NONEXIST** Attempt to yield to a non-existent thread.
- LE_ILLPRIO** Attempt to yield to thread with different priority.

lwp_sleep() will fail if one or more of the following is true:

- LE_INVALIDARG** Illegal timeout specified.

lwp_resched() will fail if one or more of the following is true:

- LE_INVALIDARG** Attempt to reschedule thread at priority greater than that of the caller.
- LE_ILLPRIO** The priority queue specified contains no threads to reschedule.

lwp_setpri() will fail if one or more of the following is true:

- LE_INVALIDARG** The priority specified is beyond the maximum available to the pod.
- LE_NONEXIST** Attempt to set priority of a non-existent thread.

lwp_join() will fail if one or more of the following are true:

- LE_NONEXIST** Attempt to join a thread that does not exist.

lwp_suspend() will fail if one or more of the following is true:

- LE_NONEXIST** Attempt to suspend a non-existent thread.

lwp_resume() will fail if one or more of the following is true:

- LE_NONEXIST** Attempt to resume a non-existent thread.

NAME

`mon_create`, `mon_destroy`, `mon_enter`, `mon_exit`, `mon_enumerate`, `mon_waiters`, `mon_cond_enter`, `mon_break`, `MONITOR`, `SAMEMON` – LWP routines to manage critical sections

SYNOPSIS

```
#include <lwp/lwp.h>

int
mon_create(mid)
mon_t *mid;

int
mon_destroy(mid)
mon_t mid;

int
mon_enter(mid)
mon_t mid;

int
mon_exit(mid)
mon_t mid;

int
mon_enumerate(vec, maxsize)
mon_t vec[ ]; /* list of all monitors */
int maxsize; /* max size of vec */

int
mon_waiters(mid, owner, vec, maxsize)
mon_t mid; /* monitor in question */
thread_t *owner; /* which thread owns the monitor */
thread_t vec[ ]; /* list of blocked threads */
int maxsize; /* max size of vec */

int
mon_cond_enter(mid)
mon_t mid;

int
mon_break(mid)
mon_t mid;

MONITOR(mid)

SAMEMON(m1, m2)
```

DESCRIPTION

Monitors are used to synchronize access to common resources. Although it is possible (on a uniprocessor) to use knowledge of how scheduling priorities work to serialize access to a resource, monitors (and condition variables) provide a general tool to provide the necessary synchronization.

`mon_create()` creates a new monitor and returns its identity in *mid*. `mon_destroy()` destroys a monitor, as well as any conditions bound to it (see `cv_create(3L)`). Because the lifetime of a monitor can transcend the lifetime of the lwp that created it, monitor destruction is not automatic upon lwp destruction.

`mon_enter()` blocks the calling thread (if the monitor is in use) until the monitor becomes free by being exited or by waiting on a condition (see `cv_create(3L)`). Threads unable to gain entry into the monitor are queued for monitor service by the priority of the thread requesting monitor access, FCFS within a priority. Monitor calls may nest. If, while holding monitor M1 a request for monitor M2 is made, M1 will be held until M2 can be acquired.

mon_cond_enter() will enter the monitor only if the monitor is not busy. Otherwise, an error is returned.

mon_enter() and **mon_cond_enter()** will allow a thread which already has the monitor to reenter the monitor. In this case, the nesting level of monitor entries is returned. Thus, the first time a monitor is entered, **mon_enter()** returns 0. The next time the monitor is entered, **mon_enter()** returns 1. **mon_exit()** frees the current monitor and allows the next thread blocked on the monitor (if any) to enter the monitor. However, if a monitor is entered more than once, **mon_exit()** returns the previous monitor nesting level without freeing the monitor to other threads. Thus, if the monitor was not reentered, **mon_exit()** returns 0.

mon_enumerate() lists all the monitors in the system. The vector supplied is filled in with the ID's of the monitors. *maxsize* is used to avoid exceeding the capacity of the list. If the number of monitors is greater than *maxsize*, only *maxsize* monitor ID's are filled in *vec*.

mon_waiters() puts the thread that currently owns the monitor in *owner* and all threads blocked on the monitor in *vec* (subject to the *maxsize* limitation), and returns the number of waiting threads.

mon_break() forces the release of a monitor lock not necessarily held by the invoking thread. This enables the next thread blocked on the monitor to enter it.

MONITOR is a macro that can be used at the start of a procedure to indicate that the procedure is a monitor. It uses the exception handling mechanism to ensure that the monitor is exited automatically when the procedure exits. Ordinarily, this single macro replaces paired **mon_enter-** **mon_exit()** calls in a monitor procedure.

SAMEMON is a convenient predicate used to compare two monitors for equality.

Monitor locks are released automatically when the lwp holding them dies. This may have implications for the validity of the monitor invariant (a condition that is always true *outside* of the monitor) if a thread unexpectedly terminates.

RETURN VALUE

mon_create() returns the ID of a new monitor.

A 0 return by **mon_destroy()** indicates success; -1 indicates error.

mon_enter() returns the nesting level of the monitor.

Upon successful completion, **mon_exit()** returns the previous nesting level. It returns -1 if there is an error.

mon_enumerate() returns the total number of monitors.

mon_waiters() returns the number of threads waiting for the monitor.

mon_cond_enter() returns the nesting level of the monitor if the monitor is not busy. It return -1 if the monitor is busy.

Upon successful completion, **mon_break()** returns 0. Otherwise, it returns -1.

ERRORS

mon_destroy() will fail if one or more of the following are true:

LE_INUSE Attempt to destroy a monitor that has threads blocked on it.

LE_NONEXIST Attempt to destroy non-existent monitor.

mon_exit() will fail if one or more of the following are true:

LE_INVALIDARG Attempt to exit a monitor that the thread does not own.

LE_NONEXIST Attempt to exit non-existent monitor.

mon_cond_enter() will fail if one or more of the following are true:

LE_INUSE The requested monitor is being used by another thread.

LE_NONEXIST Attempt to destroy non-existent monitor.

mon_break() will fail if one or more of the following are true:

LE_NOTOWNED Attempt to break a monitor lock that is not set.

LE_NONEXIST Attempt to break lock on non-existent monitor.

SEE ALSO

cv_create(3L)

BUGS

There should be language support to enforce the monitor enter-exit discipline.

NAME

`msg_send`, `msg_recv`, `msg_reply`, `MSG_RECVALL`, `msg_ensend`, `msg_enumrecv` – LWP send and receive messages

SYNOPSIS

```
#include <lwp/lwp.h>
```

```
int
```

```
msg_send(dest, arg, argsize, res, ressize)
```

```
thread_t dest; /* destination thread */
```

```
caddr_t arg; /* argument buffer */
```

```
int argsize; /* size of argument buffer */
```

```
caddr_t res; /* result buffer */
```

```
int ressize; /* size of result buffer */
```

```
int
```

```
msg_recv(sender, arg, argsize, res, ressize, timeout)
```

```
thread_t *sender; /* value-result: sending thread or agent */
```

```
caddr_t *arg; /* argument buffer */
```

```
int *argsize; /* argument size */
```

```
caddr_t *res; /* result buffer */
```

```
int *ressize; /* result size */
```

```
struct timeval *timeout; /* POLL, INFINITY, else timeout */
```

```
int
```

```
msg_reply(sender)
```

```
thread_t sender; /* agent id or thread id */
```

```
int
```

```
msg_ensend(vec, maxsize)
```

```
thread_t vec[ ]; /* list of blocked senders */
```

```
int maxsize;
```

```
int
```

```
msg_enumrecv(vec, maxsize)
```

```
thread_t vec[ ]; /* list of blocked receivers */
```

```
int maxsize;
```

```
MSG_RECVALL(sender, arg, argsize, res, ressize, timeout)
```

DESCRIPTION

Each thread queues messages addressed to it as they arrive. Threads may either specify that a particular sender's message is to be received next, or that *any* sender's message may be received next.

`msg_send()` specifies a message buffer and a reply buffer, and initiates one half of a rendezvous with the receiver. The sender will block until the receiver replies using `msg_reply`. `msg_recv()` initiates the other half of a rendezvous and blocks the invoking thread until a corresponding `msg_send()` is received. When unblocked by `msg_send`, the receiver may read the message and generate a reply by filling in the reply buffer and issuing `msg_reply`. `msg_reply()` unblocks the sender. Once a reply is sent, the receiver should no longer access either the message or reply buffer.

In `msg_send`, *argsize* specifies the size in bytes of the argument buffer *argbuf*, which is intended to be a read-only (to the receiver) buffer. *ressize* specifies the size in bytes of the result buffer *resbuf*, which is intended to be a write-only (to the receiver) buffer. *dest* is the thread that is the target of the send.

`msg_recv()` blocks the receiver until:

- A message from the agent or thread bound to *sender* has been sent to the receiver or,
- *sender* points to a THREADNULL-valued variable and *any* message has been sent to the receiver from an thread or agent, or,

- After the time specified by *timeout* elapses and no message is received.

If *timeout* is `POLL`, `msg_rcv()` returns immediately, returning success if the message expected has arrived; otherwise an error is returned. If *timeout* is `INFINITY`, `msg_rcv()` blocks forever or until the expected message arrives. If *timeout* is any other value `msg_rcv()` blocks for the time specified by *timeout* or until the expected message arrives, whichever comes first. When `msg_rcv()` returns, *sender* is filled in with the identity of the sending thread or agent, and the buffer addresses and sizes specified by the matching send are stored in *arg*, *argsize*, *res*, and *ressize*.

`msg_ensend()` and `msg_enumrcv()` are used to list all of the threads blocked on sends (awaiting a reply) and receives (awaiting a send), respectively. The value returned is the number of such blocked threads. The vector supplied by the client is filled in (subject to the *maxsize* limitation) with the ID's of the blocked threads. *maxsize* is used to avoid exceeding the capacity of the list. If the number of threads blocked on sends or receives is greater than *maxsize*, only *maxsize* thread ID's are filled in *vec*. If *maxsize* is 0, just the total number of blocked threads is returned.

sender in `msg_rcv()` is a reference parameter. If you wish to receive from *any* sender, be sure to reinitialize the thread *sender* points to as `THREADNULL` before each use (do not use the address of `THREADNULL` for the sender). Alternatively, use the `MSG_RECVALL` macro. This macro has the same parameters that `msg_rcv()` does, but will ensure that the sender is properly initialized to allow receipt from any sender. `MSG_RECVALL` returns the result from `msg_rcv`.

RETURN VALUE

Upon successful completion, `msg_send`, `msg_rcv()` and `msg_reply()` return 0. Otherwise, -1 is returned.

`msg_ensend()` returns the number of threads blocked on `msg_send`.

`msg_enumrcv()` returns the number of threads blocked on `msg_rcv`.

ERRORS

`msg_rcv()` will fail if one or more of the following is true:

<code>LE_TIMEOUT</code>	Timed out before message arrived.
<code>LE_INVALIDARG</code>	An illegal timeout was specified or the sender address is that of <code>THREADNULL</code> .
<code>LE_NONEXIST</code>	The specified thread or agent does not exist.

`msg_send()` will fail if one or more of the following is true:

<code>LE_INVALIDARG</code>	Attempt to send a message to yourself.
<code>LE_NONEXIST</code>	The specified destination thread does not exist or has terminated.

`msg_reply()` will fail if one or more of the following is true:

<code>LE_NOWAIT</code>	Attempt to reply to a sender that is not expecting a reply.
<code>LE_NONEXIST</code>	Attempt to reply to a sender that does not exist or has terminated.

NAME

`pod_setmaxpri`, `pod_getmaxpri`, `pod_getmaxsize` – control LWP scheduling priority

SYNOPSIS

```
int
pod_setmaxpri(maxprio)
int maxprio;

int
pod_getmaxpri()

int
pod_getmaxsize()
```

DESCRIPTION

The lwp library is self-initializing: the first time you use a primitive that requires threads to be supported, *main* is automatically converted into a thread. A pod will terminate when all client-created lightweight threads (including the thread bound to *main*) are dead.

By default, only a single priority (MINPRIO) is available. However, by using `pod_setmaxpri`, you can make an arbitrary number (up to the limit imposed by the implementation) of priorities available. The *main* thread will receive the highest available scheduling priority at the time of initialization. By using `pod_setmaxpri()` before any other lwp primitives, you can ensure that *main* will receive the same priority as the argument to `pod_setmaxpri`. `pod_setmaxpri()` can be called repeatedly, as long as the number of scheduling priorities (*maxprio*) increases with each call.

`pod_getmaxpri()` returns the current number of available priorities. Priorities are numbered from 1 (MINPRIO) to *maxprio*.

The implementation-dependent maximum number of priorities available can be retrieved using `pod_getmaxsize`. This value will never be less than 255.

RETURN VALUE

`pod_setmaxpri()` returns 0 if success; else -1.

`pod_getmaxsize()` returns the maximum number of priorities that your system supports.

`pod_getmaxpri()` returns the number of priority levels set by the most recent `pod_setmaxpri()` call.

ERRORS

`pod_setmaxpri()` will fail if one or more of the following are true:

LE_INVALIDARG	Attempt to allocate more priorities than supported.
LE_NOROOM	No internal memory left to create pod.

NAME

intro -- introduction to mathematical library functions and constants

SYNOPSIS

```
#include <sys/ieeefp.h>
#include <floatingpoint.h>
#include <math.h>
```

DESCRIPTION

The include file `<math.h>` contains declarations of all the functions described in Section 3M that are implemented in the math library, `libm`. C programs should be linked with the `-lm` option in order to use this library.

`<sys/ieeefp.h>` and `<floatingpoint.h>` define certain types and constants used for `libm` exception handling, conforming to ANSI/IEEE Std 754-1985, the *IEEE Standard for Binary Floating-Point Arithmetic*.

ACKNOWLEDGEMENT

The Sun version of `libm` is based upon and developed from ideas embodied and codes contained in 4.3 BSD, which may not be compatible with earlier BSD or UNIX implementations.

IEEE ENVIRONMENT

The IEEE Standard specifies modes for rounding direction, precision, and exception trapping, and status reflecting accrued exceptions. These modes and status constitute the IEEE run-time environment. On Sun-2 and Sun-3 systems without 68881 floating-point co-processors, only the default rounding direction to nearest is available, only the default non-stop exception handling is available, and accrued exception bits are not maintained.

IEEE EXCEPTION HANDLING

The IEEE Standard specifies exception handling for `aint`, `ceil`, `floor`, `rint`, `remainder`, `rint`, and `sqrt`, and suggests appropriate exception handling for `fp_class`, `copysign`, `fabs`, `finite`, `fmod`, `isinf`, `isnan`, `ilogb`, `ldexp`, `logb`, `nextafter`, `scalb`, `scalbn` and `signbit`, but does not specify exception handling for the other `libm` functions.

For these other unspecified functions the spirit of the IEEE Standard is generally followed in `libm` by handling invalid operand, singularity (division by zero), overflow, and underflow exceptions, as much as possible, in the same way they are handled for the fundamental floating-point operations such as addition and multiplication.

These unspecified functions are usually not quite correctly rounded, may not observe the optional rounding directions, and may not set the inexact exception correctly.

SYSTEM V EXCEPTION HANDLING

The *System V Interface Definition* (SVID) specifies exception handling for some `libm` functions: `j0()`, `j1()`, `jn()`, `y0()`, `y1()`, `yn()`, `exp()`, `log()`, `log10()`, `pow()`, `sqrt()`, `hypot()`, `lgamma()`, `sinh()`, `cosh()`, `sin()`, `cos()`, `tan()`, `asin()`, `acos()`, and `atan2()`. See `matherr(3M)` for a discussion of the extent to which Sun's implementation of `libm` follows the SVID when it is consistent with the IEEE Standard and with hardware efficiency.

LIST OF MATH LIBRARY FUNCTIONS

Name	Appears on Page	Description
–	bessel(3M)	Bessel functions
–	frexp(3M)	floating-point analysis
–	hyperbolic(3M)	hyperbolic functions
–	ieee_functions(3M)	IEEE classification
–	ieee_test(3M)	IEEE tests for compliance
–	ieee_values(3M)	returns double-precision IEEE infinity
–	trig(3M)	trigonometric functions
acos()	trig(3M)	inverse trigonometric functions
acosh()	hyperbolic(3M)	inverse hyperbolic function
aint()	rint(3M)	convert to integral value in floating-point format
anint()	rint(3M)	convert to integral value in floating-point format
asin()	trig(3M)	inverse trigonometric function
asinh()	hyperbolic(3M)	inverse hyperbolic function
atan()	trig(3M)	inverse trigonometric function
atan2()	trig(3M)	rectangular to polar conversion
atanh()	hyperbolic(3M)	inverse hyperbolic function
cbrt()	sqrt(3M)	cube root
ceil()	rint(3M)	ceiling function
copysign()	ieee_functions(3M)	copy sign bit
cos()	trig(3M)	trigonometric function
cosh()	hyperbolic(3M)	hyperbolic function
erf()	erf(3M)	error function
erfc()	erf(3M)	complementary error function
exp()	exp(3M)	exponential function
expm1()	exp(3M)	exp(X)-1
exp2()	exp(3M)	2**X
exp10()	exp(3M)	10**X
fabs()	ieee_functions(3M)	absolute value function
finite()	ieee_functions(3M)	test for finite number
floor()	rint(3M)	floor function
fmod()	ieee_functions(3M)	floating-point remainder
fp_class()	ieee_functions(3M)	classify operand
frexp()	frexp(3M)	floating-point analysis
hypot()	hypot(3M)	Euclidean distance
ieee_flags()	ieee_flags(3M)	IEEE modes and status
ieee_handler()	ieee_handler(3M)	IEEE trapping
ilogb()	ieee_functions(3M)	exponent extraction
infinity()	ieee_values(3M)	returns double-precision IEEE infinity
rint()	rint(3M)	convert to integral value in integer format
isinf()	ieee_functions(3M)	IEEE classification
isnan()	ieee_functions(3M)	IEEE classification
isnormal()	ieee_functions(3M)	IEEE classification
issubnormal()	ieee_functions(3M)	IEEE classification
iszero()	ieee_functions(3M)	IEEE classification
j0()	bessel(3M)	Bessel function
j1()	bessel(3M)	Bessel function
jn()	bessel(3M)	Bessel function
ldexp()	frexp(3M)	exponent adjustment
lgamma()	lgamma(3M)	log gamma function
log()	exp(3M)	natural logarithm

logb()	ieee_test(3M)	exponent extraction
log1p()	exp(3M)	log(1+X)
log2()	exp(3M)	log base 2
log10()	exp(3M)	common logarithm
matherr()	matherr(3M)	math library exception-handling routines
max_normal()	ieee_values(3M)	double-precision IEEE largest positive normalized number
max_subnormal()	ieee_values(3M)	double-precision IEEE largest positive subnormal number
min_normal()	ieee_values(3M)	double-precision IEEE smallest positive normalized number
min_subnormal()	ieee_values(3M)	double-precision IEEE smallest positive subnormal number
modf()	frexp(3M)	floating-point analysis
nextafter()	ieee_functions(3M)	IEEE nearest neighbor
nint()	rint(3M)	convert to integral value in integer format
pow()	exp(3M)	power X**Y
quiet_nan()	ieee_values(3M)	returns double-precision IEEE quiet NaN
remainder()	ieee_functions(3M)	floating-point remainder
rint()	rint(3M)	convert to integral value in floating-point format
scalb()	ieee_test(3M)	exponent adjustment
scalbn()	ieee_functions(3M)	exponent adjustment
signaling_nan()	ieee_values(3M)	returns double-precision IEEE signaling NaN
signbit()	ieee_functions(3M)	IEEE sign bit test
significand()	ieee_test(3M)	scalb(x,-ilogb(x))
sin()	trig(3M)	trigonometric function
sincos()	trig(3M)	simultaneous sin and cos
single_precision()	single_precision(3M)	single-precision libm access
sinh()	hyperbolic(3M)	hyperbolic function
sqrt()	sqrt(3M)	square root
tan()	trig(3M)	trigonometric function
tanh()	hyperbolic(3M)	hyperbolic function
y0()	bessel(3M)	Bessel function
y1()	bessel(3M)	Bessel function
yn()	bessel(3M)	Bessel function

NAME

j_0 , j_1 , j_n , y_0 , y_1 , y_n – Bessel functions

SYNOPSIS

```
#include <math.h>
```

```
double j0(x)
```

```
double x;
```

```
double j1(x)
```

```
double x;
```

```
double jn(n, x)
```

```
double x;
```

```
int n;
```

```
double y0(x)
```

```
double x;
```

```
double y1(x)
```

```
double x;
```

```
double yn(n, x)
```

```
double x;
```

```
int n;
```

DESCRIPTION

These functions calculate Bessel functions of the first and second kinds for real arguments and integer orders.

SEE ALSO

[exp\(3M\)](#)

DIAGNOSTICS

The functions y_0 , y_1 , and y_n have logarithmic singularities at the origin, so they treat zero and negative arguments the way *log* does, as described in [exp\(3M\)](#). Such arguments are unexceptional for j_0 , j_1 , and j_n .

NAME

erf, erfc – error functions

SYNOPSIS

```
#include <math.h>
```

```
double erf(x)
```

```
double x;
```

```
double erfc(x)
```

```
double x;
```

DESCRIPTION

erf(x) returns the error function of x ; where $\text{erf}(x) := (2/\sqrt{\pi}) \int_0^x \exp(-t^2) dt$.

erfc(x) returns $1.0 - \text{erf}(x)$, computed however by other methods that avoid cancellation for large x .

NAME

exp, expm1, exp2, exp10, log, log1p, log2, log10, pow – exponential, logarithm, power

SYNOPSIS

```
#include <math.h>

double exp(x)
double x;

double expm1(x)
double x;

double exp2(x)
double x;

double exp10(x)
double x;

double log(x)
double x;

double log1p(x)
double x;

double log2(x)
double x;

double log10(x)
double x;

double pow(x, y)
double x, y;
```

DESCRIPTION

exp() returns the exponential function e^{**x} .

expm1() returns $e^{**x}-1$ accurately even for tiny x .

exp2() and **exp10()** return 2^{**x} and 10^{**x} respectively.

log() returns the natural logarithm of x .

log1p() returns $\log(1+x)$ accurately even for tiny x .

log2() and **log10()** return the logarithm to base 2 and 10 respectively.

pow() returns x^{**y} . **pow(x,0.0)** is 1 for all x , in conformance with 4.3BSD, as discussed in the *Floating Point Programmers Guide*.

SEE ALSO

matherr(3M)

DIAGNOSTICS

All these functions handle exceptional arguments in the spirit of ANSI/IEEE Std 754-1985. Thus for $x == \pm 0$, **log(x)** is $-\infty$ with a division by zero exception; for $x < 0$, including $-\infty$, **log(x)** is a quiet NaN with an invalid operation exception; for $x == +\infty$ or a quiet NaN, **log(x)** is x without exception; for x a signaling NaN, **log(x)** is a quiet NaN with an invalid operation exception; for $x == 1$, **log(x)** is 0 without exception; for any other positive x , **log(x)** is a normalized number with an inexact exception.

In addition, **exp,exp2,exp10, log,log2,log10, and pow()** may also set **errno** and call **matherr(3M)**.

NAME

frexp, modf, ldexp – traditional UNIX functions

SYNOPSIS

```
#include <math.h>

double frexp(value, eptr)
double value;
int *eptr;

double ldexp(x,n)
double x;
int n;

double modf(value, iptr)
double value, *iptr;
```

DESCRIPTION

These functions are provided for compatibility with other UNIX system implementations. They are not used internally in `libm` or `libc`. Better ways to accomplish similar ends may be found in `ieee_functions(3M)` and `rint(3M)`.

`ldexp(x,n)` returns $x * 2^{**n}$ computed by exponent manipulation rather than by actually performing an exponentiation or a multiplication. Note: `ldexp(x,n)` differs from `scalbn(x,n)`, defined in `ieee_functions(3M)`, only that in the event of IEEE overflow and underflow, `ldexp(x,n)` sets `errno` to `ERANGE`.

Every non-zero number can be written uniquely as $x * 2^{**n}$, where the significand x is in the range $0.5 \leq |x| < 1.0$ and the exponent n is an integer. The function `frexp()` returns the significand of a double *value* as a double quantity, x , and stores the exponent n , indirectly through *eptr*. If *value* == 0, both results returned by `frexp()` are 0.

`modf()` returns the fractional part of *value* and stores the integral part indirectly through *iptr*. Thus the argument *value* and the returned values `modf()` and **iptr* satisfy

$$(*iptr + modf) == value$$

and both results have the same sign as *value*. The definition of `modf()` varies among UNIX system implementations, so avoid `modf()` in portable code.

The results of `frexp()` and `modf()` are not defined when *value* is an IEEE infinity or NaN.

SEE ALSO

`ieee_functions(3M)`, `rint(3M)`

NAME

`sinh`, `cosh`, `tanh`, `asinh`, `acosh`, `atanh` – hyperbolic functions

SYNOPSIS

```
#include <math.h>
```

```
double sinh(x)
```

```
double x;
```

```
double cosh(x)
```

```
double x;
```

```
double tanh(x)
```

```
double x;
```

```
double asinh(x)
```

```
double x;
```

```
double acosh(x)
```

```
double x;
```

```
double atanh(x)
```

```
double x;
```

DESCRIPTION

These functions compute the designated direct and inverse hyperbolic functions for real arguments. They inherit much of their roundoff error from `expm1()` and `log1p`, described in `exp(3M)`.

DIAGNOSTICS

These functions handle exceptional arguments in the spirit of ANSI/IEEE Std 754-1985. Thus `sinh()` and `cosh()` return $\pm\infty$ on overflow, `acosh()` returns a NaN if its argument is less than 1, and `atanh()` returns a NaN if its argument has absolute value greater than 1. In addition, `sinh`, `cosh`, and `tanh()` may also set `errno` and call `matherr(3M)`.

SEE ALSO

`exp(3M)`, `matherr(3M)`

NAME

hypot – Euclidean distance

SYNOPSIS

```
#include <math.h>
```

```
double hypot(x, y)
```

```
double x, y;
```

DESCRIPTION

hypot() returns

```
sqrt(x*x + y*y) ,
```

taking precautions against unwarranted IEEE exceptions. On IEEE overflow, **hypot()** may also set **errno** and call **matherr(3M)**. **hypot($\pm\infty$, y)** is $+\infty$ for any **y**, even a NaN, and is exceptional only for a signaling NaN.

hypot(x,y) and **atan2(3M)** convert rectangular coordinates **(x,y)** to polar **(r,θ)**; **hypot()** computes **r**, the modulus or radius.

SEE ALSO

matherr(3M)

NAME

`ieee_flags` – mode and status function for IEEE standard arithmetic

SYNOPSIS

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

int ieee_flags(action,mode,in,out)
char *action, *mode, *in, **out;
```

DESCRIPTION

This function provides easy access to the modes and status required to fully exploit ANSI/IEEE Std 754-1985 arithmetic in a C program. All arguments are pointers to strings. Results arising from invalid arguments and invalid combinations are undefined for efficiency.

There are four types of *action*: “get”, “set”, “clear”, and “clearall”. There are three valid settings for *mode*, two corresponding to modes of IEEE arithmetic:

```
“direction”,      ... current rounding direction mode
“precision”,     ... current rounding precision mode
```

and one corresponding to status of IEEE arithmetic:

```
“exception”.     ... accrued exception-occurred status
```

There are 14 types of *in* and *out* :

```
“nearest”,      ... round toward nearest
“tozero”,      ... round toward zero
“negative”,     ... round toward negative infinity
“positive”,     ... round toward positive infinity
“extended”,
“double”,
“single”,
“inexact”,
“division”,     ... division by zero exception
“underflow”,
“overflow”,
“invalid”,
“all”,         ... all five exceptions above
“common”.     ... invalid, overflow, and division exceptions
```

Note: “all” and “common” only make sense with “set” or “clear”.

For “clearall”, `ieee_flags()` returns 0 and restores all default modes and status. Nothing will be assigned to *out*. Thus

```
char *mode, *out, *in;
ieee_flags("clearall",mode, in, &out);
```

set rounding direction to “nearest”, rounding precision to “extended”, and all accrued exception-occurred status to zero.

For “clear”, `ieee_flags()` returns 0 and restores the default mode or status. Nothing will be assigned to *out*. Thus

```
char *out, *in;
ieee_flags("clear","direction", in, &out);    ... set rounding direction to round to nearest.
```

For “set”, `ieee_flags()` returns 0 if the action is successful and 1 if the corresponding required status or mode is not available (for instance, not supported in hardware). Nothing will be assigned to *out*. Thus

```
char *out, *in;
ieee_flags ("set", "direction", "tozero", &out);    ... set rounding direction to round toward zero;
```

For “get”, we have the following cases:

Case 1: *mode* is “direction”. In that case, *out* returns one of the four strings “nearest”, “tozero”, “positive”, “negative”; and `ieee_flags()` returns a value corresponding to *out* according to the enum `fp_direction_type` defined in `<sys/ieeefp.h>`.

Case 2: *mode* is “precision”. In that case, *out* returns one of the three strings “extended”, “double”, “single”; and `ieee_flags()` returns a value corresponding to *out* according to the enum `fp_precision_type` defined in `<sys/ieeefp.h>`.

Case 3: *mode* is “exception”. In that case, *out* returns

- (a) “not available” if information on exception is not available,
- (b) “no exception” if no accrued exception,
- (c) the accrued exception that has the highest priority according to the list below
 - (1) the exception named by *in*,
 - (2) “invalid”,
 - (3) “overflow”,
 - (4) “division”,
 - (5) “underflow”,
 - (6) “inexact”.

In this case `ieee_flags()` returns a five bit value where each bit (cf. enum `fp_exception_type` in `<sys/ieeefp.h>`) corresponds to an exception-occurred accrued status flag: 0 = off, 1 = on. The bit corresponding to a particular exception varies among architectures.

Example:

```
char *out; int k, ieee_flags();
ieee_flags ("clear", "exception", "all", &out);    /* clear all accrued exceptions */
...
... (code that generates three exceptions: overflow, invalid, inexact)
...
k = ieee_flags("get", "exception", "overflow", &out);
```

then `out = “overflow”`, and on a Sun-3, `k=25`.

FILES

```
/usr/include/sys/ieeefp.h
/usr/lib/libm.a
```

NAME

ieee_functions, fp_class, finite, ilogb, isinf, isnan, isnormal, issubnormal, iszero, signbit, copysign, fabs, fmod, nextafter, remainder, scalbn – appendix and related miscellaneous functions for IEEE arithmetic

SYNOPSIS

```
#include <math.h>

enum fp_class_type fp_class(x)
double x;

int finite(x)
double x;

int ilogb(x)
double x;

int isinf(x)
double x;

int isnan(x)
double x;

int isnormal(x)
double x;

int issubnormal(x)
double x;

int iszero(x)
double x;

int signbit(x)
double x;

double copysign(x,y)
double x, y;

double fabs(x)
double x;

double fmod(x,y)
double x, y;

double nextafter(x,y)
double x, y;

double remainder(x,y)
double x, y;

double scalbn(x,n)
double x; int n;
```

DESCRIPTION

Most of these functions provide capabilities required by ANSI/IEEE Std 754-1985 or suggested in its appendix.

fp_class(x) corresponds to the IEEE's `class()` and classifies *x* as zero, subnormal, normal, ∞ , or quiet or signaling *NaN*; `<floatingpoint.h>` defines `enum fp_class_type`. The following functions return 0 if the indicated condition is not satisfied:

finite(x)	returns 1 if <i>x</i> is zero, subnormal or normal
isinf(x)	returns 1 if <i>x</i> is ∞
isnan(x)	returns 1 if <i>x</i> is <i>NaN</i>
isnormal(x)	returns 1 if <i>x</i> is normal
issubnormal(x)	returns 1 if <i>x</i> is subnormal
iszero(x)	returns 1 if <i>x</i> is zero
signbit(x)	returns 1 if <i>x</i> 's sign bit is set

ilogb(x) returns the unbiased exponent of *x* in integer format. **ilogb($\pm\infty$)** = +MAXINT and **ilogb(0)** = -MAXINT; `<values.h>` defines MAXINT as the largest int. **ilogb(x)** never generates an exception. When *x* is subnormal, **ilogb(x)** returns an exponent computed as if *x* were first normalized.

copysign(x,y) returns *x* with *y*'s sign bit.

fabs(x) returns the absolute value of *x*.

nextafter(x,y) returns the next machine representable number from *x* in the direction *y*.

remainder(x,y) and **fmod(x,y)** return a remainder of *x* with respect to *y*; that is, the result *r* is one of the numbers that differ from *x* by an integral multiple of *y*. Thus $(x-r)/y$ is an integral value, even though it might exceed MAXINT if it were explicitly computed as an int. Both functions return one of the two such smallest in magnitude. **remainder(x,y)** is the operation specified in ANSI/IEEE Std 754-1985; the result of **fmod(x,y)** may differ from **remainder's** result by $\pm y$. The magnitude of **remainder's** result can not exceed half that of *y*; its sign might not agree with either *x* or *y*. The magnitude of **fmod's** result is less than that of *y*; its sign agrees with that of *x*. Neither function can generate an exception as long as both arguments are normal or subnormal. **remainder(x, 0)**, **fmod(x, 0)**, **remainder(∞ , y)**, and **fmod(∞ , y)** are invalid operations that produce a *NaN*.

scalbn(x,n) returns $x * 2^{*n}$ computed by exponent manipulation rather than by actually performing an exponentiation or a multiplication. Thus

$$1 \leq \text{scalbn}(\text{fabs}(x), -\text{ilogb}(x)) < 2$$

for every *x* except 0, ∞ , and *NaN*.

FILES

`/usr/include/floatingpoint.h`
`/usr/include/math.h`
`/usr/include/values.h`
`/usr/lib/libm.a`

SEE ALSO

floatingpoint(3), **ieee_environment(3M)**, **matherr(3M)**

NAME

`ieee_handler` – IEEE exception trap handler function

SYNOPSIS

```
#include <floatingpoint.h>

int ieee_handler(action,exception,hdl)
char action[ ], exception[ ];
sigfpe_handler_type hdl;
```

DESCRIPTION

This function provides easy exception handling to exploit ANSI/IEEE Std 754-1985 arithmetic in a C program. All arguments are pointers to strings. Results arising from invalid arguments and invalid combinations are undefined for efficiency.

There are three types of *action* : “get”, “set”, and “clear”. There are five types of *exception* :

```
“inexact”
“division”      ... division by zero exception
“underflow”
“overflow”
“invalid”
“all”           ... all five exceptions above
“common”       ... invalid, overflow, and division exceptions
```

Note: “all” and “common” only make sense with “set” or “clear”.

`hdl` contains the address of a signal-handling routine. `<floatingpoint.h>` defines `sigfpe_handler_type`.

“get” will get the location of the current handler routine for *exception* in `hdl`. “set” will set the routine pointed at by `hdl` to be the handler routine and at the same time enable the trap on *exception*, except when `hdl == SIGFPE_DEFAULT` or `SIGFPE_IGNORE`; then `ieee_handler()` will disable the trap on *exception*. When `hdl == SIGFPE_ABORT`, any trap on *exception* will dump core using `abort(3)`. “clear” “all” disables trapping on all five exceptions.

Two steps are required to intercept an IEEE-related SIGFPE code with `ieee_handler`:

- 1) Set up a handler with `ieee_handler`.
- 2) Perform a floating-point operation that generates the intended IEEE exception.

Unlike `sigfpe(3)`, `ieee_handler()` also adjusts floating-point hardware mode bits affecting IEEE trapping. For “clear”, “set” `SIGFPE_DEFAULT`, or “set” `SIGFPE_IGNORE`, the hardware trap is disabled. For any other “set”, the hardware trap is enabled.

SIGFPE signals can be handled using `sigvec(2)`, `signal(3)`, `signal(3F)`, `sigfpe(3)`, or `ieee_handler(3M)`. In a particular program, to avoid confusion, use only one of these interfaces to handle SIGFPE signals.

DIAGNOSTICS

`ieee_handler()` normally returns 0. In the case of “set”, 1 will be returned if the action is not available (for instance, not supported in hardware).

EXAMPLE

A user-specified signal handler might look like this:

```
void sample_handler( sig, code, scp, addr)
int sig ;          /* sig == SIGFPE always */
int code ;
struct sigcontext *scp ;
char *addr ;
{
    /*
     * Sample user-written sigfpe code handler.
     * Prints a message and continues.
     * struct sigcontext is defined in <signal.h>.
     */
    printf("ieee exception code %x occurred at pc %X \n",code,scp->sc_pc);
}
```

and it might be set up like this:

```
extern void sample_handler();
main()
{
    sigfpe_handler_type hdl, old_handler1, old_handler2;
    /*
     * save current overflow and invalid handlers
     */
    ieee_handler("get","overflow",old_handler1);
    ieee_handler("get","invalid", old_handler2);
    /*
     * set new overflow handler to sample_handler() and set new
     * invalid handler to SIGFPE_ABORT (abort on invalid)
     */
    hdl = (sigfpe_handler_type) sample_handler;
    if(ieee_handler("set","overflow",hdl) != 0)
        printf("ieee_handler can't set overflow \n");
    if(ieee_handler("set","invalid",SIGFPE_ABORT) != 0)
        printf("ieee_handler can't set invalid \n");
    ...
    /*
     * restore old overflow and invalid handlers
     */
    ieee_handler("set","overflow", old_handler1);
    ieee_handler("set","invalid", old_handler2);
}
```

FILES

```
/usr/include/floatingpoint.h
/usr/include/signal.h
/usr/lib/libm.a
```

SEE ALSO

sigvec(2), abort(3), floatingpoint(3), sigfpe(3), signal(3), signal(3F)

NAME

ieee_test, logb, scalb, significand – IEEE test functions for verifying standard compliance

SYNOPSIS

```
#include <math.h>

double logb(x)
double x;

double scalb(x,y)
double x; double y;

double significand(x)
double x;
```

DESCRIPTION

These functions allow users to verify compliance to ANSI/IEEE Std 754-1985 by running certain test vectors distributed by the University of California. Their use is not otherwise recommended; instead use `scalbn(x,n)` and `ilogb(x)` described in `ieee_functions(3M)`. See the *Floating Point Programmers Guide* for details.

`logb(x)` returns the unbiased exponent of x in floating-point format, for exercising the `logb(L)` test vector. `logb($\pm\infty$) = $+\infty$` ; `logb(0) = $-\infty$` with a division by zero exception. `logb(x)` differs from `ilogb(x)` in returning a result in floating-point rather than integer format, in sometimes signaling IEEE exceptions, and in not normalizing subnormal x .

`scalb(x,(double)n)` returns $x * 2^{**n}$ computed by exponent manipulation rather than by actually performing an exponentiation or a multiplication, for exercising the `scalb(S)` test vector. Thus

$$0 \leq \text{scalb}(\text{fabs}(x), -\text{logb}(x)) < 2$$

for every x except 0, ∞ and *NaN*. `scalb(x,y)` is not defined when y is not an integral value. `scalb(x,y)` differs from `scalbn(x,n)` in that the second argument is in floating-point rather than integer format.

`significand(x)` computes just

$$\text{scalb}(x, (\text{double}) -\text{ilogb}(x)),$$

for exercising the `fraction-part(F)` test vector.

FILES

```
/usr/include/math.h
/usr/lib/libm.a
```

SEE ALSO

`floatingpoint(3)`, `ieee_values(3M)`, `ieee_functions(3M)`, `matherr(3M)`

NAME

ieee_values, min_subnormal, max_subnormal, min_normal, max_normal, infinity, quiet_nan, signaling_nan, HUGE, HUGE_VAL – functions that return extreme values of IEEE arithmetic

SYNOPSIS

```
#include <math.h>

double min_subnormal()
double max_subnormal()
double min_normal()
double max_normal()
double infinity()
double quiet_nan(n)
long n;
double signaling_nan(n)
long n;
#define HUGE (infinity())
#define HUGE_VAL (infinity())
```

DESCRIPTION

These functions return special values associated with ANSI/IEEE Std 754-1985 double-precision floating-point arithmetic: the smallest and largest positive subnormal numbers, the smallest and largest positive normalized numbers, positive infinity, and a quiet and signaling NaN. The long parameters *n* to `quiet_nan(n)` and `signaling_nan(n)` are presently unused but are reserved for future use to specify the significand of the returned NaN.

None of these functions are affected by IEEE rounding or trapping modes or generate any IEEE exceptions.

The macro `HUGE` returns $+\infty$ in accordance with previous SunOS releases. The macro `HUGE_VAL` returns $+\infty$ in accordance with the System V Interface Definition.

FILES

```
/usr/include/math.h
/usr/lib/libm.a
```

SEE ALSO

ieee_functions(3M)

NAME

lgamma – log gamma function

SYNOPSIS

```
#include <math.h>
extern int signgam;
double lgamma(x)
double x;
```

DESCRIPTION

lgamma() returns

$$\ln |\Gamma(x)|$$

where

$$\Gamma(x) = \int_0^{\infty} t^{x-1} e^{-t} dt$$

for $x > 0$ and

$$\Gamma(x) = \pi / (\Gamma(1-x) \sin(\pi x))$$

for $x < 1$.

The external integer **signgam** returns the sign of $\Gamma(x)$.

IDIOSYNCRASIES

Do *not* use the expression **signgam*exp(lgamma(x))** to compute 'g := $\Gamma(x)$ '. Instead compute **lgamma()** first:

```
lg = lgamma(x); g = signgam*exp(lg);
```

only after **lgamma()** has returned can **signgam** be correct. Note: $\Gamma(x)$ must overflow when x is large enough, underflow when $-x$ is large enough, and generate a division by zero exception at the singularities x a nonpositive integer. In addition, **lgamma()** may also set **errno** and call **matherr(3M)**.

SEE ALSO

matherr(3M)

NAME

matherr – math library exception-handling function

SYNOPSIS

```
#include <math.h>

int matherr(exc)
struct exception *exc;
```

DESCRIPTION

The SVID (*System V Interface Definition*) specifies that certain **libm** functions call **matherr()** when exceptions are detected. Users may define their own mechanisms for handling exceptions, by including a function named **matherr()** in their programs. **matherr()** is of the form described above. When an exception occurs, a pointer to the exception structure *exc* will be passed to the user-supplied **matherr()** function. This structure, which is defined in the **<math.h>** header file, is as follows:

```
struct exception {
    int type;
    char *name;
    double arg1, arg2, retval;
};
```

The element **type** is an integer describing the type of exception that has occurred, from the following list of constants (defined in the header file):

DOMAIN	argument domain exception
SING	argument singularity
OVERFLOW	overflow range exception
UNDERFLOW	underflow range exception

The element **name** points to a string containing the name of the function that incurred the exception. The elements **arg1** and **arg2** are the arguments with which the function was invoked. **retval** is set to the default value that will be returned by the function unless the user's **matherr()** sets it to a different value.

If the user's **matherr()** function returns non-zero, no exception message will be printed, and **errno** will not be set.

If **matherr()** is not supplied by the user, the default **matherr** exception-handling mechanisms, summarized in the table below, will be invoked upon exception:

DOMAIN==fp_invalid

An IEEE NaN is usually returned, **errno** is set to EDOM, and a message is printed on standard error. **pow(x,0.0)** for any *x* and **atan2(0.0,0.0)** return numerical default results but set **errno** and print the message.

SING==fp_division

An IEEE ∞ of appropriate sign is returned, **errno** is set to EDOM, and a message is printed on standard error.

OVERFLOW==fp_overflow

In the default rounding direction, an IEEE ∞ of appropriate sign is returned. In optional rounding directions, \pm MAXDOUBLE, the largest finite double-precision number, is sometimes returned instead of $\pm\infty$. **errno** is set to ERANGE.

UNDERFLOW==fp_underflow

An appropriately-signed zero, subnormal number, or smallest normalized number is returned, and **errno** is set to ERANGE.

The facilities provided by **matherr()** are not available in situations such as compiling on a Sun-3 system with **/usr/lib/f68881/libm.il** or **/usr/lib/ffpa/libm.il**, in which case some **libm** functions are converted to atomic hardware operations. In these cases setting **errno** and calling **matherr()** are not worth the adverse performance impact, but regular ANSI/IEEE Std 754-1985 exception handling remains available. In any

case `errno` is not a reliable error indicator in that it may be unexpectedly set by a function in a handler for an asynchronous signal.

DEFAULT ERROR HANDLING PROCEDURES				
<i>Types of Errors</i>				
<math.h> type	DOMAIN	SING	OVERFLOW	UNDERFLOW
<code>errno</code>	EDOM	EDOM	ERANGE	ERANGE
IEEE Exception	Invalid Operation	Division by Zero	Overflow	Underflow
<floatingpoint.h> type	<code>fp_invalid</code>	<code>fp_division</code>	<code>fp_overflow</code>	<code>fp_underflow</code>
ACOS, ASIN:	M, NaN	–	–	–
ATAN2(0,0):	M, ± 0.0 or $\pm\pi$	–	–	–
BESSEL: y0, y1, yn (x < 0) y0, y1, yn (x = 0)	M, NaN –	– M, $-\infty$	– –	– –
COSH, SINH:	–	–	IEEE Overflow	–
EXP:	–	–	IEEE Overflow	IEEE Underflow
HYPOT:	–	–	IEEE Overflow	–
LGAMMA:	–	M, $+\infty$	IEEE Overflow	–
LOG, LOG10: (x < 0) (x = 0)	M, NaN –	– M, $-\infty$	– –	– –
POW: usual cases (x < 0) ** (y not an integer) 0 ** 0 0 ** (y < 0)	– M, NaN M, 1.0 –	– – – M, $\pm\infty$	IEEE Overflow – – –	IEEE Underflow – – –
SQRT:	M, NaN	–	–	–

ABBREVIATIONS

M	Message is printed (EDOM exception).
NaN	IEEE NaN result and invalid operation exception.
∞	IEEE ∞ result and division-by-zero exception.
IEEE Overflow	IEEE Overflow result and exception.
IEEE Underflow	IEEE Underflow result and exception.
π	Closest machine-representable approximation to pi.

The interaction of IEEE arithmetic and `matherr()` is not defined when executing under IEEE rounding modes other than the default round to nearest: `matherr()` may not be called on overflow or underflow, and the Sun-provided `matherr()` may return results that differ from those in this table.

EXAMPLE

```
#include <math.h>

int
matherr(x)
register struct exception *x;
{
    switch (x->type) {
        case
            DOMAIN:
                /* change sqrt to return sqrt(-arg1), not NaN */
                if (!strcmp(x->name, "sqrt")) {
                    x->retval = sqrt(-x->arg1);
                    return (0); /* print message and set errno */
                } /* fall through */
        case
            SING:
                /* all other domain or sing exceptions, print message and abort */
                fprintf(stderr, "domain exception in %s\n", x->name);
                abort();
                break;
    }
    return (0); /* all other exceptions, execute default procedure */
}
```

NAME

aint, anint, ceil, floor, rint, irint, nint – round to integral value in floating-point or integer format

SYNOPSIS

```
#include <math.h>
double aint(x)
double x;
double anint(x)
double x;
double ceil(x)
double x;
double floor(x)
double x;
double rint(x)
double x;
int irint(x)
double x;
int nint(x)
double x;
```

DESCRIPTION

aint, anint, ceil, floor, and rint() convert a double value into an integral value in double format. They vary in how they choose the result when the argument is not already an integral value. Here an “integral value” means a value of a mathematical integer, which however might be too large to fit in a particular computer’s int format. All sufficiently large values in a particular floating-point format are already integral; in IEEE double-precision format, that means all values $\geq 2^{52}$. Zeros, infinities, and quiet NaNs are treated as integral values by these functions, which always preserve their argument’s sign.

aint() returns the integral value between x and 0, nearest x . This corresponds to IEEE rounding toward zero and to the Fortran generic intrinsic function **aint**.

anint() returns the nearest integral value to x , except halfway cases are rounded to the integral value larger in magnitude. This corresponds to the Fortran generic intrinsic function **anint**.

ceil() returns the least integral value greater than or equal to x . This corresponds to IEEE rounding toward positive infinity.

floor() returns the greatest integral value less than or equal to x . This corresponds to IEEE rounding toward negative infinity.

rint() rounds x to an integral value according to the current IEEE rounding direction.

irint converts x into int format according to the current IEEE rounding direction.

nint() converts x into int format rounding to the nearest int value, except halfway cases are rounded to the int value larger in magnitude. This corresponds to the Fortran generic intrinsic function **nint**.

NAME

single_precision - Single-precision access to math library functions

SYNOPSIS

```
#include <math.h>

FLOATFUNCTIONTYPE r_acos_(x)
FLOATFUNCTIONTYPE r_acosh_(x)
FLOATFUNCTIONTYPE r_aint_(x)
FLOATFUNCTIONTYPE r_anint_(x)
FLOATFUNCTIONTYPE r_asin_(x)
FLOATFUNCTIONTYPE r_asinh_(x)
FLOATFUNCTIONTYPE r_atan_(x)
FLOATFUNCTIONTYPE r_atanh_(x)
FLOATFUNCTIONTYPE r_atan2_(x,y)
FLOATFUNCTIONTYPE r_cbrt_(x)
FLOATFUNCTIONTYPE r_ceil_(x)
enum fp_class_type ir_fp_class_(x)
FLOATFUNCTIONTYPE r_copysign_(x,y)
FLOATFUNCTIONTYPE r_cos_(x)
FLOATFUNCTIONTYPE r_cosh_(x)
FLOATFUNCTIONTYPE r_erf_(x)
FLOATFUNCTIONTYPE r_erfc_(x)
FLOATFUNCTIONTYPE r_exp_(x)
FLOATFUNCTIONTYPE r_expml_(x)
FLOATFUNCTIONTYPE r_exp2_(x)
FLOATFUNCTIONTYPE r_exp10_(x)
FLOATFUNCTIONTYPE r_fabs_(x)
int ir_finite_(x)
FLOATFUNCTIONTYPE r_floor_(x)
FLOATFUNCTIONTYPE r_fmod_(x,y)
FLOATFUNCTIONTYPE r_hypot_(x,y)
int ir_ilogb_(x)
int ir_rint_(x)
int ir_isinf_(x)
int ir_isnan_(x)
int ir_isnormal_(x)
int ir_issubnormal_(x)
int ir_iszero_(x)
int ir_nint_(x)
FLOATFUNCTIONTYPE r_infinity_( )
FLOATFUNCTIONTYPE r_j0_(x)
FLOATFUNCTIONTYPE r_j1_(x)
FLOATFUNCTIONTYPE r_jn_(n,x)
FLOATFUNCTIONTYPE r_lgamma_(x)
FLOATFUNCTIONTYPE r_logb_(x)
FLOATFUNCTIONTYPE r_log_(x)
FLOATFUNCTIONTYPE r_log1p_(x)
FLOATFUNCTIONTYPE r_log2_(x)
FLOATFUNCTIONTYPE r_log10_(x)
FLOATFUNCTIONTYPE r_max_normal_( )
FLOATFUNCTIONTYPE r_max_subnormal_( )
FLOATFUNCTIONTYPE r_min_normal_( )
FLOATFUNCTIONTYPE r_min_subnormal_( )
FLOATFUNCTIONTYPE r_nextafter_(x,y)
```

```

FLOATFUNCTIONTYPE r_pow_(x,y)
FLOATFUNCTIONTYPE r_quiet_nan_(n)
FLOATFUNCTIONTYPE r_remainder_(x,y)
FLOATFUNCTIONTYPE r_rint_(x)
FLOATFUNCTIONTYPE r_scalb_(x,y)
FLOATFUNCTIONTYPE r_scalbn_(x,n)
FLOATFUNCTIONTYPE r_signaling_nan_(n)
int ir_signbit_(x)
FLOATFUNCTIONTYPE r_significand_(x)
FLOATFUNCTIONTYPE r_sin_(x)
void r_sincos_(x,s,c)
FLOATFUNCTIONTYPE r_sinh_(x)
FLOATFUNCTIONTYPE r_sqrt_(x)
FLOATFUNCTIONTYPE r_tan_(x)
FLOATFUNCTIONTYPE r_tanh_(x)
FLOATFUNCTIONTYPE r_y0_(x)
FLOATFUNCTIONTYPE r_y1_(x)
FLOATFUNCTIONTYPE r_yn_(n,x)

float *x, *y, *s, *c
int *n
```

DESCRIPTION

These functions are single-precision versions of certain **libm** functions. Primarily for use by Fortran programmers, these functions may also be used in other languages. The single-precision floating-point results are deviously declared to avoid C's automatic type conversion to double.

FILES

/usr/lib/libm.a

NAME

sqrt, cbrt – cube root, square root

SYNOPSIS

```
#include <math.h>
```

```
double cbrt(x)
```

```
double x;
```

```
double sqrt(x)
```

```
double x;
```

DESCRIPTION

sqrt(x) returns the square root of x , correctly rounded according to ANSI/IEEE 754-1985. In addition, **sqrt()** may also set **errno** and call **matherr(3M)**.

cbrt(x) returns the cube root of x . **cbrt()** is accurate to within 0.7 *ulps*.

SEE ALSO

matherr(3M)

NAME

sin, cos, tan, asin, acos, atan, atan2 – trigonometric functions

SYNOPSIS

```
#include <math.h>
```

```
double sin(x)
```

```
double x;
```

```
double cos(x)
```

```
double x;
```

```
void sincos(x, s, c)
```

```
double x, *s, *c;
```

```
double tan(x)
```

```
double x;
```

```
double asin(x)
```

```
double x;
```

```
double acos(x)
```

```
double x;
```

```
double atan(x)
```

```
double x;
```

```
double atan2(y, x)
```

```
double y, x;
```

DESCRIPTION

sin, cos, sincos, and tan() return trigonometric functions of radian arguments. The values of trigonometric functions of arguments exceeding $\pi/4$ in magnitude are affected by the precision of the approximation to $\pi/2$ used to reduce those arguments to the range $-\pi/4$ to $\pi/4$. Argument reduction may occur in hardware or software; if in software, the variable `fp_pi` defined in `<math.h>` allows changing that precision at run time. Trigonometric argument reduction is discussed in the *Floating Point Programmers Guide*. Note that `sincos(x,s,c)` allows simultaneous computation of `*s = sin(x)` and `*c = cos(x)`.

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(y,x) and hypot(3M) convert rectangular coordinates (x,y) to polar (r,θ) ; atan2() computes θ , the argument or phase, by computing an arc tangent of y/x in the range $-\pi$ to π . atan2(0.0,0.0) is ± 0.0 or $\pm\pi$, in conformance with 4.3BSD, as discussed in the *Floating Point Programmers Guide*.

DIAGNOSTICS

These functions handle exceptional arguments in the spirit of ANSI/IEEE Std 754-1985. sin($\pm\infty$), cos($\pm\infty$), tan($\pm\infty$), or asin(x) or acos(x) with $|x|>1$, return NaN. In addition, asin, acos, and atan2() may also set `errno` and call `matherr(3M)`.

SEE ALSO

hypot(3M), matherr(3M)

NAME

intro – introduction to RPC service library functions and protocols

DESCRIPTION

These functions constitute the RPC service library. Most of these describe RPC protocols. The **PROTOCOL** section describes how to access the protocol description file. This file may be compiled with **rpcgen(1)** to produce data definitions and XDR routines. Procompiled versions of header files sometimes exist as **<rpcsvc/*.h>** and precompiled XDR routines and programming interfaces to the protocols sometimes exist in *librpcsvc*. **Warning:** some of these header files and XDR routines were hand-written because they existed before *rpcgen*. They do not correspond to their protocol description file. In order to get the link editor to load this library, use the **-lrpcsvc** option of **cc(1V)**. Information about the availability of programming interfaces to these protocols is available under **PROGRAMMING** section of each manual page.

Some routines in the *librpcsvc* library do not correspond to protocols, but are useful utilities for RPC programming. These are distinguished by the presence of the **SYNOPSIS** section instead of the usual **PROTOCOL** section.

LIST OF STANDARD RPC SERVICES

Name	Appears on Page	Description
bootparam()	bootparam(3R)	bootparam protocol
ether()	ether(3R)	monitor traffic on the Ethernet
get()	publickey(3R)	get secret key
getrpcport()	getrpcport(3R)	get RPC port number
getsecretkey()	publickey(3R)	get secret key
ipalloc()	ipalloc(3R)	determine or temporarily allocate IP address
key()	publickey(3R)	get secret key
klm_prot()	klm_prot(3R)	protocol between kernel and local lock manager
mount()	mount(3R)	keep track of remotely mounted filesystems
nlm_prot()	nlm_prot(3R)	protocol between local and remote network lock managers
pnp()	pnp(3R)	Automated network installer
public()	publickey(3R)	get secret key
rex()	rex(3R)	remote execution protocol
rnusers()	rnusers(3R)	return information about users on remote machines
rquota()	rquota(3R)	implement quotas on remote machines
rstat()	rstat(3R)	get performance data from remote kernel
rusers()	rnusers(3R)	return information about users on remote machines
rwall()	rwall(3R)	write to specified remote machines
secret()	publickey(3R)	get secret key
sm_inter()	sm_inter(3R)	status monitor protocol
spray()	spray(3R)	scatter data in order to check the network
xcrypt()	xcrypt(3R)	hex encryption and utility routines
yp()	yp(3R)	Yellow Pages protocol
yppasswd()	yppasswd(3R)	update user password in Yellow Pages

NAME

bootparam – bootparam protocol

PROTOCOL

`/usr/include/rpcsvc/bootparam_prot.x`

DESCRIPTION

The bootparam protocol is used for providing information to the diskless clients necessary for booting.

PROGRAMMING

`#include <rpcsvc/bootparam.h>`

XDR Routines

The following XDR routines are available in `librpcsvc`:

`xdr_bp_whoami_arg`
`xdr_bp_whoami_res`
`xdr_bp_getfile_arg`
`xdr_bp_getfile_res`

SEE ALSO

`bootparams(5)`, `bootparamd(8)`

NAME

ether – monitor traffic on the Ethernet

PROTOCOL

`/usr/include/rpcsvc/ether.x`

DESCRIPTION

The ether protocol is used for monitoring traffic on the ethernet.

PROGRAMMING

`#include <rpcsvc/ether.h>`

The following XDR routines are available in `librpcsvc`:

- `xdr_etherstat`
- `xdr_etheraddrs`
- `xdr_etherhtable`
- `xdr_etherhmem`
- `xdr_addrmask`

SEE ALSO

`traffic(1C)`, `etherfind(8C)`, `etherd(8C)`

NAME

getrpcport – get RPC port number

SYNOPSIS

```
int getrpcport(host, prognum, versnum, proto)
    char *host;
    int prognum, versnum, proto;
```

DESCRIPTION

getrpcport() returns the port number for version *versnum* of the RPC program *prognum* running on *host* and using protocol *proto*. It returns 0 if it cannot contact the portmapper, or if *prognum* is not registered. If *prognum* is registered but not with version *versnum*, it will still return a port number (for some version of the program) indicating that the program is indeed registered. The version mismatch will be detected upon the first call to the service.

NAME

klm_prot – protocol between kernel and local lock manager

PROTOCOL

/usr/include/rpcsvc/klm_prot.x

DESCRIPTION

The protocol is used for communication between kernel and local lock manager.

PROGRAMMING

#include <rpcsvc/klm_prot.h>

XDR Routines

The following XDR routines are available in **librpcsvc**:

xdr_klm_testargs
xdr_klm_testrply
xdr_klm_lockargs
xdr_klm_unlockargs
xdr_klm_stat

SEE ALSO

lockd(8C)

NAME

mount – keep track of remotely mounted filesystems

PROTOCOL

`/usr/include/rpcsvc/mount.x`

DESCRIPTION

The mount protocol is separate from, but related to, the NFS protocol. It provides all of the operating system specific services to get the NFS off the ground — looking up path names, validating user identity, and checking access permissions. Clients use the mount protocol to get the first file handle, which allows them entry into a remote filesystem.

The mount protocol is kept separate from the NFS protocol to make it easy to plug in new access checking and validation methods without changing the NFS server protocol.

Note: the protocol definition implies stateful servers because the server maintains a list of client's mount requests. The mount list information is not critical for the correct functioning of either the client or the server. It is intended for advisory use only, for example, to warn people when a server is going down.

PROGRAMMING

`#include <rpcsvc/mount.h>`

The following XDR routines are available in `librpcsvc`:

`xdr_exportbody`
`xdr_exports`
`xdr_fhandle`
`xdr_fhstatus`
`xdr_groups`
`xdr_mountbody`
`xdr_mountlist`
`xdr_path`

SEE ALSO

`mount(8)`, `mountd(8C)`, `showmount(8)`

NFS Protocol Spec, in *Network Programming*

NAME

ipalloc - determine or temporarily allocate IP address

PROTOCOL

/usr/include/rpcsvc/ipalloc.x

AVAILABILITY

Sun386i systems only.

DESCRIPTION

ipalloc() is the protocol for allocating the IP address that a system should use.

PROGRAMMING

#include <rpcsvc/ipalloc.h>

The following RPC calls are available in version 2 of this protocol:

NULLPROC

This is a standard null entry, used to ping a service to measure overhead or to discover servers.

IP_ALLOC

Returns an IP address corresponding to a given Ethernet address, if possible. This RPC must be called using DES authentication, from a client authorized to allocate IP addresses. A cache of allocated addresses is maintained.

The first action taken on receipt of this RPC is to verify that no existing mapping between the *etheraddr* and the *netnum* exists in the YP database. If one is found, then that is returned. Otherwise, an internal cache is checked, and if an entry is found there for the given *etheraddr* on the right network, that entry is used. If no address was found either in the YP database or in the cache, a new one may be allocated and returned, and the *ip_success* status is returned.

If an unusable entry was found in the cache, this RPC returns **ip_failure** status.

IP_TONAME

Used to determine whether a given IP address is known to the YP service, since YP allows a delay between the posting of an address and its availability in some locations on the network.

IP_FREE

This RPC is used to delete *ipaddr* entries from the cache when they are no longer needed there. It requires the same protections as the **IP_ALLOC** RPC.

SEE ALSO

ipallocald(8C), **pnpboot(8C)**

NAME

nlm_prot – protocol between local and remote network lock managers

PROTOCOL

/usr/include/rpcsvc/nlm_prot.x

DESCRIPTION

The network lock manager protocol is used for communication between local and remote lock managers.

PROGRAMMING

#include <rpcsvc/nlm_prot.h>

XDR Routines

The following XDR routines are available in librpcsvc:

- xdr_nlm_testargs
- xdr_nlm_testres
- xdr_nlm_lockargs
- xdr_nlm_cancargs
- xdr_nlm_unlockargs
- xdr_nlm_res

SEE ALSO

lockd(8C)

NAME

pnp - automatic network installation

PROTOCOL

/usr/include/rpcsvc/pnprpc.x

AVAILABILITY

Sun386i systems only.

DESCRIPTION

pnp() is used during unattended network installation, and routine booting, of Sun386i systems on a Sun386i network. Each network cable (subnetwork or full network) must have at least one **pnpd(8C)** server running on it to support PNP.

PROGRAMMING

#include <rpcsvc/pnprpc.h>

The following RPC calls are available in version 2 of the PNP protocol:

NULLPROC

Finds a PNP daemon on the local network. Used with **clntudp_broadcast()**, often to measure network overhead.

PNP_WHOAMI

Used early in the boot process to acquire network configuration information about a system, or to determine that a system is not known by the network.

PNP_ACQUIRE

Used to acquire a server willing to configure a new system after a **PNP_WHOAMI** request fails. This RPC is typically broadcast; any successful reply may be used.

PNP_SETUP

Requests a network configuration from a PNP daemon that has responded to a previous **PNP_ACQUIRE** RPC.

PNP_POLL

After a **PNP_SETUP** request, if the status is **in_progress**, the procedure is to wait 20 seconds, and issue a **PNP_POLL** request, and then check the status again. Once the status is **success**, the system will be configured for the network. Entries in the yp database may be added or old ones deleted, and file storage may be assigned, according to the architecture and boot type.

If the server misses 5 **PNP_POLL** requests, it will assume that the client system crashed and back out of the procedure. Similarly, if the client system does not receive responses from the server for **PNP_MISSEDPOLLS** consecutive requests, it should assume the server crashed and begin its PNP sequence again.

SEE ALSO

pnplib(8C), **pnpd(8C)**

NAME

publickey, getpublickey, getsecretkey – get public or secret key

SYNOPSIS

```
#include <rpc/rpc.h>  
#include <rpc/key_prot.h>  
  
getpublickey(netname, publickey)  
    char netname[MAXNETNAMELEN+1];  
    char publickey[HEXKEYBYTES+1];  
  
getsecretkey(netname, secretkey, passwd)  
    char netname[MAXNETNAMELEN+1];  
    char secretkey[HEXKEYBYTES+1];  
    char *passwd;
```

DESCRIPTION

These routines are used to get public and secret keys from the YP database. **getsecretkey()** has an extra argument, *passwd*, which is used to decrypt the encrypted secret key stored in the database. Both routines return 1 if they are successful in finding the key, 0 otherwise. The keys are returned as NULL-terminated, hexadecimal strings. If the password supplied to **getsecretkey()** fails to decrypt the secret key, the routine will return 1 but the *secretkey* argument will be a NULL string (“”).

SEE ALSO

publickey(5)

RPC Programmer's Manual in Network Programming

NAME

rex – remote execution protocol

PROTOCOL

/usr/include/rpcsvc/rex.x

DESCRIPTION

This server will execute commands remotely. The working directory and environment of the command can be specified, and the standard input and output of the command can be arbitrarily redirected. An option is provided for interactive I/O for programs that expect to be running on terminals. Note: this service is only provided with the TCP transport.

PROGRAMMING

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

```
#include <rpcsvc/rex.h> /* not compiled with rpgen
```

The following XDR routines are available in librpcsvc:

```
xdr_rex_start
```

```
xdr_rex_result
```

```
xdr_rex_ttymode
```

```
xdr_rex_ttysize
```

SEE ALSO

on(1C), rexd(8C)

NAME

rnusers, rusers – return information about users on remote machines

PROTOCOL

/usr/include/rpcsvc/rnusers.x

PROGRAMMING

```
#include <rpcsvc/rusers.h>
```

```
rnusers(host)
```

```
    char *host
```

```
rusers(host, up)
```

```
    char *host
```

```
    struct utmpidlearr *up;
```

rnusers() returns the number of users logged on to *host* (–1 if it cannot determine that number). **rusers()** fills the **utmpidlearr** structure with data about *host*, and returns 0 if successful.

The following XDR routines are also available:

```
xdr_utmpidle
```

```
xdr_utmpidlearr
```

SEE ALSO

rusers(1C)

NAME

rquota – implement quotas on remote machines

PROTOCOL

/usr/include/rpcsvc/rquota.x

DESCRIPTION

The **rquota()** protocol inquires about quotas on remote machines. It is used in conjunction with NFS, since NFS itself does not implement quotas.

PROGRAMMING

#include <rpcsvc/rquota.h>

The following XDR routines are available in **librpcsvc**:

xdr_getquota_arg

xdr_getquota_rslt

xdr_rquota

SEE ALSO

quota(1), quotactl(2)

NAME

rstat – get performance data from remote kernel

PROTOCOL

/usr/include/rpcsvc/rstat.x

DESCRIPTION

The **rstat()** protocol is used to gather statistics from remote kernel. Statistics are available on items such as paging, swapping and cpu utilization.

PROGRAMMING

#include <rpcsvc/rstat.h>

havedisk(host)

char *host;

rstat(host, statp)

char *host;

struct statstime *statp;

havedisk() returns 1 if *host* has a disk, 0 if it does not, and -1 if this cannot be determined. **rstat()** fills in the **statstime** structure for *host*, and returns 0 if it was successful.

The following XDR routines are available in **librpcsvc**:

xdr_statstime

xdr_statsswtch

xdr_stats

SEE ALSO

perfmeter(1), rup(1C), rstatd(8C)

NAME

rwall – write to specified remote machines

SYNOPSIS

```
#include <rpcsvc/rwall.h>

rwall(host, msg);
char *host, *msg;
```

DESCRIPTION

host prints the string *msg* to all its users. It returns 0 if successful.

RPC INFO

program number:

WALLPROG

procs:

WALLPROC_WALL

Takes string as argument (*wrapstring*), returns no arguments.

Executes *wall* on remote host with string.

versions:

RSTATVERS_ORIG

SEE ALSO

rwall(1C), rwalld(8C), shutdown(8)

NAME

sm_inter – status monitor protocol

PROTOCOL

/usr/include/rpcsvc/sm_inter.x

DESCRIPTION

The status monitor protocol is used for monitoring the status of remote hosts.

PROGRAMMING

#include <rpcsvc/sm_inter.h>

XDR Routines

The following XDR routines are available in `librpcsvc`:

- xdr_sm_name
- xdr_mon
- xdr_mon_id
- xdr_sm_stat_res
- xdr_sm_stat

SEE ALSO

statd(8C)

NAME

spray – scatter data in order to check the network

PROTOCOL

/usr/include/rpcsvc/spray.x

DESCRIPTION

The spray protocol sends packets to a given machine to test the speed and reliability of it.

PROGRAMMING

#include <rpcsvc/spray.h>

The following XDR routines are available in **librpcsvc**:

xdr_sprayarr

xdr_spraycumul

SEE ALSO

spray(8C), sprayd(8C)

NAME

xcrypt, xdecrypt, passwd2des – hex encryption and utility routines

SYNOPSIS

xencrypt(data, key)

char *data;

char *key;

xdecrypt(data, key)

char *data;

char *key;

passwd2des(pass, key)

char *pass;

char *key;

DESCRIPTION

The routines **xencrypt** and **xdecrypt** take NULL-terminated hexadecimal strings as arguments, and encrypt them using the 8-byte *key* as input to the DES algorithm. The input strings must have a length that is a multiple on 16 hex digits (64 bits is the DES block size).

passwd2des converts a password, of arbitrary length, into an 8-byte DES key, with odd-parity set in the low bit of each byte. The high-order bit of each input byte is ignored.

These routines are used by the DES authentication subsystem for encrypting and decrypting the secret keys stored in the publickey database.

SEE ALSO

des_crypt(3), publickey(5)

NAME

yp – Yellow Pages protocol

PROTOCOL

`/usr/include/rpcsvc/yp.x`

DESCRIPTION

The Yellow Pages Service is used for the administration of network-wide databases. The service is composed mainly of two programs: `YPBINDPROG` for finding a YP server and `YPPROG` for accessing the YP databases.

PROGRAMMING

Refer to `ypclnt(3N)` for information on the programmatic interface to YP servers and databases.

SEE ALSO

`ypclnt(3N)`, `yppasswd(3R)`

NAME

yppasswd – update user password in Yellow Pages

PROTOCOL

/usr/include/rpcsvc/yppasswd.x

DESCRIPTION

The yppasswd protocol is used to change a user's password entry in the YP password database.

PROGRAMMING

```
#include <rpcsvc/yppasswd.h>
```

```
yppasswd(oldpass, newpw)  
char *oldpass  
struct passwd *newpw;
```

If *oldpass* is indeed the old user password, this routine replaces the password entry with *newpw*. It returns 0 if successful.

SEE ALSO

yppasswd(1), yppasswdd(8C)

NAME

intro – introduction to System V functions

SYNOPSIS

/usr/5bin/cc

DESCRIPTION

These functions are contained in the System V library, */usr/5lib/libc.a*. They are automatically linked when you compile a C program with the C compiler in */usr/5bin/cc*.

LIST OF SYSTEM V LIBRARY FUNCTIONS

Name	Appears on Page	Description
<code>_tolower()</code>	<code>ctype(3V)</code>	character classification and conversion macros and functions
<code>_toupper()</code>	<code>ctype(3V)</code>	character classification and conversion macros and functions
<code>addch()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>addstr()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>asctime()</code>	<code>ctime(3V)</code>	convert date and time
<code>assert()</code>	<code>assert(3V)</code>	verify program assertion
<code>attroff()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>attron()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>attrset()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>baudrate()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>beep()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>box()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>cbreak()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>clear()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>clearerr()</code>	<code>ferror(3V)</code>	stream status inquiries
<code>clearok()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>clrtoebot()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>clrtoeol()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>ctime()</code>	<code>ctime(3V)</code>	convert date and time
<code>ctype()</code>	<code>ctype(3V)</code>	character classification and conversion macros and functions
<code>curses()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>curses()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>delay_output()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>delch()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>deleteln()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>delwin()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>doupdate()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>echo()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>endpwent()</code>	<code>getpwent(3V)</code>	get password file entry
<code>endwin()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>erase()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>erasechar()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>fdopen()</code>	<code>fopen(3V)</code>	open a stream
<code>feof()</code>	<code>ferror(3V)</code>	stream status inquiries
<code>ferror()</code>	<code>ferror(3V)</code>	stream status inquiries
<code>fgetc()</code>	<code>getc(3V)</code>	get character or integer from stream
<code>fgetpwent()</code>	<code>getpwent(3V)</code>	get password file entry
<code>fileno()</code>	<code>ferror(3V)</code>	stream status inquiries
<code>fixterm()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>flash()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>flushinp()</code>	<code>curses(3V)</code>	System V cursor addressing and screen display library
<code>fopen()</code>	<code>fopen(3V)</code>	open a stream

fprintf()	printf(3V)	formatted output conversion
freopen()	fopen(3V)	open a stream
fscanf()	scanf(3V)	formatted input conversion
getc()	getc(3V)	get character or integer from stream
getch()	curses(3V)	System V cursor addressing and screen display library
getchar()	getc(3V)	get character or integer from stream
getpass()	getpass(3V)	read a password
getpwent()	getpwent(3V)	get password file entry
getpwnam()	getpwent(3V)	get password file entry
getpwuid()	getpwent(3V)	get password file entry
getstr()	curses(3V)	System V cursor addressing and screen display library
gettmode()	curses(3V)	System V cursor addressing and screen display library
getw()	getc(3V)	get character or integer from stream
getyx()	curses(3V)	System V cursor addressing and screen display library
gmtime()	ctime(3V)	convert date and time
has_ic()	curses(3V)	System V cursor addressing and screen display library
has_il()	curses(3V)	System V cursor addressing and screen display library
idlok()	curses(3V)	System V cursor addressing and screen display library
inch()	curses(3V)	System V cursor addressing and screen display library
initscr()	curses(3V)	System V cursor addressing and screen display library
insch()	curses(3V)	System V cursor addressing and screen display library
insertln()	curses(3V)	System V cursor addressing and screen display library
intrflush()	curses(3V)	System V cursor addressing and screen display library
isalnum()	ctype(3V)	character classification and conversion macros and functions
isalpha()	ctype(3V)	character classification and conversion macros and functions
isascii()	ctype(3V)	character classification and conversion macros and functions
iscntrl()	ctype(3V)	character classification and conversion macros and functions
isdigit()	ctype(3V)	character classification and conversion macros and functions
isgraph()	ctype(3V)	character classification and conversion macros and functions
islower()	ctype(3V)	character classification and conversion macros and functions
isprint()	ctype(3V)	character classification and conversion macros and functions
ispunct()	ctype(3V)	character classification and conversion macros and functions
isspace()	ctype(3V)	character classification and conversion macros and functions
isupper()	ctype(3V)	character classification and conversion macros and functions
isxdigit()	ctype(3V)	character classification and conversion macros and functions
keypad()	curses(3V)	System V cursor addressing and screen display library
killchar()	curses(3V)	System V cursor addressing and screen display library
leaveok()	curses(3V)	System V cursor addressing and screen display library
localtime()	ctime(3V)	convert date and time
longname()	curses(3V)	System V cursor addressing and screen display library
meta()	curses(3V)	System V cursor addressing and screen display library
move()	curses(3V)	System V cursor addressing and screen display library
mvaddch()	curses(3V)	System V cursor addressing and screen display library
mvaddstr()	curses(3V)	System V cursor addressing and screen display library
mvcur()	curses(3V)	System V cursor addressing and screen display library
mvdelch()	curses(3V)	System V cursor addressing and screen display library
mvgetch()	curses(3V)	System V cursor addressing and screen display library
mvgetstr()	curses(3V)	System V cursor addressing and screen display library
mvinch()	curses(3V)	System V cursor addressing and screen display library
mvinsch()	curses(3V)	System V cursor addressing and screen display library
mvprintw()	curses(3V)	System V cursor addressing and screen display library
mvscanw()	curses(3V)	System V cursor addressing and screen display library
mvwaddch()	curses(3V)	System V cursor addressing and screen display library

mvwaddstr()	curses(3V)	System V cursor addressing and screen display library
mvwdelch()	curses(3V)	System V cursor addressing and screen display library
mvwgetch()	curses(3V)	System V cursor addressing and screen display library
mvwgetstr()	curses(3V)	System V cursor addressing and screen display library
mvwin()	curses(3V)	System V cursor addressing and screen display library
mvwinch()	curses(3V)	System V cursor addressing and screen display library
mvwinsch()	curses(3V)	System V cursor addressing and screen display library
mvwprintw()	curses(3V)	System V cursor addressing and screen display library
mvwscanw()	curses(3V)	System V cursor addressing and screen display library
newpad()	curses(3V)	System V cursor addressing and screen display library
newterm()	curses(3V)	System V cursor addressing and screen display library
newwin()	curses(3V)	System V cursor addressing and screen display library
nice()	nice(3V)	change priority of a process
nl()	curses(3V)	System V cursor addressing and screen display library
nocbreak()	curses(3V)	System V cursor addressing and screen display library
nodelay()	curses(3V)	System V cursor addressing and screen display library
noecho()	curses(3V)	System V cursor addressing and screen display library
nonl()	curses(3V)	System V cursor addressing and screen display library
noraw()	curses(3V)	System V cursor addressing and screen display library
overlay()	curses(3V)	System V cursor addressing and screen display library
overwrite()	curses(3V)	System V cursor addressing and screen display library
pnoutrefresh()	curses(3V)	System V cursor addressing and screen display library
printf()	printf(3V)	formatted output conversion
rand()	rand(3V)	simple random number generator
raw()	curses(3V)	System V cursor addressing and screen display library
refresh()	curses(3V)	System V cursor addressing and screen display library
resetterm()	curses(3V)	System V cursor addressing and screen display library
resetty()	curses(3V)	System V cursor addressing and screen display library
saveterm()	curses(3V)	System V cursor addressing and screen display library
savetty()	curses(3V)	System V cursor addressing and screen display library
scanf()	scanf(3V)	formatted input conversion
scanw()	curses(3V)	System V cursor addressing and screen display library
scroll()	curses(3V)	System V cursor addressing and screen display library
scrollok()	curses(3V)	System V cursor addressing and screen display library
set_term()	curses(3V)	System V cursor addressing and screen display library
setbuf()	setbuf(3V)	assign buffering to a stream
setbuffer()	setbuf(3V)	assign buffering to a stream
setgid()	setuid(3V)	set user and group IDs
setlinebuf()	setbuf(3V)	assign buffering to a stream
setpwent()	getpwent(3V)	get password file entry
setpwfile()	getpwent(3V)	get password file entry
setscreg()	curses(3V)	System V cursor addressing and screen display library
setterm()	curses(3V)	System V cursor addressing and screen display library
setuid()	setuid(3V)	set user and group IDs
setupterm()	curses(3V)	System V cursor addressing and screen display library
setvbuf()	setbuf(3V)	assign buffering to a stream
signal()	signal(3V)	simplified software signal facilities
sleep()	sleep(3V)	suspend execution for interval
sprintf()	printf(3V)	formatted output conversion
srand()	rand(3V)	simple random number generator
sscanf()	scanf(3V)	formatted input conversion
standend()	curses(3V)	System V cursor addressing and screen display library
standout()	curses(3V)	System V cursor addressing and screen display library

subwin()	curses(3V)	System V cursor addressing and screen display library
timegm()	ctime(3V)	convert date and time
timelocal()	ctime(3V)	convert date and time
times()	times(3V)	get process and child process times
toascii()	ctype(3V)	character classification and conversion macros and functions
tolower()	ctype(3V)	character classification and conversion macros and functions
touchwin()	curses(3V)	System V cursor addressing and screen display library
toupper()	ctype(3V)	character classification and conversion macros and functions
traceoff()	curses(3V)	System V cursor addressing and screen display library
traceon()	curses(3V)	System V cursor addressing and screen display library
ttyslot()	ttyslot(3V)	find the slot in the utmp file of the current process
typeahead()	curses(3V)	System V cursor addressing and screen display library
tzset()	ctime(3V)	convert date and time
tzsetwall()	ctime(3V)	convert date and time
unctrl()	curses(3V)	System V cursor addressing and screen display library
waddch()	curses(3V)	System V cursor addressing and screen display library
waddstr()	curses(3V)	System V cursor addressing and screen display library
wattroff()	curses(3V)	System V cursor addressing and screen display library
wattron()	curses(3V)	System V cursor addressing and screen display library
wattrset()	curses(3V)	System V cursor addressing and screen display library
wclear()	curses(3V)	System V cursor addressing and screen display library
wclrtobot()	curses(3V)	System V cursor addressing and screen display library
wclrtoeol()	curses(3V)	System V cursor addressing and screen display library
wdelch()	curses(3V)	System V cursor addressing and screen display library
wdeleteln()	curses(3V)	System V cursor addressing and screen display library
werase()	curses(3V)	System V cursor addressing and screen display library
wgetch()	curses(3V)	System V cursor addressing and screen display library
wgetstr()	curses(3V)	System V cursor addressing and screen display library
winch()	curses(3V)	System V cursor addressing and screen display library
winsch()	curses(3V)	System V cursor addressing and screen display library
winsertln()	curses(3V)	System V cursor addressing and screen display library
wmove()	curses(3V)	System V cursor addressing and screen display library
wnoutrefresh()	curses(3V)	System V cursor addressing and screen display library
wprintw()	curses(3V)	System V cursor addressing and screen display library
wrefresh()	curses(3V)	System V cursor addressing and screen display library
wscanw()	curses(3V)	System V cursor addressing and screen display library
wsetscrreg()	curses(3V)	System V cursor addressing and screen display library
wstandend()	curses(3V)	System V cursor addressing and screen display library
wstandout()	curses(3V)	System V cursor addressing and screen display library

NAME

assert – verify program assertion

SYNOPSIS

```
#include <assert.h>
```

```
assert(expression)
```

```
int expression;
```

DESCRIPTION

assert() is a macro that indicates *expression* is expected to be true at this point in the program. When it is executed, if *expression* is false (zero), assert() prints

```
'Assertion failed: expression, file xyz, line nnn'
```

on the standard error output and aborts. In the error message, *xyz* is the name of the source file and *nnn* the source line number of the assert() statement.

Compiling with the cc(1V) option -DNDEBUG, or with the preprocessor control statement '#define NDEBUG' ahead of the '#include <assert.h>' statement, will stop assertions from being compiled into the program.

SEE ALSO

cc(1V), abort(3)

NAME

`ctime`, `localtime`, `gmtime`, `asctime`, `timelocal`, `timegm`, `tzset`, `tzsetwall` – convert date and time

SYNOPSIS

```
#include <time.h>

struct tm *localtime(clock)
long *clock;

struct tm *gmtime(clock)
long *clock;

char *asctime(tm)
struct tm *tm;

char *ctime(clock)
long *clock;

time_t timelocal(tm)
struct tm *tm;

time_t timegm(tm)
struct tm *tm;

void tzset()

void tzsetwall()

extern long timezone;

extern int daylight;

extern char *tzname[2];
```

DESCRIPTION

`localtime()` and `gmtime()` return pointers to structures containing the time, broken down into various components of that time represented in a particular time zone. `localtime()` breaks down a time specified by the `clock()` argument, correcting for the time zone and any time zone adjustments (such as Daylight Savings Time). Before doing so, `localtime()` calls `tzset()` (if `tzset()` has not been called in the current process). `gmtime()` breaks down a time specified by the `clock()` argument into GMT, which is the time the system uses.

`asctime()` converts a time value contained in a “tm” structure to a 26-character string of the form:

```
Sun Sep 16 01:03:52 1973\n\0
```

Each field has a constant width. `asctime()` returns a pointer to the string.

`ctime()` converts a long integer, pointed to by `clock`, to a 26-character string of the form produced by `asctime()`. It first breaks down `clock()` to a `struct tm` by calling `localtime`, and then calls `asctime()` to convert that `struct tm` to a string.

`timelocal()` and `timegm()` convert the time specified by the `tm` argument to a time value that represents that time expressed as the number of seconds since Jan. 1, 1970, 00:00, Greenwich Mean Time. `timelocal()` converts a `struct tm` that represents local time, correcting for the time zone and any time zone adjustments (such as Daylight Savings Time). Before doing so, `timelocal()` calls `tzset()` (if `tzset()` has not been called in the current process). `timegm()` converts a `struct tm` that represents GMT.

`tzset()` uses the value of the environment variable `TZ` to set time conversion information used by `localtime`. If `TZ` is absent from the environment, the best available approximation to local wall clock time is used by `localtime`. If `TZ` appears in the environment but its value is a NULL string, Greenwich Mean Time is used; if `TZ` appears and begins with a slash, it is used as the absolute pathname of the `tzfile-format` (see `tzfile(5)`) file from which to read the time conversion information; if `TZ` appears and begins with a character other than a slash, it is used as a pathname relative to a system time conversion information directory.

tzsetwall() sets things up so that **localtime()** returns the best available approximation of local wall clock time.

Declarations of all the functions and externals, and the “tm” structure, are in the **<time.h>** header file. The structure (of type) **struct tm** includes the following fields:

```

int tm_sec;      /* seconds (0 - 59) */
int tm_min;     /* minutes (0 - 59) */
int tm_hour;    /* hours (0 - 23) */
int tm_mday;    /* day of month (1 - 31) */
int tm_mon;     /* month of year (0 - 11) */
int tm_year;    /* year - 1900 */
int tm_wday;    /* day of week (Sunday = 0) */
int tm_yday;    /* day of year (0 - 365) */
int tm_isdst;   /* 1 if DST in effect */
char *tm_zone;  /* abbreviation of timezone name */
long tm_gmtoff; /* offset from GMT in seconds */

```

tm_isdst is non-zero if Daylight Savings Time is in effect. **tm_zone** points to a string that is the name used for the local time zone at the time being converted. **tm_gmtoff** is the offset (in seconds) of the time represented from GMT, with positive values indicating East of Greenwich.

The external **long** variable *timezone* contains the difference, in seconds, between GMT and local standard time (in PST, *timezone* is 8*60*60). If this difference is not a constant, *timezone* will contain the value of the offset on January 1, 1970 at 00:00 GMT. Since this is not necessarily the same as the value at some particular time, the time in question should be converted to a “struct tm” using **localtime** (see **ctime(3)**) and the **tm_gmtoff** field of that structure should be used. The external variable *daylight* is non-zero if and only if Daylight Savings Time would be in effect within the current time zone at some time; it does not indicate whether Daylight Savings Time is currently in effect.

The external variable *tzname* is an array of two **char *** pointers. The first pointer points to a character string that is the name of the current time zone when Daylight Savings Time is not in effect; the second one, if Daylight Savings Time conversion should be applied, points to a character string that is the name of the current time zone when Daylight Savings Time is in effect. These strings are updated by **localtime()** whenever a time is converted. If Daylight Savings Time is in effect at the time being converted, the second pointer is set to point to the name of the current time zone at that time, otherwise the first pointer is so set.

timezone, *daylight*, and *tzname* are retained for compatibility with existing programs.

FILES

/usr/share/lib/zoneinfo standard time conversion information directory
/usr/share/lib/zoneinfo/localtime
 local time zone file

SEE ALSO

gettimeofday(2), **ctime(3)**, **getenv(3)**, **time(3C)**, **environ(5V)**, **tzfile(5)**

BUGS

The return values point to static data, whose contents are overwritten by each call. The **tm_zone** field of a returned **struct tm** points to a static array of characters, which will also be overwritten at the next call (and by calls to **tzset()** or **tzsetwall()**).

NAME

`ctype`, `isalpha`, `isupper`, `islower`, `isdigit`, `isxdigit`, `isalnum`, `isspace`, `ispunct`, `isprint`, `isctrl`, `isascii`, `isgraph`, `toupper`, `tolower`, `toascii`, `_toupper`, `_tolower` – character classification and conversion macros and functions

SYNOPSIS

```
#include <ctype.h>
```

```
isalpha(c)
```

```
...
```

CHARACTER CLASSIFICATION MACROS

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(c)` is true and on the single non-ASCII value EOF (see `stdio(3V)`).

<code>isalpha(c)</code>	<code>c</code> is a letter
<code>isupper(c)</code>	<code>c</code> is an upper case letter
<code>islower(c)</code>	<code>c</code> is a lower case letter
<code>isdigit(c)</code>	<code>c</code> is a digit [0-9].
<code>isxdigit(c)</code>	<code>c</code> is a hexadecimal digit [0-9], [A-F], or [a-f].
<code>isalnum(c)</code>	<code>c</code> is an alphanumeric character, that is, <code>c</code> is a letter or a digit
<code>isspace(c)</code>	<code>c</code> is a space, tab, carriage return, newline, vertical tab, or formfeed
<code>ispunct(c)</code>	<code>c</code> is a punctuation character (neither control nor alphanumeric)
<code>isprint(c)</code>	<code>c</code> is a printing character, code 040(8) (space) through 0176 (tilde)
<code>isctrl(c)</code>	<code>c</code> is a delete character (0177) or ordinary control character (less than 040).
<code>isascii(c)</code>	<code>c</code> is an ASCII character, code less than 0200
<code>isgraph(c)</code>	<code>c</code> is a visible graphic character, code 041 (exclamation mark) through 0176 (tilde).

CHARACTER CONVERSION MACROS AND FUNCTIONS

`toupper` and `tolower` are functions, rather than macros, and work correctly on all characters. The macros `_toupper` and `_tolower` are faster than the equivalent functions (`toupper` and `tolower`) but only work properly on a restricted range of characters.

These functions perform simple conversions on single characters.

<code>toupper(c)</code>	converts <code>c</code> to its upper-case equivalent. If <code>c</code> is not a lower-case letter, it is returned unchanged.
<code>tolower(c)</code>	converts <code>c</code> to its lower-case equivalent. If <code>c</code> is not an upper-case letter, it is returned unchanged.
<code>toascii(c)</code>	masks <code>c</code> with the correct value so that <code>c</code> is guaranteed to be an ASCII character in the range 0 thru 0x7f.

These macros perform simple conversions on single characters.

<code>_toupper(c)</code>	converts <code>c</code> to its upper-case equivalent. Note that this <i>only</i> works where <code>c</code> is known to be a lower-case character to start with (presumably checked using <code>islower</code>).
<code>_tolower(c)</code>	converts <code>c</code> to its lower-case equivalent. Note: this <i>only</i> works where <code>c</code> is known to be an upper-case character to start with (presumably checked using <code>isupper</code>).

DIAGNOSTICS

If the argument to any of these macros is not in the domain of the function, the result is undefined.

SEE ALSO

ctype(3), stdio(3V), ascii(7)

NAME

curses – System V terminal screen handling and optimization package

SYNOPSIS

The **curses** manual page is organized as follows:

In SYNOPSIS

- compiling information
- summary of parameters used by **curses** routines
- alphabetical list of **curses** routines, showing their parameters

In DESCRIPTION:

- An overview of how **curses** routines should be used

In **ROUTINES**, descriptions of **curses** routines are grouped under the appropriate topics:

- Overall Screen Manipulation
- Window and Pad Manipulation
- Output
- Input
- Output Options Setting
- Input Options Setting
- Environment Queries
- Soft Labels
- Low-level Curses Access
- Terminfo-Level Manipulations
- Termcap Emulation
- Miscellaneous
- Use of **curscr**

Then come sections on:

- ATTRIBUTES
- FUNCTION CALLS
- LINE GRAPHICS

`/usr/5bin/cc [flag ...] file ... -lcurses [library ...]`

`#include <curses.h>` (automatically includes `<stdio.h>`, `<termio.h>`, and `<unctrl.h>`).

The parameters in the following list are not global variables; this is a summary of the parameters used by the **curses** library routines. All routines return the **int** values **ERR** or **OK** unless otherwise noted. Routines that return pointers always return **NULL** on error. (**ERR**, **OK**, and **NULL** are all defined in `<curses.h>`.) Routines that return integers are not listed in the parameter list below.

bool *bf*

char ***area, *boolnames[], *boolcodes[], *boolfnames[], *bp*
char **cap, *capname, codename[2], erasechar, *filename, *fmt*
char **keyname, killchar, *label, *longname*
char **name, *numnames[], *numcodes[], *numfnames[]*
char **slk_label, *str, *strnames[], *strcodes[], *strfnames[]*
char **term, *tgetstr, *tigetstr, *tgoto, *tparm, *type*
chtype *attrs, ch, horch, vertch*

FILE **infd, *outfd*

int *begin_x, begin_y, begline, bot, c, col, count*
int *dmaxcol, dmaxrow, dmincol, dminrow, *errret, fildes*
int *(*init()), labfmt, labnum, line*
int *ms, ncol, new, newcol, newrow, nlines, numlines*

int oldcol, oldrow, overlay
int p1, p2, p9, pmincol, pminrow, (*putc()), row
int smaxcol, smaxrow, smincol, sminrow, start
int tenths, top, visibility, x, y
SCREEN *new, *newterm, *set_term
TERMINAL *cur_term, *nterm, *oterm
va_list varglist
WINDOW *curscr, *dstwin, *initscr, *newpad, *newwin, *orig
WINDOW *pad, *srcwin, *stdscr, *subpad, *subwin, *win
addch (*ch*)
addstr (*str*)
attroff (*attrs*)
attron (*attrs*)
attrset (*attrs*)
baudrate()
beep()
box (*win, vertch, horch*)
cbreak()
clear()
clearok (*win, bf*)
clrtoebot()
clrtoeol()
copywin (*srcwin, dstwin, sminrow, smincol, dminrow, dmincol, dmaxrow, dmaxcol, overlay*)
curs_set (*visibility*)
def_prog_mode()
def_shell_mode()
del_curterm (*oterm*)
delay_output (*ms*)
delch()
deleteln()
delwin (*win*)
doupdate()
draino (*ms*)
echo()
echochar (*ch*)
endwin()
erase()
erasechar()
filter()
flash()
flushinp()
garbagedlines (*win, begline, numlines*)
getbegyx (*win, y, x*)
getch()
getmaxyx (*win, y, x*)
getstr (*str*)
getsyx (*y, x*)
getyx (*win, y, x*)
halfdelay (*tenths*)
has_ic()
has_il()
idlok (*win, bf*)

inch()
initscr()
insch (*ch*)
insertln()
intrflush (*win, bf*)
isendwin()
keyname (*c*)
keypad (*win, bf*)
killchar()
leaveok (*win, bf*)
longname()
meta (*win, bf*)
move (*y, x*)
mvaddch (*y, x, ch*)
mvaddstr (*y, x, str*)
mvcur (*oldrow, oldcol, newrow, newcol*)
mvdclch (*y, x*)
mvgetch (*y, x*)
mvgetstr (*y, x, str*)
mvinch (*y, x*)
mvinsch (*y, x, ch*)
mvprintw (*y, x, fmt* [, *arg ...*])
mvscanw (*y, x, fmt* [, *arg ...*])
mvwaddch (*win, y, x, ch*)
mvwaddstr (*win, y, x, str*)
mvwclch (*win, y, x*)
mvwgetch (*win, y, x*)
mvwgetstr (*win, y, x, str*)
mvwin (*win, y, x*)
mvwinch (*win, y, x*)
mvwinsch (*win, y, x, ch*)
mvwprintw (*win, y, x, fmt* [, *arg ...*])
mvwscanw (*win, y, x, fmt* [, *arg ...*])
napms (*ms*)
newpad (*nlines, ncols*)
newterm (*type, outfd, infd*)
newwin (*nlines, ncols, begin_y, begin_x*)
nl()
nocbreak()
nodelay (*win, bf*)
noecho()
nonl()
noraw()
notimeout (*win, bf*)
overlay (*srcwin, dstwin*)
overwrite (*srcwin, dstwin*)
pechochar (*pad, ch*)
pnoutrefresh (*pad, pminrow, pmincol, sminrow, smincol, smaxrow, smaxcol*)
prefresh (*pad, pminrow, pmincol, sminrow, smincol, smaxrow, smaxcol*)
printw (*fmt* [, *arg ...*])
putp (*str*)
raw()
refresh()

reset_prog_mode()
reset_shell_mode()
resetty()
restartterm (*term, fildes, errret*)
ripoffline (*line, init*)
savetty()
scanw (*fmt* [, *arg...*])
scr_dump (*filename*)
scr_init (*filename*)
scr_restore (*filename*)
scroll (*win*)
scrollok (*win, bf*)
set_curterm (*nterm*)
set_term (*new*)
setscrreg (*top, bot*)
setsyx (*y, x*)
setupterm (*term, fildes, errret*)
slk_clear()
slk_init (*fmt*)
slk_label (*labnum*)
slk_noutrefresh()
slk_refresh()
slk_restore()
slk_set (*labnum, label, fmt*)
slk_touch()
standend()
standout()
subpad (*orig, nlines, ncols, begin_y, begin_x*)
subwin (*orig, nlines, ncols, begin_y, begin_x*)
tgetent (*bp, name*)
tgetflag (*codename*)
tgetnum (*codename*)
tgetstr (*codename, area*)
tgoto (*cap, col, row*)
tigetflag (*capname*)
tigetnum (*capname*)
tigetstr (*capname*)
touchline (*win, start, count*)
touchwin (*win*)
tparm (*str, p1, p2, ..., p9*)
tputs (*str, count, putc*)
traceoff()
traceon()
typeahead (*fildes*)
unctrl (*c*)
ungetch (*c*)
vidattr (*attrs*)
vidputs (*attrs, putc*)
vwprintw (*win, fmt, varglist*)
vwscanw (*win, fmt, varglist*)
waddch (*win, ch*)
waddstr (*win, str*)
wattroff (*win, attrs*)

wattron (*win, attrs*)
wattrset (*win, attrs*)
wclear (*win*)
wclrtoebot (*win*)
wclrtoeol (*win*)
wdelch (*win*)
wdeleteln (*win*)
wechochar (*win, ch*)
werase (*win*)
wgetch (*win*)
wgetstr (*win, str*)
winch (*win*)
winsch (*win, ch*)
winsertln (*win*)
wmove (*win, y, x*)
wnoutrefresh (*win*)
wprintw (*win, fmt [, arg ...]*)
wrefresh (*win*)
wscanw (*win, fmt [, arg ...]*)
wsetscrreg (*win, top, bot*)
wstandend (*win*)
wstandout (*win*)

DESCRIPTION

The **curses** routines give the user a terminal-independent method of updating screens with reasonable optimization.

In order to initialize the routines, the routine **initscr()** or **newterm()** must be called before any of the other routines that deal with windows and screens are used. (Three exceptions are noted where they apply.) The routine **endwin()** must be called before exiting. To get character-at-a-time input without echoing, (most interactive, screen oriented programs want this) after calling **initscr()** you should call **'cbreak (); noecho ();'** Most programs would additionally call **'nonl (); intrflush(stdscr, FALSE); keypad(stdscr, TRUE);'**

Before a **curses** program is run, a terminal's TAB stops should be set and its initialization strings, if defined, must be output. This can be done by executing the **tset** command in your **.profile** or **.login** file. For further details, see **tset(1)** and the **Tabs and Initialization** subsection of **terminfo(5V)**.

The **curses** library contains routines that manipulate data structures called *windows* that can be thought of as two-dimensional arrays of characters representing all or part of a terminal screen. A default window called **stdscr** is supplied, which is the size of the terminal screen. Others may be created with **newwin()**. Windows are referred to by variables declared as **WINDOW ***; the type **WINDOW** is defined in **<curses.h>** to be a C structure. These data structures are manipulated with routines described below, among which the most basic are **move()** and **addch()**. (More general versions of these routines are included with names beginning with **w**, allowing you to specify a window. The routines not beginning with **w** usually affect **stdscr**.) Then **refresh()** is called, telling the routines to make the user's terminal screen look like **stdscr**. The characters in a window are actually of type **chtype**, so that other information about the character may also be stored with each character.

Special windows called *pads* may also be manipulated. These are windows that are not constrained to the size of the screen and whose contents need not be displayed completely. See the description of **newpad()** under **Window and Pad Manipulation** for more information.

In addition to drawing characters on the screen, video attributes may be included that cause the characters to show up in modes such as underlined or in reverse video on terminals that support such display enhancements. Line drawing characters may be specified to be output. On input, **curses** is also able to translate arrow and function keys that transmit escape sequences into single values. The video attributes, line drawing characters, and input values use names, defined in **<curses.h>**, such as **A_REVERSE**, **ACS_HLINE**, and

KEY_LEFT.

curses also defines the **WINDOW *** variable, **curscr**, which is used only for certain low-level operations like clearing and redrawing a garbaged screen. **curscr** can be used in only a few routines. If the window argument to **clearok()** is **curscr**, the next call to **wrefresh()** with any window will clear and repaint the screen from scratch. If the window argument to **wrefresh()** is **curscr**, the screen is immediately cleared and repainted from scratch. This is how most programs would implement a "repaint-screen" function. More information on using **curscr** is provided where its use is appropriate.

The environment variables **LINES** and **COLUMNS** may be set to override **curses**'s idea of how large a screen is.

If the environment variable **TERMINFO** is defined, any program using **curses** will check for a local terminal definition before checking in the standard place. For example, if the environment variable **TERM** is set to **sun**, then the compiled terminal definition is found in **/usr/share/lib/terminfo/s/sun**. (The **s** is copied from the first letter of **sun** to avoid creation of huge directories.) However, if **TERMINFO** is set to **\$HOME/myterms**, **curses** will first check **\$HOME/myterms/s/sun**, and, if that fails, will then check **/usr/share/lib/terminfo/s/sun**. This is useful for developing experimental definitions or when write permission on **/usr/share/lib/terminfo** is not available.

The integer variables **LINES** and **COLS** are defined in **<curses.h>**, and will be filled in by **initscr()** with the size of the screen. (For more information, see the subsection **Terminfo-Level Manipulations**.) The constants **TRUE** and **FALSE** have the values **1** and **0**, respectively. The constants **ERR** and **OK** are returned by routines to indicate whether the routine successfully completed. These constants are also defined in **<curses.h>**.

ROUTINES

Many of the following routines have two or more versions. The routines prefixed with **w** require a *window* argument. The routines prefixed with **p** require a *pad* argument. Those without a prefix generally use **stdscr**.

The routines prefixed with **mv** require **y** and **x** coordinates to move to before performing the appropriate action. The **mv** routines imply a call to **move()** before the call to the other routine. The window argument is always specified before the coordinates. **y** always refers to the row (of the window), and **x** always refers to the column. The upper left corner is always **(0,0)**, not **(1,1)**. The routines prefixed with **mvw** take both a *window* argument and **y** and **x** coordinates.

In each case, *win* is the window affected and *pad* is the pad affected. (*win* and *pad* are always of type **WINDOW ***.) Option-setting routines require a boolean flag *bf* with the value **TRUE** or **FALSE**. (*bf* is always of type **bool**.) The types **WINDOW**, **bool**, and **chtype** are defined in **<curses.h>**. See the **SYNOPSIS** for a summary of what types all variables are.

All routines return either the integer **ERR** or the integer **OK**, unless otherwise noted. Routines that return pointers always return **NULL** on error.

Overall Screen Manipulation

WINDOW *initscr() The first routine called should almost always be **initscr()**. (The exceptions are **slk_init()**, **filter()**, and **ripoffline()**.) This will determine the terminal type and initialize all **curses** data structures. **initscr()** also arranges that the first call to **refresh()** will clear the screen. If errors occur, **initscr()** will write an appropriate error message to standard error and exit; otherwise, a pointer to **stdscr** is returned. If the program wants an indication of error conditions, **newterm()** should be used instead of **initscr()**. **initscr()** should only be called once per application.

endwin() A program should always call **endwin()** before exiting or escaping from **curses** mode temporarily, to do a shell escape or **system(3)** call, for example. This routine will restore **termio(4)** modes, move the cursor to the lower left corner of the screen and reset the terminal into the proper non-visual mode. To resume after a temporary escape, call **wrefresh()** or **doupdate()**.

isendwin() Returns TRUE if **endwin()** has been called without any subsequent calls to **wrefresh()**.

SCREEN *newterm(*type, outfd, infd*)

A program that outputs to more than one terminal must use **newterm()** for each terminal instead of **initscr()**. A program that wants an indication of error conditions, so that it may continue to run in a line-oriented mode if the terminal cannot support a screen-oriented program, must also use this routine. **newterm()** should be called once for each terminal. It returns a variable of type **SCREEN*** that should be saved as a reference to that terminal. The arguments are the *type* of the terminal to be used in place of the environment variable **TERM**; *outfd*, a **stdio(3V)** file pointer for output to the terminal; and *infd*, another file pointer for input from the terminal. When it is done running, the program must also call **endwin()** for each terminal being used. If **newterm()** is called more than once for the same terminal, the first terminal referred to must be the last one for which **endwin()** is called.

SCREEN *set_term (*new*)

This routine is used to switch between different terminals. The screen reference *new* becomes the new current terminal. A pointer to the screen of the previous terminal is returned by the routine. This is the only routine that manipulates **SCREEN** pointers; all other routines affect only the current terminal.

Window and Pad Manipulation

refresh()

wrefresh (*win*)

These routines (or **prefresh()**, **pnoutrefresh()**, **wnoutrefresh()**, or **doupdate()**) must be called to write output to the terminal, as most other routines merely manipulate data structures. **wrefresh()** copies the named window to the physical terminal screen, taking into account what is already there in order to minimize the amount of information that's sent to the terminal (called optimization). **refresh()** does the same thing, except it uses **stdscr** as a default window. Unless **leaveok()** has been enabled, the physical cursor of the terminal is left at the location of the window's cursor. The number of characters output to the terminal is returned.

Note: **refresh()** is a macro.

wnoutrefresh (*win*)

doupdate()

These two routines allow multiple updates to the physical terminal screen with more efficiency than **wrefresh()** alone. How this is accomplished is described in the next paragraph.

curses keeps two data structures representing the terminal screen: a *physical* terminal screen, describing what is actually on the screen, and a *virtual* terminal screen, describing what the programmer wants to have on the screen. **wrefresh()** works by first calling **wnoutrefresh()**, which copies the named window to the virtual screen, and then by calling **doupdate()**, which compares the virtual screen to the physical screen and does the actual update. If the programmer wishes to output several windows at once, a series of calls to **wrefresh()** will result in alternating calls to **wnoutrefresh()** and **doupdate()**, causing several bursts of output to the screen. By first calling **wnoutrefresh()** for each window, it is then possible to call **doupdate()** once, resulting in only one burst of output, with probably fewer total characters transmitted and certainly less processor time used.

WINDOW *newwin (*nlines, ncols, begin_y, begin_x*)

Create and return a pointer to a new window with the given number of lines (or rows), *nlines*, and columns, *ncols*. The upper left corner of the window is at line *begin_y*, column *begin_x*. If either *nlines* or *ncols* is 0, they will be set to the

value of *lines*–*begin_y* and *cols*–*begin_x*. A new full-screen window is created by calling `newwin(0,0,0,0)`.

mvwin (*win*, *y*, *x*) Move the window so that the upper left corner will be at position (*y*, *x*). If the move would cause the window to be off the screen, it is an error and the window is not moved.

WINDOW *subwin (*orig*, *nlines*, *ncols*, *begin_y*, *begin_x*) Create and return a pointer to a new window with the given number of lines (or rows), *nlines*, and columns, *ncols*. The window is at position (*begin_y*, *begin_x*) on the screen. (This position is relative to the screen, and not to the window *orig*.) The window is made in the middle of the window *orig*, so that changes made to one window will affect both windows. When using this routine, often it will be necessary to call `touchwin()` or `touchline()` on *orig* before calling `wrefresh`.

delwin (*win*) Delete the named window, freeing up all memory associated with it. In the case of overlapping windows, subwindows should be deleted before the main window.

WINDOW *newpad (*nlines*, *ncols*) Create and return a pointer to a new pad data structure with the given number of lines (or rows), *nlines*, and columns, *ncols*. A pad is a window that is not restricted by the screen size and is not necessarily associated with a particular part of the screen. Pads can be used when a large window is needed, and only a part of the window will be on the screen at one time. Automatic refreshes of pads (for example, from scrolling or echoing of input) do not occur. It is not legal to call `wrefresh()` with a pad as an argument; the routines `prefresh()` or `pnoutrefresh()` should be called instead. Note: these routines require additional parameters to specify the part of the pad to be displayed and the location on the screen to be used for display.

WINDOW *subpad (*orig*, *nlines*, *ncols*, *begin_y*, *begin_x*) Create and return a pointer to a subwindow within a pad with the given number of lines (or rows), *nlines*, and columns, *ncols*. Unlike `subwin()`, which uses screen coordinates, the window is at position (*begin_y*, *begin_x*) on the pad. The window is made in the middle of the window *orig*, so that changes made to one window will affect both windows. When using this routine, often it will be necessary to call `touchwin()` or `touchline()` on *orig* before calling `prefresh()`.

prefresh (*pad*, *pminrow*, *pmincol*, *sminrow*, *smincol*, *smaxrow*, *smaxcol*)

pnoutrefresh (*pad*, *pminrow*, *pmincol*, *sminrow*, *smincol*, *smaxrow*, *smaxcol*)

These routines are analogous to `wrefresh()` and `wnoutrefresh()` except that pads, instead of windows, are involved. The additional parameters are needed to indicate what part of the pad and screen are involved. *pminrow* and *pmincol* specify the upper left corner, in the pad, of the rectangle to be displayed. *sminrow*, *smincol*, *smaxrow*, and *smaxcol* specify the edges, on the screen, of the rectangle to be displayed in. The lower right corner in the pad of the rectangle to be displayed is calculated from the screen coordinates, since the rectangles must be the same size. Both rectangles must be entirely contained within their respective structures. Negative values of *pminrow*, *pmincol*, *sminrow*, or *smincol* are treated as if they were zero.

Output

These routines are used to “draw” text on windows.

addch (*ch*)
waddch (*win, ch*)
mvaddch (*y, x, ch*)
mvwaddch (*win, y, x, ch*)

The character *ch* is put into the window at the current cursor position of the window and the position of the window cursor is advanced. Its function is similar to that of **putchar()** (see **putc(3S)**). At the right margin, an automatic newline is performed. At the bottom of the scrolling region, if **scrollok()** is enabled, the scrolling region will be scrolled up one line.

If *ch* is a TAB, NEWLINE, or backspace, the cursor will be moved appropriately within the window. A NEWLINE also does a **clrtoeol()** before moving. TAB characters are considered to be at every eighth column. If *ch* is another control character, it will be drawn in the CTRL-X notation. (Calling **winch()** after adding a control character will not return the control character, but instead will return the representation of the control character.)

Video attributes can be combined with a character by or-ing them into the parameter. This will result in these attributes also being set. (The intent here is that text, including attributes, can be copied from one place to another using **inch()** and **addch()**.) See **standout()**, below.

Note: *ch* is actually of type **chtype**, not a character.

Note: **addch()**, **mvaddch()**, and **mvwaddch()** are macros.

echochar (*ch*)
wechochar (*win, ch*)
pechochar (*pad, ch*)

These routines are functionally equivalent to a call to **addch** (*ch*) followed by a call to **refresh()**, a call to **waddch** (*win, ch*) followed by a call to **wrefresh** (*win*), or a call to **waddch** (*pad, ch*) followed by a call to **prefresh** (*pad*). The knowledge that only a single character is being output is taken into consideration and, for non-control characters, a considerable performance gain can be seen by using these routines instead of their equivalents. In the case of **pechochar()**, the last location of the pad on the screen is reused for the arguments to **prefresh()**.

Note: *ch* is actually of type **chtype**, not a character.

Note: **echochar()** is a macro.

addstr (*str*)
waddstr (*win, str*)
mvwaddstr (*win, y, x, str*)
mvaddstr (*y, x, str*)

These routines write all the characters of the null-terminated character string *str* on the given window. This is equivalent to calling **waddch()** once for each character in the string.

Note: **addstr()**, **mvaddstr()**, and **mvwaddstr()** are macros.

attroff (*attrs*)
wattroff (*win, attrs*)
attron (*attrs*)
wattron (*win, attrs*)
attrset (*attrs*)
wattrset (*win, attrs*)
standend ()
wstandend (*win*)
standout ()
wstandout (*win*)

These routines manipulate the current attributes of the named window. These

attributes can be any combination of `A_STANDOUT`, `A_REVERSE`, `A_BOLD`, `A_DIM`, `A_BLINK`, `A_UNDERLINE`, and `A_ALTCHARSET`. These constants are defined in `<curses.h>` and can be combined with the C logical OR (`|`) operator.

The current attributes of a window are applied to all characters that are written into the window with `waddch()`. Attributes are a property of the character, and move with the character through any scrolling and insert/delete line/character operations. To the extent possible on the particular terminal, they will be displayed as the graphic rendition of the characters put on the screen.

`attrset(attrs)` sets the current attributes of the given window to *attrs*. `attroff(attrs)` turns off the named attributes without turning on or off any other attributes. `attron(attrs)` turns on the named attributes without affecting any others. `standout()` is the same as `attron(A_STANDOUT)`. `standend()` is the same as `attrset(0)`, that is, it turns off all attributes.

Note: *attrs* is actually of type `chtype`, not a character.

Note: `attroff()`, `attron()`, `attrset()`, `standend()`, and `standout()` are macros.

beep()

flash()

These routines are used to signal the terminal user. `beep()` will sound the audible alarm on the terminal, if possible, and if not, will flash the screen (visible bell), if that is possible. `flash()` will flash the screen, and if that is not possible, will sound the audible signal. If neither signal is possible, nothing will happen. Nearly all terminals have an audible signal (bell or beep) but only some can flash the screen.

box(win, vertch, horch)

A box is drawn around the edge of the window, *win*. *vertch* and *horch* are the characters the box is to be drawn with. If *vertch* and *horch* are `0`, then appropriate default characters, `ACS_VLINE` and `ACS_HLINE`, will be used.

Note: *vertch* and *horch* are actually of type `chtype`, not characters.

erase()

werase(win)

These routines copy blanks to every position in the window.

Note: `erase()` is a macro.

clear()

wclear(win)

These routines are like `erase()` and `werase()`, but they also call `clearok()`, arranging that the screen will be cleared completely on the next call to `wrefresh()` for that window, and repainted from scratch.

Note: `clear()` is a macro.

clrtoobot()

wclrtoobot(win)

All lines below the cursor in this window are erased. Also, the current line to the right of the cursor, inclusive, is erased.

Note: `clrtoobot()` is a macro.

clrtoeol()

wclrtoeol(win)

The current line to the right of the cursor, inclusive, is erased.

Note: `clrtoeol()` is a macro.

delay_output(ms)

Insert a *ms* millisecond pause in the output. It is not recommended that this routine be used extensively, because padding characters are used rather than a processor pause.

delch()

wdelch(win)

mvdelch(y, x)

mvwdelch(win, y, x)

The character under the cursor in the window is deleted. All characters to the right on the same line are moved to the left one position and the last character on the line is filled with a blank. The cursor position does not change (after moving

to (y, x) , if specified). (This does not imply use of the hardware “delete-character” feature.)

Note: `delch()`, `mvdclch()`, and `mvwdclch()` are macros.

deleteln()

wdeleteln (*win*)

The line under the cursor in the window is deleted. All lines below the current line are moved up one line. The bottom line of the window is cleared. The cursor position does not change. (This does not imply use of the hardware “delete-line” feature.)

Note: `deleteln()` is a macro.

getyx (*win, y, x*)

The cursor position of the window is placed in the two integer variables y and x . This is implemented as a macro, so no ‘&’ is necessary before the variables.

getbegyx (*win, y, x*)

getmaxyx (*win, y, x*)

Like `getyx()`, these routines store the current beginning coordinates and size of the specified window.

Note: `getbegyx()` and `getmaxyx()` are macros.

insch (*ch*)

winsch (*win, ch*)

mvwinsch (*win, y, x, ch*)

mvinsch (*y, x, ch*)

The character *ch* is inserted before the character under the cursor. All characters to the right are moved one SPACE to the right, possibly losing the rightmost character of the line. The cursor position does not change (after moving to (y, x) , if specified). (This does not imply use of the hardware “insert-character” feature.)

Note: *ch* is actually of type `chtype`, not a character.

Note: `insch()`, `mvinsch()`, and `mvwinsch()` are macros.

insertln()

winsertln (*win*)

A blank line is inserted above the current line and the bottom line is lost. (This does not imply use of the hardware “insert-line” feature.)

Note: `insertln()` is a macro.

move (*y, x*)

wmove (*win, y, x*)

The cursor associated with the window is moved to line (row) y , column x . This does not move the physical cursor of the terminal until `refresh()` is called. The position specified is relative to the upper left corner of the window, which is $(0, 0)$.

Note: `move()` is a macro.

overlay (*srcwin, dstwin*)

overwrite (*srcwin, dstwin*)

These routines overlay *srcwin* on top of *dstwin*; that is, all text in *srcwin* is copied into *dstwin*. *srcwin* and *dstwin* need not be the same size; only text where the two windows overlap is copied. The difference is that `overlay()` is non-destructive (blanks are not copied), while `overwrite()` is destructive.

copywin (*srcwin, dstwin, sminrow, smincol, dminrow, dmincol, dmaxrow, dmaxcol, overlay*)

This routine provides a finer grain of control over the `overlay()` and `overwrite()` routines. Like in the `prefresh()` routine, a rectangle is specified in the destination window, $(dminrow, dmincol)$ and $(dmaxrow, dmaxcol)$, and the upper-left-corner coordinates of the source window, $(sminrow, smincol)$. If the argument *overlay* is true, then copying is non-destructive, as in `overlay()`.

printw (*fmt* [, *arg* ...])

wprintw (*win, fmt* [, *arg* ...])

mvprintw (*y, x, fmt* [, *arg* ...])

mvwprintw (*win, y, x, fmt* [, *arg* ...])

These routines are analogous to `printf(3S)`. The string that would be output by `printf(3S)` is instead output using `waddstr()` on the given window.

vwprintw (*win, fmt, varglist*)

This routine corresponds to `vprintf(3V)`. It performs a `wprintw()` using a variable argument list. The third argument is a `va_list`, a pointer to a list of arguments, as defined in `<varargs.h>`. See the `vprintf(3V)` and `varargs(3)` manual pages for a detailed description on how to use variable argument lists.

scroll (*win*)

The window is scrolled up one line. This involves moving the lines in the window data structure. As an optimization, if the window is `stdscr` and the scrolling region is the entire window, the physical screen will be scrolled at the same time.

touchwin (*win*)

touchline (*win, start, count*)

Throw away all optimization information about which parts of the window have been touched, by pretending that the entire window has been drawn on. This is sometimes necessary when using overlapping windows, since a change to one window will affect the other window, but the records of which lines have been changed in the other window will not reflect the change. `touchline()` only pretends that `count` lines have been changed, beginning with line `start`.

Input

getch()

wgetch (*win*)

mvgetch (*y, x*)

mvwgetch (*win, y, x*)

A character is read from the terminal associated with the window. In `NODELAY` mode, if there is no input waiting, the value `ERR` is returned. In `DELAY` mode, the program will hang until the system passes text through to the program. Depending on the setting of `cbreak()`, this will be after one character (`CBREAK` mode), or after the first newline (`NOCBREAK` mode). In `HALF-DELAY` mode, the program will hang until a character is typed or the specified timeout has been reached. Unless `noecho()` has been set, the character will also be echoed into the designated window. No `refresh()` will occur between the `move()` and the `getch()` done within the routines `mvgetch()` and `mvwgetch()`.

When using `getch()`, `wgetch()`, `mvgetch()`, or `mvwgetch()`, do not set both `NOCBREAK` mode (`nocbreak()`) and `ECHO` mode (`echo()`) at the same time. Depending on the state of the terminal driver when each character is typed, the program may produce undesirable results.

If `keypad(win, TRUE)` has been called, and a function key is pressed, the token for that function key will be returned instead of the raw characters. (See `keypad()` under **Input Options Setting**.) Possible function keys are defined in `<curses.h>` with integers beginning with `0401`, whose names begin with `KEY_`. If a character is received that could be the beginning of a function key (such as escape), `curses` will set a timer. If the remainder of the sequence is not received within the designated time, the character will be passed through, otherwise the function key value will be returned. For this reason, on many terminals, there will be a delay after a user presses the escape key before the escape is returned to the program. (Use by a programmer of the escape key for a single character routine is discouraged. Also see `notimeout()` below.)

Note: `getch()`, `mvgetch()`, and `mvwgetch()` are macros.

getstr (*str*)

wgetstr (*win, str*)

mvgetstr (*y, x, str*)

mvwgetstr (*win, y, x, str*)

A series of calls to `getch()` is made, until a newline, carriage return, or enter key

is received. The resulting value is placed in the area pointed at by the character pointer *str*. The user's erase and kill characters are interpreted. As in `mvgetch()`, no `refresh()` is done between the `move()` and `getstr()` within the routines `mvgetstr()` and `mvwgetstr()`.

Note: `getstr()`, `mvgetstr()`, and `mvwgetstr()` are macros.

flushinp() Throws away any typeahead that has been typed by the user and has not yet been read by the program.

ungetch (c) Place *c* back onto the input queue to be returned by the next call to `wgetch()`.

inch()

winch (win)

mvinch (y, x)

mvwinch (win, y, x)

The character, of type `chtype`, at the current position in the named window is returned. If any attributes are set for that position, their values will be OR'ed into the value returned. The predefined constants `A_CHARTEXT` and `A_ATTRIBUTES`, defined in `< curses.h >`, can be used with the C logical AND (`&`) operator to extract the character or attributes alone.

Note: `inch()`, `winch()`, `mvinch()`, and `mvwinch()` are macros.

scanw (fmt [, arg ...])

wscanw (win, fmt [, arg ...])

mvscanw (y, x, fmt [, arg ...])

mvwscanw (win, y, x, fmt [, arg ...])

These routines correspond to `scanf(3V)`, as do their arguments and return values. `wgetstr()` is called on the window, and the resulting line is used as input for the scan.

vwscanw (win, fmt, ap) This routine is similar to `vwprintw()` above in that performs a `wscanw()` using a variable argument list. The third argument is a `va_list`, a pointer to a list of arguments, as defined in `< varargs.h >`. See the `vprintf(3V)` and `varargs(3)` manual pages for a detailed description on how to use variable argument lists.

Output Options Setting

These routines set options within `curses` that deal with output. All options are initially `FALSE`, unless otherwise stated. It is not necessary to turn these options off before calling `endwin()`.

clearok (win, bf) If enabled (*bf* is `TRUE`), the next call to `wrefresh()` with this window will clear the screen completely and redraw the entire screen from scratch. This is useful when the contents of the screen are uncertain, or in some cases for a more pleasing visual effect.

idlok (win, bf) If enabled (*bf* is `TRUE`), `curses` will consider using the hardware "insert/delete-line" feature of terminals so equipped. If disabled (*bf* is `FALSE`), `curses` will very seldom use this feature. (The "insert/delete-character" feature is always considered.) This option should be enabled only if your application needs "insert/delete-line", for example, for a screen editor. It is disabled by default because "insert/delete-line" tends to be visually annoying when used in applications where it is not really needed. If "insert/delete-line" cannot be used, `curses` will redraw the changed portions of all lines.

leaveok (win, bf) Normally, the hardware cursor is left at the location of the window cursor being refreshed. This option allows the cursor to be left wherever the update happens to leave it. It is useful for applications where the cursor is not used, since it reduces the need for cursor motions. If possible, the cursor is made invisible when this option is enabled.

setscrreg (*top, bot*)

wsetscrreg (*win, top, bot*)

These routines allow the user to set a software scrolling region in a window. *top* and *bot* are the line numbers of the top and bottom margin of the scrolling region. (Line 0 is the top line of the window.) If this option and **scrollok()** are enabled, an attempt to move off the bottom margin line will cause all lines in the scrolling region to scroll up one line. (Note: this has nothing to do with use of a physical scrolling region capability in the terminal, like that in the DEC VT100. Only the text of the window is scrolled; if **idlok()** is enabled and the terminal has either a scrolling region or "insert/delete-line" capability, they will probably be used by the output routines.)

Note: **setscrreg()** and **wsetscrreg()** are macros.

scrollok (*win, bf*)

This option controls what happens when the cursor of a window is moved off the edge of the window or scrolling region, either from a newline on the bottom line, or typing the last character of the last line. If disabled (*bf* is FALSE), the cursor is left on the bottom line at the location where the offending character was entered. If enabled (*bf* is TRUE), **wrefresh()** is called on the window, and then the physical terminal and window are scrolled up one line. (Note: in order to get the physical scrolling effect on the terminal, it is also necessary to call **idlok()**.)

nl()

nonl()

These routines control whether NEWLINE is translated into RETURN and LINEFEED on output, and whether RETURN is translated into NEWLINE on input. Initially, the translations do occur. By disabling these translations using **nonl()**, **curses** is able to make better use of the linefeed capability, resulting in faster cursor motion.

Input Options Setting

These routines set options within **curses** that deal with input. The options involve using **ioctl(2)** and therefore interact with **curses** routines. It is not necessary to turn these options off before calling **endwin()**.

For more information on these options, refer to *Programming Utilities and Libraries*.

cbreak()

nocbreak()

These two routines put the terminal into and out of CBREAK mode, respectively. In CBREAK mode, characters typed by the user are immediately available to the program and erase/kill character processing is not performed. When in NOCBREAK mode, the tty driver will buffer characters typed until a NEWLINE or RETURN is typed. Interrupt and flow-control characters are unaffected by this mode (see **termio(4)**). Initially the terminal may or may not be in CBREAK mode, as it is inherited, therefore, a program should call **cbreak()** or **nocbreak()** explicitly. Most interactive programs using **curses** will set CBREAK mode.

Note: **cbreak()** overrides **raw()**. See **getch()** under **Input** for a discussion of how these routines interact with **echo()** and **noecho()**.

echo()

noecho()

These routines control whether characters typed by the user are echoed by **getch()** as they are typed. Echoing by the tty driver is always disabled, but initially **getch()** is in ECHO mode, so characters typed are echoed. Authors of most interactive programs prefer to do their own echoing in a controlled area of the screen, or not to echo at all, so they disable echoing by calling **noecho()**. See **getch()** under **Input** for a discussion of how these routines interact with **cbreak()** and **nocbreak()**.

halfdelay (*tenths*)

Half-delay mode is similar to CBREAK mode in that characters typed by the user are immediately available to the program. However, after blocking for *tenths* tenths of seconds, ERR will be returned if nothing has been typed. *tenths* must be a number between 1 and 255. Use **nocbreak()** to leave half-delay mode.

- intrflush** (*win, bf*) If this option is enabled, when an interrupt key is pressed on the keyboard (interrupt, break, quit) all output in the tty driver queue will be flushed, giving the effect of faster response to the interrupt, but causing curses to have the wrong idea of what is on the screen. Disabling the option prevents the flush. The default for the option is inherited from the tty driver settings. The window argument is ignored.
- keypad** (*win, bf*) This option enables the keypad of the user's terminal. If enabled, the user can press a function key (such as an arrow key) and **wgetch()** will return a single value representing the function key, as in **KEY_LEFT**. If disabled, curses will not treat function keys specially and the program would have to interpret the escape sequences itself. If the keypad in the terminal can be turned on (made to transmit) and off (made to work locally), turning on this option will cause the terminal keypad to be turned on when **wgetch()** is called.
- meta** (*win, bf*) If enabled, characters returned by **wgetch()** are transmitted with all 8 bits, instead of with the highest bit stripped. In order for **meta()** to work correctly, the **km** (**has_meta_key**) capability has to be specified in the terminal's **terminfo(5V)** entry.
- nodelay** (*win, bf*) This option causes **wgetch()** to be a non-blocking call. If no input is ready, **wgetch()** will return **ERR**. If disabled, **wgetch()** will hang until a key is pressed.
- notimeout** (*win, bf*) While interpreting an input escape sequence, **wgetch()** will set a timer while waiting for the next character. If **notimeout** (*win, TRUE*) is called, then **wgetch()** will not set a timer. The purpose of the timeout is to differentiate between sequences received from a function key and those typed by a user.
- raw()**
noraw() The terminal is placed into or out of RAW mode. RAW mode is similar to CBREAK mode, in that characters typed are immediately passed through to the user program. The differences are that in RAW mode, the interrupt, quit, suspend, and flow control characters are passed through uninterpreted, instead of generating a signal. RAW mode also causes 8-bit input and output. The behavior of the BREAK key depends on other bits in the terminal driver that are not set by curses.
- typeahead** (*fildef*) curses does "line-breakout optimization" by looking for typeahead periodically while updating the screen. If input is found, and it is coming from a tty, the current update will be postponed until **refresh()** or **doupdate()** is called again. This allows faster response to commands typed in advance. Normally, the file descriptor for the input FILE pointer passed to **newterm()**, or **stdin** in the case that **initscr()** was used, will be used to do this typeahead checking. The **typeahead()** routine specifies that the file descriptor *fildef* is to be used to check for typeahead instead. If *fildef* is **-1**, then no typeahead checking will be done.
- Note: *fildef* is a file descriptor, not a **<stdio.h>** FILE pointer.
- Environment Queries**
- baudrate()** Returns the output speed of the terminal. The number returned is in bits per second, for example, 9600, and is an integer.
- char erasechar()** The user's current erase character is returned.
- has_ic()** True if the terminal has insert- and delete-character capabilities.
- has_il()** True if the terminal has insert- and delete-line capabilities, or can simulate them using scrolling regions. This might be used to check to see if it would be appropriate to turn on physical scrolling using **scrollok()**.
- char killchar()** The user's current line-kill character is returned.
- char *longname()** This routine returns a pointer to a static area containing a verbose description of

the current terminal. The maximum length of a verbose description is 128 characters. It is defined only after the call to `initscr()` or `newterm()`. The area is overwritten by each call to `newterm()` and is not restored by `set_term()`, so the value should be saved between calls to `newterm()` if `longname()` is going to be used with multiple terminals.

Soft Labels

If desired, `curses` will manipulate the set of soft function-key labels that exist on many terminals. For those terminals that do not have soft labels, if you want to simulate them, `curses` will take over the bottom line of `stdscr`, reducing the size of `stdscr` and the variable `LINES`. `curses` standardizes on 8 labels of 8 characters each.

slk_init(*labfmt*) In order to use soft labels, this routine must be called before `initscr()` or `newterm()` is called. If `initscr()` winds up using a line from `stdscr` to emulate the soft labels, then *labfmt* determines how the labels are arranged on the screen. Setting *labfmt* to 0 indicates that the labels are to be arranged in a 3-2-3 arrangement; 1 asks for a 4-4 arrangement.

slk_set(*labnum*, *label*, *labfmt*) *labnum* is the label number, from 1 to 8. *label* is the string to be put on the label, up to 8 characters in length. A NULL string or a NULL pointer will put up a blank label. *labfmt* is one of 0, 1 or 2, to indicate whether the label is to be left-justified, centered, or right-justified within the label.

slk_refresh()
slk_noutrefresh() These routines correspond to the routines `wrefresh()` and `wnoutrefresh()`. Most applications would use `slk_noutrefresh()` because a `wrefresh()` will most likely soon follow.

char *slk_label(*labnum*) The current label for label number *labnum*, with leading and trailing blanks stripped, is returned.

slk_clear() The soft labels are cleared from the screen.

slk_restore() The soft labels are restored to the screen after a `slk_clear()`.

slk_touch() All of the soft labels are forced to be output the next time a `slk_noutrefresh()` is performed.

Low-Level curses Access

The following routines give low-level access to various `curses` functionality. These routines typically would be used inside of library routines.

def_prog_mode()
def_shell_mode() Save the current terminal modes as the "program" (in `curses`) or "shell" (not in `curses`) state for use by the `reset_prog_mode()` and `reset_shell_mode()` routines. This is done automatically by `initscr()`.

reset_prog_mode()
reset_shell_mode() Restore the terminal to "program" (in `curses`) or "shell" (out of `curses`) state. These are done automatically by `endwin()` and `doupdate()` after an `endwin()`, so they normally would not be called.

resetty()
savetty() These routines save and restore the state of the terminal modes. `savetty()` saves the current state of the terminal in a buffer and `resetty()` restores the state to what it was at the last call to `savetty()`.

getsyx(*y*, *x*) The current coordinates of the virtual screen cursor are returned in *y* and *x*. Like `getyx()`, the variables *y* and *x* do not take an `&` before them. If `leaveok()` is

currently TRUE, then `-1, -1` will be returned. If lines may have been removed from the top of the screen using `ripoffline()` and the values are to be used beyond just passing them on to `setsyx()`, the value `y+stdscr->_yoffset` should be used for those other uses.

Note: `getsyx()` is a macro.

- setsyx** (*y, x*) The virtual screen cursor is set to *y, x*. If *y* and *x* are both `-1`, then `leaveok()` will be set. The two routines `getsyx()` and `setsyx()` are designed to be used by a library routine that manipulates curses windows but does not want to mess up the current position of the program's cursor. The library routine would call `getsyx()` at the beginning, do its manipulation of its own windows, do a `wnoutrefresh()` on its windows, call `setsyx()`, and then call `doupdate()`.
- ripoffline** (*line, init*) This routine provides access to the same facility that `slk_init()` uses to reduce the size of the screen. `ripoffline()` must be called before `initscr()` or `newterm()` is called. If *line* is positive, a line will be removed from the top of `stdscr`; if negative, a line will be removed from the bottom. When this is done inside `initscr()`, the routine *init* is called with two arguments: a window pointer to the 1-line window that has been allocated and an integer with the number of columns in the window. Inside this initialization routine, the integer variables `LINES` and `COLS` (defined in `<curses.h>`) are not guaranteed to be accurate and `wrefresh()` or `doupdate()` must not be called. It is allowable to call `wnoutrefresh()` during the initialization routine.
- `ripoffline()` can be called up to five times before calling `initscr()` or `newterm()`.
- scr_dump** (*filename*) The current contents of the virtual screen are written to the file *filename*.
- scr_restore** (*filename*) The virtual screen is set to the contents of *filename*, which must have been written using `scr_dump()`. The next call to `doupdate()` will restore the screen to what it looked like in the dump file.
- scr_init** (*filename*) The contents of *filename* are read in and used to initialize the curses data structures about what the terminal currently has on its screen. If the data is determined to be valid, curses will base its next update of the screen on this information rather than clearing the screen and starting from scratch. `scr_init()` would be used after `initscr()` or a `system(3)` call to share the screen with another process that has done a `scr_dump()` after its `endwin()` call. The data will be declared invalid if the time-stamp of the tty is old or the `terminfo(5V)` capability `nrrmc` is true.
- curs_set** (*visibility*) The cursor is set to invisible, normal, or very visible for *visibility* equal to `0`, `1` or `2`.
- draino** (*ms*) Wait until the output has drained enough that it will only take *ms* more milliseconds to drain completely.
- garbagedlines** (*win, begline, numlines*) This routine indicates to curses that a screen line is garbaged and should be thrown away before having anything written over the top of it. It could be used for programs such as editors that want a command to redraw just a single line. Such a command could be used in cases where there is a noisy communications line and redrawing the entire screen would be subject to even more communication noise. Just redrawing the single line gives some semblance of hope that it would show up unblemished. The current location of the window is used to determine which lines are to be redrawn.
- napms** (*ms*) Sleep for *ms* milliseconds.

Terminfo-Level Manipulations

These low-level routines must be called by programs that need to deal directly with the `terminfo(5V)` database to handle certain terminal capabilities, such as programming function keys. For all other functionality, `curses` routines are more suitable and their use is recommended.

Initially, `setupterm()` should be called. (Note: `setupterm()` is automatically called by `initscr()` and `newterm()`.) This will define the set of terminal-dependent variables defined in the `terminfo(5V)` database. The `terminfo(5V)` variables `lines` and `columns` (see `terminfo(5V)`) are initialized by `setupterm()` as follows: if the environment variables `LINES` and `COLUMNS` exist, their values are used. If the above environment variables do not exist, and the window sizes in rows and columns as returned by the `TIOCGWINSZ` ioctl are non-zero, those sizes are used. Otherwise, the values for `lines` and `columns` specified in the `terminfo(5V)` database are used.

The header files `<curses.h>` and `<term.h>` should be included, in this order, to get the definitions for these strings, numbers, and flags. Parameterized strings should be passed through `tparm()` to instantiate them. All `terminfo(5V)` strings (including the output of `tparm()`) should be printed with `tputs()` or `putp()`. Before exiting, `reset_shell_mode()` should be called to restore the tty modes. Programs that use cursor addressing should output `enter_ca_mode` upon startup and should output `exit_ca_mode` before exiting (see `terminfo(5V)`). (Programs desiring shell escapes should call `reset_shell_mode()` and output `exit_ca_mode` before the shell is called and should output `enter_ca_mode` and call `reset_prog_mode()` after returning from the shell. Note: this is different from the `curses` routines (see `endwin()`)

`setupterm (term, fildes, errret)`

Reads in the `terminfo(5V)` database, initializing the `terminfo(5V)` structures, but does not set up the output virtualization structures used by `curses`. The terminal type is in the character string `term`; if `term` is NULL, the environment variable `TERM` will be used. All output is to the file descriptor `fildes`. If `errret` is not NULL, then `setupterm()` will return OK or ERR and store a status value in the integer pointed to by `errret`. A status of 1 in `errret` is normal, 0 means that the terminal could not be found, and -1 means that the `terminfo(5V)` database could not be found. If `errret` is NULL, `setupterm()` will print an error message upon finding an error and exit. Thus, the simplest call is `'setupterm ((char *)0, 1, (int *)0)'`, which uses all the defaults.

The `terminfo(5V)` boolean, numeric and string variables are stored in a structure of type `TERMINAL`. After `setupterm()` returns successfully, the variable `cur_term` (of type `TERMINAL *`) is initialized with all of the information that the `terminfo(5V)` boolean, numeric and string variables refer to. The pointer may be saved before calling `setupterm()` again. Further calls to `setupterm()` will allocate new space rather than reuse the space pointed to by `cur_term`.

`set_curterm (nterm)` `nterm` is of type `TERMINAL *`. `set_curterm()` sets the variable `cur_term` to `nterm`, and makes all of the `terminfo(5V)` boolean, numeric and string variables use the values from `nterm`.

`del_curterm (oterm)` `oterm` is of type `TERMINAL *`. `del_curterm()` frees the space pointed to by `oterm` and makes it available for further use. If `oterm` is the same as `cur_term`, then references to any of the `terminfo(5V)` boolean, numeric and string variables thereafter may refer to invalid memory locations until another `setupterm()` has been called.

`restartterm (term, fildes, errret)`

Like `setupterm()` after a memory restore.

`char *tparm (str, p1, p2, ..., pg)`

Instantiate the string `str` with parms `p1`. A pointer is returned to the result of `str` with the parameters applied.

`tputs (str, count, putc)` Apply padding to the string `str` and output it. `str` must be a `terminfo(5V)` string

variable or the return value from `tparm()`, `tgetstr()`, `tigetstr()` or `tgoto()`. `count` is the number of lines affected, or 1 if not applicable. `putchar()` is a `putc(3S)`-like routine to which the characters are passed, one at a time.

- putp** (*str*) A routine that calls `tputs()` (*str*, 1, `putchar()`).
- vidputs** (*attrs*, *putc*) Output a string that puts the terminal in the video attribute mode *attrs*, which is any combination of the attributes listed below. The characters are passed to the `putc(3S)`-like routine `putchar()`.
- vidattr** (*attrs*) Like `vidputs()`, except that it outputs through `putchar(3S)`.
- mvcur** (*oldrow*, *oldcol*, *newrow*, *newcol*)
Low-level cursor motion.

The following routines return the value of the capability corresponding to the `terminfo(5V)` *capname* passed to them, such as `xenl`.

- tigetflag** (*capname*) The value `-1` is returned if *capname* is not a boolean capability.
- tigetnum** (*capname*) The value `-2` is returned if *capname* is not a numeric capability.
- tigetstr** (*capname*) The value (`char *`) `-1` is returned if *capname* is not a string capability.

`char *boolnames[], *boolcodes[], *boolfnames[]`
`char *numnames[], *numcodes[], *numfnames[]`
`char *strnames[], *strcodes[], *strfnames[]`

These null-terminated arrays contain the *capnames*, the `termcap(5)` codes, and the full C names, for each of the `terminfo(5V)` variables.

Termcap Emulation

These routines are included as a conversion aid for programs that use the `termcap(3X)` library. Their parameters are the same and the routines are emulated using the `terminfo(5V)` database.

- tgetent** (*bp*, *name*) Look up `termcap` entry for *name*. The emulation ignores the buffer pointer *bp*.
- tgetflag** (*codename*) Get the boolean entry for *codename*.
- tgetnum** (*codes*) Get numeric entry for *codename*.
- char *tgetstr** (*codename*, *area*)
Return the string entry for *codename*. If *area* is not `NULL`, then also store it in the buffer pointed to by *area* and advance *area*. `tputs()` should be used to output the returned string.
- char *tgoto** (*cap*, *col*, *row*)
Instantiate the parameters into the given capability. The output from this routine is to be passed to `tputs()`.
- tputs** (*str*, *affcnt*, *putc*) See `tputs()` above, under **Terminfo-Level Manipulations**.

Miscellaneous

- unctrl** (*c*) This macro expands to a character string which is a printable representation of the character *c*. Control characters are displayed in the `^X` notation. Printing characters are displayed as is.

`unctrl()` is a macro, defined in `<unctrl.h>`, which is automatically included by `<curses.h>`.
- char *keyname** (*c*) A character string corresponding to the key *c* is returned.
- filter()** This routine is one of the few that is to be called before `initscr()` or `newterm()` is called. It arranges things so that `curses` thinks that there is a 1-line screen. `curses` will not use any terminal capabilities that assume that they know what line on the screen the cursor is on.

Use of curscr

The special window `curscr` can be used in only a few routines. If the window argument to `clearok()` is `curscr`, the next call to `wrefresh()` with any window will cause the screen to be cleared and repainted from scratch. If the window argument to `wrefresh()` is `curscr`, the screen is immediately cleared and repainted from scratch. (This is how most programs would implement a "repaint-screen" routine.) The source window argument to `overlay()`, `overwrite()`, and `copywin` may be `curscr`, in which case the current contents of the virtual terminal screen will be accessed.

Obsolete Calls

Various routines are provided to maintain compatibility in programs written for older versions of the curses library. These routines are all emulated as indicated below.

- `crmode()` Replaced by `cbreak()`.
- `fixterm()` Replaced by `reset_prog_mode()`.
- `gettmode()` A no-op.
- `nocrmode()` Replaced by `nocbreak()`.
- `resetterm()` Replaced by `reset_shell_mode()`.
- `saveterm()` Replaced by `def_prog_mode()`.
- `setterm()` Replaced by `setupterm()`.

ATTRIBUTES

The following video attributes, defined in `<curses.h>`, can be passed to the routines `attron()`, `attroff()`, and `attrset()`, or OR'ed with the characters passed to `addch()`.

- `A_STANDOUT` Terminal's best highlighting mode
- `A_UNDERLINE` Underlining
- `A_REVERSE` Reverse video
- `A_BLINK` Blinking
- `A_DIM` Half bright
- `A_BOLD` Extra bright or bold
- `A_ALTCHARSET` Alternate character set

- `A_CHARTEXT` Bit-mask to extract character (described under `winch`)
- `A_ATTRIBUTES` Bit-mask to extract attributes (described under `winch`)
- `A_NORMAL` Bit mask to reset all attributes off
(for example: `'attrset (A_NORMAL)'`)

FUNCTION-KEYS

The following function keys, defined in `<curses.h>`, might be returned by `getch()` if `keypad()` has been enabled. Note: not all of these may be supported on a particular terminal if the terminal does not transmit a unique code when the key is pressed or the definition for the key is not present in the `terminfo(5V)` database.

<i>Name</i>	<i>Value</i>	<i>Key name</i>
<code>KEY_BREAK</code>	0401	break key (unreliable)
<code>KEY_DOWN</code>	0402	The four arrow keys ...
<code>KEY_UP</code>	0403	
<code>KEY_LEFT</code>	0404	
<code>KEY_RIGHT</code>	0405	...
<code>KEY_HOME</code>	0406	Home key (upward+left arrow)
<code>KEY_BACKSPACE</code>	0407	backspace (unreliable)
<code>KEY_F0</code>	0410	Function keys. Space for 64 keys is reserved.
<code>KEY_F(n)</code>	<code>(KEY_F0+(n))</code>	Formula for f_n
<code>KEY_DL</code>	0510	Delete line
<code>KEY_IL</code>	0511	Insert line
<code>KEY_DC</code>	0512	Delete character

KEY_IC	0513	Insert char or enter insert mode
KEY_EIC	0514	Exit insert char mode
KEY_CLEAR	0515	Clear screen
KEY_EOS	0516	Clear to end of screen
KEY_EOL	0517	Clear to end of line
KEY_SF	0520	Scroll 1 line forward
KEY_SR	0521	Scroll 1 line backwards (reverse)
KEY_NPAGE	0522	Next page
KEY_PPAGE	0523	Previous page
KEY_STAB	0524	Set TAB
KEY_CTAB	0525	Clear TAB
KEY_CATAB	0526	Clear all TAB characters
KEY_ENTER	0527	Enter or send
KEY_SRESET	0530	soft (partial) reset
KEY_RESET	0531	reset or hard reset
KEY_PRINT	0532	print or copy
KEY_LL	0533	home down or bottom (lower left) keypad is arranged like this: A1 up A3 left B2 right C1 down C3
KEY_A1	0534	Upper left of keypad
KEY_A3	0535	Upper right of keypad
KEY_B2	0536	Center of keypad
KEY_C1	0537	Lower left of keypad
KEY_C3	0540	Lower right of keypad
KEY_BTAB	0541	Back TAB key
KEY_BEG	0542	beg(inning) key
KEY_CANCEL	0543	cancel key
KEY_CLOSE	0544	close key
KEY_COMMAND	0545	cmd (command) key
KEY_COPY	0546	copy key
KEY_CREATE	0547	create key
KEY_END	0550	end key
KEY_EXIT	0551	exit key
KEY_FIND	0552	find key
KEY_HELP	0553	help key
KEY_MARK	0554	mark key
KEY_MESSAGE	0555	message key
KEY_MOVE	0556	move key
KEY_NEXT	0557	next object key
KEY_OPEN	0560	open key
KEY_OPTIONS	0561	options key
KEY_PREVIOUS	0562	previous object key
KEY_REDO	0563	redo key
KEY_REFERENCE	0564	ref(erence) key
KEY_REFRESH	0565	refresh key
KEY_REPLACE	0566	replace key
KEY_RESTART	0567	restart key
KEY_RESUME	0570	resume key
KEY_SAVE	0571	save key
KEY_SBEG	0572	shifted beginning key
KEY_SCANCEL	0573	shifted cancel key

KEY_SCOMMAND	0574	shifted command key
KEY_SCOPY	0575	shifted copy key
KEY_SCREATE	0576	shifted create key
KEY_SDC	0577	shifted delete char key
KEY_SDL	0600	shifted delete line key
KEY_SELECT	0601	select key
KEY_SEND	0602	shifted end key
KEY_SEOL	0603	shifted clear line key
KEY_SEXIT	0604	shifted exit key
KEY_SFIND	0605	shifted find key
KEY_SHELP	0606	shifted help key
KEY_SHOME	0607	shifted home key
KEY_SIC	0610	shifted input key
KEY_SLEFT	0611	shifted left arrow key
KEY_SMESSAGE	0612	shifted message key
KEY_SMOVE	0613	shifted move key
KEY_SNEXT	0614	shifted next key
KEY_SOPTIONS	0615	shifted options key
KEY_SPREVIOUS	0616	shifted prev key
KEY_SPRINT	0617	shifted print key
KEY_SREDO	0620	shifted redo key
KEY_SREPLACE	0621	shifted replace key
KEY_SRIGHT	0622	shifted right arrow
KEY_SRSUME	0623	shifted resume key
KEY_SSAVE	0624	shifted save key
KEY_SSUSPEND	0625	shifted suspend key
KEY_SUNDO	0626	shifted undo key
KEY_SUSPEND	0627	suspend key
KEY_UNDO	0630	undo key

LINE GRAPHICS

The following variables may be used to add line-drawing characters to the screen with **waddce**. When defined for the terminal, the variable will have the **A_ALTCHARSET** bit turned on. Otherwise, the default character listed below will be stored in the variable. The names were chosen to be consistent with the DEC VT100 nomenclature.

<i>Name</i>	<i>Default</i>	<i>Glyph Description</i>
ACS_ULCORNER	+	upper left corner
ACS_LLCORNER	+	lower left corner
ACS_URCORNER	+	upper right corner
ACS_LRCORNER	+	lower right corner
ACS_RTEE	+	right tee (┘)
ACS_LTEE	+	left tee (└)
ACS_BTEE	+	bottom tee (└┘)
ACS_TTEE	+	top tee (┐)
ACS_HLINE	-	horizontal line
ACS_VLINE		vertical line
ACS_PLUS	+	plus
ACS_S1	-	scan line 1
ACS_S9	-	scan line 9
ACS_DIAMOND	+	diamond
ACS_CKBOARD	:	checker board (stipple)
ACS_DEGREE	'	degree symbol
ACS_PLMINUS	#	plus/minus

ACS_BULLET	o	bullet
ACS_LARROW	<	arrow pointing left
ACS_RARROW	>	arrow pointing right
ACS_DARROW	v	arrow pointing down
ACS_UARROW	^	arrow pointing up
ACS_BOARD	#	board of squares
ACS_LANTERN	#	lantern symbol
ACS_BLOCK	#	solid square block

RETURN VALUES

All routines return the integer OK upon successful completion and the integer ERR upon failure, unless otherwise noted in the preceding routine descriptions.

All macros return the value of their w version, except `setscrreg()`, `wsetscrreg()`, `getsyx()`, `getyx()`, `getbegy()`, `getmaxyx()`. For these macros, no useful value is returned.

Routines that return pointers always return *(type *)* NULL "on error.

FILES

`/usr/share/lib/terminfo`
`.login`
`.profile`

SEE ALSO

`cc(1V)`, `ld(1)`, `ioctl(2)`, `plot(3X)`, `printf(3S)`, `putc(3S)`, `scanf(3V)`, `stdio(3V)`, `system(3)`, `varargs(3)`, `vprintf(3V)`, `termio(4)`, `term(5V)`, `terminfo(5V)`

WARNINGS

The plotting library `plot(3X)` and the curses library `curses(3V)` both use the names `erase()` and `move()`. The `curses` versions are macros. If you need both libraries, put the `plot(3X)` code in a different source file than the `curses(3V)` code, and/or `#undef move` and `#undef erase` in the `plot(3X)` code.

Between the time a call to `initscr()` and `endwin()` has been issued, use only the routines in the `curses` library to generate output. Using system calls or the "standard I/O package" (see `stdio(3V)`) for output during that time can cause unpredictable results.

INT C

VT - 100

→ = 405	PF1 = 411	ENTER = 527
← = 405	PF2 = 412	— Key Pad —
↓ = 402	PF3 = 413	1 = 534 — = 155
↑ = 403	PF4 = 414	2 = 536) = 420
		3 = 535
		4 = 415
		5 = 416
		6 = 417
		7 = 421
		8 = 422
		9 = 410
		0 = 537

NAME

ferror, feof, clearerr, fileno – stream status inquiries

SYNOPSIS

#include <stdio.h>

ferror(stream)

FILE *stream;

feof(stream)

FILE *stream;

clearerr(stream)

FILE *stream;

fileno(stream)

FILE *stream;

DESCRIPTION

ferror() returns non-zero when an error has occurred reading from or writing to the named stream, otherwise zero. Unless cleared by **clearerr**, the error indication lasts until the stream is closed.

feof() returns non-zero when EOF has previously been detected reading the named input stream, otherwise zero. Unless cleared by **clearerr**, the EOF indication lasts until the stream is closed; however, operations which attempt to read from the stream will ignore the current state of the EOF indication and attempt to read from the file descriptor associated with the stream.

clearerr() resets the error indication and EOF indication to zero on the named stream.

fileno() returns the integer file descriptor associated with the stream; see **open(2V)**.

NOTE

All these functions are implemented as macros; they cannot be redeclared.

SEE ALSO

open(2V), fopen(3S)

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. If the open succeeds, `fopen()` returns a pointer to be used to identify the stream in subsequent operations.

filename points to a character string that contains the name of the file to be opened.

type is a character string having one of the following values:

<code>r</code>	open for reading
<code>w</code>	truncate or create for writing
<code>a</code>	append: open for writing at end of file, or create for writing
<code>r+</code>	open for update (reading and writing)
<code>w+</code>	truncate or create for update
<code>a+</code>	append; open or create for update at EOF

`freopen()` opens the file named by *filename* and associates the stream pointed to by *stream* with it. The *type* argument is used just as in `fopen`. The original stream is closed, regardless of whether the open ultimately succeeds. If the open succeeds, `freopen()` returns the original value of *stream*.

`freopen()` is typically used to attach the preopened streams associated with `stdin`, `stdout`, and `stderr` to other files.

`fdopen()` associates a stream with the file descriptor *fildes*. File descriptors are obtained from calls like `open`, `dup`, `creat`, or `pipe(2)`, which open files but do not return streams. Streams are necessary input for many of the Section 3S library routines. The *type* of the stream must agree with the mode of the open file.

When a file is opened for update, both input and output may be done on the resulting stream. However, output may not be directly followed by input without an intervening `fseek()` or `rewind`, and input may not be directly followed by output without an intervening `fseek`, `rewind`, or an input operation which encounters end-of-file.

When a file is opened for append (that is, when *type* is `a` or `a+`), it is impossible to overwrite information already in the file. `fseek()` may be used to reposition the file pointer to any position in the file, but when output is written to the file, the current file pointer is disregarded. All output is written at the end of the file and causes the file pointer to be repositioned at the end of the output. If two separate processes open the same file for append, each process may write freely to the file without fear of destroying output being written by the other. The output from the two processes will be intermixed in the file in the order in which it is written.

SEE ALSO

`open(2V)`, `pipe(2)`, `fclose(3S)`, `fopen(3S)`, `fseek(3S)`

DIAGNOSTICS

`fopen`, `freopen`, and `fdopen()` return a NULL pointer on failure.

BUGS

In order to support the same number of open files that the system does, **fopen()** must allocate additional memory for data structures using **calloc()** after 64 files have been opened. This confuses some programs which use their own memory allocators.

NAME

`getc`, `getchar`, `fgetc`, `getw` – get character or integer 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 (that is, byte) from the named input stream, as an integer. It also moves the file pointer, if defined, ahead one character in stream. `getchar()` is defined as `getc(stdin)`. `getc` and `getchar` are macros.

`fgetc()` behaves like `getc`, but is a function rather than a macro. `fgetc()` runs more slowly than `getc`, but it takes less space per invocation and its name can be passed as an argument to a function.

`getw()` returns the next C int (*word*) from the named input stream. `getw()` increments the associated file pointer, if defined, to point to the next word. The size of a word is the size of an integer and varies from machine to machine. `getw()` assumes no special alignment in the file.

SEE ALSO

`ferror(3S)`, `fopen(3S)`, `fread(3S)`, `gets(3S)`, `putc(3S)`, `scanf(3S)`, `ungetc(3S)`

DIAGNOSTICS

These functions return the integer constant EOF at EOF or upon an error. Because EOF is a valid integer, `ferror(3S)` should be used to detect `getw()` errors.

WARNING

If the integer value returned by `getc`, `getchar`, or `fgetc` is stored into a character variable and then compared against the integer constant EOF, the comparison may never succeed, because sign-extension of a character on widening to integer is machine-dependent.

BUGS

Because it is implemented as a macro, `getc()` treats a stream argument with side effects incorrectly. In particular, `getc(*f++)` does not work sensibly. `fgetc()` should be used instead.

Because of possible differences in word length and byte ordering, files written using `putw()` are machine-dependent, and may not be readable using `getw()` on a different processor.

NAME

getpass – read a password

SYNOPSIS

```
char *getpass(prompt)  
char *prompt;
```

DESCRIPTION

getpass() reads up to a NEWLINE or EOF from the file **/dev/tty**, after prompting with the NULL-terminated string *prompt* and disabling echoing. A pointer is returned to a NULL-terminated string of at most 8 characters. An interrupt will terminate input and send an interrupt signal to the calling program before returning. If **/dev/tty** cannot be opened, a NULL pointer is returned; the standard input is not read.

FILES

/dev/tty

SEE ALSO

crypt(3), **getpass(3)**

WARNING

The above routine uses **<stdio.h>**, which increases the size of programs not otherwise using standard I/O, more than might be expected.

BUGS

The return value points to static data whose content is overwritten by each call.

NAME

getpwent, getpwuid, getpwnam, setpwent, endpwent, setpwfile, fgetpwent – 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()

setpwfile(name)
char *name;

struct passwd *fgetpwent(f)
FILE *f;
```

DESCRIPTION

getpwent, getpwuid() and getpwnam() each return a pointer to an object with the following structure containing the broken-out fields of a line in the password file. Each line in the file contains a “passwd” structure, declared in the <pwd.h> header file:

```
struct passwd { /* see getpwent(3) */
    char    *pw_name;
    char    *pw_passwd;
    int     pw_uid;
    int     pw_gid;
    char    *pw_age;
    char    *pw_comment;
    char    *pw_gecos;
    char    *pw_dir;
    char    *pw_shell;
};

struct passwd *getpwent(), *getpwuid(), *getpwnam();
```

This structure is declared in <pwd.h> so it is not necessary to redeclare it.

The field `pw_comment` is unused; the others have meanings described in `passwd(5)`. When first called, `getpwent()` returns a pointer to the first `passwd` structure in the file; thereafter, it returns a pointer to the next `passwd` structure in the file; so successive calls can be used to search the entire file. `getpwuid()` searches from the beginning of the file until a numerical user ID matching `uid` is found and returns a pointer to the particular structure in which it was found. `getpwnam()` searches from the beginning of the file until a login name matching `name` is found, and returns a pointer to the particular structure in which it was found. If an end-of-file or an error is encountered on reading, these functions return a NULL pointer.

A call to `getpwent()` has the effect of rewinding the password file to allow repeated searches. `endpwent()` may be called to close the password file when processing is complete.

`setpwfile()` changes the default password file to `name` thus allowing alternate password files to be used. Note: it does *not* close the previous file. If this is desired, `endpwent()` should be called prior to it.

`fgetpwent()` returns a pointer to the next `passwd` structure in the stream `f`, which matches the format of the password file `/etc/passwd`.

The field `pw_age` is used to hold a value for “password aging” on some systems; “password aging” is not supported on Sun systems. As such, it is effectively not used.

FILES

`/etc/passwd`
`/var/yp/domainname/passwd.byname`
`/var/yp/domainname/passwd.byuid`

SEE ALSO

`getgrent(3)`, `getlogin(3)`, `getpwent(3)`, `passwd(5)`, `ypserv(8)`

DIAGNOSTICS

A NULL pointer is returned on end-of-file or error.

WARNING

The above routines use the standard I/O library, which causes them to increase the size of programs, not otherwise using standard I/O, more than might be expected.

BUGS

All information is contained in a static area, so it must be copied if it is to be saved.

NAME

`nice` – change priority of a process

SYNOPSIS

`int nice(incr)`

DESCRIPTION

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 undue impact on system performance.

Negative increments are illegal, except when specified by the super-user. The priority is limited to the range -20 (most urgent) to 19 (least). Requests for values above or below these limits result in the scheduling priority being set to the corresponding limit.

The priority of a process is passed to a child process by `fork(2)`.

RETURN VALUE

Upon successful completion, `nice()` returns the new scheduling priority. Otherwise, a value of -1 is returned and `errno` is set to indicate the error.

ERRORS

The priority is not changed if:

EPERM The value of *incr* specified was negative, or greater than 40, and the effective user ID is not super-user.

SEE ALSO

`nice(1)`, `fork(2)`, `getpriority(2)`, `renice(8)`

NAME

nlist – get entries from symbol table

SYNOPSIS

```
#include <nlist.h>

int nlist(filename, nl)
char *filename;
struct nlist *nl;
```

DESCRIPTION

nlist() examines the symbol table from the executable image whose name is pointed to by *filename*, and selectively extracts a list of values and puts them in the array of **nlist()** structures pointed to by *nl*. The name list pointed to by **nl()** consists of an array of structures containing names, types and values. The *n_name* field of each such structure is taken to be a pointer to a character string representing a symbol name. The list is terminated by an entry with a NULL pointer (or a pointer to a NULL string) in the *n_name* field. For each entry in *nl*, if the named symbol is present in the executable image's symbol table, its value and type are placed in the *n_value* and *n_type* fields. If a symbol cannot be located, the corresponding *n_type* field of **nl()** is set to zero.

RETURN VALUE

Upon normal completion, **nlist()** returns 0. If an error occurs, **nlist()** returns -1 and sets all of the *n_type* fields in members of the array pointed to by **nl()** to zero.

SEE ALSO

a.out(5)

NAME

printf, fprintf, sprintf – formatted output conversion

SYNOPSIS

```
#include <stdio.h>
int printf(format [ , arg ] ... )
char *format;

int fprintf(stream, format [ , arg ] ... )
FILE *stream;
char *format;

int sprintf(s, format [ , arg ] ... )
char *s, *format;

#include <varargs.h>
int _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”, followed by the NULL character (0), in consecutive bytes starting at **s*; it is the user’s responsibility to ensure that enough storage is available. **printf**, **fprintf()** and **sprintf()** return the number of characters transmitted (excluding the NULL character in the case of **sprintf**).

If an output error is encountered **printf**, **fprintf()** and **sprintf()** return EOF.

Each of these functions converts, formats, and prints its *args* under control of the *format*. The *format* 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 zero or more *args*. The results are undefined if there are insufficient *args* for the format. If the format is exhausted while *args* remain, the excess *args* are simply ignored.

Each conversion specification is introduced by the character **%**. After the **%**, the following appear in sequence:

Zero or more *flags*, which modify the meaning of the conversion specification.

An optional decimal digit string specifying a minimum *field width*. If the converted value has fewer characters than the field width, it will be padded on the left (or right, if the left-adjustment flag ‘-’, described below, has been given) to the field width. The padding is with blanks unless the field width digit string starts with a zero, in which case the padding is with zeros.

A *precision* that gives the minimum number of digits to appear for the **d**, **i**, **o**, **u**, **x**, or **X** conversions, the number of digits to appear after the decimal point for the **e**, **E**, and **f** conversions, the maximum number of significant digits for the **g** and **G** conversion, or the maximum number of characters to be printed from a string in **s** conversion. The precision takes the form of a period (.) followed by a decimal digit string; a NULL digit string is treated as zero. Padding specified by the precision overrides the padding specified by the field width.

An optional **l** (ell) specifying that a following **d**, **i**, **o**, **u**, **x**, or **X** conversion character applies to a long integer *arg*. An **l** before any other conversion character is ignored.

A character that indicates the type of conversion to be applied.

A field width or precision or both may be indicated by an asterisk (*) instead of a digit string. In this case, an integer *arg* supplies the field width or precision. The *arg* that is actually converted is not fetched until the conversion letter is seen, so the *args* specifying field width or precision must appear *before* the *arg* (if any) to be converted. A negative field width argument is taken as a ‘-’ flag followed by a positive field width. If the precision argument is negative, it will be changed to zero.

The flag characters and their meanings are:

- The result of the conversion will be left-justified within the field.
- + The result of a signed conversion will always begin with a sign (+ or -).
- blank If the first character of a signed conversion is not a sign, a blank will be prefixed to the result. This implies that if the blank and + flags both appear, the blank flag will be ignored.
- # This flag specifies that the value is to be converted to an "alternate form." For c, d, i, s, and u conversions, the flag has no effect. For o conversion, it increases the precision to force the first digit of the result to be a zero. For x or X conversion, a non-zero result will have 0x or 0X prefixed to it. For e, E, f, g, and G conversions, the result will always contain a decimal point, even if no digits follow the point (normally, a decimal point appears in the result of these conversions only if a digit follows it). For g and G conversions, trailing zeroes will *not* be removed from the result (which they normally are).

The conversion characters and their meanings are:

- d,i,o,u,x,X The integer *arg* is converted to signed decimal (d or i), unsigned octal (o), unsigned decimal (u), or unsigned hexadecimal notation (x and X), respectively; the letters abcdef are used for x conversion and the letters ABCDEF for X conversion. The precision specifies the minimum number of digits to appear; if the value being converted can be represented in fewer digits, it will be expanded with leading zeroes. (For compatibility with older versions, padding with leading zeroes may alternatively be specified by prepending a zero to the field width. This does not imply an octal value for the field width.) The default precision is 1. The result of converting a zero value with a precision of zero is a NULL string.
- f The float or double *arg* is converted to decimal notation in the style "[-.]ddd.ddd" where the number of digits after the decimal point is equal to the precision specification. If the precision is missing, 6 digits are given; if the precision is explicitly 0, no digits and no decimal point are printed.
- e,E The float or double *arg* is converted in the style "[-.]d.ddde±ddd," where there is one digit before the decimal point and the number of digits after it is equal to the precision; when the precision is missing, 6 digits are produced; if the precision is zero, no decimal point appears. The E format code will produce a number with E instead of e introducing the exponent. The exponent always contains at least two digits.
- g,G The float or double *arg* is printed in style f or e (or in style E in the case of a G format code), with the precision specifying the number of significant digits. The style used depends on the value converted: style e or E will be used only if the exponent resulting from the conversion is less than -4 or greater than the precision. Trailing zeroes are removed from the result; a decimal point appears only if it is followed by a digit.

The e, E, f, g, and G formats print IEEE indeterminate values (infinity or not-a-number) as "Infinity" or "NaN" respectively.

- c The character *arg* is printed.
- s The *arg* is taken to be a string (character pointer) and characters from the string are printed until a NULL character (0) is encountered or until the number of characters indicated by the precision specification is reached. If the precision is missing, it is taken to be infinite, so all characters up to the first NULL character are printed. A NULL value for *arg* will yield undefined results.
- % Print a %; no argument is converted.

In no case does a non-existent or small field width cause truncation of a field; if the result of a conversion is wider than the field width, the field is simply expanded to contain the conversion result. Padding takes place only if the specified field width exceeds the actual width. Characters generated by `printf()` and `fprintf()` are printed as if `putc(3S)` had been called.

EXAMPLES

To print a date and time in the form "Sunday, July 3, 10:02," where *weekday* and *month* are pointers to NULL-terminated strings:

```
printf("%s, %s %i, %d:%.2d", weekday, month, day, hour, min);
```

To print π to 5 decimal places:

```
printf("pi = %.5f", 4 * atan(1.0));
```

NOTE

These routines call `_doprnt`, which is an implementation-dependent routine. Each uses the variable-length argument facilities of `varargs(3)`. Although it is possible to use `_doprnt` to take a list of arguments and pass them on to a routine like `printf`, not all implementations have such a routine. We strongly recommend that you use the routines described in `vprintf(3S)` instead.

SEE ALSO

`econvert(3)`, `printf(3S)`, `putc(3S)`, `scanf(3V)`, `varargs(3)`, `vprintf(3S)`

BUGS

Very wide fields (>128 characters) fail.

NAME

rand, srand – simple random number generator

SYNOPSIS

srand(seed)

int seed;

rand()

DESCRIPTION

rand() uses a multiplicative congruential random number generator with period 2^{32} to return successive pseudo-random numbers in the range from 0 to $2^{15}-1$.

srand() can be called at any time to reset the random-number generator to a random starting point. The generator is initially seeded with a value of 1.

NOTE

The spectral properties of **rand()** leave a great deal to be desired. **drand48(3)** and **random(3)** provide much better, though more elaborate, random-number generators.

SEE ALSO

drand48(3), random(3), rand(3C)

BUGS

The low bits of the numbers generated are not very random; use the middle bits. In particular the lowest bit alternates between 0 and 1.

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 results are undefined in there are insufficient *args* for the format. If the format is exhausted while *args* remain, the excess *args* are simply ignored.

The control string usually contains conversion specifications, which are used to direct interpretation of input sequences. The control string may contain:

1. White-space characters (SPACE, TAB, NEWLINE, or FORMFEED) which, except in two cases described below, cause input to be read up to the next non-white-space character.
2. An ordinary character (not '%'), which must match the next character of the input stream.
3. Conversion specifications, consisting of the character '%', an optional assignment suppressing character '*', an optional numerical maximum field width, an optional l (ell) or h indicating the size of the receiving variable, and a conversion code.

A conversion specification directs the conversion of the next input field; the result is placed in the variable pointed to by the corresponding argument, unless assignment suppression was indicated by '*'. The suppression of assignment provides a way of describing an input field which is to be skipped. An input field is defined as a string of non-space characters; it extends to the next inappropriate character or until the field width, if specified, is exhausted. For all descriptors except '[' and 'c', white space leading an input field is ignored.

The conversion character indicates the interpretation of the input field; the corresponding pointer argument must usually be of a restricted type. For a suppressed field, no pointer argument is given. 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.
u	An unsigned decimal integer is expected; the corresponding argument should be an unsigned integer pointer.
o	An octal integer is expected; the corresponding argument should be an integer pointer.
x	A hexadecimal integer is expected; the corresponding argument should be an integer pointer.
i	An integer is expected; the corresponding argument should be an integer pointer. It will store the value of the next input item interpreted according to C conventions: a leading "0" implies octal; a leading "0x" implies hexadecimal; otherwise, decimal.
n	Stores in an integer argument the total number of characters (including white space) that have been scanned so far since the function call. No input is consumed.
e,f,g	A floating point number is expected; the next field is converted accordingly and stored through the corresponding argument, which should be a pointer to a <i>float</i> . The input format for floating point numbers is as described for <i>string_to_decimal(3)</i> , with

- fortran_exponent* zero.
- 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 automatically. The input field is terminated by a white space character.
 - c A character is expected; the corresponding argument should be a character pointer. The normal skip over white space is suppressed in this case; to read the next non-space character, use `%1s`. If a field width is given, the corresponding argument should refer to a character array, and the indicated number of characters is read.
 - [Indicates string data; the normal skip over leading white space is suppressed. The left bracket is followed by a set of characters, which we will call the *scanset*, and a right bracket; the input field is the maximal sequence of input characters consisting entirely of characters in the scanset. The circumflex (`^`), when it appears as the first character in the scanset, serves as a complement operator and redefines the scanset as the set of all characters *not* contained in the remainder of the scanset string. There are some conventions used in the construction of the scanset. A range of characters may be represented by the construct *first-last*, thus `[0123456789]` may be expressed `[0-9]`. Using this convention, *first* must be lexically less than or equal to *last*, or else the dash will stand for itself. The dash will also stand for itself whenever it is the first or the last character in the scanset. To include the right square bracket as an element of the scanset, it must appear as the first character (possibly preceded by a circumflex) of the scanset, and in this case it will not be syntactically interpreted as the closing bracket. The corresponding argument must point to a character array large enough to hold the data field and the terminating `\0`, which will be added automatically. At least one character must match for this conversion to be considered successful.

The conversion characters `d`, `u`, `o`, `x`, and `i` may be preceded by `l` or `h` to indicate that a pointer to `long` or `short` rather than to `int` is in the argument list. Similarly, the conversion characters `e`, `f`, and `g` may be preceded by `l` to indicate that a pointer to `double` rather than to `float` is in the argument list. The `l` or `h` modifier is ignored for other conversion characters.

Avoid this common error: because `printf(3V)` does not require that the lengths of conversion descriptors and actual parameters match, coders sometimes are careless with the `scanf()` functions. But converting `%f` to `&double` or `%lf` to `&float` *does not work*; the results are quite incorrect.

`scanf()` conversion terminates at EOF, at the end of the control string, or when an input character conflicts with the control string. In the latter case, the offending character is left unread in the input stream.

`scanf()` returns the number of successfully matched and assigned input items; this number can be zero in the event of an early conflict between an input character and the control string. The constant EOF is returned upon end of input. Note: 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.

If the input ends before the first conflict or conversion, EOF is returned. If the input ends after the first conflict or conversion, the number of successfully matched items is returned.

EXAMPLES

The call:

```
int i, n; float x; char name[50];
n = scanf("%d%f%s", &i, &x, name);
```

with the input line:

```
25 54.32E-1 thompson
```

will assign to `n` the value 3, to `i` the value 25, to `x` the value 5.432, and `name` will contain `thompson\0`. Or:

```
int i, j; float x; char name[50];
(void) scanf("%i%2d%f%*d %[0-9]", &j, &i, &x, name);
```

with input:

```
011 56789 0123 56a72
```

will assign 9 to *j*, 56 to *i*, 789.0 to *x*, skip 0123, and place the string 56\0 in *name*. The next call to `getchar()` (see `getc(3S)`) will return a. Or:

```
int i, j, s, e; char name[50];
```

```
(void) scanf("%i %i %n %s %n", &i, &j, &s, name, &e);
```

with input:

```
0x11 0xy johnson
```

will assign 17 to *i*, 0 to *j*, 6 to *s*, will place the string xy\0 in *name*, and will assign 8 to *e*. Thus, the length of *name* is $e - s = 2$. The next call to `getchar()` (see `getc(3S)`) will return a SPACE.

SEE ALSO

`getc(3S)`, `printf(3V)`, `stdio(3V)`, `string_to_decimal(3)`, `strtol(3)`, `scanf(3S)`

DIAGNOSTICS

These functions return EOF on end of input, and a short count for missing or illegal data items.

BUGS

The success of literal matches and suppressed assignments is not directly determinable.

CAVEATS

Trailing white space (including a NEWLINE) is left unread unless matched in the control string.

NAME

setbuf, setbuffer, setlinebuf, setvbuf – 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;

int setvbuf(stream, buf, type, size)
FILE *stream;
char *buf;
int type, size;
```

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 any line buffered input stream. `fflush()` (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.

By default, output to a terminal is line buffered, except for output to the standard stream `stderr` which is unbuffered, and all other input/output is fully buffered.

`setbuf()` can be used after a stream has been opened but before it is read or written. It causes the array pointed to by `buf` to be used instead of an automatically allocated buffer. If `buf` is the NULL pointer, input/output will be completely unbuffered. A manifest constant `BUFSIZ`, defined in the `<stdio.h>` header file, tells how big an array is needed:

```
char buf[BUFSIZ];
```

`setbuffer`, an alternate form of `setbuf`, can be used after a stream has been opened but before it is read or written. It uses the character array `buf` whose size is determined by the `size` argument instead of an automatically allocated buffer. If `buf` is the NULL pointer, input/output will be completely unbuffered.

`setvbuf()` can be used after a stream has been opened but before it is read or written. `type` determines how stream will be buffered. Legal values for `type` (defined in `<stdio.h>`) are:

```
_IOFBF    fully buffers the input/output.
_IOLBF    line buffers the output; the buffer will be flushed when a NEWLINE is written, the buffer is full, or input is requested.
_IONBF    completely unbuffers the input/output.
```

If `buf` is not the NULL pointer, the array it points to will be used for buffering, instead of an automatically allocated buffer. `size` specifies the size of the buffer to be used.

`setlinebuf()` is used to change the buffering on a stream from block buffered or unbuffered to line buffered. Unlike `setbuf`, `setbuffer`, and `setvbuf`, 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.

NOTE

A common source of error is allocating buffer space as an “automatic” variable in a code block, and then failing to close the stream in the same block.

SEE ALSO

fclose(3S), fopen(3V), fread(3S), getc(3S), malloc(3), printf(3V), putc(3S), puts(3S), setbuf(3S)

DIAGNOSTICS

If an illegal value for *type* or *size* is provided, **setvbuf()** returns a non-zero value. Otherwise, the value returned will be zero.

NAME

setjmp, longjmp, sigsetjmp, siglongjmp – non-local goto

SYNOPSIS

```
#include <setjmp.h>

int setjmp(env)
jmp_buf env;

longjmp(env, val)
jmp_buf env;
int val;

int _setjmp(env)
jmp_buf env;

_longjmp(env, val)
jmp_buf env;
int val;

int sigsetjmp(env, savemask)
sigjmp_buf env;
int savemask;

siglongjmp(env, val)
sigjmp_buf env;
int val;
```

DESCRIPTION

setjmp() and longjmp() 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. A normal call to setjmp() returns zero. setjmp() also saves the register environment. If a longjmp() call will be made, the routine which called setjmp() should not return until after the longjmp() has returned control (see below).

longjmp() restores the environment saved by the last call of setjmp, and 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; however, if *val* were zero, execution would continue as if the call of setjmp() had returned one. This ensures that a “return” from setjmp() caused by a call to longjmp() can be distinguished from a regular return from setjmp. The calling function must not itself have returned in the interim, otherwise longjmp() will be returning control to a possibly non-existent environment. All memory-bound data have values as of the time longjmp() was called. The CPU and floating-point data registers are restored to the values they had at the time that setjmp() was called. But, because the register storage class is only a hint to the C compiler, variables declared as register variables may not necessarily be assigned to machine registers, so their values are unpredictable after a longjmp. This is especially a problem for programmers trying to write machine-independent C routines.

setjmp() and longjmp() manipulate only the C stack and registers; they do not save or restore the signal mask. _setjmp behaves identically to setjmp, and _longjmp behaves identically to longjmp. If the *savemask* flag to sigsetjmp is non-zero, the signal mask (see sigsetmask(2)) is saved, and a subsequent siglongjmp using the same *env* will restore the signal mask. If the *savemask* flag is zero, the signal mask is not saved, and a subsequent siglongjmp using the same *env* will not restore the signal mask. In all other ways, sigsetjmp functions in the same way that setjmp() does, and siglongjmp functions in the same way that longjmp() does.

None of these functions save or restore any floating-point status or control registers, in particular the MC68881 *fpsr*, *fpcr*, or *fpiar*, the Sun-3 FPA *fpamode* or *fpastatus*, and the Sun-4 *%fsr*. See *ieee_flags(3M)* to save and restore floating-point status or control information.

EXAMPLE

The following code fragment indicates the flow of control of the `setjmp()` and `longjmp()` combination:

```
function declaration
...
    jmp_buf my_environment;
    ...
    if (setjmp(my_environment)) {
        /* register variables have unpredictable values
           code after the return from longjmp
           ...
        } else {
            /* do not modify register vars
               this is the return from setjmp
               ...
            }
        }
```

SEE ALSO

`cc(1V)`, `sigsetmask(2)`, `sigvec(2)`, `ieee_flags(3M)`, `signal(3V)`, `setjmp(3)`

BUGS

`setjmp()` does not save the 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.

On Sun-2 and Sun-3 systems `setjmp()` also saves the register environment. Therefore, all data that are bound to registers are restored to the values they had at the time that `setjmp()` was called. All memory-bound data have values as of the time `longjmp()` was called. However, because the `register` storage class is only a hint to the C compiler, variables declared as `register` variables may not necessarily be assigned to machine registers, so their values are unpredictable after a `longjmp`. When using compiler options that specify automatic register allocation (see `cc(1V)`), the compiler will not attempt to assign variables to registers in routines that call `setjmp`.

`longjmp()` never causes `setjmp()` to return zero in the Sun implementation; this is also true of many other implementations, including all System V implementations, so programmers should not depend on `longjmp()` being able to cause `setjmp()` to return zero.

NAME

setuid, setgid – set user and group IDs

SYNOPSIS

setuid(uid)

setgid(gid)

DESCRIPTION

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

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

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

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

SEE ALSO

execve(2), **getgid(2)**, **getuid(2)**, **setregid(2)**, **setreuid(2)**,

DIAGNOSTICS

Zero is returned if the user (group) ID is set; -1 is returned otherwise, with the global variable **errno** set as for **setreuid()** (**setregid**).

NAME

signal – simplified software signal facilities

SYNOPSIS

```
#include <signal.h>

void (*signal(sig, func))()
void (*func)();
```

DESCRIPTION

signal() is a simplified interface to the more general **sigvec(2)** facility. Programs that use **signal()** in preference to **sigvec()** are more likely to be portable to all systems.

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 **termio(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 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
SIGABRT	6*	abort (generated by abort(3) routine)
SIGEMT	7*	emulator trap
SIGFPE	8*	arithmetic exception
SIGKILL	9	kill (cannot be caught, blocked, or ignored)
SIGBUS	10*	bus error
SIGSEGV	11*	segmentation violation
SIGSYS	12*	bad argument to system call
SIGPIPE	13	write on a pipe or other socket with no one to read it
SIGALRM	14	alarm clock
SIGTERM	15	software termination signal
SIGURG	16●	urgent condition present on socket
SIGSTOP	17†	stop (cannot be caught, blocked, or ignored)
SIGTSTP	18†	stop signal generated from keyboard
SIGCONT	19●	continue after stop (cannot be blocked)
SIGCHLD	20●	child status has changed
SIGTTIN	21†	background read attempted from control terminal
SIGTTOU	22†	background write attempted to control terminal
SIGIO	23●	I/O is possible on a descriptor (see fcntl(2V))
SIGXCPU	24	cpu time limit exceeded (see getrlimit(2))
SIGXFSZ	25	file size limit exceeded (see getrlimit(2))
SIGVTALRM	26	virtual time alarm (see getitimer(2))
SIGPROF	27	profiling timer alarm (see getitimer(2))
SIGWINCH	28●	window changed (see termio(4) and win(4S))
SIGLOST	29*	resource lost (see lockd(8C))
SIGUSR1	30	user-defined signal 1
SIGUSR2	31	user-defined signal 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 *func* is called. The value of *func* for the caught signal is reset to SIG_DFL before *func* is called, unless the signal is SIGILL or SIGTRAP

A return from the function continues the process at the point it was interrupted.

If a caught signal occurs during certain system calls, causing the call to terminate prematurely, the call is interrupted. In particular this can occur during a read(2V) or write(2V) on a slow device (such as a terminal; but not a file) and during a wait(2). After the signal catching function returns, the interrupted system call may return a -1 to the calling process with *errno* set to EINTR.

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. An execve(2) resets all caught signals to the default action; ignored signals remain ignored.

NOTES

The handler routine can be declared:

```
void handler(sig, code, scp, addr)
int sig, code;
struct sigcontext *scp;
char *addr;
```

Here *sig* is the signal number; *code* is a parameter of certain signals that provides additional detail; *scp* is a pointer to the sigcontext structure (defined in <signal.h>), used to restore the context from before the signal; and *addr* is additional address information. See sigvec(2) for more details.

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	<i>sig</i> 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), execve(2), fork(2), getitimer(2), getrlimit(2), kill(2V), ptrace(2), read(2V), sigblock(2), sig-pause(2), sigsetmask(2), sigstack(2), sigvec(2), vfork(2), wait(2), write(2V), setjmp(3), termio(4)

NAME

sleep – suspend execution for interval

SYNOPSIS

unsigned sleep(seconds)
unsigned seconds;

DESCRIPTION

sleep() suspends the current process from execution for the number of seconds specified by the argument. The actual suspension time may be less than that requested for two reasons: (1) because scheduled wake-ups occur at fixed 1-second intervals and (2) because any caught signal will terminate the **sleep()** following execution of that signal's catching routine. Also, the suspension time may be an arbitrary amount longer than requested because of other activity in the system. The value returned by **sleep()** will be the "unslept" amount (the requested time minus the time actually slept) in case the caller had an alarm set to go off earlier than the end of the requested **sleep()** time, or premature arousal due to another caught signal.

sleep() is implemented by setting an interval timer and pausing until it expires. The previous state of this timer is saved and restored. If the sleep time exceeds the time to the expiration of the previous value of the timer, the process sleeps only until the timer would have expired, and the signal which occurs with the expiration of the timer is sent one second later.

SEE ALSO

setitimer(2), **sigpause(2)**, **usleep(3)**

NAME

stdio – standard buffered input/output package

SYNOPSIS

```
#include <stdio.h>
```

```
FILE *stdin;
```

```
FILE *stdout;
```

```
FILE *stderr;
```

DESCRIPTION

The functions described in sections 3V and 3S constitute a user-level I/O buffering scheme. The in-line macros **getc(3V)** and **putc(3S)** handle characters quickly. The macros **getchar** and **putchar**, and the higher level routines **fgetc**, **getw**, **gets**, **fgets**, **scanf**, **fscanf**, **fread**, **fputc**, **putw**, **puts**, **fputs**, **printf**, **fprintf**, **fwrite** all use or act as if they 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(3V)** creates certain descriptive data for a stream and returns a pointer to designate the stream in all further transactions. Normally, there are three open streams with constant pointers declared in the **<stdio.h>** include file and associated with the standard open files:

```
stdin      standard input file
stdout     standard output file
stderr     standard error file
```

A constant **NULL** (0) designates a nonexistent pointer.

An integer constant **EOF** (-1) is returned upon end-of-file or error by most integer functions that deal with streams (see the individual descriptions for details).

Any module that uses this package must include the header file of pertinent macro definitions, as follows:

```
#include <stdio.h>
```

The functions and constants mentioned in sections labeled 3V and 3S of this manual are declared in that header file and need no further declaration. The constants and the following ‘functions’ are implemented as macros; redeclaration of these names is perilous: **getc**, **getchar**, **putc**, **putchar**, **feof**, **ferror**, **fileno**, and **clearerr**.

Output streams, with the exception of the standard error stream **stderr**, are by default buffered if the output refers to a file and line-buffered if the output refers to a terminal. The standard error output stream **stderr** is by default unbuffered, but use of **fopen(3V)** will cause it to become buffered or line-buffered. When an output stream is unbuffered, information is written to the destination file or terminal as soon as it is output to the stream; when it is buffered, many characters are saved up and written as a block. When it is line-buffered, each line of output is written to the destination file or terminal as soon as the line is completed (that is, as soon as a **NEWLINE** character is output or as soon as input is read from a line-buffered stream). **setbuf(3V)**, **setbuffer**, **setlinebuf**, or **setvbuf** can be used to change the stream’s buffering strategy.

SEE ALSO

open(2V), **close(2)**, **lseek(2)**, **pipe(2)**, **read(2V)**, **vfork(2)**, **write(2V)**, **ctermid(3S)**, **cuserid(3S)**, **fclose(3S)**, **ferror(3V)**, **fopen(3V)**, **fread(3S)**, **fseek(3S)**, **getc(3V)**, **gets(3S)**, **popen(3S)**, **printf(3V)**, **putc(3S)**, **puts(3S)**, **scanf(3V)**, **setbuf(3V)**, **system(3)**, **tmpfile(3S)**, **tmpnam(3S)**, **ungetc(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.

BUGS

The standard buffered functions do not interact well with certain other library and system functions, especially **vfork(2)**.

NOTES

The line buffering of output to terminals is almost always transparent, but may cause confusion or malfunctioning of programs which use standard I/O routines but use `read(2V)` to read from the standard input, as calls to `read()` do not cause output to line-buffered streams to be flushed.

Output saved up on *all* line-buffered streams is written when input is read from *any* line-buffered stream. Input read from a stream that is not line-buffered does not flush output on line-buffered streams.

In cases where a large amount of computation is done after printing part of a line on an output terminal, it is necessary to call `fflush` (see `fclose(3S)`) on the standard output before performing the computation so that the output will appear.

NAME

times – get process and child process times

SYNOPSIS

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

long times(buffer)
struct tms *buffer;
```

DESCRIPTION

times() 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**:

```
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 */
};
```

This information comes from the calling process and each of its terminated child processes for which it has executed a **wait**.

tms_utime is the CPU time used while executing instructions in the user space of the calling process.

tms_stime is the CPU time used by the system on behalf of the calling process.

tms_cutime is the sum of the **tms_utimes** and **tms_cutimes** of the child processes.

tms_cstime is the sum of the **tms_stimes** and **tms_cstimes** of the child processes.

RETURN VALUE

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

SEE ALSO

time(1V), **getrusage(2)**, **wait(2)**, **time(3C)**

NAME

ttyslot – find the slot in the utmp file of the current process

SYNOPSIS

ttyslot()

DESCRIPTION

ttyslot() returns the index of the current user's entry in the **/etc/utmp** file. This is accomplished by actually scanning the file **/etc/ttys** for the name of the terminal associated with the standard input, the standard output, or the error output (0, 1 or 2).

FILES

/etc/ttys
/etc/utmp

DIAGNOSTICS

A value of **-1** is returned if an error was encountered while searching for the terminal name or if none of the above file descriptors is associated with a terminal device.

NAME

vprintf, vfprintf, vsprintf – print formatted output of a varargs argument list

SYNOPSIS

```
#include <stdio.h>
#include <varargs.h>

int vprintf(format, ap)
char *format;
va_list ap;

int vfprintf(stream, format, ap)
FILE *stream;
char *format;
va_list ap;

int vsprintf(s, format, ap)
char *s, *format;
va_list ap;
```

DESCRIPTION

vprintf, vfprintf, and vsprintf() are the same as printf(3V), fprintf, and sprintf respectively, except that instead of being called with a variable number of arguments, they are called with an argument list as defined by varargs(3).

EXAMPLE

The following demonstrates how vfprintf() could be used to write an error routine.

```
#include <stdio.h>
#include <varargs.h>
...
/* error should be called like:
 *      error(function_name, format, arg1, arg2...);
 * Note that function_name and format cannot be declared
 * separately because of the definition of varargs.
 */

/*VARARGS0*/
void
error (va_alist
      va_dcl;
{
    va_list args;
    char *fmt;

    va_start(args);
    /* print name of function causing error */
    (void) fprintf(stderr, "ERROR in %s: ", va_arg(args, char *));
    fmt = va_arg(args, char *);
    /* print out remainder of message */
    (void) vfprintf(stderr, fmt, args);
    va_end(args);
    (void) abort();
}
```

SEE ALSO

printf(3V), varargs(3)

NAME

intro – introduction to device drivers, protocols, and network interfaces

DESCRIPTION

This section describes device drivers, high-speed network interfaces, and protocols available under SunOS. The system provides drivers for a variety of hardware devices, such as disks, magnetic tapes, serial communication lines, mice and frame buffers, as well as virtual devices such as pseudo-terminals and windows. SunOS provides hardware support and a network interface for the 10-Megabit Ethernet, along with interfaces for the IP protocol family and a STREAMS-based Network Interface Tap (NIT) facility.

In addition to describing device drivers that are supported by the 4.3BSD operating system, this section contains subsections that describe:

- SunOS-specific device drivers, under '4S'.
- Protocol families, under '4F'.
- Protocols and raw interfaces, under '4P'.
- STREAMS modules, under '4M'.
- Network interfaces, under '4N'.

Configuration

The SunOS kernel can be configured to include or omit many of the device drivers described in this section. The CONFIG section of the manual page gives the line(s) to include in the kernel configuration file for each machine architecture on which a device is supported. If no specific architectures are indicated, the configuration syntax applies to all Sun systems.

The GENERIC kernel is the default configuration for SunOS. It contains all of the optional drivers for a given machine architecture. See `config(8)`, for details on configuring a new SunOS kernel.

The manual page for a device driver may also include a DIAGNOSTICS section, listing error messages that the driver might produce. Normally, these messages are logged to the appropriate system log using the kernel's standard message-buffering mechanism (see `syslogd(8)`); they may also appear on the system console.

Ioctls

Various special functions, such as querying or altering the operating characteristics of a device, are performed by supplying appropriate parameters to the `ioctl(2)` system call. These parameters are often referred to as "ioctls." Ioctls for a specific device are presented in the manual page for that device. Ioctls that pertain to a class of devices are listed in a manual page with a name that suggests the class of device, and ending in 'io', such as `mtio(4)` for magnetic tape devices, or `dkio(4S)` for disk controllers. In addition, some ioctls operate directly on higher-level objects such as files, terminals, sockets, and streams:

- Ioctls that operate directly on files, file descriptors, and sockets are described in `filio(4)`. Note: the `fcntl(2)` system call is the primary method for operating on file descriptors as such, rather than on the underlying files. Also note that the `setsockopt` system call (see `getsockopt(2)`) is the primary method for operating on sockets as such, rather than on the underlying protocol or network interface. Ioctls for a specific network interface are documented in the manual page for that interface.
- Ioctls for terminals, including pseudo-terminals, are described in `termio(4)`. This manual page includes information about both the BSD `termios` structure, as well as the System V `termio` structure.
- Ioctls for STREAMS are described in `streamio(4)`.

Devices Always Present

Device drivers present in every kernel include:

- The paging device; see `drum(4)`.
- Drivers for accessing physical, virtual, and I/O space in memory; see `mem(4S)`.
- The data sink; see `null(4)`.

Terminals and Serial Communications Devices

Serial communication lines are normally supported by the terminal driver; see `tty(4)`. This driver manages serial lines provided by communications drivers, such as those described in `mti(4S)` and `zs(4S)`. The terminal driver also handles serial lines provided by virtual terminals, such as the Sun console monitor described in `console(4S)`, and true pseudo-terminals, described in `pty(4)`.

Disk Devices

Drivers for the following disk controllers provide standard block and raw interfaces under SunOS;

- SCSI controllers, in `sd(4S)`,
- Xylogics 450 and 451 SMD controllers, in `xy(4S)`,
- Xylogics 7053 SMD controllers, in `xd(4S)`.

Ioctls to query or set a disk's geometry and partitioning are described in `dkio(4S)`.

Magnetic Tape Devices

Magnetic tape devices supported by SunOS include those described in `ar(4S)`, `tm(4S)`, `st(4S)`, and `xt(4S)`. Ioctls for all tape-device drivers are described in `mtio(4S)`.

Frame Buffers

Frame buffer devices include color frame buffers described in the `cg*(4S)` manual pages, monochrome frame buffers described in the `bw*(4S)` manual pages, graphics processor interfaces described in the `gp*(4S)` manual pages, and an indirect device for the console frame buffer described in `fb(4S)`. Ioctls for all frame-buffer devices are described in `fbio(4S)`.

Miscellaneous Devices

Miscellaneous devices include the console keyboard described in `kbd(4S)`, the console mouse described in `mouse(4S)`, window devices described in `win(4S)`, and the DES encryption-chip interface described in `des(4S)`.

Network-Interface Devices

SunOS supports the 10-Megabit Ethernet as its primary network interface; see `ec(4S)`, `ie(4S)`, and `le(4S)` for details. However, a software loopback interface, `lo(4)` is also supported. General properties of these network interfaces are described in `if(4N)`, along with the ioctls that operate on them.

Support for network routing is described in `routing(4N)`.

Protocols and Protocol Families

SunOS supports both socket-based and STREAMS-based network communications. The Internet protocol family, described in `inet(4F)`, is the primary protocol family primary supported by SunOS, although the system can support a number of others. The raw interface provides low-level services, such as packet fragmentation and reassembly, routing, addressing, and basic transport for socket-based implementations. Facilities for communicating using an Internet-family protocol are generally accessed by specifying the `AF_INET` address family when binding a socket; see `socket(2)` for details.

Major protocols in the Internet family include:

- The Internet Protocol (IP) itself, which supports the universal datagram format, as described in `ip(4P)`. This is the default protocol for `SOCK_RAW` type sockets within the `AF_INET` domain.
- The Transmission Control Protocol (TCP); see `tcp(4P)`. This is the default protocol for `SOCK_STREAM` type sockets.
- The User Datagram Protocol (UDP); see `udp(4P)`. This is the default protocol for `SOCK_DGRAM` type sockets.
- The Address Resolution Protocol (ARP); see `arp(4P)`.
- The Internet Control Message Protocol (ICMP); see `icmp(4P)`.

The Network Interface Tap (NIT) protocol, described in `nit(4P)`, is a STREAMS-based facility for accessing the network at the link level.

SEE ALSO

`fcntl(2)`, `getsockopt(2)`, `ioctl(2)`, `socket(2)`, `ar(4S)`, `arp(4P)`, `dkio(4S)`, `drum(4)`, `ec(4S)`, `fb(4S)`, `fbio(4S)`, `filio(4)`, `icmp(4P)`, `if(4N)`, `inet(4F)`, `ip(4P)`, `kbd(4S)`, `le(4)`, `lo(4)`, `mbio(4S)`, `mem(4S)`, `mti(4)`, `mtio(4)`, `nit(4P)`, `null(4)`, `pty(4)`, `routing(4N)`, `sd(4S)`, `st(4S)`, `streamio(4)`, `tcp(4P)`, `termio(4)`, `tm(4S)`, `tty(4)`, `udp(4P)`, `win(4S)`, `xd(4S)`, `xy(4S)`, `zs(4S)`

LIST OF DEVICES, INTERFACES AND PROTOCOLS

Name	Appears on Page	Description
<code>alm</code>	<code>mcp(4S)</code>	Asynchronous Line Multiplexer
<code>ar</code>	<code>ar(4S)</code>	Archive 1/4 inch Streaming Tape Drive
<code>arp</code>	<code>arp(4P)</code>	Address Resolution Protocol
<code>bk</code>	<code>bk(4)</code>	line discipline for machine-machine communication
<code>bwone</code>	<code>bwone(4S)</code>	Sun-1 black and white frame buffer
<code>bwtwo</code>	<code>bwtwo(4S)</code>	Sun-3/Sun-2 black and white frame buffer
<code>cgfour</code>	<code>cgfour(4S)</code>	Sun-3 color memory frame buffer
<code>cgone</code>	<code>cgone(4S)</code>	Sun-1 color graphics interface
<code>cgthree</code>	<code>cgthree(4S)</code>	Sun386i color memory frame buffer
<code>cgtwo</code>	<code>cgtwo(4S)</code>	Sun-3/Sun-2 color graphics interface
<code>clone</code>	<code>clone(4)</code>	open any minor device on a STREAMS driver
<code>console</code>	<code>console(4S)</code>	console driver and terminal emulator for the Sun workstation
<code>des</code>	<code>des(4S)</code>	DES encryption chip interface
<code>dkio</code>	<code>dkio(4S)</code>	generic disk control operations
<code>drum</code>	<code>drum(4)</code>	paging device
<code>ec</code>	<code>ec(4S)</code>	3Com 10 Mb/s Ethernet interface
<code>fb</code>	<code>fb(4S)</code>	driver for Sun console frame buffer
<code>fbio</code>	<code>fbio(4S)</code>	general properties of frame buffers
<code>fd</code>	<code>fd(4S)</code>	Disk driver for Floppy Disk Controllers
<code>filio</code>	<code>filio(4)</code>	ioctl's that operate directly on files, file descriptors, and sockets
<code>fpa</code>	<code>fpa(4S)</code>	Sun-3 floating point accelerator
<code>gpone</code>	<code>gpone(4S)</code>	Sun-3/Sun-2 graphics processor
<code>icmp</code>	<code>icmp(4P)</code>	Internet Control Message Protocol
<code>ie</code>	<code>ie(4S)</code>	Intel 10 Mb/s Ethernet interface
<code>if</code>	<code>if(4N)</code>	general properties of network interfaces
<code>inet</code>	<code>inet(4F)</code>	Internet protocol family
<code>ip</code>	<code>ip(4P)</code>	Internet Protocol
<code>kb</code>	<code>kb(4M)</code>	Sun keyboard STREAMS module
<code>kbd</code>	<code>kbd(4S)</code>	Sun keyboard
<code>kmem</code>	<code>mem(4S)</code>	main memory and bus I/O space
<code>ldterm</code>	<code>ldterm(4M)</code>	standard terminal STREAMS module
<code>le</code>	<code>le(4S)</code>	Sun-3/50, Sun-3/60 10MB Ethernet interface
<code>lo</code>	<code>lo(4)</code>	software loopback network interface
<code>lofs</code>	<code>lofs(4S)</code>	loopback virtual file system
<code>mbio</code>	<code>mem(4S)</code>	main memory and bus I/O space
<code>mbmem</code>	<code>mem(4S)</code>	main memory and bus I/O space
<code>mcp</code>	<code>mcp(4S)</code>	MCP Multiprotocol Communications Processor
<code>mem</code>	<code>mem(4S)</code>	main memory and bus I/O space
<code>mouse</code>	<code>mouse(4S)</code>	Sun mouse
<code>ms3</code>	<code>mouse(4S)</code>	Sun mouse
<code>ms</code>	<code>ms(4M)</code>	Sun mouse STREAMS module
<code>mti</code>	<code>mti(4S)</code>	System MTI-800/1600 multi-terminal interface
<code>mtio</code>	<code>mtio(4)</code>	UNIX system magnetic tape interface

NFS	nfs(4P)	network file system
nif_pf	nit_pf(4M)	streams NIT packet filtering module
nit	nit(4P)	Network Interface Tap facility
nit_buf	nit_buf(4M)	streams NIT buffering module
nit_if	nit_if(4M)	streams NIT device interface module
null	null(4)	data sink
pp	pp(4)	Centronics-compatible parallel printer port
pty	pty(4)	pseudo terminal driver
root	root(4S)	pseudo-driver for Sun root disk
routing	routing(4N)	system supporting for local network packet routing
sd	sd(4S)	Disk driver for SCSI Disk Controllers
st	st(4S)	Sysgen SC 4000 and Emulex MT-02 Tape Controller
streamio	streamio(4)	STREAMS ioctl commands
tcp	tcp(4P)	Transmission Control Protocol
termio	termio(4)	general terminal interface
tm	tm(4S)	tapemaster 1/2 inch tape drive
ttcompat	ttcompat(4M)	V7/4BSD compatibility STREAMS module
tty	tty(4)	controlling terminal interface
udp	udp(4P)	User Datagram Protocol
vme16d16	mem(4S)	main memory and bus I/O space
vme16d32	mem(4S)	main memory and bus I/O space
vme24d16	mem(4S)	main memory and bus I/O space
vme24d32	mem(4S)	main memory and bus I/O space
vme32d16	mem(4S)	main memory and bus I/O space
vme32d32	mem(4S)	main memory and bus I/O space
vp	vp(4S)	Ikon 10071-5 Versatec parallel printer interface
vpc	vpc(4S)	Systech VPC-2200 Versatec plotter and Centronics printer
win	win(4S)	Sun window system
xd	xd(4S)	Disk driver for Xylogics 7053 SMD Disk Controller
xt	xt(4S)	Xylogics 472 1/2 inch tape controller
xy	xy(4S)	Disk driver for Xylogics SMD Disk Controllers
zero	zero(4S)	source of zeroes
zs	zs(4S)	Zilog 8530 SCC serial communications driver

NAME

ar – Archive 1/4 inch Streaming Tape Drive

CONFIG — SUN-2 SYSTEM

device ar0 at mbio ? csr 0x200 priority 3

device ar1 at mbio ? csr 0x208 priority 3

AVAILABILITY

Sun-2, Sun-3, and Sun-4 systems only.

DESCRIPTION

The Archive tape controller is a Sun 'QIC-II' interface to an Archive streaming tape drive. It provides a standard tape interface to the device, see `mtio(4)`, with some deficiencies listed under **BUGS** below.

The maximum blocksize for the raw device is limited only by available memory.

FILES

`/dev/rar*`

`/dev/nrar*` non-rewinding

SEE ALSO

`mtio(4)`

DIAGNOSTICS

ar*: would not initialize

ar*: already open

The tape can be open by only one process at a time

ar*: no such drive

ar*: no cartridge in drive

ar*: cartridge is write protected

ar: interrupt from uninitialized controller %x

ar*: many retries, consider retiring this

ar*: %b error at block #

ar*: %b error at block #

ar: giving up on Rdy, try

BUGS

The tape cannot reverse direction so the `BSF` and `BSR` `ioctl`s are not supported.

The `FSR` `ioctl` is not supported.

The system will hang if the tape is removed while running.

When using the raw device, the number of bytes in any given transfer must be a multiple of 512 bytes. If it is not, the device driver returns an error.

The driver will only write an EOF mark on close if the last operation was a write, without regard for the mode used when opening the file. This delete empty files on a raw tape copy operation.

NAME

arp – Address Resolution Protocol

CONFIG

pseudo-device ether

SYNOPSIS

```
#include <sys/socket.h>
#include <net/if_arp.h>
#include <netinet/in.h>

s = socket(AF_INET, SOCK_DGRAM, 0);
```

DESCRIPTION

ARP is a protocol used to dynamically map between DARPA Internet Protocol (IP) and 10Mb/s Ethernet addresses. It is used by all the 10Mb/s Ethernet interface drivers. It is not specific to the Internet Protocol or to the 10Mb/s Ethernet, but this implementation currently supports only that combination.

ARP caches IP-to-Ethernet address mappings. When an interface requests a mapping for an address not in the cache, ARP queues the message which requires the mapping and broadcasts a message on the associated network requesting the address mapping. If a response is provided, the new mapping is cached and any pending message is transmitted. ARP will queue at most one packet while waiting for a mapping request to be responded to; only the most recently “transmitted” packet is kept.

To facilitate communications with systems which do not use ARP, `ioctl`s are provided to enter and delete entries in the IP-to-Ethernet tables.

USAGE

```
#include <sys/sockio.h>
#include <sys/socket.h>
#include <net/if.h>
#include <net/if_arp.h>
struct arpreq arpreq;
ioctl(s, SIOCSARP, (caddr_t)&arpreq);
ioctl(s, SIOCGARP, (caddr_t)&arpreq);
ioctl(s, SIOCDEARP, (caddr_t)&arpreq);
```

Each `ioctl` takes the same structure as an argument. `SIOCSARP` sets an ARP entry, `SIOCGARP` gets an ARP entry, and `SIOCDEARP` deletes an ARP entry. These `ioctl`s may be applied to any socket descriptor `s`, but only by the super-user. The `arpreq` structure contains:

```
/*
 * ARP ioctl request
 */
struct arpreq {
    struct sockaddr arp_pa;    /* protocol address */
    struct sockaddr arp_ha;    /* hardware address */
    int    arp_flags;         /* flags */
};
/* arp_flags field values */
#define ATF_COM                0x2    /* completed entry (arp_ha valid) */
#define ATF_PERM              0x4    /* permanent entry */
#define ATF_PUBL              0x8    /* publish (respond for other host) */
#define ATF_USETRAILERS       0x10   /* send trailer packets to host */
```

The address family for the `arp_pa` `sockaddr` must be `AF_INET`; for the `arp_ha` `sockaddr` it must be `AF_UNSPEC`. The only flag bits which may be written are `ATF_PERM`, `ATF_PUBL` and `ATF_USETRAILERS`. `ATF_PERM` makes the entry permanent if the `ioctl` call succeeds. The peculiar nature of the ARP tables may cause the `ioctl` to fail if more than 6 (permanent) IP addresses hash to the same slot. `ATF_PUBL` specifies that the ARP code should respond to ARP requests for the indicated host

coming from other machines. This allows a host to act as an "ARP server" which may be useful in convincing an ARP-only machine to talk to a non-ARP machine.

ARP is also used to negotiate the use of trailer IP encapsulations; trailers are an alternate encapsulation used to allow efficient packet alignment for large packets despite variable-sized headers. Hosts which wish to receive trailer encapsulations so indicate by sending gratuitous ARP translation replies along with replies to IP requests; they are also sent in reply to IP translation replies. The negotiation is thus fully symmetrical, in that either or both hosts may request trailers. The ATF_USETRAILERS flag is used to record the receipt of such a reply, and enables the transmission of trailer packets to that host.

ARP watches passively for hosts impersonating the local host (that is, a host which responds to an ARP mapping request for the local host's address).

SEE ALSO

ec(4S), ie(4S), inet(4F), arp(8C), ifconfig(8C)

Plummer, Dave, "*An Ethernet Address Resolution Protocol -or- Converting Network Protocol Addresses to 48.bit Ethernet Addresses for Transmission on Ethernet Hardware,*" RFC 826, Network Information Center, SRI International, Menlo Park, Calif., November 1982. (Sun 800-1059-10)

Leffler, Sam, and Michael Karels, "*Trailer Encapsulations,*" RFC 893, Network Information Center, SRI International, Menlo Park, Calif., April 1984.

DIAGNOSTICS

duplicate IP address!! sent from ethernet address: %x:%x:%x:%x:%x:%x. ARP has discovered another host on the local network which responds to mapping requests for its own Internet address.

BUGS

ARP packets on the Ethernet use only 42 bytes of data, however, the smallest legal Ethernet packet is 60 bytes (not including CRC). Some systems may not enforce the minimum packet size, others will.

NAME

bk – line discipline for machine-machine communication

SYNOPSIS

pseudo-device bk

DESCRIPTION

This line discipline provides a replacement for the tty driver `tty(4)` when high speed output to and especially input from another machine is to be transmitted over an asynchronous communications line. The discipline was designed for use by a (now obsolete) store-and-forward local network running over serial lines. It may be suitable for uploading of data from microprocessors into the system. If you are going to send data over asynchronous communications lines at high speed into the system, you must use this discipline, as the system otherwise may detect high input data rates on terminal lines and disable the lines; in any case the processing of such data when normal terminal mechanisms are involved saturates the system.

The line discipline is enabled by a sequence:

```
#include <sgtty.h>
int ldisc = NETLDISC, fildes; ...
ioctl(fildes, TIOCSETD, &ldisc);
```

A typical application program then reads a sequence of lines from the terminal port, checking header and sequencing information on each line and acknowledging receipt of each line to the sender, who then transmits another line of data. Typically several hundred bytes of data and a smaller amount of control information will be received on each handshake.

The old standard teletype discipline can be restored by doing:

```
ldisc = OTTYDISC;
ioctl(fildes, TIOCSETD, &ldisc);
```

While in networked mode, normal teletype output functions take place. Thus, if an 8 bit output data path is desired, it is necessary to prepare the output line by putting it into RAW mode using `ioctl(2)`. This must be done *before* changing the discipline with `TIOCSETD`, as most `ioctl(2)` calls are disabled while in network line-discipline mode.

When in network mode, input processing is very limited to reduce overhead. Currently the input path is only 7 bits wide, with newline the only character terminating an input record. Each input record must be read and acknowledged before the next input is read as the system refuses to accept any new data when there is a record in the buffer. The buffer is limited in length, but the system guarantees to always be willing to accept input resulting in 512 data characters and then the terminating newline.

User level programs should provide sequencing and checksums on the information to guarantee accurate data transfer.

SEE ALSO

`ioctl(2)`, `tty(4)`

NAME

bwone – Sun-1 black and white frame buffer

CONFIG — SUN-2 SYSTEM

device bwone0 at mbmem ? csr 0xc0000 priority 3

DESCRIPTION

The **bwone** interface provides access to Sun-1 system black and white graphics controller boards. It supports the ioctls described in **fbio(4S)**.

FILES

/dev/bwone[0-9]

SEE ALSO

mmap(2), fb(4S), fbio(4S)

BUGS

Use of vertical-retrace interrupts is not supported.

The video state returned by the **FBIOSVIDEO** ioctl may be incorrect. It is not possible for the driver to determine the state of the hardware video enable bit, so it reports the last state stored by the **FBIOSVIDEO** ioctl. User processes which map the frame buffer can directly enable or disable the video, unknown to the driver.

NAME

bwtwo – Sun-3/Sun-2 black and white frame buffer

CONFIG — SUN-3 SYSTEM

device bwtwo0 at obmem 1 csr 0xff000000 priority 4

device bwtwo0 at obmem 2 csr 0x100000 priority 4

device bwtwo0 at obmem 3 csr 0xff000000 priority 4

device bwtwo0 at obmem 4 csr 0xff000000

device bwtwo0 at obmem 7 csr 0xff000000 priority 4

The first synopsis line given above should be used to generate a kernel for a Sun-3/75, Sun-3/140 or Sun-3/160 system; the second, for a Sun-3/50 system; the third, for a Sun-3/260 system; the fourth, for a Sun-3/110 system; and the fifth, for a Sun-3/60 system.

CONFIG — SUN-2 SYSTEM

device bwtwo0 at obmem 1 csr 0x700000 priority 4

device bwtwo0 at obio 2 csr 0x0 priority 4

The first synopsis line given above should be used to generate a kernel for a Sun-2/120 or Sun-2/170 system; the second, for a Sun-2/50 or Sun-2/160 system.

CONFIG — Sun386i SYSTEM

device bwtwo0 at obmem ? csr 0xA0200000

DESCRIPTION

The **bwtwo** interface provides access to Sun monochrome memory frame buffers. It supports the ioctl's described in **fbio(4S)**.

If **flags 0x1** is specified, frame buffer write operations are buffered through regular high-speed RAM. This “copy memory” mode of operation speeds frame buffer accesses, but consumes an extra 128K bytes of memory. Only Sun-2, Sun-3/75, and Sun-3/160 systems support copy memory; on other systems a warning message is printed and the flag is ignored.

Reading or writing to the frame buffer is not allowed — you must use the **mmap(2)** system call to map the board into your address space.

FILES

/dev/bwtwo[0-9]

SEE ALSO

mmap(2), **cgfour(4S)**, **fb(4S)**, **fbio(4S)**

BUGS

Use of vertical-retrace interrupts is not supported.

NAME

cgfour – Sun-3 color memory frame buffer

CONFIG — SUN-3 SYSTEM

device cgfour0 at obmem 4 csr 0xff000000 priority 4

device cgfour0 at obmem 7 csr 0xff000000 priority 4

The first synopsis line given should be used to generate a kernel for the Sun-3/110 system; the second, for a Sun-3/60 system.

DESCRIPTION

The **cgfour** is a color memory frame buffer with a monochrome overlay plane and an overlay enable plane implemented on the Sun-3/110 system and some Sun-3/60 system models. It provides the standard frame buffer interface as defined in **fbio(4S)**.

In addition to the **ioctl**s described under **fbio(4S)**, the **cgfour** interface responds to two **cgfour**-specific colormap **ioctl**s, **FBIOPUTCMAP** and **FBIOGETCMAP**. **FBIOPUTCMAP** returns no information other than success/failure using the **ioctl** return value. **FBIOGETCMAP** returns its information in the arrays pointed to by the red, green, and blue members of its **fbcmmap** structure argument; **fbcmmap** is defined in `<sun/fbio.h>` as:

```

struct fbcmmap {
    int         index;          /* first element (0 origin) */
    int         count;         /* number of elements */
    unsigned char *red;        /* red color map elements */
    unsigned char *green;     /* green color map elements */
    unsigned char *blue;      /* blue color map elements */
};

```

The driver uses color board vertical-retrace interrupts to load the colormap.

The Sun-3/60 system has an overlay plane colormap, which is accessed by encoding the plane group into the index value with the **PIX_GROUP** macro (see `<pixrect/pr_planegroups.h>`).

FILES

/dev/cgfour0

SEE ALSO

mmap(2), **fbio(4S)**

NAME

cgone – Sun-1 color graphics interface

CONFIG — SUN-2 SYSTEM

device cgone0 at mbmem ? csr 0xec000 priority 3

DESCRIPTION

The **cgone** interface provides access to the Sun-1 system color graphics controller board, which is normally supplied with a 13'' or 19'' RS170 color monitor. It provides the standard frame buffer interface as defined in **fbio(4S)**.

It supports the **FBIOGPIXRECT** ioctl which allows SunView to be run on it; see **fbio(4S)**

The hardware consumes 16 kilobytes of Multibus memory space. The board starts at standard addresses 0xE8000 or 0xEC000. The board must be configured for interrupt level 3.

FILES

/dev/cgone[0-9]

SEE ALSO

mmap(2), **fbio(4S)**

BUGS

Use of color board vertical-retrace interrupts is not supported.

NAME

cgthree – Sun386i color memory frame buffer

CONFIG

device cgthree0 at obmem ? csr 0xA0400000

AVAILABILITY

Sun386i systems only.

DESCRIPTION

cgthree is a color memory frame buffer. It provides the standard frame buffer interface as defined in fbio(4S).

In addition to the ioctls described under fbio(4S), the cgthree interface responds to two cgthree-specific colormap ioctl(2) parameters, FBIOPUTCMAP and .SB FBIOGETCMAP. FBIOPUTCMAP returns no information other than success/failure via the ioctl return value. FBIOGETCMAP returns its information in the arrays pointed to by the red, green, and blue members of its fbcmap structure argument; fbcmap is defined in <sun/fbio.h> as:

```

    struct fbcmap {
        int         index;           /* first element (0 origin) */
        int         count;          /* number of elements */
        unsigned char *red;         /* red color map elements */
        unsigned char *green;      /* green color map elements */
        unsigned char *blue;       /* blue color map elements */
    };

```

FILES

/dev/cgthree0

SEE ALSO

mmap(2), fbio(4S)

NAME

cgtwo – Sun-3/Sun-2 color graphics interface

CONFIG — SUN-3 SYSTEM

cgtwo0 at vme24d16 ? csr 0x400000

CONFIG — SUN-2 SYSTEM

cgtwo0 at vme24 ? csr 0x400000

DESCRIPTION

The **cgtwo** interface provides access to the Sun-3/Sun-2 system color graphics controller board, which is normally supplied with a 19" 66 Hz non-interlaced color monitor. It provides the standard frame buffer interface as defined in **fbio(4S)**.

The hardware consumes 4 megabytes of VME bus address space. The board starts at standard address 0x400000. The board must be configured for interrupt level 3.

FILES

/dev/cgtwo[0-9]

SEE ALSO

mmap(2), **fbio(4S)**

NAME

clone – open any minor device on a STREAMS driver

DESCRIPTION

clone is a STREAMS software driver that finds and opens an unused minor device on another STREAMS driver. The minor device passed to **clone** during the open operation is interpreted as the major device number of another STREAMS driver for which an unused minor device is to be obtained. Each such open results in a separate stream to a previously unused minor device.

The **clone** driver supports only an **open(2V)** function. This open function performs all of the necessary work so that subsequent system calls (including **close(2)**) require no further involvement of the **clone** driver.

ERRORS

clone generates an ENXIO error, without opening the device, if the minor device number provided does not correspond to a valid major device, or if the driver indicated is not a STREAMS driver.

CAVEATS

Multiple opens of the same minor device are not supported through the **clone** interface. Executing **stat(2)** on the file system node for a cloned device yields a different result than does executing **fstat** using a file descriptor obtained from opening that node.

SEE ALSO

close(2), **open(2V)**, **stat(2)**

NAME

console – console driver and terminal emulator for the Sun workstation

CONFIG

None; included in standard system.

SYNOPSIS

```
#include <fcntl.h>
#include <sys/termios.h>
open("/dev/console", mode);
```

DESCRIPTION

console is an indirect driver for the Sun console terminal. On a Sun workstation, this driver refers to the workstation console driver, which implements a standard UNIX system terminal. On a Sun server without a keyboard or a frame buffer, this driver refers to the CPU serial port driver (**zs(4S)**); a terminal is normally connected to this port.

The workstation console does not support any of the **termio(4)** device control functions specified by flags in the **c_cflag** word of the **termios** structure or by the **IGNBRK**, **IGNPAR**, **PARMRK**, or **INPCK** flags in the **c_iflag** word of the **termios** structure, as these functions apply only to asynchronous serial ports. All other **termio(4)** functions must be performed by **STREAMS** modules pushed atop the driver; when a slave device is opened, the **ldterm(4M)** and **ttcompat(4M)** **STREAMS** modules are automatically pushed on top of the stream, providing the standard **termio(4)** interface.

The workstation console driver calls the PROM resident monitor to output data to the console frame buffer. Keystrokes from the CPU serial port to which the keyboard is connected are routed through the keyboard **STREAMS** module (**kb(4M)**) and treated as input.

When the Sun window system **win(4S)** is active, console input is directed through the window system rather than being treated as input by the workstation console driver.

IOCTLS

An ioctl **TIOCCONS** can be applied to pseudo-terminals (**pty(4)**) to route output that would normally appear on the console to the pseudo-terminal instead. Thus, the window system does a **TIOCCONS** on a pseudo-terminal so that the system will route console output to the window to which that pseudo-terminal is connected, rather than routing output through the PROM monitor to the screen, since routing output through the PROM monitor destroys the integrity of the screen. Note: when you use **TIOCCONS** in this way, the console *input* is routed from the pseudo-terminal as well.

If a **TIOCCONS** is performed on **/dev/console**, or the pseudo-terminal to which console output is being routed is closed, output to the console will again be routed to the workstation console driver.

ANSI STANDARD TERMINAL EMULATION

The Sun Workstation's PROM monitor provides routines that emulates a standard ANSI X3.64 terminal.

Note: the VT100 also follows the ANSI X3.64 standard but both the Sun and the VT100 have nonstandard extensions to the ANSI X3.64 standard. The Sun terminal emulator and the VT100 are *not* compatible in any true sense.

The Sun console displays 34 lines of 80 ASCII characters per line, with scrolling, (x, y) cursor addressability, and a number of other control functions.

The Sun console displays a non-blinking block cursor which marks the current line and character position on the screen. ASCII characters between 0x20 (space) and 0x7E (tilde) inclusive are printing characters — when one is written to the Sun console (and is not part of an escape sequence), it is displayed at the current cursor position and the cursor moves one position to the right on the current line. If the cursor is already at the right edge of the screen, it moves to the first character position on the next line. If the cursor is already at the right edge of the screen on the bottom line, the Line-feed function is performed (see control-J below), which scrolls the screen up by one or more lines or wraps around, before moving the cursor to the first character position on the next line.

Control Sequence Syntax

The Sun console defines a number of control sequences which may occur in its input. When such a sequence is written to the Sun console, it is not displayed on the screen, but effects some control function as described below, for example, moves the cursor or sets a display mode.

Some of the control sequences consist of a single character. The notation **control-X**

for some character *X*, represents a control character.

Other ANSI control sequences are of the form

ESC [<params> <char>

Spaces are included only for readability; these characters must occur in the given sequence without the intervening spaces.

ESC represents the ASCII escape character (ESC, control-[, 0x1B).

[The next character is a left square bracket '[' (0x5B).

<params>

are a sequence of zero or more decimal numbers made up of digits between 0 and 9, separated by semicolons.

<char> represents a function character, which is different for each control sequence.

Some examples of syntactically valid escape sequences are (again, ESC represent the single ASCII character 'Escape'):

ESC[m	<i>select graphic rendition with default parameter</i>
ESC[7m	<i>select graphic rendition with reverse image</i>
ESC[33;54H	<i>set cursor position</i>
ESC[123;456;0;;3;B	<i>move cursor down</i>

Syntactically valid ANSI escape sequences which are not currently interpreted by the Sun console are ignored. Control characters which are not currently interpreted by the Sun console are also ignored.

Each control function requires a specified number of parameters, as noted below. If fewer parameters are supplied, the remaining parameters default to 1, except as noted in the descriptions below.

If more than the required number of parameters is supplied, only the last *n* are used, where *n* is the number required by that particular command character. Also, parameters which are omitted or set to zero are reset to the default value of 1 (except as noted below).

Consider, for example, the command character M which requires one parameter. ESC[;M and ESC[0M and ESC[M and ESC[23;15;32;1M are all equivalent to ESC[1M and provide a parameter value of 1. Note: ESC[;5M (interpreted as 'ESC[5M') is *not* equivalent to ESC[5;M (interpreted as 'ESC[5;1M') which is ultimately interpreted as 'ESC[1M').

In the syntax descriptions below, parameters are represented as '#' or '#1;#2'.

ANSI Control Functions

The following paragraphs specify the ANSI control functions implemented by the Sun console. Each description gives:

- the control sequence syntax
- the hex equivalent of control characters where applicable
- the control function name and ANSI or Sun abbreviation (if any).
- description of parameters required, if any
- description of the control function
- for functions which set a mode, the initial setting of the mode. The initial settings can be restored with the SUNRESET escape sequence.

Control Character Functions**control-G (0x7) Bell (BEL)**

The Sun Workstation Model 100 and 100U is not equipped with an audible bell. It 'rings the bell' by flashing the entire screen. The Sun-2 models have an audible bell which beeps. The window system flashes the window.

control-H (0x8) Backspace (BS)

The cursor moves one position to the left on the current line. If it is already at the left edge of the screen, nothing happens.

control-I (0x9) Tab (TAB)

The cursor moves right on the current line to the next tab stop. The tab stops are fixed at every multiple of 8 columns. If the cursor is already at the right edge of the screen, nothing happens; otherwise the cursor moves right a minimum of one and a maximum of eight character positions.

control-J (0xA) Line-feed (LF)

The cursor moves down one line, remaining at the same character position on the line. If the cursor is already at the bottom line, the screen either scrolls up or "wraps around" depending on the setting of an internal variable *S* (initially 1) which can be changed by the ESC[*r* control sequence. If *S* is greater than zero, the entire screen (including the cursor) is scrolled up by *S* lines before executing the line-feed. The top *S* lines scroll off the screen and are lost. *S* new blank lines scroll onto the bottom of the screen. After scrolling, the line-feed is executed by moving the cursor down one line.

If *S* is zero, 'wrap-around' mode is entered. 'ESC [1 r' exits back to scroll mode. If a line-feed occurs on the bottom line in wrap mode, the cursor goes to the same character position in the top line of the screen. When any line-feed occurs, the line that the cursor moves to is cleared. This means that no scrolling occurs. Wrap-around mode is not implemented in the window system.

The screen scrolls as fast as possible depending on how much data is backed up waiting to be printed. Whenever a scroll must take place and the console is in normal scroll mode ('ESC [1 r'), it scans the rest of the data awaiting printing to see how many line-feeds occur in it. This scan stops when any control character from the set {VT, FF, SO, SI, DLE, DC1, DC2, DC3, DC4, NAK, SYN, ETB, CAN, EM, SUB, ESC, FS, GS, RS, US} is found. At that point, the screen is scrolled by *N* lines ($N \geq 1$) and processing continues. The scanned text is still processed normally to fill in the newly created lines. This results in much faster scrolling with scrolling as long as no escape codes or other control characters are intermixed with the text.

See also the discussion of the 'Set scrolling' (ESC[*r*) control function below.

control-K (0xB) Reverse Line-feed

The cursor moves up one line, remaining at the same character position on the line. If the cursor is already at the top line, nothing happens.

control-L (0xC) Form-feed (FF)

The cursor is positioned to the Home position (upper-left corner) and the entire screen is cleared.

control-M (0xD) Return (CR)

The cursor moves to the leftmost character position on the current line.

Escape Sequence Functions**control-[(0x1B) Escape (ESC)**

This is the escape character. Escape initiates a multi-character control sequence.

ESC[#@ Insert Character (ICH)

Takes one parameter, # (default 1). Inserts # spaces at the current cursor position. The tail of the current line starting at the current cursor position inclusive is shifted to the right by # character positions to make room for the spaces. The rightmost # character positions shift off the line and are lost. The position of the cursor is unchanged.

- ESC[#A** **Cursor Up (CUU)**
Takes one parameter, # (default 1). Moves the cursor up # lines. If the cursor is fewer than # lines from the top of the screen, moves the cursor to the topmost line on the screen. The character position of the cursor on the line is unchanged.
- ESC[#B** **Cursor Down (CUD)**
Takes one parameter, # (default 1). Moves the cursor down # lines. If the cursor is fewer than # lines from the bottom of the screen, move the cursor to the last line on the screen. The character position of the cursor on the line is unchanged.
- ESC[#C** **Cursor Forward (CUF)**
Takes one parameter, # (default 1). Moves the cursor to the right by # character positions on the current line. If the cursor is fewer than # positions from the right edge of the screen, moves the cursor to the rightmost position on the current line.
- ESC[#D** **Cursor Backward (CUB)**
Takes one parameter, # (default 1). Moves the cursor to the left by # character positions on the current line. If the cursor is fewer than # positions from the left edge of the screen, moves the cursor to the leftmost position on the current line.
- ESC[#E** **Cursor Next Line (CNL)**
Takes one parameter, # (default 1). Positions the cursor at the leftmost character position on the #-th line below the current line. If the current line is less than # lines from the bottom of the screen, positions the cursor at the leftmost character position on the bottom line.
- ESC[#1;#2f** **Horizontal And Vertical Position (HVP)**
or
ESC[#1;#2H **Cursor Position (CUP)**
Takes two parameters, #1 and #2 (default 1, 1). Moves the cursor to the #2-th character position on the #1-th line. Character positions are numbered from 1 at the left edge of the screen; line positions are numbered from 1 at the top of the screen. Hence, if both parameters are omitted, the default action moves the cursor to the home position (upper left corner). If only one parameter is supplied, the cursor moves to column 1 of the specified line.
- ESC[J** **Erase in Display (ED)**
Takes no parameters. Erases from the current cursor position inclusive to the end of the screen. In other words, erases from the current cursor position inclusive to the end of the current line and all lines below the current line. The cursor position is unchanged.
- ESC[K** **Erase in Line (EL)**
Takes no parameters. Erases from the current cursor position inclusive to the end of the current line. The cursor position is unchanged.
- ESC[#L** **Insert Line (IL)**
Takes one parameter, # (default 1). Makes room for # new lines starting at the current line by scrolling down by # lines the portion of the screen from the current line inclusive to the bottom. The # new lines at the cursor are filled with spaces; the bottom # lines shift off the bottom of the screen and are lost. The position of the cursor on the screen is unchanged.
- ESC[#M** **Delete Line (DL)**
Takes one parameter, # (default 1). Deletes # lines beginning with the current line. The portion of the screen from the current line inclusive to the bottom is scrolled upward by # lines. The # new lines scrolling onto the bottom of the screen are filled with spaces; the # old lines beginning at the cursor line are deleted. The position of the cursor on the screen is unchanged.
- ESC[#P** **Delete Character (DCH)**
Takes one parameter, # (default 1). Deletes # characters starting with the current cursor position. Shifts to the left by # character positions the tail of the current line from the current cursor position inclusive to the end of the line. Blanks are shifted into the rightmost # character positions. The position of the cursor on the screen is unchanged.

ESC[#m **Select Graphic Rendition (SGR)**

Takes one parameter, # (default 0). Note: unlike most escape sequences, the parameter defaults to zero if omitted. Invokes the graphic rendition specified by the parameter. All following printing characters in the data stream are rendered according to the parameter until the next occurrence of this escape sequence in the data stream. Currently only two graphic renditions are defined:

- 0 Normal rendition.
- 7 Negative (reverse) image.

Negative image displays characters as white-on-black if the screen mode is currently black-on-white, and vice-versa. Any non-zero value of # is currently equivalent to 7 and selects the negative image rendition.

ESC[p **Black On White (SUNBOW)**

Takes no parameters. Sets the screen mode to black-on-white. If the screen mode is already black-on-white, has no effect. In this mode spaces display as solid white, other characters as black-on-white. The cursor is a solid black block. Characters displayed in negative image rendition (see 'Select Graphic Rendition' above) is white-on-black in this mode. This is the initial setting of the screen mode on reset.

ESC[q **White On Black (SUNWOB)**

Takes no parameters. Sets the screen mode to white-on-black. If the screen mode is already white-on-black, has no effect. In this mode spaces display as solid black, other characters as white-on-black. The cursor is a solid white block. Characters displayed in negative image rendition (see 'Select Graphic Rendition' above) is black-on-white in this mode. The initial setting of the screen mode on reset is the alternative mode, black on white.

ESC[#r **Set scrolling (SUNSCRL)**

Takes one parameter, # (default 0). Sets to # an internal register which determines how many lines the screen scrolls up when a line-feed function is performed with the cursor on the bottom line. A parameter of 2 or 3 introduces a small amount of "jump" when a scroll occurs. A parameter of 34 clears the screen rather than scrolling. The initial setting is 1 on reset.

A parameter of zero initiates "wrap mode" instead of scrolling. In wrap mode, if a linefeed occurs on the bottom line, the cursor goes to the same character position in the top line of the screen. When any linefeed occurs, the line that the cursor moves to is cleared. This means that no scrolling ever occurs. 'ESC [1 r' exits back to scroll mode.

For more information, see the description of the Line-feed (CTRL-J) control function above.

ESC[s **Reset terminal emulator (SUNRESET)**

Takes no parameters. Resets all modes to default, restores current font from PROM. Screen and cursor position are

4014 TERMINAL EMULATION

The PROM monitor for Sun models 100U and 150U provides the Sun Workstation with the capability to emulate a subset of the Tektronix 4014 terminal. This feature does not exist in other Sun PROMs and will be removed from models 100U and 150U in future Sun releases. **tektool(1)** provides Tektronix 4014 terminal emulation and should be used instead of relying on the capabilities of the PROM monitor.

FILES

/dev/console

SEE ALSO

tektool(1) **kb(4M)**, **pty(4)**, **termio(4)**, **ttcompat(4M)**, **ldterm(4M)**, **win(4S)**, **zs(4S)**

ANSI Standard X3.64, "Additional Controls for Use with ASCII", Secretariat: CBEMA, 1828 L St., N.W., Washington, D.C. 20036.

BUGS

TIOCCONS should be restricted to the owner of /dev/console.

NAME

des – DES encryption chip interface

CONFIG — SUN-3 SYSTEM

des0 at obio ? csr 0x1c0000

#include <sys/des.h>

CONFIG — SUN-2 SYSTEM

des0 at virtual ? csr 0xee1800

#include <sys/des.h>

DESCRIPTION

The **des** driver provides a high level interface to the AmZ8068 Data Ciphering Processor, a hardware implementation of the NBS Data Encryption Standard.

The high level interface provided by this driver is hardware independent and could be shared by future drivers in other systems.

The interface allows access to two modes of the DES algorithm: Electronic Code Book (ECB) and Cipher Block Chaining (CBC). All access to the DES driver is through **ioctl(2)** calls rather than through reads and writes; all encryption is done in-place in the user's buffers.

IOCTLS

The **ioctl**s provided are:

DESIOCBLOCK

This call encrypts/decrypts an entire buffer of data, whose address and length are passed in the 'struct **desparams**' addressed by the argument. The length must be a multiple of 8 bytes.

DESIOCQUICK

This call encrypts/decrypts a small amount of data quickly. The data is limited to **DES_QUICKLEN** bytes, and must be a multiple of 8 bytes. Rather than being addresses, the data is passed directly in the 'struct **desparams**' argument.

FILES

/dev/des

SEE ALSO

des(1), **des_crypt(3)**

Federal Information Processing Standards Publication 46

AmZ8068 DCP Product Description, Advanced Micro Devices

NAME

dkio – generic disk control operations

DESCRIPTION

All Sun disk drivers support a set of ioctl's for disk formatting and labeling operations. Basic to these ioctl's are the definitions in <sun/dkio.h>:

```

/*
 * Structures and definitions for disk io control commands
 */
/* Disk identification */
struct dk_info {
    int     dki_ctlr;           /* controller address */
    short   dki_unit;          /* unit (slave) address */
    short   dki_ctype;         /* controller type */
    short   dki_flags;         /* flags */
};
/* controller types */
#define DKC_UNKNOWN      0
#define DKC_SMD2180     1
#define DKC_DSD5215     5
#define DKC_XY450       6
#define DKC_ACB4000     7
#define DKC_MD21        8
#define DKC_CSS         12
#define DKC_NEC765     13    /* floppy on Sun386i */
/* flags */
#define DKI_BAD144      0x01  /* use DEC std 144 bad sector fwding */
#define DKI_MAPTRK     0x02  /* controller does track mapping */
#define DKI_FMTTRK     0x04  /* formats only full track at a time */
#define DKI_FMTVOL     0x08  /* formats only full volume at a time */
/* Definition of a disk's geometry */
struct dk_geom {
    unsigned short  dkg_ncyl;    /* # of data cylinders */
    unsigned short  dkg_acyl;    /* # of alternate cylinders */
    unsigned short  dkg_bcyl;    /* cyl offset (for fixed head area) */
    unsigned short  dkg_nhead;   /* # of heads */
    unsigned short  dkg_bhead;   /* head offset (for Larks, etc.) */
    unsigned short  dkg_nsect;   /* # of sectors per track */
    unsigned short  dkg_intrlv;  /* interleave factor */
    unsigned short  dkg_gap1;    /* gap 1 size */
    unsigned short  dkg_gap2;    /* gap 2 size */
    unsigned short  dkg_apc;     /* alternates per cyl (SCSI only) */
    unsigned short  dkg_extra[9]; /* for compatible expansion */
};
/* disk io control commands */
#define DKIOCGGEOM   _IOR(d, 2, struct dk_geom)    /* Get geometry */
#define DKIOCSGEOM   _IOW(d, 3, struct dk_geom)    /* Set geometry */
#define DKIOCGPART   _IOR(d, 4, struct dk_map)     /* Get partition info */
#define DKIOCSPART   _IOW(d, 5, struct dk_map)     /* Set partition info */
#define DKIOCINFO    _IOR(d, 8, struct dk_info)    /* Get info */

```

The DKIOCINFO ioctl returns a `dk_info` structure which tells the type of the controller and attributes about how bad-block processing is done on the controller. The DKIOCGPART and DKIOCSPART get and set the controller's current notion of the partition table for the disk (without changing the partition table on the

disk itself), while the **DKIOCGGGEOM** and **DKIOCSGGEOM** ioctl's do similar things for the per-drive geometry information.

SEE ALSO

ip(4P), sd(4S), xy(4S)

NAME

drum – paging device

CONFIG

None; included with standard system.

SYNOPSIS

```
#include <fcntl.h>
```

```
open("/dev/drum", mode);
```

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

CONFIG — SUN-2 SYSTEM

device ec0 at mbmem ? csr 0xe0000 priority 3

device ec1 at mbmem ? csr 0xe2000 priority 3

DESCRIPTION

The **ec** interface provides access to a 10 Mb/s Ethernet network through a 3COM controller. For a general description of network interfaces see **if(4N)**.

The hardware consumes 8 kilobytes of Multibus memory space. This memory is used for internal buffering by the board. The board starts at standard addresses 0xE0000 or 0xE2000. The board must be configured for interrupt level 3.

The interface software implements an exponential backoff algorithm when notified of a collision on the cable.

The interface handles the Internet protocol family, with the interface address maintained in Internet format. The Address Resolution Protocol **arp(4P)** is used to map 32-bit Internet addresses used in **inet(4F)** to the 48-bit addresses used on the Ethernet.

DIAGNOSTICS

ec%d: Ethernet jammed

After 16 failed transmissions and backoffs using the exponential backoff algorithm, the packet was dropped.

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

arp(4P), **if(4N)**, **inet(4F)**

BUGS

The interface hardware is not capable of talking to itself, making diagnosis more difficult.

NAME

fb – driver for Sun console frame buffer

CONFIG

None; included in standard system.

DESCRIPTION

The **fb** driver provides indirect access to a Sun graphics controller board. It is an indirect driver for the Sun workstation console's frame buffer. At boot time, the workstation's frame buffer device is determined from information from the PROM monitor and set to be the one that **fb** will indirect to. The device driver for the console's frame buffer must be configured into the kernel so that this indirect driver can access it.

The idea behind this driver is that user programs can open a known device, query its characteristics and access it in a device dependent way, depending on the type. **fb** redirects **open(2V)**, **close(2)**, **ioctl(2)**, and **mmap(2)** calls to the real frame buffer. All of the Sun frame buffers support the same general interface; see **fbio(4S)**

FILES

/dev/fb

SEE ALSO

close(2), **ioctl(2)**, **mmap(2)**, **open(2V)**, **bwone(4S)**, **bwtwo(4S)**, **cgone(4S)**, **cgtwo(4S)**, **fbio(4S)**, **gpone(4S)**

NAME

fbio – general properties of frame buffers

DESCRIPTION

All of the Sun frame buffers support the same general interface. Each responds to a **FBIOTYPE** ioctl which returns information in a structure defined in `<sun/fbio.h>`:

```

struct fbtype {
    int    fb_type;        /* as defined below */
    int    fb_height;     /* in pixels */
    int    fb_width;      /* in pixels */
    int    fb_depth;      /* bits per pixel */
    int    fb_cmsize;     /* size of color map (entries) */
    int    fb_size;       /* total size in bytes */
};

#define FBTYPE_SUN1BW      0
#define FBTYPE_SUN1COLOR  1
#define FBTYPE_SUN2BW      2
#define FBTYPE_SUN2COLOR  3
#define FBTYPE_SUN2GP      4
#define FBTYPE_SUN3COLOR  6
#define FBTYPE_SUN4COLOR  8

```

Each device has an **FBTYPE** which is used by higher-level software to determine how to perform raster-op and other functions. Each device is used by opening it, doing an **FBIOTYPE** ioctl to see which frame buffer type is present, and thereby selecting the appropriate device-management routines.

Full-fledged frame buffers (that is, those that run SunView) implement an **FBIOPIXRECT** ioctl, which returns a **pixrect**. This call is made only from inside the kernel. The returned **pixrect** is used by **win(4S)** for cursor tracking and colormap loading.

FBIOSVIDEO and **FBIOGVIDEO** are general-purpose ioctls for controlling possible video features of frame buffers. They are defined in `<sun/fbio.h>`. These ioctls either set or return the value of a flags integer. At this point, only the **FBVIDEO_ON** option is available, controlled by **FBIOSVIDEO**. **FBIOGVIDEO** returns the current video state.

The **FBIOSATTR** and **FBIOGATTR** ioctls allow access to special features of newer frame buffers. They use the following structures as defined in `<sun/fbio.h>`:

```

#define FB_ATTR_NDEVSPECIFIC  8    /* no. of device specific values */
#define FB_ATTR_NEMUTYPES  4    /* no. of emulation types */
struct fbsattr {
    int    flags;          /* misc flags */
#define FB_ATTR_AUTOINIT      1    /* emulation auto init flag */
#define FB_ATTR_DEVSPECIFIC  2    /* dev. specific stuff valid flag */
    int    emu_type;       /* emulation type (-1 if unused) */
    int    dev_specific[FB_ATTR_NDEVSPECIFIC]; /* catchall */
};
struct fbgattr {
    int    real_type;      /* real device type */
    int    owner;          /* PID of owner, 0 if myself */
    struct fbtype fbtype;  /* fbtype info for real device */
    struct fbsattr sattr; /* see above */
    int    emu_types[FB_ATTR_NEMUTYPES]; /* possible emulations */
                                           /* (-1 if unused) */
};

```

SEE ALSO

mmap(2), bwone(4S), bwtwo(4S), cgfour(4S), cgone(4S), cgtwo(4S), fb(4S), gpone(4S), win(4S)

BUGS

FBIOSATTR and FBIOGATTR are only supported by the cgfour(4S) frame buffer.

The FBVIDEO_ON flag may be incorrect for Sun-1 system black and white frame buffers; see bwone(4S).

NAME

fd – Disk driver for Floppy Disk Controllers

CONFIG

controller fdc0 at atmem ? csr 0x1000 dmachan 2 irq 6 priority 2
disk fd0 at fdc0 drive 0 flags 0

AVAILABILITY

Sun386i systems only.

DESCRIPTION

The block-files access the disk using the system's normal buffering mechanism and may be read and written without regard to physical disk records. There is also a 'raw' interface that provides for direct transmission between the disk and the user's read or write buffer. A single read or write call usually results in one I/O operation; therefore raw I/O is considerably more efficient when many words are transmitted. The names of the raw files conventionally begin with an extra 'r.'

Disk Support

The **fd0** partition on a floppy disk is normally used for the file system.

FILES

1.44 MB Floppy Disk Drives:

/dev/fd0a block file
/dev/fd0c block file /dev/rfd0a raw file /dev/rfd0c raw file

720 K Floppy Disk Drives:

/dev/fd10a block file
/dev/fd10c block file
/dev/rfd10a raw file
/dev/rfd10c raw file

SEE ALSO

dkio(4S)

DIAGNOSTICS

fd drv %d, trk %d: %s

A command such as read or write encountered a format-related error condition. The value of %s is derived from the error number given by the controller, indicating the nature of the error. The track number is relative to the beginning of the partition involved.

fd drv %d, blk %d: %s

A command such as read or write encountered an error condition related to I/O. The value of %s is derived from the error number returned by the controller and indicates the nature of the error. The block number is relative to the start of the partition involved.

fd controller: %s

An error occurred in the controller. The value of %s is derived from the status returned by the controller and specifies the error encountered.

fd(%d): %s please insert

I/O was attempted while the floppy drive door was not latched. The value of %s indicates which disk was expected to be in the drive.

NAME

filio – ioctls that operate directly on files, file descriptors, and sockets

SYNOPSIS

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

DESCRIPTION

The IOCTL's listed in this manual page apply directly to files, file descriptors, and sockets, independent of any underlying device or protocol.

Note: the `fcntl(2V)` system call is the primary method for operating on file descriptors as such, rather than on the underlying files.

IOCTLS for File Descriptors

FIOCLEX The argument is ignored. Set the close-on-exec flag for the file descriptor passed to `ioctl`. This flag is also manipulated by the `F_SETFD` command of `fcntl(2V)`.

FIONCLEX The argument is ignored. Clear the close-on-exec flag for the file descriptor passed to `ioctl`.

IOCTLS for Files

FIONREAD The argument is a pointer to a `long`. Set the value of that `long` to the number of immediately readable characters from whatever the descriptor passed to `ioctl` refers to. This works for files, pipes, sockets, and terminals.

FIONBIO The argument is a pointer to an `int`. Set or clear non-blocking I/O. If the value of that `int` is a 1 (one) the descriptor is set for non-blocking I/O. If the value of that `int` is a 0 (zero) the descriptor is cleared for non-blocking I/O.

FIOASYNC The argument is a pointer to an `int`. Set or clear asynchronous I/O. If the value of that `int` is a 1 (one) the descriptor is set for asynchronous I/O. If the value of that `int` is a 0 (zero) the descriptor is cleared for asynchronous I/O.

FIOSETOWN The argument is a pointer to an `int`. Set the process-group ID that will subsequently receive `SIGIO` or `SIGURG` signals for the object referred to by the descriptor passed to `ioctl` to the value of that `int`.

FIOGETOWN The argument is a pointer to an `int`. Set the value of that `int` to the process-group ID that is receiving `SIGIO` or `SIGURG` signals for the object referred to by the descriptor passed to `ioctl`.

SEE ALSO

`ioctl(2)`, `fcntl(2V)`, `getsockopt(2)`, `sockio(4)`

NAME

fpa – Sun-3 floating-point accelerator

CONFIG — SUN-3 SYSTEM

fpa0 at virtual ? csr 0xe0000000

DESCRIPTION

The **fpa** is a floating point accelerator available for Sun-3 systems.

The **fpa** device driver manipulates the 32 contexts supported by the floating point accelerator hardware.

The **open(2V)**, **close(2)**, and **ioctl(2)** system calls generally produce errors when applied to this device. However, since the 32 **fpa** contexts are allocated and deallocated automatically, and no user program needs to access **fpa** registers explicitly, such calls to the device are generally unnecessary. Access to the device is normally provided at compile time by a compiler option, such as the **-ffpa** option to **cc(1V)**.

FILES

/dev/fpa

SEE ALSO

cc(1V), **close(2)**, **ioctl(2)**, **open(2V)**

NAME

gpone – Sun-3/Sun-2 graphics processor

CONFIG — SUN-3 SYSTEM

device gpone0 at vme24d16 ? csr

CONFIG — SUN-2 SYSTEM

device gpone0 at vme24 ? csr

DESCRIPTION

The **gpone** interface provides access to the optional Graphics Processor Board (GP).

The hardware consumes 64 kilobytes of VME bus address space. The GP board starts at standard address 0x210000 and must be configured for interrupt level 3.

IOCTLS

The graphics processor responds to a number of ioctl calls as described here. One of the calls uses a **gp1fbinfo** structure that looks like this:

```

struct gp1fbinfo {
    int          fb_vmeaddr; /* physical color board address */
    int          fb_hwwidth; /* fb board width */
    int          fb_hwheight; /* fb board height */
    int          addrdelta; /* phys addr diff between fb and gp */
    caddr_t     fb_ropaddr; /* cg2 va thru kernelmap */
    int          fbunit; /* fb unit to use for a,b,c,d */
};

```

The ioctl call looks like this:

```

ioctl(file, request, argp)
int file, request;

```

argp is defined differently for each GP ioctl request and is specified in the descriptions below.

The following ioctl commands provide for transferring data between the graphics processor and color boards and processes.

GP1IO_PUT_INFO

Passes information about the frame buffer into driver. **argp** points to a **struct gp1fbinfo** which is passed to the driver.

GP1IO_GET_STATIC_BLOCK

Hands out a static block from the GP. **argp** points to an **int** which is returned from the driver.

GP1IO_FREE_STATIC_BLOCK

Frees a static block from the GP. **argp** points to an **int** which is passed to the driver.

GP1IO_GET_GBUFFER_STATE

Checks to see if there is a buffer present on the GP. **argp** points to an **int** which is returned from the driver.

GP1IO_CHK_GP

Restarts the GP if necessary. **argp** points to an **int** which is passed to the driver.

GP1IO_GET_RESTART_COUNT

Returns the number of restarts of a GP since power on. Needed to differentiate SIGXCPU calls in user processes. **argp** points to an **int** which is returned from the driver.

GP1IO_REDIRECT_DEVFB

Configures **/dev/fb** to talk to a graphics processor device. **argp** points to an **int** which is passed to the driver.

GP1IO_GET_REQDEV

Returns the requested minor device. **argp** points to a **dev_t** which is returned from the driver.

GPIO_GET_TRUMINORDEV

Returns the true minor device. `argp` points to a `char` which is returned from the driver.

The graphics processor driver also responds to the `FBIOGTYPE`, `ioctl` which a program can use to inquire as to the characteristics of the display device, the `FBIOGINFO`, `ioctl` for passing generic information, and the `FBIOGPIXRECT` `ioctl` so that SunWindows can run on it. See `fbio(4S)`.

FILES

`/dev/gpone[0-3][abcd]`
`/usr/include/sun/gpio.h`
`/usr/include/pixrect/{gp1cmds.h,gp1reg.h,gp1var.h}`
`/dev/fb`

SEE ALSO

`fbio(4S)`, `mmap(2)`, `gpconfig(8)`
SunCGI Reference Manual

DIAGNOSTICS

The Graphics Processor has been restarted. You may see display garbage as a result.

NAME

icmp – Internet Control Message Protocol

SYNOPSIS

```
#include <sys/socket.h>
#include <netinet/in.h>
#include <netinet/ip_icmp.h>

s = socket(AF_INET, SOCK_RAW, proto);
```

DESCRIPTION

ICMP is the error and control message protocol used by the Internet protocol family. It is used by the kernel to handle and report errors in protocol processing. It may also be accessed through a “raw socket” for network monitoring and diagnostic functions. The protocol number for ICMP, used in the *proto* parameter to the socket call, can be obtained from `getprotobyname` (see `getprotoent(3N)`). ICMP 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(2V)` or `recv(2)` and `write(2V)` or `send(2)` system calls may be used).

Outgoing packets automatically have an Internet Protocol (IP) header prepended to them. Incoming packets are provided to the holder of a raw socket with the IP header and options intact.

ICMP is an unreliable datagram protocol layered above IP. It is used internally by the protocol code for various purposes including routing, fault isolation, and congestion control. Receipt of an ICMP “redirect” message will add a new entry in the routing table, or modify an existing one. ICMP messages are routinely sent by the protocol code. Received ICMP messages may be reflected back to users of higher-level protocols such as TCP or UDP as error returns from system calls. A copy of all ICMP message received by the system is provided using the ICMP raw socket.

ERRORS

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

`connect(2)`, `read(2V)`, `recv(2)`, `send(2)`, `write(2V)`, `getprotoent(3N)`, `inet(4F)`, `ip(4P)`, `routing(4N)`

Postel, Jon, *Internet Control Message Protocol — DARPA Internet Program Protocol Specification*, RFC 792, Network Information Center, SRI International, Menlo Park, Calif., September 1981. (Sun 800-1064-01)

BUGS

Replies to ICMP “echo” messages which are source routed are not sent back using inverted source routes, but rather go back through the normal routing mechanisms.

NAME

ie – Intel 10 Mb/s Ethernet interface

CONFIG — SUN-3 SYSTEM

device ie0 at obio ? csr 0xc0000 priority 3
 device ie1 at vme24d16 ? csr 0xe88000 priority 3 vector ieintr 0x75
 device ie1 at vme24d16 ? csr 0x31ff02 priority 3 vector ieintr 0x75

CONFIG — SUN-2 SYSTEM

device ie0 at obio 2 csr 0x7f0800 priority 3
 device ie1 at vme24 ? csr 0xe88000 priority 3 vector ieintr 0x75
 device ie0 at mbmem ? csr 0x88000 priority 3
 device ie1 at mbmem ? csr 0x8c000 flags 2 priority 3

CONFIG — Sun386i SYSTEM

device ie0 at obmem ? csr 0xD0000000 irq 21 priority 3

DESCRIPTION

The **ie** interface provides access to a 10 Mb/s Ethernet network through the Intel 82586 controller chip. For a general description of network interfaces see **if(4N)**.

In the Sun-3 lines above, the first line specifies the CPU-board-resident Intel Ethernet interface. The second line specifies a Multibus Intel Ethernet interface for use with a VME adapter. The third line specifies the Intel Ethernet interface present on a Sun-3 Eurocard board.

In the Sun-2 lines above, the first line specifies the CPU-board-resident Intel Ethernet interface on a Sun-2/50 or Sun-2/160 system. The second line specifies a Multibus Intel Ethernet controller for use with a VME adapter on these systems. The third line specifies the first Multibus Intel Ethernet controller for a Sun-2/120 or Sun-2/170 system. The fourth line specifies the second such controller for these systems.

The Sun386i line above specifies the CPU-board-resident Intel Ethernet interface.

SEE ALSO

if(4N)

DIAGNOSTICS

There are too many driver messages to list them all individually here. Some of the more common messages and their meanings follow.

ie%d: Ethernet jammed

Network activity has become so intense that sixteen successive transmission attempts failed, and the 82586 gave up on the current packet. Another possible cause of this message is a noise source somewhere in the network, such as a loose transceiver connection.

ie%d: no carrier

The 82586 has lost input to its carrier detect pin while trying to transmit a packet, causing the packet to be dropped. Possible causes include an open circuit somewhere in the network and noise on the carrier detect line from the transceiver.

ie%d: lost interrupt: resetting

The driver and 82586 chip have lost synchronization with each other. The driver recovers by resetting itself and the chip.

ie%d: iebark reset

The 82586 failed to complete a watchdog timeout command in the allotted time. The driver recovers by resetting itself and the chip.

ie%d: WARNING: requeueing

The driver has run out of resources while getting a packet ready to transmit. The packet is put back on the output queue for retransmission after more resources become available.

ie%d: panic: scb overwritten

The driver has discovered that memory that should remain unchanged after initialization has become corrupted. This error usually is a symptom of a bad 82586 chip.

NAME

if – general properties of network interfaces

DESCRIPTION

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 with 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 can 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 (for example, 10Mb/s Ethernets using arp(4P)), the entire address specified in the ioctl is used.

The following ioctl calls may be used to manipulate network interfaces. Unless specified otherwise, the request takes an ifreq structure as its parameter. This structure has the form

```

struct ifreq {
    char    ifr_name[16];          /* name of interface (e.g. "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 */
};

```

SIOCADDMULTI Enable a multicast address for the interface. A maximum of 64 multicast addresses may be enabled for any given interface.

SIOCDELMULTI Disable a previously set multicast address.

SIOCSPROMISC Toggle promiscuous mode.

SEE ALSO

arp(4P), ec(4S), lo(4)

NAME

inet – Internet protocol family

SYNOPSIS

options INET

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

```
#include <netinet/in.h>
```

DESCRIPTION

The Internet protocol family implements a collection of protocols which are centered around the *Internet Protocol* (IP) and which share a common address format. The Internet family provides protocol support for the **SOCK_STREAM**, **SOCK_DGRAM**, and **SOCK_RAW** socket types.

PROTOCOLS

The Internet protocol family is comprised of the Internet Protocol (IP), the Address Resolution Protocol (ARP), the Internet Control Message Protocol (ICMP), the Transmission Control Protocol (TCP), and the User Datagram Protocol (UDP).

TCP is used to support the **SOCK_STREAM** abstraction while UDP is used to support the **SOCK_DGRAM** abstraction; see **tcp(4P)** and **udp(4P)**. A raw interface to IP is available by creating an Internet socket of type **SOCK_RAW**; see **ip(4P)**. ICMP is used by the kernel to handle and report errors in protocol processing. It is also accessible to user programs; see **icmp(4P)**. ARP is used to translate 32-bit IP addresses into 48-bit Ethernet addresses; see **arp(4P)**.

The 32-bit IP address is divided into network number and host number parts. It is frequency-encoded; the most-significant bit is zero in Class A addresses, in which the high-order 8 bits are the network number. Class B addresses have their high order two bits set to 10 and use the high-order 16 bits as the network number field. Class C addresses have a 24-bit network number part of which the high order three bits are 110. Sites with a cluster of local networks may chose to use a single network number for the cluster; this is done by using subnet addressing. The local (host) portion of the address is further subdivided into subnet number and host number parts. Within a subnet, each subnet appears to be an individual network; externally, the entire cluster appears to be a single, uniform network requiring only a single routing entry. Subnet addressing is enabled and examined by the following **ioctl(2)** commands on a datagram socket in the Internet domain; they have the same form as the **SIOCIFADDR** command (see **intro(4N)**).

SIOCSIFNETMASK Set interface network mask. The network mask defines the network part of the address; if it contains more of the address than the address type would indicate, then subnets are in use.

SIOCGIFNETMASK Get interface network mask.

ADDRESSING

IP addresses are four byte quantities, stored in network byte order (on Sun386i systems these are word and byte reversed).

Sockets in the Internet protocol family use the following addressing structure:

```
struct sockaddr_in {
    short  sin_family;
    u_short sin_port;
    struct in_addr sin_addr;
    char   sin_zero[8];
};
```

Library routines are provided to manipulate structures of this form; see **intro(3N)**.

The **sin_addr** field of the **sockaddr_in** structure specifies a local or remote IP address. Each network interface has its own unique IP address. The special value **INADDR_ANY** may be used in this field to effect “wildcard” matching. Given in a **bind(2)** call, this value leaves the local IP address of the socket unspecified, so that the socket will receive connections or messages directed at any of the valid IP addresses of the system. This can prove useful when a process neither knows nor cares what the local IP

address is or when a process wishes to receive requests using all of its network interfaces. The `sockaddr_in` structure given in the `bind(2)` call must specify an `in_addr` value of either `IPADDR_ANY` or one of the system's valid IP addresses. Requests to bind any other address will elicit the error `EADDRNOTAVAIL`. When a `connect(2)` call is made for a socket that has a wildcard local address, the system sets the `sin_addr` field of the socket to the IP address of the network interface that the packets for that connection are routed via.

The `sin_port` field of the `sockaddr_in` structure specifies a port number used by TCP or UDP. The local port address specified in a `bind(2)` call is restricted to be greater than `IPPORT_RESERVED` (defined in `<netinet/in.h>`) unless the creating process is running as the super-user, providing a space of protected port numbers. In addition, the local port address must not be in use by any socket of same address family and type. Requests to bind sockets to port numbers being used by other sockets return the error `EADDRINUSE`. If the local port address is specified as 0, then the system picks a unique port address greater than `IPPORT_RESERVED`. A unique local port address is also picked when a socket which is not bound is used in a `connect(2)` or `sendto` (see `send(2)`) call. This allows programs which do not care which local port number is used to set up TCP connections by simply calling `socket(2)` and then `connect(2)`, and to send UDP datagrams with a `socket(2)` call followed by a `sendto(2)` call.

Although this implementation restricts sockets to unique local port numbers, TCP allows multiple simultaneous connections involving the same local port number so long as the remote IP addresses or port numbers are different for each connection. Programs may explicitly override the socket restriction by setting the `SO_REUSEADDR` socket option with `setsockopt` (see `getsockopt(2)`).

SEE ALSO

`bind(2)`, `connect(2)`, `getsockopt(2)`, `ioctl(2)`, `sendto(2)`, `socket(2)`, `byteorder(3N)`, `gethostent(3N)`, `getnetent(3N)`, `getprotoent(3N)`, `getservent(3N)`, `inet(3N)`, `intro(3N)`, `arp(4P)`, `icmp(4P)`, `intro(4N)`, `ip(4P)`, `tcp(4P)`, `udp(4P)`,

Network Information Center, *DDN Protocol Handbook* (3 vols.), Network Information Center, SRI International, Menlo Park, Calif., 1985.

A 4.2BSD Interprocess Communication Primer

CAVEAT

The Internet protocol support is subject to change as the Internet protocols develop. Users should not depend on details of the current implementation, but rather the services exported.

NAME

`ip` – Internet Protocol

SYNOPSIS

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

```
#include <netinet/in.h>
```

```
s = socket(AF_INET, SOCK_RAW, proto);
```

DESCRIPTION

IP is the internetwork datagram delivery protocol that is central to the Internet protocol family. Programs may use IP through higher-level protocols such as the Transmission Control Protocol (TCP) or the User Datagram Protocol (UDP), or may interface directly using a “raw socket.” See `tcp(4P)` and `udp(4P)`. The protocol options defined in the IP specification may be set in outgoing datagrams.

Raw IP sockets are connectionless and are normally used with the `sendto` and `recvfrom` calls, (see `send(2)` and `recv(2)`) although the `connect(2)` call may also be used to fix the destination for future datagrams (in which case the `read(2V)` or `recv(2)` and `write(2V)` or `send(2)` calls may be used). If `proto` is zero, the default protocol, `IPPROTO_RAW`, is used. If `proto` is non-zero, that protocol number will be set in outgoing datagrams and will be used to filter incoming datagrams. An IP header will be generated and prepended to each outgoing datagram; Received datagrams are returned with the IP header and options intact.

A single socket option, `IP_OPTIONS`, is supported at the IP level. This socket option may be used to set IP options to be included in each outgoing datagram. IP options to be sent are set with `setsockopt` (see `getsockopt(2)`). The `getsockopt(2)` call returns the IP options set in the last `setsockopt` call. IP options on received datagrams are visible to user programs only using raw IP sockets. The format of IP options given in `setsockopt` matches those defined in the IP specification with one exception: the list of addresses for the source routing options must include the first-hop gateway at the beginning of the list of gateways. The first-hop gateway address will be extracted from the option list and the size adjusted accordingly before use. IP options may be used with any socket type in the Internet family.

At the socket level, the socket option `SO_DONTROUTE` may be applied. This option forces datagrams being sent to bypass the routing step in output. Normally, IP selects a network interface to send the datagram via, and possibly an intermediate gateway, based on an entry in the routing table. See `routing(4N)`. When `SO_DONTROUTE` is set, the datagram will be sent via the interface whose network number or full IP address matches the destination address. If no interface matches, the error `ENETUNRCH` will be returned.

Datagrams flow through the IP layer in two directions: from the network `ip` to user processes and from user processes *down* to the network. Using this orientation, IP is layered *above* the network interface drivers and *below* the transport protocols such as UDP and TCP. The Internet Control Message Protocol (ICMP) is logically a part of IP. See `icmp(4P)`.

IP provides for a checksum of the header part, but not the data part of the datagram. The checksum value is computed and set in the process of sending datagrams and checked when receiving datagrams. IP header checksumming may be disabled for debugging purposes by patching the kernel variable `ipcksum` to have the value zero.

IP options in received datagrams are processed in the IP layer according to the protocol specification. Currently recognized IP options include: security, loose source and record route (LSRR), strict source and record route (SSRR), record route, stream identifier, and internet timestamp.

The IP layer will normally forward received datagrams that are not addressed to it. Forwarding is under the control of the kernel variable `ipforwarding`: if `ipforwarding` is zero, IP datagrams will not be forwarded; if `ipforwarding` is one, IP datagrams will be forwarded. `ipforwarding` is usually set to one only in machines with more than one network interface (internetwork routers). This kernel variable can be patched to enable or disable forwarding.

The IP layer will send an ICMP message back to the source host in many cases when it receives a datagram that can not be handled. A "time exceeded" ICMP message will be sent if the "time to live" field in the IP header drops to zero in the process of forwarding a datagram. A "destination unreachable" message will be sent if a datagram can not be forwarded because there is no route to the final destination, or if it can not be fragmented. If the datagram is addressed to the local host but is destined for a protocol that is not supported or a port that is not in use, a destination unreachable message will also be sent. The IP layer may send an ICMP "source quench" message if it is receiving datagrams too quickly. ICMP messages are only sent for the first fragment of a fragmented datagram and are never returned in response to errors in other ICMP messages.

The IP layer supports fragmentation and reassembly. Datagrams are fragmented on output if the datagram is larger than the maximum transmission unit (MTU) of the network interface. Fragments of received datagrams are dropped from the reassembly queues if the complete datagram is not reconstructed within a short time period.

Errors in sending discovered at the network interface driver layer are passed by IP back up to the user process.

ERRORS

A socket operation may fail with one of the following errors returned:

EACCESS	when specifying an IP broadcast destination address if the caller is not the super-user;
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;
EMSGSIZE	when sending datagram that is too large for an interface, but is not allowed be fragmented (such as broadcasts);
ENETUNREACH	when trying to establish a connection or send a datagram, if there is no matching entry in the routing table, or if an ICMP "destination unreachable" message is received.
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 fragmentation buffers or other internal data structure;
EADDRNOTAVAIL	when an attempt is made to create a socket with a local address that matches no network interface, or when specifying an IP broadcast destination address and the network interface does not support broadcast;

The following errors may occur when setting or getting IP options:

EINVAL	An unknown socket option name was given.
EINVAL	The IP option field was improperly formed; an option field was shorter than the minimum value or longer than the option buffer provided.

SEE ALSO

connect(2), getsockopt(2), read(2V), recv(2), send(2), write(2V), icmp(4P), inet(4F) routing(4N), tcp(4P), udp(4P),

Postel, Jon, "*Internet Protocol - DARPA Internet Program Protocol Specification*," RFC 791, Network Information Center, SRI International, Menlo Park, Calif., September 1981. (Sun 800-1063-01)

BUGS

Raw sockets should receive ICMP error packets relating to the protocol; currently such packets are simply discarded.

Users of higher-level protocols such as TCP and UDP should be able to see received IP options.

NAME

kb – Sun keyboard STREAMS module

CONFIG

pseudo-device *kbnumber*

SYNOPSIS

```
#include <sys/stream.h>
#include <sys/stropt.h>
#include <sundev/vuid_event.h>
#include <sundev/kbio.h>
#include <sundev/kbd.h>

ioctl(fd, I_PUSH, "kb");
```

DESCRIPTION

The **kb** STREAMS module processes byte streams generated by Sun keyboards attached to a CPU serial or parallel port. Definitions for altering keyboard translation, and reading events from the keyboard, are in `<sundev/kbio.h>` and `<sundev/kbd.h>`. *number* specifies the maximum number of keyboards supported by the system.

kb recognizes which keys have been typed using a set of tables for each known type of keyboard. Each translation table is an array of 128 bytes (unsigned characters). If a character value is less than 0x80, it is treated as an ASCII character (perhaps with the META bit included). Higher values indicate special characters that invoke more complicated actions.

Keyboard Translation State

The keyboard can be in one of the following translation modes:

TR_NONE	Keyboard translation is turned off and up/down key codes are reported.
TR_ASCII	ASCII codes are reported.
TR_EVENT	firm_events (see <i>The SunView System Programmer's Guide — Appendix: Writing a Virtual User Input Device Driver</i>) are reported.
TR_UNTRANS_EVENT	firm_events containing unencoded keystation codes are reported for all input events within the window system.

Keyboard Translation-Table Entries

All instances of the **kb** module share five translation tables used to convert raw keystation codes to event values. The tables are:

Unshifted	Used when a key is depressed and no shifts are in effect.
Shifted	Used when a key is depressed and a Shift key is being held down.
Caps Lock	Used when a key is depressed and Caps Lock is in effect.
Controlled	Used when a key is depressed and the Control key is being held down (regardless of whether a Shift key is being held down or Caps Lock is in effect).
Key Up	Used when a key is released.

Each key on the keyboard has a "key station" code which is a number from 0 to 127. This number is used as an index into the translation table that is currently in effect. If the corresponding entry in that translation table is a value from 0 to 127, this value is treated as an ASCII character, and that character is the result of the translation.

If the entry is a value from 128 to 255, it is a “special” entry. Special entry values are classified according to the value of the high-order bits. The high-order value for each class is defined as a constant, as shown in the list below. The value of the low-order bits, when added to this constant, distinguishes between keys within each class:

SHIFTKEYS 0x80	A shift key. The value of the particular shift key is added to determine which shift mask to apply:
CAPSLCK 0	“Caps Lock” key.
SHIFTLOCK 1	“Shift Lock” key.
LEFTSHIFT 2	Left-hand “Shift” key.
RIGHTSHIFT 3	Right-hand “Shift” key.
LEFTCTRL 4	Left-hand (or only) “Control” key.
RIGHTCTRL 5	Right-hand “Control” key.
BUCKYBITS 0x90	Used to toggle mode-key-up/down status without altering the value of an accompanying ASCII character. (The actual bit-position value, minus 7, is added.)
METABIT 0	The “Meta” key was pressed along with the key. This is the only user-accessible bucky bit.
SYSTEMBIT 1	The “System” key was pressed. This is a place holder to indicate which key is the system-abort key.
FUNNY 0xA0	Performs various functions depending on the value of the low 4 bits:
NOP 0xA0	Does nothing.
OOPS 0xA1	Exists, but is undefined.
HOLE 0xA2	There is no key in this position on the keyboard, and the position-code should not be used.
NOSCROLL 0xA3	Alternately sends ^S and ^Q.
CTRLS 0xA4	Sends ^S and toggles NOScroll key.
CTRLQ 0xA5	Sends ^Q and toggles NOScroll key.
RESET 0xA6	Keyboard reset.
ERROR 0xA7	The keyboard driver detected an internal error.
IDLE 0xA8	The keyboard is idle (no keys down).
0xA9 — 0xAF	Reserved for nonparameterized functions.
STRING 0xB0	The low-order bits index a table of strings. When a key with a STRING entry is depressed, the characters in the null-terminated string for that key are sent, character by character. The maximum length is defined as:
	KTAB_STRLLEN 10
	Individual string numbers are defined as:
	HOMEARROW 0x00
	UPARROW 0x01
	DOWNARROW 0x02
	LEFTARROW 0x03
	RIGHTARROW 0x04
	String numbers 0x05 — 0x0F are available for custom entries.

LEFTFUNC 0xC0
 RIGHTFUNC 0xD0
 TOPFUNC 0xE0
 BOTTOMFUNC 0xF0

Function keys. The low 4 bits indicate the function key number within the group:

LF(<i>n</i>)	(LEFTFUNC+(<i>n</i>)-1)
RF(<i>n</i>)	(RIGHTFUNC+(<i>n</i>)-1)
TF(<i>n</i>)	(TOPFUNC+(<i>n</i>)-1)
BF(<i>n</i>)	(BOTTOMFUNC+(<i>n</i>)-1)

There are 64 keys reserved for function keys. The actual positions may not be on left/right/top/bottom of the keyboard, although they usually are.

In TR_ASCII mode, when a function key is pressed, the following escape sequence is sent:

<ESC>[0...9z

where <ESC> is a single escape character and “0...9” indicates the decimal representation of the function-key value. For example, function key R1 sends the sequence:

<ESC>[208z

because the decimal value of RF(1) is 208. In TR_EVENT mode, if there is a VUID event code for the function key in question, an event with that event code is generated; otherwise, individual events for the characters of the escape sequence are generated.

IOCTLS

Two `ioctl`s set and retrieve the current translation mode of a keyboard:

KIOCTRANS The argument is a pointer to an `int`. The translation mode is set to the value in the `int` pointed to by the argument.

KIOCGTRANS The argument is a pointer to an `int`. The current translation mode is stored in the `int` pointed to by the argument.

`ioctl`s for changing and retrieving entries from the keyboard translation table use the `kiockey` structure:

```
struct kiockey {
    int    kio_tablemask; /* Translation table (one of: 0, CAPSMASK,
                          SHIFTMASK, CTRLMASK, UPMASK) */
#define KIOCABORT1  -1 /* Special “mask”: abort1 keystation */
#define KIOCABORT2  -2 /* Special “mask”: abort2 keystation */
    u_char kio_station; /* Physical keyboard key station (0-127) */
    u_char kio_entry; /* Translation table station’s entry */
    char   kio_string[10]; /* Value for STRING entries (null terminated) */
};
```

KIOCSETKEY The argument is a pointer to a `kiockey` structure. The translation table entry referred to by the values in that structure is changed.

`kio_tablemask` specifies which of the five translation tables contains the entry to be modified:

UPMASK 0x0080	“Key Up” translation table.
CTRLMASK 0x0030	“Controlled” translation table.
SHIFTMASK 0x000E	“Shifted” translation table.
CAPSMASK 0x0001	“Caps Lock” translation table.
(No shift keys pressed or locked)	“Unshifted” translation table.

kio_station specifies the keystation code for the entry to be modified. The value of **kio_entry** is stored in the entry in question. If **kio_entry** is between **STRING** and **STRING+15**, the string contained in **kio_string** is copied to the appropriate string table entry. This call may return **EINVAL** if there are invalid arguments.

There are a couple special values of **kio_tablemask** that affect the two step “break to the PROM monitor” sequence. The usual sequence is **SETUP-a** or **L1-a**. If **kio_tablemask** is **KIOCABORT1** then the value of **kio_station** is set to be the first keystation in the sequence. If **kio_tablemask** is **KIOCABORT2** then the value of **kio_station** is set to be the second keystation in the sequence.

KIOCGTKEY The argument is a pointer to a **kiockey** structure. The current value of the keyboard translation table entry specified by **kio_tablemask** and **kio_station** is stored in the structure pointed to by the argument. This call may return **EINVAL** if there are invalid arguments.

KIOCTYPE The argument is a pointer to an **int**. A code indicating the type of the keyboard is stored in the **int** pointed to by the argument:

KB_KLUNK	Micro Switch 103SD32-2
KB_VT100	Keytronics VT100 compatible
KB_SUN2	Sun-2 keyboard
KB_SUN3	Type 3 keyboard
KB_SUN4	Type 4 keyboard
KB_ASCII	ASCII terminal masquerading as keyboard

-1 is stored in the **int** pointed to by the argument if the keyboard type is unknown.

KIOCCMD The argument is a pointer to an **int**. The command specified by the value of the **int** pointed to by the argument is sent to the keyboard. The commands that can be sent are:

Commands to the Sun-2, Type 3, and Type 4 keyboard:

KBD_CMD_RESET	Reset keyboard as if power-up.
KBD_CMD_BELL	Turn on the bell.
KBD_CMD_NOBELL	Turn off the bell

Commands to the Type 3 and Type 4 keyboard:

KBD_CMD_CLICK	Turn on the click annunciator.
KBD_CMD_NOCLICK	Turn off the click annunciator.

Inappropriate commands for particular keyboard types are ignored. Since there is no reliable way to get the state of the bell or click (because we cannot query the keyboard, and also because a process could do writes to the appropriate serial driver — thus going around this **ioctl**) we do not provide an equivalent **ioctl** to query its state.

KIOCSDIRECT

KIOCGDIRECT These **ioctls** are supported for compatibility with the system keyboard device **/dev/kbd**. **KIOCSDIRECT** has no effect, and **KIOCGDIRECT** always returns 1.

INDEX STRUCTURES

There is a hierarchy of structures for accessing keyboard translation data. The array *keytables* contains pointers to the translation data for each of the known keyboard types:

```
struct keyboard *keytables[] = {
    &keyindex_ms,
    &keyindex_vt,
    &keyindex_s2,
    &keyindex_s3,
};
```

Each keyboard type is described by a **struct keyboard** that contains pointers to the five translation-tables (“Unshifted”, “Shifted”, “Caps Locked”, “Controlled”, and “Key Up”) associated with that type, plus bit-masks that indicate what state can persist with no keys pressed, and the key-pair used as the abort

sequence for the system.

An array `keystringtab` contains the strings sent by various keys, and can be accessed by any translation:

```
#define kstescinit(c) {'\033', '[', 'c', '\0'}
char keystringtab[16][KTAB_STRLEN] = {
    kstescinit(H),          /*home*/
    kstescinit(A),          /*up*/
    kstescinit(B),          /*down*/
    kstescinit(D),          /*left*/
    kstescinit(C),          /*right*/
};
```

Index Structure for the Type 4 Keyboard

```
/* Index to keymaps for Type 4 keyboard */
static struct keyboard keyindex_s4 = {
    &keytab_s4_lc,
    &keytab_s4_uc,
    &keytab_s4_cl,
    &keytab_s4_ct,
    &keytab_s4_up,
    0x0000,          /* Shift bits which stay on with idle keyboard */
    0x0000,          /* Bucky bits which stay on with idle keyboard */
    1,              /* 77, /* abort keys */
    CAPSMASK,       /* Shift bits which toggle on down event */
};
```

Index Structure for the Type 3 Keyboard

```
static struct keyboard keyindex_s3 = {
    &keytab_s3_lc,
    &keytab_s3_uc,
    &keytab_s3_cl,
    &keytab_s3_ct,
    &keytab_s3_up,
    0x0000,          /* Shift bits that stay on with idle keyboard */
    0x0000,          /* Bucky bits that stay on with idle keyboard */
    0x01, 0x4d,     /* Abort sequence L1-A */
    CAPSMASK,       /* Shift bits that toggle on down event */
};
```

Index Structure for the Sun-2 Keyboard

```
static struct keyboard keyindex_s2 = {
    &keytab_s2_lc,
    &keytab_s2_uc,
    &keytab_s2_cl,
    &keytab_s2_ct,
    &keytab_s2_up,
    CAPSMASK,       /* Shift bits that stay on with idle keyboard */
    0x0000,          /* Bucky bits that stay on with idle keyboard */
    0x01, 0x4d,     /* Abort sequence L1-A */
    0x0000,          /* Shift bits that toggle on down event */
};
```

Index Structure for the Micro Switch 103SD32-2 Keyboard

```
static struct keyboard keyindex_ms = {
    &keytab_ms_lc,
    &keytab_ms_uc,
    &keytab_ms_cl,
    &keytab_ms_ct,
    &keytab_ms_up,
    CTLSMASK, /* Shift bits that stay on with idle keyboard */
    0x0000, /* Bucky bits that stay on with idle keyboard */
    0x01, 0x4d, /* Abort sequence L1-A */
    0x0000, /* Shift bits that toggle on down event */
};
```

Index Structure for the VT100-Style Keyboard

```
static struct keyboard keyindex_vt = {
    &keytab_vt_lc,
    &keytab_vt_uc,
    &keytab_vt_cl,
    &keytab_vt_ct,
    &keytab_vt_up,
    CAPSMASK+CTLSMASK, /* Shift keys that stay on with idle keyboard */
    0x0000, /* Bucky bits that stay on with idle keyboard */
    0x01, 0x3b, /* Abort sequence SETUP-A */
    0x0000, /* Shift bits that toggle on down event */
};
```

DEFAULT TRANSLATION TABLES

Type 4 Keyboard

Unshifted

Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value
00 HOLE	01 BUCKYBITS+	02 HOLE	03 LF(2)	04 HOLE	05 TF(1)	06 TF(2)	07 TF(10)
	SYSTEMBIT						
08 TF(3)	09 TF(11)	0A TF(4)	0B TF(12)	0C TF(5)	0D HOLE	0E TF(6)	0F HOLE
10 TF(7)	11 TF(8)	12 TF(9)	13 ALT	14 HOLE	15 RF(1)	16 RF(2)	17 RF(3)
18 HOLE	19 LF(3)	1A LF(4)	1B HOLE	1C HOLE	1D c ('[')	1E '1'	1F '2'
20 '3'	21 '4'	22 '5'	23 '6'	24 '7'	25 '8'	26 '9'	27 '0'
28 '.'	29 '='	2A ''	2B '^b'	2C HOLE	2D RF(4)	2E RF(5)	2F RF(6)
30 BF(13)	31 LF(5)	32 BF(10)	33 LF(6)	34 HOLE	35 '^c'	36 'q'	37 'w'
38 'e'	39 'r'	3A 't'	3B 'y'	3C 'u'	3D 'i'	3E 'o'	3F 'p'
40 '['	41 ']'	42 0x7F	43 OOPS	44 RF(7)	45 STRING+	46 RF(9)	47 BF(15)
			(temp)		UPARROW		
48 LF(7)	49 LF(8)	4A HOLE	4B HOLE	4C SHIFTKEYS+	4D 'a'	4E 's'	4F 'd'
				LEFTCTRL			
50 'f'	51 'g'	52 'h'	53 'j'	54 'k'	55 'l'	56 ';' :	57 '^'
58 '\'	59 '^r'	5A BF(11)	5B STRING+	5C RF(11)	5D STRING+	5E BF(8)	5F LF(9)
			LEFTARROW		RIGHTARROW		
60 HOLE	61 LF(10)	62 BF(16)	63 SHIFTKEYS+	64 'z'	65 'x'	66 'c'	67 'v'
			LEFTSHIFT				
68 'b'	69 'n'	6A 'm'	6B ','	6C '.'	6D '/'	6E SHIFTKEYS+	6F '^n'
						RIGHTSHIFT	
70 RF(13)	71 STRING+	72 RF(15)	73 HOLE	74 HOLE	75 HOLE	76 LF(16)	77 SHIFTKEYS+
	DOWNARROW						CAPSLOCK
78 BUCKYBITS+	79 ' ' METABIT	7A BUCKYBITS+	7B HOLE	7C HOLE	7D BF(14)	7E ERROR	7F IDLE
		METABIT					

**Type 4 Keyboard
Shifted**

Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value
00 HOLE	01 BUCKYBITS+ SYSTEMBIT	02 HOLE	03 LF(2)	04 HOLE	05 TF(1)	06 TF(2)	07 TF (10)
08 TF(3)	09 TF(11)	0A TF(4)	0B TF(12)	0C TF(5)	0D HOLE	0E TF (6)	0F HOLE
10 TF(7)	11 TF(8)	12 TF(9)	13 ALT	14 HOLE	15 RF(1)	16 RF(2)	17 RF (3)
18 HOLE	19 LF(3)	1A LF(4)	1B HOLE	1C HOLE	1D c('[')	1E ']'	1F '@'
20 '#'	21 '\$'	22 '%'	23 ''	24 '&'	25 '*'	26 '('	27 ')'
28 '.'	29 '+'	2A ''	2B '^b'	2C HOLE	2D RF(4)	2E RF(5)	2F RF (6)
30 BF(13)	31 LF(5)	32 BF(10)	33 LF (6)	34 HOLE	35 '^c'	36 'Q'	37 'W'
38 'E'	39 'R'	3A 'T'	3B 'Y'	3C 'U'	3D 'I'	3E 'O'	3F 'P'
40 '['	41 ']'	42 0x7F	43 OOPS (temp)	44 RF(7)	45 STRING+ UPARROW	46 RF(9)	47 BF (15)
48 LF(7)	49 LF (8)	4A HOLE	4B HOLE	4C SHIFTKEYS+ LEFTCTRL	4D 'A'	4E 'S'	4F 'D'
50 'F'	51 'G'	52 'H'	53 'J'	54 'K'	55 'L'	56 ';'	57 ''
58 'I'	59 '^c'	5A BF(11)	5B STRING+ LEFTARROW	5C RF(11)	5D STRING+ RIGHTARROW	5E BF(8)	5F LF (9)
60 HOLE	61 LF(10)	62 BF (16)	63 SHIFTKEYS+ LEFTSHIFT	64 'Z'	65 'X'	66 'C'	67 'V'
68 'B'	69 'N'	6A 'M'	6B '<'	6C '>'	6D '?'	6E SHIFTKEYS+ RIGHTSHIFT	6F '^n'
70 RF(13)	71 STRING+ DOWNARROW	72 RF(15)	73 HOLE	74 HOLE	75 HOLE	76 LF (16)	77 SHIFTKEYS+ CAPSLOCK
78 BUCKYBITS+ METABIT	79 ' '	7A BUCKYBITS+ METABIT	7B HOLE	7C HOLE	7D BF (14)	7E ERROR	7F IDLE

**Type 4 Keyboard
Caps Locked**

Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value
00 HOLE	01 BUCKYBITS+ SYSTEMBIT	02 HOLE	03 LF(2)	04 HOLE	05 TF(1)	06 TF(2)	07 TF (10)
08 TF(3)	09 TF(11)	0A TF(4)	0B TF(12)	0C TF(5)	0D HOLE	0E TF (6)	0F HOLE
10 TF(7)	11 TF(8)	12 TF(9)	13 ALT	14 HOLE	15 RF(1)	16 RF(2)	17 RF (3)
18 HOLE	19 LF(3)	1A LF(4)	1B HOLE	1C HOLE	1D c('[')	1E ']'	1F '2'
21 '3'	21 '4'	22 '5'	23 '6'	24 '7'	25 '8'	26 '9'	27 '0'
28 '.'	29 '='	2A ''	2B '^b'	2C HOLE	2D RF(4)	2E RF(5)	2F RF (6)
30 BF(23)	32 LF(5)	32 BF(20)	33 LF (6)	34 HOLE	35 '^c'	36 'Q'	37 'W'
38 'E'	39 'R'	3A 'T'	3B 'Y'	3C 'U'	3D 'I'	3E 'O'	3F 'P'
40 '['	41 ']'	43 0x7F	43 OOPS (temp)	44 RF(7)	45 STRING+ UPARROW	46 RF(9)	47 BF (15)
48 LF(7)	49 LF (8)	4A HOLE	4B HOLE	4C SHIFTKEYS+ LEFTCTRL	4D 'A'	4E 'S'	4F 'D'
50 'F'	51 'G'	52 'H'	53 'J'	54 'K'	55 'L'	56 ';'	57 '^'
58 '^N'	59 '^c'	5A BF(11)	5B STRING+ LEFTARROW	5C RF(11)	5D STRING+ RIGHTARROW	5E BF(8)	5F LF (9)
60 HOLE	61 LF(10)	62 BF (16)	63 SHIFTKEYS+ LEFTSHIFT	64 'Z'	65 'X'	66 'C'	67 'V'
68 'B'	69 'N'	6A 'M'	6B '^'	6C '^'	6D '/'	6E SHIFTKEYS+ RIGHTSHIFT	6F '^n'
70 RF(13)	71 STRING+ DOWNARROW	72 RF(15)	73 HOLE	74 HOLE	75 HOLE	76 LF (16)	77 SHIFTKEYS+ CAPSLOCK
78 BUCKYBITS+ METABIT	79 ' '	7A BUCKYBITS+ METABIT	7B HOLE	7C HOLE	7D BF (14)	7E ERROR	7F IDLE

**Type 4 Keyboard
Controlled**

Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value
00 HOLE	01 BUCKYBITS+ SYSTEMBIT	02 HOLE	03 LF(2)	04 HOLE	05 TF(1)	06 TF(2)	07 TF(10)
08 TF(3)	09 TF(11)	0A TF(4)	0B TF(12)	0C TF(5)	0D HOLE	0E TF(6)	0F HOLE
10 TF(7)	11 TF(8)	12 TF(9)	13 ALT	14 HOLE	15 RF(1)	16 RF(2)	17 RF(3)
18 HOLE	19 LF(3)	1A LF(4)	1B HOLE	1C HOLE	1D c('I')	1E '1'	1F c('@')
20 '3'	21 '4'	22 '5'	23 c('"'	24 '7'	25 '8'	26 '9'	27 '0'
28 c('_')	29 '='	2A c('"'	2B '\b'	2C HOLE	2D RF(4)	2E RF(5)	2F RF(6)
30 BF(13)	31 LF(5)	32 BF(10)	33 LF(6)	34 HOLE	35 '\t'	36 c('q')	37 c('w')
38 c('e')	39 c('r')	3A c('t')	3B c('y')	3C c('u')	3D c('i')	3E c('o')	3F c('p')
40 c('I')	41 c('J')	42 0x7F	43 OOPS (temp)	44 RF(7)	45 STRING+ UPARROW	46 RF(9)	47 BF(15)
48 LF(7)	49 LF(8)	4A HOLE	4B HOLE	4C SHIFTKEYS+ LEFTCTRL	4D c('a')	4E c('s')	4F c('d')
50 c('f')	51 c('g')	52 c('h')	53 c('j')	54 c('k')	55 c('l')	56 ';' :	57 '^'
58 c('\N')	59 '\e'	5A BF(11)	5B STRING+ LEFTARROW	5C RF(11)	5D STRING+ RIGHTARROW	5E BF(8)	5F LF(9)
60 HOLE	61 LF(10)	62 BF(16)	63 SHIFTKEYS+ LEFTSHIFT	64 c('z')	65 c('x')	66 c('c')	67 c('v')
68 c('b')	69 c('n')	6A c('m')	6B ';' :	6C ';' :	6D c(' _')	6E SHIFTKEYS+ RIGHTSHIFT	6F '\n'
70 RF(13)	71 STRING+ DOWNARROW	72 RF(15)	73 HOLE	74 HOLE	75 HOLE	76 LF(16)	77 SHIFTKEYS+ CAPSLOCK
78 BUCKYBITS+ METABIT	79 c(' ')	7A BUCKYBITS+ METABIT	7B HOLE	7C HOLE	7D BF(14)	7E ERROR	7F IDLE

**Type 4 Keyboard
Key Up**

Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value
00 HOLE	01 BUCKYBITS+ SYSTEMBIT	02 HOLE	03 OOPS	04 HOLE	05 OOPS	06 OOPS	07 HOLE
08 OOPS	09 HOLE	0A OOPS	0B HOLE	0C OOPS	0D HOLE	0E OOPS	0F HOLE
10 OOPS	11 OOPS	12 OOPS	13 OOPS	14 HOLE	15 OOPS	16 OOPS	17 NOP
18 HOLE	19 OOPS	1A OOPS	1B HOLE	1C HOLE	1D NOP	1E NOP	1F NOP
20 NOP	21 NOP	22 NOP	23 NOP	24 NOP	25 NOP	26 NOP	27 NOP
28 NOP	29 NOP	2A NOP	2B NOP	2C HOLE	2D OOPS	2E OOPS	2F NOP
30 HOLE	31 OOPS	32 HOLE	33 OOPS	34 HOLE	35 NOP	36 NOP	37 NOP
38 NOP	39 NOP	3A NOP	3B NOP	3C NOP	3D NOP	3E NOP	3F NOP
40 NOP	41 NOP	42 NOP	43 HOLE	44 OOPS	45 OOPS	46 NOP	47 HOLE
48 OOPS	49 OOPS	4A HOLE	4B HOLE	4C SHIFTKEYS+ LEFTCTRL	4D NOP	4E NOP	4F NOP
50 NOP	51 NOP	52 NOP	53 NOP	54 NOP	55 NOP	56 NOP	57 NOP
58 NOP	59 NOP	5A HOLE	5B OOPS	5C OOPS	5D NOP	5E HOLE	5F OOPS
60 OOPS	61 OOPS	62 HOLE	63 SHIFTKEYS+ LEFTSHIFT	64 NOP	65 NOP	66 NOP	67 NOP
68 NOP	69 NOP	6A NOP	6B NOP	6C NOP	6D NOP	6E SHIFTKEYS+ RIGHTSHIFT	6F NOP
70 OOPS	71 OOPS	72 NOP	73 HOLE	74 HOLE	75 HOLE	76 HOLE	77 NOP
78 BUCKYBITS+ METABIT	79 NOP	7A BUCKYBITS+ METABIT	7B HOLE	7C HOLE	7D HOLE	7E HOLE	7F RESET

Type 3 Keyboard

Unshifted

Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value
00 HOLE	01 BUCKYBITS+ SYSTEMBIT	02 HOLE	03 LF(2)	04 HOLE	05 TF(1)	06 TF (2)	07 HOLE
08 TF(3)	09 HOLE	0A TF(4)	0B HOLE	0C TF(5)	0D HOLE	0E TF (6)	0F HOLE
10 TF(7)	11 TF(8)	12 TF(9)	13 ALT	14 HOLE	15 RF(1)	16 RF(2)	17 RF (3)
18 HOLE	19 LF(3)	1A LF(4)	1B HOLE	1C HOLE	1D c ('[')	1E '1'	1F '2'
20 '3'	21 '4'	22 '5'	23 '6'	24 '7'	25 '8'	26 '9'	27 '0'
28 '-'	29 '='	2A ''	2B '^b'	2C HOLE	2D RF(4)	2E RF(5)	2F RF (6)
30 HOLE	31 LF(5)	32 HOLE	33 LF (6)	34 HOLE	35 '^t'	36 'q'	37 'w'
38 'e'	39 'r'	3A 't'	3B 'y'	3C 'u'	3D 'i'	3E 'o'	3F 'p'
40 '['	41 ']'	42 0x7F	43 HOLE	45 RF(7)	45 STRING+ UPARROW	46 RF (9)	47 HOLE
48 LF(7)	49 LF(8)	4A LF (40)	4B HOLE	4C SHIFTKEYS+ LEFTCTRL	4D 'a'	4E 's'	4F 'd'
50 'f'	51 'g'	52 'h'	53 'j'	54 'k'	55 'l'	56 ':'	57 '^'
58 '^	59 '^r'	5A HOLE	5B STRING+ LEFTARROW	5C RF(11)	5D STRING+ RIGHTARROW	5E HOLE	5F LF (9)
60 LF(15)	61 LF (10)	62 HOLE	63 SHIFTKEYS+ LEFTSHIFT	64 'z'	65 'x'	66 'c'	67 'v'
68 'b'	69 'n'	6A 'm'	6B ';'	6C '.'	6D '/'	6E SHIFTKEYS+ RIGHTSHIFT	6F '^n'
70 RF(13)	71 STRING+ DOWNARROW	72 RF (15)	73 HOLE	74 HOLE	75 HOLE	76 HOLE	77 SHIFTKEYS+ CAPSLOCK
78 BUCKYBITS+ METABIT	79 ''	7A BUCKYBITS+ METABIT	7B HOLE	7C HOLE	7D HOLE	7E ERROR	7F IDLE

Shifted

Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value
00 HOLE	01 BUCKYBITS+ SYSTEMBIT	02 HOLE	03 LF(2)	04 HOLE	05 TF(1)	06 TF (2)	07 HOLE
08 TF(3)	09 HOLE	0A TF(4)	0B HOLE	0C TF(5)	0D HOLE	0E TF (6)	0F HOLE
10 TF(7)	11 TF(8)	12 TF(9)	13 ALT	14 HOLE	15 RF(1)	16 RF(2)	17 RF (3)
18 HOLE	19 LF(3)	1A LF(4)	1B HOLE	1C HOLE	1D c ('[')	1E '1'	1F '@'
20 '#'	21 '\$'	22 '%'	23 ''	24 '&'	25 '*'	26 '('	27 ')'
28 '-'	29 '+'	2A ''	2B '^b'	2C HOLE	2D RF(4)	2E RF(5)	2F RF (6)
30 HOLE	31 LF(5)	32 HOLE	33 LF (6)	34 HOLE	35 '^t'	36 'Q'	37 'W'
38 'E'	39 'R'	3A 'T'	3B 'Y'	3C 'U'	3D 'I'	3E 'O'	3F 'P'
40 '['	41 ']'	42 0x7F	43 HOLE	44 RF(7)	45 STRING+ UPARROW	46 RF (9)	47 HOLE
48 LF(7)	49 LF (8)	4A HOLE	4B HOLE	4C SHIFTKEYS+ LEFTCTRL	4D 'A'	4E 'S'	4F 'D'
50 'F'	51 'G'	52 'H'	53 'J'	54 'K'	55 'L'	56 ':'	57 ''
58 'I'	59 '^r'	5A HOLE	5B STRING+ LEFTARROW	5C RF(11)	5D STRING+ RIGHTARROW	5E HOLE	5F LF (9)
60 LF(15)	61 LF (10)	62 HOLE	63 SHIFTKEYS+ LEFTSHIFT	64 'Z'	65 'X'	66 'C'	67 'V'
68 'B'	69 'N'	6A 'M'	6B '<'	6C '>'	6D '?'	6E SHIFTKEYS+ RIGHTSHIFT	6F '^n'
70 RF(13)	71 STRING+ DOWNARROW	72 RF (15)	73 HOLE	74 HOLE	75 HOLE	76 HOLE	77 SHIFTKEYS+ CAPSLOCK
78 BUCKYBITS+ METABIT	79 ''	7A BUCKYBITS+ METABIT	7B HOLE	7C HOLE	7D HOLE	7E ERROR	7F IDLE

Type 3 Keyboard

Caps Locked

Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value
00 HOLE	01 BUCKYBITS+ SYSTEMBIT	02 HOLE	03 LF(2)	04 HOLE	05 TF(1)	06 TF (2)	07 HOLE
08 TF(3)	09 HOLE	0A TF(4)	0B HOLE	0C TF(5)	0D HOLE	0E TF (6)	0F HOLE
10 TF(7)	11 TF(8)	12 TF(9)	13 ALT	14 HOLE	15 RF(1)	16 RF(2)	17 RF (3)
18 HOLE	19 LF(3)	1A LF(4)	1B HOLE	1C HOLE	1D c('I')	1E '1'	1F '2'
20 '3'	21 '4'	22 '5'	23 '6'	24 '7'	25 '8'	26 '9'	27 '0'
28 '-'	29 '='	2A ''	2B `b`	2C HOLE	2D RF(4)	2E RF(5)	2F RF (6)
30 HOLE	31 LF(5)	32 HOLE	33 LF (6)	34 HOLE	35 `x`	36 'Q'	37 'W'
38 'E'	39 'R'	3A 'T'	3B 'Y'	3C 'U'	3D 'I'	3E 'O'	3F 'P'
40 'I'	41 'J'	42 0x7F	43 HOLE	44 RF(7)	45 STRING+ UPARROW	46 RF (9)	47 HOLE
48 LF(7)	49 LF (8)	4A HOLE	4B HOLE	4C SHIFTKEYS+ LEFTCTRL	4D 'A'	4E 'S'	4F 'D'
50 'F'	51 'G'	52 'H'	53 'J'	54 'K'	55 'L'	56 ';' :	57 `\"`
58 `\"`	59 `x`	5A HOLE	5B STRING+ LEFTARROW	5C RF(11)	5D STRING+ RIGHTARROW	5E HOLE	5F LF (9)
60 LF(15)	61 LF (10)	62 HOLE	63 SHIFTKEYS+ LEFTSHIFT	64 'Z'	65 'X'	66 'C'	67 'V'
68 'B'	69 'N'	6A 'M'	6B `;`	6C `.`	6D `/`	6E SHIFTKEYS+ RIGHTSHIFT	6F `n`
70 RF(13)	71 STRING+ DOWNARROW	72 RF (15)	73 HOLE	74 HOLE	75 HOLE	76 HOLE	77 SHIFTKEYS+ CAPSLOCK
78 BUCKYBITS+ METABIT	79 ` `	7A BUCKYBITS+ METABIT	7B HOLE	7C HOLE	7D HOLE	7E ERROR	7F IDLE

Controlled

Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value
00 HOLE	01 BUCKYBITS+ SYSTEMBIT	02 HOLE	03 LF(2)	04 HOLE	05 TF(1)	06 TF (2)	07 HOLE
08 TF(3)	09 HOLE	0A TF(4)	0B HOLE	0C TF(5)	0D HOLE	0E TF (6)	0F HOLE
10 TF(7)	11 TF(8)	12 TF(9)	13 ALT	14 HOLE	15 RF(1)	16 RF(2)	17 RF (3)
18 HOLE	19 LF(3)	1A LF(4)	1B HOLE	1C HOLE	1D c('I')	1E '1'	1F c('@')
20 '3'	21 '4'	22 '5'	23 c('')	24 '7'	25 '8'	26 '9'	27 '0'
28 c('_')	29 '='	2A c('')	2B `b`	2C HOLE	2D RF(4)	2E RF(5)	2F RF (6)
30 HOLE	31 LF(5)	32 HOLE	33 LF(6)	34 HOLE	35 `x`	36 c('q')	37 c('w')
38 c('e')	39 c('r')	3A c('t')	3B c('y')	3C c('u')	3D c('i')	3E c('o')	3F c('p')
40 c('I')	41 c('J')	42 0x7F	43 HOLE	44 RF(7)	45 STRING+ UPARROW	46 RF (9)	47 HOLE
48 LF(7)	49 LF (8)	4A HOLE	4B HOLE	4C SHIFTKEYS+ LEFTCTRL	4D c('a')	4E c('s')	4F c('d')
50 c('f')	51 c('g')	52 c('h')	53 c('j')	54 c('k')	55 c('l')	56 ';' :	57 `\"`
58 c(`\"`)	59 `x`	5A HOLE	5B STRING+ LEFTARROW	5C RF(11)	5D STRING+ RIGHTARROW	5E HOLE	5F LF (9)
60 LF(15)	61 LF (10)	62 HOLE	63 SHIFTKEYS+ LEFTSHIFT	64 c('z')	65 c('x')	66 c('c')	67 c('v')
68 c('b')	69 c('n')	6A c('m')	6B `;`	6C `.`	6D c('_')	6E SHIFTKEYS+ RIGHTSHIFT	6F `n`
70 RF(13)	71 STRING+ DOWNARROW	72 RF (15)	73 HOLE	74 HOLE	75 HOLE	76 HOLE	77 SHIFTKEYS+ CAPSLOCK
78 BUCKYBITS+ METABIT	79 c(' ')	7A BUCKYBITS+ METABIT	7B HOLE	7C HOLE	7D HOLE	7E ERROR	7F IDLE

Type 3 Keyboard

Key Up

Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value
00 HOLE	01 BUCKYBITS+ SYSTEMBIT	02 HOLE	03 OOPS	04 HOLE	05 OOPS	06 OOPS	07 HOLE	
08 OOPS	09 HOLE	0A OOPS	0B HOLE	0C OOPS	0D HOLE	0E OOPS	0F HOLE	
10 OOPS	11 OOPS	12 OOPS	13 OOPS	14 HOLE	15 OOPS	16 OOPS	17 NOP	
18 HOLE	19 OOPS	1A OOPS	1B HOLE	1C HOLE	1D NOP	1E NOP	1F NOP	
20 NOP	21 NOP	22 NOP	23 NOP	24 NOP	25 NOP	26 NOP	27 NOP	
28 NOP	29 NOP	2A NOP	2B NOP	2C HOLE	2D OOPS	2E OOPS	2F NOP	
30 HOLE	31 OOPS	32 HOLE	33 OOPS	34 HOLE	35 NOP	36 NOP	37 NOP	
38 NOP	39 NOP	3A NOP	3B NOP	3C NOP	3D NOP	3E NOP	3F NOP	
40 NOP	41 NOP	42 NOP	43 HOLE	44 OOPS	45 OOPS	46 NOP	47 HOLE	
48 OOPS	49 OOPS	4A HOLE	4B HOLE	4C SHIFTKEYS+ LEFTCTRL	4D NOP	4E NOP	4F NOP	
50 NOP	51 NOP	52 NOP	53 NOP	54 NOP	55 NOP	56 NOP	57 NOP	
58 NOP	59 NOP	5A HOLE	5B OOPS	5C OOPS	5D NOP	5E HOLE	5F OOPS	
60 OOPS	61 OOPS	62 HOLE	63 SHIFTKEYS+ LEFTSHIFT	64 NOP	65 NOP	66 NOP	67 NOP	
68 NOP	69 NOP	6A NOP	6B NOP	6C NOP	6D NOP	6E SHIFTKEYS+ RIGHTSHIFT	6F NOP	
70 OOPS	71 OOPS	72 NOP	73 HOLE	74 HOLE	75 HOLE	76 HOLE	77 NOP	
78 BUCKYBITS+ METABIT	79 NOP	7A BUCKYBITS+ METABIT	7B HOLE	7C HOLE	7D HOLE	7E HOLE	7F RESET	

Sun-2 Keyboard

Unshifted

Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value
00 HOLE	01 BUCKYBITS+ SYSTEMBIT	02 LF(11)	03 LF(2)	04 HOLE	05 TF(1)	06 TF(2)	07 TF(11)	
08 TF(3)	09 TF(12)	0A TF(4)	0B TF(13)	0C TF(5)	0D TF(14)	0E TF(6)	0F TF(15)	
10 TF(7)	11 TF(8)	12 TF(9)	13 TF(10)	14 HOLE	15 RF(1)	16 RF(2)	17 RF(3)	
18 HOLE	19 LF(3)	1A LF(4)	1B LF(12)	1C HOLE	1D c('[')	1E '1'	1F '2'	
20 '3'	21 '4'	22 '5'	23 '6'	24 '7'	25 '8'	26 '9'	27 '0'	
28 '-'	29 '='	2A ''	2B '\b'	2C HOLE	2D RF(4)	2E RF(5)	2F RF(6)	
30 HOLE	31 LF(5)	32 LF(13)	33 LF(6)	34 HOLE	35 '\n'	36 'q'	37 'w'	
38 'e'	39 'r'	3A 't'	3B 'y'	3C 'u'	3D 'i'	3E 'o'	3F 'p'	
40 '['	41 ']'	42 0x7F	43 HOLE	44 RF(7)	45 STRING+ UPARROW	46 RF(9)	47 HOLE	
48 LF(7)	49 LF(8)	4A LF(14)	4B HOLE	4C SHIFTKEYS+ LEFTCTRL	4D 'a'	4E 's'	4F 'd'	
50 'f'	51 'g'	52 'h'	53 'j'	54 'k'	55 'l'	56 ';'	57 '\'	
58 '\'	59 '\v'	5A HOLE	5B STRING+ LEFTARROW	5C RF(11)	5D STRING+ RIGHTARROW	5E HOLE	5F LF(9)	
60 LF(15)	61 LF(10)	62 HOLE	63 SHIFTKEYS+ LEFTSHIFT	64 'z'	65 'x'	66 'c'	67 'v'	
68 'b'	69 'n'	6A 'm'	6B ','	6C '.'	6D '/'	6E SHIFTKEYS+ RIGHTSHIFT	6F '\u'	
70 RF(13)	71 STRING+ DOWNARROW	72 RF(15)	73 HOLE	74 HOLE	75 HOLE	76 HOLE	77 HOLE	
70 BUCKYBITS+ METABIT	71 ''	72 BUCKYBITS+ METABIT	73 HOLE	74 HOLE	75 HOLE	76 ERROR	77 IDLE	

Sun-2 Keyboard
Shifted

Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value
00 HOLE	01 BUCKYBITS+ SYSTEMBIT	02 LF(11)	03 LF(2)	04 HOLE	05 TF(1)	06 TF(2)	07 TF(11)
08 TF(3)	00 TF(12)	0A TF(4)	0B TF(13)	0C TF(5)	0D TF(14)	0E TF(6)	0F TF(15)
10 TF(7)	11 TF(8)	12 TF(9)	13 TF(10)	14 HOLE	15 RF(1)	16 RF(2)	17 RF(3)
18 HOLE	19 LF(3)	1A LF(4)	1B LF(12)	1C HOLE	1D c('')	1E '!	1F '@'
20 '#'	21 '\$'	22 '%'	23 ''	24 '&'	25 '*'	26 '('	27 ')'
28 '_'	29 '+'	2A ''	2B `b'	2C HOLE	2D RF(4)	2E RF(5)	2F RF(6)
30 HOLE	31 LF(5)	32 LF(13)	33 LF(6)	34 HOLE	35 `v'	36 'Q'	37 'W'
38 'E'	39 'R'	3A 'T'	3B 'Y'	3C 'U'	3D 'I'	3E 'O'	3F 'P'
40 '['	41 'J'	42 0x7F	43 HOLE	44 RF(7)	45 STRING+ UPARROW	46 RF(9)	47 HOLE
48 LF(7)	49 LF(8)	4A LF(14)	4B HOLE	4C SHIFTKEYS+ LEFTCTRL	4D 'A'	4E 'S'	4F 'D'
50 'F'	51 'G'	52 'H'	53 'J'	54 'K'	55 'L'	56 ';'	57 ''
58 'I'	59 `v'	5A HOLE	5B STRING+ LEFTARROW	5C RF(11)	5D STRING+ RIGHTARROW	5E HOLE	5F LF(9)
60 LF(15)	61 LF(10)	62 HOLE	63 SHIFTKEYS+ LEFTSHIFT	64 'Z'	65 'X'	66 'C'	67 'V'
68 'B'	69 'N'	6A 'M'	6B '<'	6C '>'	6D '?'	6E SHIFTKEYS+ RIGHTSHIFT	6F `n'
70 RF(13)	71 STRING+ DOWNARROW	72 RF(15)	73 HOLE	74 HOLE	75 HOLE	76 HOLE	77 HOLE
78 BUCKYBITS+ METABIT	79 ' '	7A BUCKYBITS+ METABIT	7B HOLE	7C HOLE	7D HOLE	7E ERROR	7F IDLE

Caps Locked

Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value
00 HOLE	01 BUCKYBITS+ SYSTEMBIT	02 LF(11)	03 LF(2)	04 HOLE	05 TF(1)	06 TF(2)	07 TF(11)
08 TF(3)	09 TF(12)	0A TF(4)	0B TF(13)	0C TF(5)	0D TF(14)	0E TF(6)	0F TF(15)
10 TF(7)	11 TF(8)	12 TF(9)	13 TF(10)	14 HOLE	15 RF(1)	16 RF(2)	17 RF(3)
18 HOLE	19 LF(3)	1A LF(4)	1B LF(12)	1C HOLE	1D c('')	1E '1'	1F '2'
20 '3'	21 '4'	22 '5'	23 '6'	24 '7'	25 '8'	26 '9'	27 '0'
28 '-'	29 '='	2A ''	2B `b'	2C HOLE	2D RF(4)	2E RF(5)	2F RF(6)
30 HOLE	31 LF(5)	32 LF(13)	33 LF(6)	34 HOLE	35 `v'	36 'Q'	37 'W'
38 'E'	39 'R'	3A 'T'	3B 'Y'	3C 'U'	3D 'I'	3E 'O'	3F 'P'
40 '['	41 'J'	42 0x7F	43 HOLE	44 RF(7)	45 STRING+ UPARROW	46 RF(9)	47 HOLE
48 LF(7)	49 LF(8)	4A LF(14)	4B HOLE	4C SHIFTKEYS+ LEFTCTRL	4D 'A'	4E 'S'	4F 'D'
50 'F'	51 'G'	52 'H'	53 'J'	54 'K'	55 'L'	56 ';'	57 `n'
58 `v'	59 `v'	5A HOLE	5B STRING+ LEFTARROW	5C RF(11)	5D STRING+ RIGHTARROW	5E HOLE	5F LF(9)
60 LF(15)	61 LF(10)	62 HOLE	63 SHIFTKEYS+ LEFTSHIFT	64 'Z'	65 'X'	66 'C'	67 'V'
68 'B'	69 'N'	6A 'M'	6B ';'	6C 'I'	6D 'I'	6E SHIFTKEYS+ RIGHTSHIFT	6F `n'
70 RF(13)	71 STRING+ DOWNARROW	72 RF(15)	73 HOLE	74 HOLE	75 HOLE	76 HOLE	77 HOLE
78 BUCKYBITS+ METABIT	79 ' '	7A BUCKYBITS+ METABIT	7B HOLE	7C HOLE	7D HOLE	7E ERROR	7F IDLE

Sun-2 Keyboard

Controlled

Key	Value	Key	Value	Key	Value	Key	Value	Key	Value
00	HOLE	01	BUCKYBITS+	02	LF(11)	03	LF(2)	04	HOLE
08	TF(3)	09	TF(12)	0A	TF(4)	0B	TF(13)	0C	TF(5)
10	TF(7)	11	TF(8)	12	TF(9)	13	TF(10)	14	HOLE
18	HOLE	19	LF(3)	1A	LF(4)	1B	LF(12)	1C	HOLE
20	'3'	21	'4'	22	'5'	23	c('')	24	'7'
28	c('_')	29	'='	2A	c('')	2B	'b'	2C	HOLE
30	HOLE	31	LF(5)	32	LF(13)	33	LF(6)	34	HOLE
38	c('e')	39	c('r')	3A	c('t')	3B	c('y')	3C	c('u')
40	c('l')	41	c('l')	42	0x7F	43	HOLE	44	RF(7)
48	LF(7)	49	LF(8)	4A	LF(14)	4B	HOLE	4C	SHIFTKEYS+4D c('a')
50	c('f')	51	c('g')	52	c('h')	53	c('j')	54	c('k')
58	c('\')	59	'r'	5A	HOLE	5B	STRING+	5C	RF(11)
60	LF(15)	61	LF(10)	62	HOLE	63	SHIFTKEYS+ LEFTSHIFT	64	c('z')
68	c('b')	69	c('n')	6A	c('m')	6B	'.'	6C	'.'
70	RF(13)	71	STRING+ DOWNARROW	72	RF(15)	73	HOLE	74	HOLE
78	BUCKYBITS+ METABIT	79	c(' ')	7A	BUCKYBITS+ METABIT	7B	HOLE	7C	HOLE
								7D	HOLE
								7E	ERROR
								7F	IDLE

Key Up

Key	Value	Key	Value	Key	Value	Key	Value	Key	Value
00	HOLE	01	BUCKYBITS+ SYSTEMBIT	02	OOPS	03	OOPS	04	HOLE
08	OOPS	09	OOPS	0A	OOPS	0B	OOPS	0C	OOPS
10	OOPS	11	OOPS	12	OOPS	13	OOPS	14	HOLE
18	HOLE	19	OOPS	1A	OOPS	1B	OOPS	1C	HOLE
20	NOP	21	NOP	22	NOP	23	NOP	24	NOP
28	NOP	29	NOP	2A	NOP	2B	NOP	2C	HOLE
30	HOLE	31	OOPS	32	OOPS	33	OOPS	34	HOLE
38	NOP	39	NOP	3A	NOP	3B	NOP	3C	NOP
40	NOP	41	NOP	42	NOP	43	HOLE	44	OOPS
48	OOPS	49	OOPS	4A	OOPS	4B	HOLE	4C	SHIFTKEYS+ LEFTCTRL
50	NOP	51	NOP	52	NOP	53	NOP	54	NOP
58	NOP	59	NOP	5A	HOLE	5B	OOPS	5C	OOPS
60	OOPS	61	OOPS	62	HOLE	63	SHIFTKEYS+ LEFTSHIFT	64	NOP
68	NOP	69	NOP	6A	NOP	6B	NOP	6C	NOP
70	OOPS	71	OOPS	72	NOP	73	HOLE	74	HOLE
78	BUCKYBITS+ METABIT	79	NOP	7A	BUCKYBITS+ METABIT	7B	HOLE	7C	HOLE
								7D	HOLE
								7E	HOLE
								7F	RESET

Micro Switch 103SD32-2 Keyboard

Unshifted

Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value
00 HOLE	01 BUCKYBITS+ SYSTEMBIT	02 LF(2)	03 LF(3)	04 HOLE	05 TF(1)	06 TF(2)	07 TF(3)	
08 TF(4)	09 TF(5)	0A TF(6)	0B TF(7)	0C TF(8)	0D TF(9)	0E TF(10)	0F TF(11)	
10 TF(12)	11 TF(13)	12 TF(14)	13 c(')	14 HOLE	15 RF(1)	16 '+'	17 '-'	
18 HOLE	19 LF(4)	1A ^	1B LF(6)	1C HOLE	1D SHIFTKEYS+ CAPSLOCK	1E '1'	1F '2'	
20 '3'	21 '4'	22 '5'	23 '6'	24 '7'	25 '8'	26 '9'	27 '0'	
28 '-'	29 '~	2A ''	2B ^b'	2C HOLE	2D '7'	2E '8'	2F '9'	
30 HOLE	31 LF(7)	32 STRING+ UPARROW	33 LF(9)	34 HOLE	35 ^	36 'q'	37 'w'	
38 'e'	39 'r'	3A 't'	3B 'y'	3C 'u'	3D 'i'	3E 'o'	3F 'p'	
40 '{'	41 '}'	42 '_'	43 HOLE	44 '4'	45 '5'	46 '6'	47 HOLE	
48 STRING+ LEFTARROW	49 STRING+ HOMEARROW	4A STRING+ RIGHTARROW	4B HOLE	4C SHIFTKEYS+ SHIFTLOCK	4D 'a'	4E 's'	4F 'd'	
50 'f'	51 'g'	52 'h'	53 'j'	54 'k'	55 'l'	56 ';' :	57 '.'	
58 'i'	59 ^c'	5A HOLE	5B '1'	5C '2'	5D '3'	5E HOLE	5F NOScroll	
60 STRING+ DOWNARROW	61 LF(15)	62 HOLE	63 HOLE	64 SHIFTKEYS+ LEFTSHIFT	65 'z'	66 'x'	67 'c'	
68 'v'	69 'b'	6A 'n'	6B 'm'	6C '^	6D '^	6E '^/	6F SHIFTKEYS+ RIGHTSHIFT	
70 NOP	71 0x7F	72 '0'	73 NOP	74 '^	75 HOLE	76 HOLE	77 HOLE	
78 HOLE	79 HOLE	7A SHIFTKEYS+ LEFTCTRL	7B '^	7C SHIFTKEYS+ RIGHTCTRL	7D HOLE	7E HOLE	7F IDLE	

Shifted

Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value
00 HOLE	01 BUCKYBITS SYSTEMBIT	02 LF(2)	03 LF(3)	04 HOLE	05 TF(1)	06 TF(2)	07 TF(3)	
08 TF(4)	09 TF(5)	0A TF(6)	0B TF(7)	0C TF(8)	0D TF(9)	0E TF(10)	0F TF(11)	
10 TF(12)	11 TF(13)	12 TF(14)	13 c(')	14 HOLE	15 RF(1)	16 '+'	17 '-'	
18 HOLE	19 LF(4)	1A ^	1B LF(6)	1C HOLE	1D SHIFTKEYS+ CAPSLOCK	1E '1'	1F ''	
20 '#'	21 '\$'	22 '%'	23 '&'	24 '^	25 '('	26 ')' :	27 '0'	
28 '=	29 '~	2A '@'	2B ^b'	2C HOLE	2D '7'	2E '8'	2F '9'	
30 HOLE	31 LF(7)	32 STRING+ UPARROW	33 LF(9)	34 HOLE	35 ^	36 'Q'	37 'W'	
38 'E'	39 'R'	3A 'T'	3B 'Y'	3C 'U'	3D 'I'	3E 'O'	3F 'P'	
40 '['	41 ']'	42 '_'	43 HOLE	44 '4'	45 '5'	46 '6'	47 HOLE	
48 STRING+ LEFTARROW	49 STRING+ HOMEARROW	4A STRING+ RIGHTARROW	4B HOLE	4C SHIFTKEYS+ SHIFTLOCK	4D 'A'	4E 'S'	4F 'D'	
50 'F'	51 'G'	52 'H'	53 'J'	54 'K'	55 'L'	56 '+'	57 '^	
58 ^	59 ^c'	5A HOLE	5B '1'	5C '2'	5D '3'	5E HOLE	5F NOScroll	
60 STRING+ DOWNARROW	61 LF(15)	62 HOLE	63 HOLE	64 SHIFTKEYS+ LEFTSHIFT	65 'Z'	66 'X'	67 'C'	
68 'V'	69 'B'	6A 'N'	6B 'M'	6C '<'	6D '>'	6E '^?	6F SHIFTKEYS+ RIGHTSHIFT	
70 NOP	71 0x7F	72 '0'	73 NOP	74 '^	75 HOLE	76 HOLE	77 HOLE	
78 HOLE	79 HOLE	7A SHIFTKEYS+ RIGHTSHIFT	7B '^	7C SHIFTKEYS+ LEFTCTRL	7D HOLE	7E HOLE	7F IDLE	

Micro Switch 103SD32-2 Keyboard

Caps Locked

Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value
00 HOLE	01 BUCKYBITS+ SYSTEMBIT	02 LF(2)	03 LF(3)	04 HOLE	05 TF(1)	06 TF(2)	07 TF(3)
08 TF(4)	09 TF(5)	0A TF(6)	0B TF(7)	0C TF(8)	0D TF(9)	0E TF(10)	0F TF(11)
10 TF(12)	11 TF(13)	12 TF(14)	13 c('I')	14 HOLE	15 RF(1)	16 '+'	17 '-'
18 HOLE	19 LF(4)	1A '^'	1B LF(6)	1C HOLE	1D SHIFTKEYS+ CAPSLOCK	1E '1'	1F '2'
20 '3'	21 '4'	22 '5'	23 '6'	24 '7'	25 '8'	26 '9'	27 '0'
28 '-'	29 '~'	2A ''	2B '^b'	2C HOLE	2D '7'	2E '8'	2F '9'
30 HOLE	31 LF(7)	32 STRING+ UPARROW	33 LF(9)	34 HOLE	35 '^'	36 'Q'	37 'W'
38 'E'	39 'R'	3A 'T'	3B 'Y'	3C 'U'	3D 'I'	3E 'O'	3F 'P'
40 '['	41 ']'	42 '_'	43 HOLE	44 '4'	45 '5'	46 '6'	47 HOLE
48 STRING+ LEFTARROW	49 STRING+ HOMEARROW	4A STRING+ RIGHTARROW	4B HOLE	4C SHIFTKEYS+ SHIFTLOCK	4D 'A'	4E 'S'	4F 'D'
50 'F'	51 'G'	52 'H'	53 'J'	54 'K'	55 'L'	56 ';' :	57 ':'
58 'I'	59 '^'	5A HOLE	5B '1'	5C '2'	5D '3'	5E HOLE	5F NOSCROLL
60 STRING+ DOWNARROW	61 LF(15)	62 HOLE	63 HOLE	64 SHIFTKEYS+ LEFTSHIFT	65 'Z'	66 'X'	67 'C'
68 'V'	69 'B'	6A 'N'	6B 'M'	6C ','	6D '.'	6E '/'	6F SHIFTKEYS+ RIGHTSHIFT
70 NOP	71 0x7F	72 'O'	73 NOP	74 '.'	75 HOLE	76 HOLE	77 HOLE
78 HOLE	79 HOLE	7A SHIFTKEYS+ LEFTCTRL	7B ''	7C SHIFTKEYS+ RIGHTCTRL	7D HOLE	7E HOLE	7F IDLE

Controlled

Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value
00 HOLE	01 BUCKYBITS+ SYSTEMBIT	02 LF(2)	03 LF(3)	04 HOLE	05 TF(1)	06 TF(2)	07 TF(3)
08 TF(4)	09 TF(5)	0A TF(6)	0B TF(7)	0C TF(8)	0D TF(9)	0E TF(10)	0F TF(11)
10 TF(02)	11 TF(03)	12 TF(04)	13 c('I')	14 HOLE	15 RF(0)	16 OOPS	17 OOPS
18 HOLE	19 LF(4)	1A '^'	1B LF(6)	1C HOLE	1D SHIFTKEYS+ CAPSLOCK	1E OOPS	1F OOPS
20 OOPS	21 OOPS	22 OOPS	23 OOPS	24 OOPS	25 OOPS	26 OOPS	27 OOPS
28 OOPS	29 c('')	2A c('@')	2B '^b'	2C HOLE	2D OOPS	2E OOPS	2F OOPS
30 HOLE	31 LF(7)	32 STRING+ UPARROW	33 F(9)	34 HOLE	35 '^'	36 CTRLQ	37 c('W')
38 c('E')	39 c('R')	3A c('T')	3B c('Y')	3C c('U')	3D c('I')	3E c('O')	3F c('P')
40 c('I')	41 c('J')	42 c('_')	43 HOLE	44 OOPS	45 OOPS	46 OOPS	47 HOLE
48 STRING+ LEFTARROW	49 STRING+ HOMEARROW	4A STRING+ RIGHTARROW	4B HOLE	4C SHIFTKEYS+ SHIFTLOCK	4D c('A')	4E CTRLS	4F c('D')
50 c('F')	51 c('G')	52 c('H')	53 c('J')	54 c('K')	55 c('L')	56 OOPS	57 OOPS
58 c('V')	59 '^'	5A HOLE	5B OOPS	5C OOPS	5D OOPS	5E HOLE	5F NOSCROLL
60 STRING+ DOWNARROW	61 LF(15)	62 HOLE	63 HOLE	64 SHIFTKEYS+ LEFTSHIFT	65 c('Z')	66 c('X')	67 c('C')
68 c('V')	69 c('B')	6A c('N')	6B c('M')	6C OOPS	6D OOPS	6E OOPS	6F SHIFTKEYS+ RIGHTSHIFT
70 NOP	71 0x7F	72 OOPS	73 NOP	74 OOPS	75 HOLE	76 HOLE	77 HOLE
78 HOLE	79 HOLE	7A SHIFTKEYS+ LEFTCTRL	7B '^O'	7C SHIFTKEYS+ RIGHTCTRL	7D HOLE	7E HOLE	7F IDLE

Micro Switch 103SD32-2 Keyboard

Key Up

Key	Value	Key	Value	Key	Value	Key	Value	Key	Value	Key	Value	Key	Value		
00	HOLE	01	BUCKYBITS+ SYSTEMBIT	02	OOPS	03	OOPS	04	HOLE	05	OOPS	06	OOPS	07	OOPS
08	OOPS	09	OOPS	0A	OOPS	0B	OOPS	0C	OOPS	0D	OOPS	0E	OOPS	0F	OOPS
10	OOPS	11	OOPS	12	OOPS	13	NOP	14	HOLE	15	OOPS	16	NOP	17	NOP
18	HOLE	19	OOPS	1A	NOP	1B	OOPS	1C	HOLE	1D	SHIFTKEYS+ CAPSLOCK	1E	NOP	1F	NOP
20	NOP	21	NOP	22	NOP	23	NOP	24	NOP	25	NOP	26	NOP	27	NOP
28	NOP	29	NOP	2A	NOP	2B	NOP	2C	HOLE	2D	NOP	2E	NOP	2F	NOP
30	HOLE	31	OOPS	32	NOP	33	OOPS	34	HOLE	35	NOP	36	NOP	37	NOP
38	NOP	39	NOP	3A	NOP	3B	NOP	3C	NOP	3D	NOP	3E	NOP	3F	NOP
40	NOP	41	NOP	42	NOP	43	HOLE	44	NOP	45	NOP	46	NOP	47	HOLE
48	NOP	49	NOP	4A	NOP	4B	HOLE	4C	SHIFTKEYS+ SHIFTLOCK	4D	NOP	4E	NOP	4F	NOP
50	NOP	51	NOP	52	NOP	53	NOP	54	NOP	55	NOP	56	NOP	57	NOP
58	NOP	59	NOP	5A	HOLE	5B	NOP	5C	NOP	5D	NOP	5E	HOLE	5F	NOP
60	NOP	61	OOPS	62	HOLE	63	HOLE	64	SHIFTKEYS+ LEFTSHIFT	65	NOP	66	NOP	67	NOP
68	NOP	69	NOP	6A	NOP	6B	NOP	6C	NOP	6D	NOP	6E	NOP	6F	SHIFTKEYS+ RIGHTSHIFT
70	NOP	71	NOP	72	NOP	73	NOP	74	NOP	75	HOLE	76	HOLE	77	HOLE
78	HOLE	79	HOLE	7A	SHIFTKEYS+ LEFTCTRL	7B	NOP	7C	SHIFTKEYS+ RIGHTCTRL	7D	HOLE	7E	HOLE	7F	RESET

VT100-Style Keyboard

Unshifted

Key	Value	Key	Value	Key	Value	Key	Value	Key	Value	Key	Value	Key	Value		
00	HOLE	01	BUCKYBITS+ SYSTEMBIT	02	HOLE	03	HOLE	04	HOLE	05	HOLE	06	HOLE	07	HOLE
08	HOLE	09	HOLE	0A	STRING+ UPARROW	0B	STRING+ DOWNARROW	0C	STRING+ LEFTARROW	0D	STRING+ RIGHTARROW	0E	HOLE	0F	TF(1)
10	TF(2)	11	TF(3)	12	TF(4)	13	c ('[')	14	'1'	15	'2'	16	'3'	17	'4'
18	'5'	19	'6'	1A	'7'	1B	'8'	1C	'9'	1D	'0'	1E	'_'	1F	'='
20	''	21	c ('H')	22	BUCKYBITS+ METABIT	23	'?'	24	'8'	25	'9'	26	'-'	27	'\x'
28	'q'	29	'w'	2A	'e'	2B	'r'	2C	't'	2D	'y'	2E	'u'	2F	'i'
30	'o'	31	'p'	32	'l'	33	'j'	34	0x7F	35	'4'	36	'5'	37	'6'
38	'.'	39	SHIFTKEYS+ LEFTCTRL	3A	SHIFTKEYS+ CAPSLOCK	3B	'a'	3C	's'	3D	'd'	3E	'f'	3F	'g'
40	'h'	41	'j'	42	'k'	43	'l'	44	','	45	'\`'	46	'\x'	47	'\`'
48	'1'	49	'2'	4A	'3'	4B	NOP	4C	NOSCROLL	4D	SHIFTKEYS+ LEFTSHIFT	4E	'z'	4F	'x'
50	'c'	51	'v'	52	'b'	53	'n'	54	'm'	55	';'	56	'.'	57	'/'
58	SHIFTKEYS+ RIGHTSHIFT	59	'n'	5A	'0'	5B	HOLE	5C	'.'	5D	'\x'	5E	HOLE	5F	HOLE
60	HOLE	61	HOLE	62	' '	63	HOLE	64	HOLE	65	HOLE	66	HOLE	67	HOLE
68	HOLE	69	HOLE	6A	HOLE	6B	HOLE	6C	HOLE	6D	HOLE	6E	HOLE	6F	HOLE
70	HOLE	71	HOLE	72	HOLE	73	HOLE	74	HOLE	75	HOLE	76	HOLE	77	HOLE
78	HOLE	79	HOLE	7A	HOLE	7B	HOLE	7C	HOLE	7D	HOLE	7E	HOLE	7F	IDLE

VT100-Style Keyboard

Shifted

Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value
00 HOLE	01 BUCKYBITS+ SYSTEMBIT	02 HOLE	03 HOLE	04 HOLE	05 HOLE	06 HOLE	07 HOLE	
08 HOLE	09 HOLE	0A STRING+ UPARROW	0B STRING+ DOWNARROW	0C STRING+ LEFTARROW	0D STRING+ RIGHTARROW	0E HOLE	0F TF(1)	
10 TF(2)	11 TF(3)	12 TF(4)	13 c('[')	14 '1'	15 '@'	16 '#'	17 '\$'	
18 '%'	19 ''	1A '&'	1B '*'	1C '('	1D ')'	1E '_'	1F '+'	
20 ''	21 c('H')	22 BUCKYBITS+ METABIT	23 '7'	24 '8'	25 '9'	26 '-'	27 '^'	
28 'Q'	29 'W'	2A 'E'	2B 'R'	2C 'T'	2D 'Y'	2E 'U'	2F 'I'	
30 'O'	31 'P'	32 '['	33 ']'	34 0x7F	35 '4'	36 '5'	37 '6'	
38 ','	39 SHIFTKEYS+ LEFTCTRL	3A SHIFTKEYS+ CAPSLOCK	3B 'A'	3C 'S'	3D 'D'	3E 'F'	3F 'G'	
40 'H'	41 'J'	42 'K'	43 'L'	44 ';' :	45 ''	46 '^'	47 'I'	
48 '1'	49 '2'	4A '3'	4B NOP	4C NOSCROLL	4D SHIFTKEYS+ LEFTSHIFT	4E 'Z'	4F 'X'	
50 'C'	51 'V'	52 'B'	53 'N'	54 'M'	55 '<'	56 '>'	57 '?'	
58 SHIFTKEYS+ RIGHTSHIFT	59 '\n'	5A '0'	5B HOLE	5C '.'	5D '^'	5E HOLE	5F HOLE	
60 HOLE	61 HOLE	62 ' '	63 HOLE	64 HOLE	65 HOLE	66 HOLE	67 HOLE	
68 HOLE	69 HOLE	6A HOLE	6B HOLE	6C HOLE	6D HOLE	6E HOLE	6F HOLE	
70 HOLE	71 HOLE	72 HOLE	73 HOLE	74 HOLE	75 HOLE	76 HOLE	77 HOLE	
78 HOLE	79 HOLE	7A HOLE	7B HOLE	7C HOLE	7D HOLE	7E HOLE	7F IDLE	

Caps Locked

Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value
00 HOLE	01 BUCKYBITS+ SYSTEMBIT	02 HOLE	03 HOLE	04 HOLE	05 HOLE	06 HOLE	07 HOLE	
08 HOLE	09 HOLE	0A STRING+ UPARROW	0B STRING+ DOWNARROW	0C STRING+ LEFTARROW	0D STRING+ RIGHTARROW	0E HOLE	0F TF(1)	
10 TF(2)	11 TF(3)	12 TF(4)	13 c('[')	14 '1'	15 '2'	16 '3'	17 '4'	
18 '5'	19 '6'	1A '7'	1B '8'	1C '9'	1D '0'	1E '-'	1F '='	
20 ''	21 c('H')	22 BUCKYBITS+ METABIT	23 '7'	24 '8'	25 '9'	26 '-'	27 '^'	
28 'Q'	29 'W'	2A 'E'	2B 'R'	2C 'T'	2D 'Y'	2E 'U'	2F 'I'	
30 'O'	31 'P'	32 '['	33 ']'	34 0x7F	35 '4'	36 '5'	37 '6'	
38 ','	39 SHIFTKEYS+ LEFTCTRL	3A SHIFTKEYS+ CAPSLOCK	3B 'A'	3C 'S'	3D 'D'	3E 'F'	3F 'G'	
40 'H'	41 'J'	42 'K'	43 'L'	44 ';' :	45 '^'	46 '^'	47 '^'	
48 '1'	49 '2'	4A '3'	4B NOP	4C NOSCROLL	4D SHIFTKEYS+ LEFTSHIFT	4E 'Z'	4F 'X'	
50 'C'	51 'V'	52 'B'	53 'N'	54 'M'	55 ';' :	56 '.'	57 '/'	
58 SHIFTKEYS+ RIGHTSHIFT	59 '\n'	5A '0'	5B HOLE	5C '.'	5D '^'	5E HOLE	5F HOLE	
60 HOLE	61 HOLE	62 ' '	63 HOLE	64 HOLE	65 HOLE	66 HOLE	67 HOLE	
68 HOLE	69 HOLE	6A HOLE	6B HOLE	6C HOLE	6D HOLE	6E HOLE	6F HOLE	
70 HOLE	71 HOLE	72 HOLE	73 HOLE	74 HOLE	75 HOLE	76 HOLE	77 HOLE	
78 HOLE	79 HOLE	7A HOLE	7B HOLE	7C HOLE	7D HOLE	7E HOLE	7F IDLE	

**VT100-Style Keyboard
Controlled**

Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value
00 HOLE	01 BUCKYBITS+ SYSTEMBIT	02 HOLE	03 HOLE	04 HOLE	05 HOLE	06 HOLE	07 HOLE	
08 HOLE	09 HOLE	0A STRING+ UPARROW	0B STRING+ DOWNARROW	0C STRING+ LEFTARROW	0D STRING+ RIGHTARROW	0E HOLE	0F TF(1)	
10 TF(2)	11 TF(3)	12 TF(4)	13 c('I')	14 '1'	15 c('@')	16 '3'	17 '4'	
18 '5'	19 c('')	1A '7'	1B '8'	1C '9'	1D '0'	1E c(' _')	1F '='	
20 c('')	21 c('H')	22 BUCKYBITS+ METABIT	23 '7'	24 '8'	25 '9'	26 ' _'	27 'x'	
28 CTRLQ	29 c('W')	2A c('E')	2B c('R')	2C c('T')	2D c('Y')	2E c('U')	2F c('I')	
30 c('O')	31 c('P')	32 c('I')	33 c('J')	34 0x7F	35 '4'	36 '5'	37 '6'	
38 ','	39 SHIFTKEYS+ LEFTCTRL	3A SHIFTKEYS+ CAPSLOCK	3B c('A')	3C CTRLS	3D c('D')	3E c('F')	3F c('G')	
40 c('H')	41 c('J')	42 c('K')	43 c('L')	44 ' _'	45 ' _'	46 'x'	47 c('\')	
48 '1'	49 '2'	4A '3'	4B NOP	4C NOSCROLL	4D SHIFTKEYS+ LEFTSHIFT	4E c('Z')	4F c('X')	
50 c('C')	51 c('V')	52 c('B')	53 c('N')	54 c('M')	55 ' _'	56 ' _'	57 c(' _')	
58 SHIFTKEYS+ RIGHTSHIFT	59 'n'	5A '0'	5B HOLE	5C ' _'	5D HOLE	5E HOLE	5F HOLE	
60 HOLE	61 HOLE	62 c(' ')	63 HOLE	64 HOLE	65 HOLE	66 HOLE	67 HOLE	
68 HOLE	69 HOLE	6A HOLE	6B HOLE	6C HOLE	6D HOLE	6E HOLE	6F HOLE	
70 HOLE	71 HOLE	72 HOLE	73 HOLE	74 HOLE	75 HOLE	76 HOLE	77 HOLE	
78 HOLE	79 HOLE	7A HOLE	7B HOLE	7C HOLE	7D HOLE	7E HOLE	7F IDLE	

Key Up

Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value	Key Value
00 HOLE	01 BUCKYBITS+ SYSTEMBIT	02 HOLE	03 HOLE	04 HOLE	05 HOLE	06 HOLE	07 HOLE		
08 HOLE	09 HOLE	0A NOP	0B NOP	0C NOP	0D NOP	0E HOLE	0F OOPS		
10 OOPS	11 OOPS	12 OOPS	13 NOP	14 NOP	15 NOP	16 NOP	17 NOP		
18 NOP	19 NOP	1A NOP	1B NOP	1C NOP	1D NOP	1E NOP	1F NOP		
20 NOP	21 NOP	22 BUCKYBITS+ METABIT	23 NOP	24 NOP	25 NOP	26 NOP	27 NOP		
28 NOP	29 NOP	2A NOP	2B NOP	2C NOP	2D NOP	2E NOP	2F NOP		
30 NOP	31 NOP	32 NOP	33 NOP	34 NOP	35 NOP	36 NOP	37 NOP		
38 NOP	39 SHIFTKEYS+ LEFTCTRL	3A SHIFTKEYS+ CAPSLOCK	3B NOP	3C NOP	3D NOP	3E NOP	3F NOP		
40 NOP	41 NOP	42 NOP	43 NOP	44 NOP	45 NOP	46 NOP	47 NOP		
48 NOP	49 NOP	4A NOP	4B NOP	4C NOP	4D SHIFTKEYS+ LEFTSHIFT	4E NOP	4F NOP		
50 NOP	51 NOP	52 NOP	53 NOP	54 NOP	55 NOP	56 NOP	57 NOP		
58 SHIFTKEYS+ RIGHTSHIFT	59 NOP	5A NOP	5B HOLE	5C NOP	5D NOP	5E HOLE	5F HOLE		
60 HOLE	61 HOLE	62 NOP	63 HOLE	64 HOLE	65 HOLE	66 HOLE	67 HOLE		
68 HOLE	69 HOLE	6A HOLE	6B HOLE	6C HOLE	6D HOLE	6E HOLE	6F HOLE		
70 HOLE	71 HOLE	72 HOLE	73 HOLE	74 HOLE	75 HOLE	76 HOLE	77 HOLE		
78 HOLE	79 HOLE	7A HOLE	7B HOLE	7C HOLE	7D HOLE	7E HOLE	7F RESET		

SEE ALSO

click(1), oldsetkeys(1), kbd(4S), termio(4), win(4S)

The SunView System Programmer's Guide— Appendix: Writing a Virtual User Input Device Driver
(describes firm_event format)

NAME

kbd – Sun keyboard

CONFIG

None; included in standard system.

DESCRIPTION

The **kbd** device provides access to the Sun Workstation keyboard. When opened, it provides access to the standard keyboard device for the workstation (attached either to a CPU serial or parallel port). It is a multiplexing driver; a stream referring to the standard keyboard device, with the **kb(4M)** and **ttcompat(4M)** STREAMS modules pushed on top of that device, is linked below it. Normally, this device passes input to the “workstation console” driver, which is linked above a special minor device of **kbd**, so that keystrokes appear as input on **/dev/console**; the **KIOCSDIRECT ioctl** must be used to direct input towards or away from the **/dev/kbd** device.

IOCTLS

KIOCSDIRECT The argument is a pointer to an **int**. If the value in the **int** pointed to by the argument is 1, subsequent keystrokes typed on the system keyboard will be sent to **/dev/kbd**; if it is 0, subsequent keystrokes will be sent to the “workstation console” device. When the last process that has **/dev/kbd** open closes it, if keystrokes had been sent to **/dev/kbd** they are redirected back to the “workstation console” device.

KIOCGDIRECT The argument is a pointer to an **int**. If keystrokes are currently being sent to **/dev/kbd**, 1 is stored in the **int** pointed to by the argument; if keystrokes are currently being sent to the “workstation console” device, 0 is stored there.

FILES

/dev/kbd

SEE ALSO

console(4S), **kb(4M)**, **ttcompat(4M)**, **win(4S)**, **zs(4S)**

NAME

ldterm – standard terminal STREAMS module

CONFIG

None; included by default.

SYNOPSIS

```
#include <sys/stream.h>
#include <sys/stropt.h>
```

```
ioctl(fd, I_PUSH, "ldterm");
```

DESCRIPTION

ldterm is a STREAMS module that provides most of the **termio** (4) terminal interface. This module does not perform the low-level device control functions specified by flags in the **c_cflag** word of the **termios** structure or by the **IGNBRK**, **IGNPAR**, **PARMRK**, or **INPCK** flags in the **c_iflag** word of the **termios** structure; those functions must be performed by the driver or by modules pushed below the **ldterm** module. All other **termio** functions are performed by **ldterm**; some of them, however, require the cooperation of the driver or modules pushed below **ldterm**, and may not be performed in some cases. These include the **IXOFF** flag in the **c_iflag** word and the delays specified in the **c_oflag** word.

Read-side Behavior

Various types of STREAMS messages are processed as follows:

- M_BREAK** When this message is received, either an interrupt signal is generated, or the message is treated as if it were an **M_DATA** message containing a single ASCII NUL character, depending on the state of the **BRKINT** flag.
- M_DATA** These messages are normally processed using the standard **termio** input processing. If the **ICANON** flag is set, a single input record (“line”) is accumulated in an internal buffer, and sent upstream when a line-terminating character is received. If the **ICANON** flag is not set, other input processing is performed and the processed data is passed upstream.
- If output is to be stopped or started as a result of the arrival of characters, **M_STOP** and **M_START** messages are sent downstream, respectively. If the **IXOFF** flag is set, and input is to be stopped or started as a result of flow-control considerations, **M_STOPI** and **M_STARTI** messages are sent downstream, respectively.
- M_DATA** messages are sent downstream, as necessary, to perform echoing.
- If a signal is to be generated, a **M_FLUSH** message with a flag byte of **FLUSHR** is placed on the read queue, and if the signal is also to flush output a **M_FLUSH** message with a flag byte of **FLUSHW** is sent downstream.
- M_CTL** If the first byte of the message is **MC_NOCANON**, the input processing normally performed on **M_DATA** messages is disabled, and those messages are passed upstream unmodified; this is for the use of modules or drivers that perform their own input processing, such as a pseudo-terminal in **TIOCREMOTE** mode connected to a program that performs this processing. If the first byte of the message is **MC_DOCANON**, the input processing is enabled. Otherwise, the message is ignored; in any case, the message is passed upstream.
- M_FLUSH** The read queue of the module is flushed of all its data messages, and all data in the record being accumulated is also flushed. The message is passed upstream.
- M_HANGUP** Data is flushed as it is for a **M_FLUSH** message, and **M_FLUSH** messages with a flag byte of **FLUSHRW** are sent upstream and downstream. Then an **M_PCSIG** message is sent upstream with a signal of **SIGCONT**, followed by the **M_HANGUP** message.
- M_IOCACK** The data contained within the message, which is to be returned to the process, is augmented if necessary, and the message is passed upstream.

All other messages are passed upstream unchanged.

Write-side behavior

Various types of STREAMS messages are processed as follows:

- M_FLUSH** The write queue of the module is flushed of all its data messages, and the message is passed downstream.
- M_IOCTL** The function to be performed for this `ioctl` by the `ldterm` module is performed, and the message is passed downstream in most cases. The `TCFLSH` and `TCXONC` `ioctls` can be performed entirely in this module, so the reply is sent upstream and the message is not passed downstream.
- M_DATA** If the `OPOST` flag is set, or both the `XCASE` and `ICANON` flags are set, output processing is performed and the processed message is passed downstream, along with any `M_DELAY` messages generated. Otherwise, the message is passed downstream without change.

All other messages are passed downstream unchanged.

IOCTLS

The following `ioctls` are processed by the `ldterm` module. All others are passed downstream.

TCGETS

TCGETA The message is passed downstream; if an acknowledgment is seen, the data provided by the driver and modules downstream is augmented and the acknowledgement is passed upstream.

TCSETS

TCSETSW

TCSETSF

TCSETA

TCSETAW

TCSETAF

The parameters that control the behavior of the `ldterm` module are changed. If a mode change requires options at the stream head to be changed, a `M_SETOPT` message is sent upstream. If the `ICANON` flag is turned on or off, the read mode at the stream head is changed to message-nondiscard or byte-stream mode, respectively. If it is turned on, the `vmin` and `vtime` values at the stream head are set to 1 and 0, respectively; if it is turned off, they are set to the values specified by the `ioctl`. The `vmin` and `vtime` values are also set if `ICANON` is off and the values are changed by the `ioctl`. If the `TOSTOP` flag is turned on or off, the `tostop` mode at the stream head is turned on or off, respectively.

TCFLSH

If the argument is 0, an `M_FLUSH` message with a flag byte of `FLUSHR` is sent downstream and placed on the read queue. If the argument is 1, the write queue is flushed of all its data messages and a `M_FLUSH` message with a flag byte of `FLUSHW` is sent upstream and downstream. If the argument is 2, the write queue is flushed of all its data messages and a `M_FLUSH` message with a flag byte of `FLUSHRW` is sent downstream and placed on the read queue.

TCXONC

If the argument is 0, and output is not already stopped, an `M_STOP` message is sent downstream. If the argument is 1, and output is stopped, an `M_START` message is sent downstream. If the argument is 2, and input is not already stopped, an `M_STOPI` message is sent downstream. If the argument is 3, and input is stopped, an `M_STARTI` message is sent downstream.

SEE ALSO

`console(4S)`, `mcp(4S)`, `mti(4S)`, `pty(4)`, `termio(4)`, `ttcompat(4M)`, `zs(4S)`

NAME

le – Sun-3/50, Sun-3/60 10MB Ethernet interface

CONFIG

device le0 at obio ? csr

DESCRIPTION

The le interface provides access to a 10 Mb/s Ethernet network through a Sun-3 controller using the AMD LANCE (Local Area Network Controller for Ethernet) Am7990 chip. For a general description of network interfaces see if(4N).

The synopsis line above specifies the first and only Ethernet controller on a Sun-3/50 system.

SEE ALSO

if(4N), kb(4S), tty_compact(4)

DIAGNOSTICS**le%d: transmitter frozen — resetting**

A bug in the LANCE chip has stopped the chip's transmitter section. The driver has detected this condition and reinitialized the chip.

le%d: out of mbufs: output packet dropped

The driver has run out of memory to use to buffer packets on output. The packet being transmitted at the time of occurrence is lost. This error is usually symptomatic of trouble elsewhere in the kernel.

le%d: stray transmitter interrupt

The LANCE chip has signalled that it completed transmitting a packet but the driver has sent no such packet.

le%d: LANCE Rev C/D Extra Byte(s) bug; Packet dropped

The LANCE chip's internal silo pointers have become misaligned. This error arises from a chip bug.

le%d: trailer error

An incoming packet claimed to have a trailing header but did not.

le%d: runt packet

An incoming packet's size was below the Ethernet minimum transmission size.

le%d: Receive buffer error - BUFF bit set in rmd

This error "should never happen," as it occurs only in conjunction with a LANCE feature that the driver does not use.

le%d: Received packet with STP bit in rmd cleared

The driver has received a packet that straddles multiple receive buffers and therefore consumes more than one of the LANCE chip's receive descriptors. Provided that all stations on the Ethernet are operating according to the Ethernet specification, this error "should never happen," since the driver allocates its receive buffers to be large enough to hold packets of the largest permitted size. Most likely, some other station on the net is transmitting packets whose lengths exceed the maximum permitted for Ethernet.

le%d: Received packet with ENP bit in rmd cleared

The driver has received a packet that straddles multiple receive buffers and therefore consumes more than one of the LANCE chip's receive descriptors. Provided that all stations on the Ethernet are operating according to the Ethernet specification, this error "should never happen," since the driver allocates its receive buffers to be large enough to hold packets of the largest permitted size. The most likely cause of the message is that some other station on the net is transmitting packets whose lengths exceed the maximum permitted for Ethernet.

le%d: Transmit buffer error - BUFF bit set in tmd

Excessive bus contention has prevented the LANCE chip from gathering packet contents quickly enough to sustain the packet's transmission over the Ethernet. The affected packet is lost.

le%d: Transmit late collision - Net problem?

A packet collision has occurred after the channel's slot time has elapsed. This error usually indicates faulty hardware elsewhere on the net.

le%d: No carrier - transceiver cable problem?

The LANCE chip has lost input to its carrier detect pin while trying to transmit a packet.

le%d: Transmit retried more than 16 times - net jammed

Network activity has become so intense that sixteen successive transmission attempts failed, the LANCE chip gave up on the current packet.

le%d: missed packet

The driver has dropped an incoming packet because it had no buffer space for it.

le%d: Babble error - sent a packet longer than the maximum length

While transmitting a packet, the LANCE chip has noticed that the packet's length exceeds the maximum allowed for Ethernet. This error indicates a kernel bug.

le%d: Memory Error! Ethernet chip memory access timed out

The LANCE chip timed out while trying to acquire the bus for a DVMA transfer.

le%d: Reception stopped

Because of some other error, the receive section of the LANCE chip shut down and had to be restarted.

le%d: Transmission stopped

Because of some other error, the transmit section of the LANCE chip shut down and had to be restarted.

NAME

lo – software loopback network interface

SYNOPSIS

pseudo-device loop

DESCRIPTION

The **loop** device is a software loopback network interface; see **if(4N)** for a general description of network interfaces.

The **loop** interface is used for performance analysis and software testing, and to provide guaranteed access to Internet protocols on machines with no local network interfaces. A typical application is the **comsat(8C)** server which accepts notification of mail delivery through a particular port on the loopback interface.

By default, the loopback interface is accessible at Internet address 127.0.0.1 (non-standard); this address may be changed with the **SIOCSIFADDR** ioctl.

SEE ALSO

if(4N), **inet(4F)**, **comsat(8C)**

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.

BUGS

It should handle all address and protocol families. An approved network address should be reserved for this interface.

NAME

lofs – loopback virtual file system

CONFIG

options LOFS

SYNOPSIS

```
#include <sys/mount.h>
mount(MOUNT_LOFS, virtual, flags, dir);
```

virtual is the mount point for the virtual file system. *dir* is the pathname of the existing file system. *flags* is either 0 or M_RDONLY. The M_RDONLY flag forces all accesses in the new name space to be read-only; without it, accesses are the same as for the underlying file system. All other `mount(2)` flags are preserved from the underlying file systems.

DESCRIPTION

The loopback filesystem device allows new, virtual, file systems to be created, which provide access to existing files using alternate pathnames. Once the virtual file system is created, other file systems can be mounted within it without affecting the original file system. File systems that are subsequently mounted onto the original filesystem, however, *are* visible to the virtual file system, unless or until the corresponding mount point in the virtual file system is covered by a file system mounted there.

For instance, a loopback mount of / onto `/tmp/newroot` allows the entire filesystem hierarchy to appear as if it were duplicated under `/tmp/newroot`, including any file systems mounted from remote NFS servers. All files would then be accessible either from a pathname relative to /, or from a pathname relative to `/tmp/newroot` until such time as a file system is mounted in `/tmp/newroot`, or any of its subdirectories.

Loopback mounts of / can be performed in conjunction with the `chroot(2)` system call, to provide a complete virtual filesystem to a process or family of processes.

Recursive traversal of loopback mount points is not allowed; after the loopback mount of `/tmp/newroot`, the file `/tmp/newroot/tmp/newroot` does not contain yet another filesystem hierarchy; rather, it appears just as `/tmp/newroot` did before the loopback mount was performed (say, as an empty directory).

SEE ALSO

`chroot(2)`, `mount(2)`

BUGS

Because only directories can be mounted or mounted on, the structure of a virtual file system can only be modified at directories.

Loopback mounts must be used with care; the potential for confusing users and applications is enormous.

NAME

mcp, alm – Sun MCP Multiprotocol Communications Processor/ALM-2 Asynchronous Line Multiplexer

CONFIG — SUN-3 SYSTEM

MCP

```
device mcp0 at vme32d32 ? csr 0x1000000 flags 0x1ffff priority 4 vector mcpintr 0x8b
device mcp1 at vme32d32 ? csr 0x1010000 flags 0x1ffff priority 4 vector mcpintr 0x8a
device mcp2 at vme32d32 ? csr 0x1020000 flags 0x1ffff priority 4 vector mcpintr 0x89
device mcp3 at vme32d32 ? csr 0x1030000 flags 0x1ffff priority 4 vector mcpintr 0x88
```

ALM-2

```
pseudo-device mcpa64
```

SYNOPSIS

```
#include <fcntl.h>
#include <sys/termios.h>
open(“/dev/ttyxy”, mode);
open(“/dev/ttydn”, mode);
open(“/dev/cuan”, mode);
```

DESCRIPTION (MCP)

The Sun MCP (Multiprotocol Communications Processor) supports up to four synchronous serial lines in conjunction with SunLink™ Multiple Communication Protocol products.

DESCRIPTION (ALM-2)

The Sun ALM-2 Asynchronous Line Multiplexer provides 16 asynchronous serial communication lines with modem control and one Centronics-compatible parallel printer port.

Each port supports those `termio(4)` device control functions specified by flags in the `c_cflag` word of the `termios` structure and by the `IGNBRK`, `IGNPAR`, `PARMRK`, or `INPCK` flags in the `c_iflag` word of the `termios` structure are performed by the `mcp` driver. All other `termio(4)` functions must be performed by STREAMS modules pushed atop the driver; when a device is opened, the `ldterm(4M)` and `ttcompat(4M)` STREAMS modules are automatically pushed on top of the stream, providing the standard `termio(4)` interface.

Bit *i* of flags may be specified to say that a line is not properly connected, and that the line *i* should be treated as hard-wired with carrier always present. Thus specifying flags `0x0004` in the specification of `mcp0` would treat line `/dev/ttyh2` in this way.

Minor device numbers in the range 0 – 63 correspond directly to the normal tty lines and are named `/dev/ttyXY`, where *X* represents the physical board as one of the characters `h`, `i`, `j`, or `k`, and *Y* is the line number on the board as a single hexadecimal digit. (Thus the first line on the first board is `/dev/ttyh0`, and the sixteenth line on the third board is `/dev/ttyjf`.)

To allow a single tty line to be connected to a modem and used for both incoming and outgoing calls, a special feature, controlled by the minor device number, has been added. Minor device numbers in the range 128 – 191 correspond to the same physical lines as those above (that is, the same line as the minor device number minus 128).

A dial-in line has a minor device in the range 0 – 63 and is conventionally renamed `/dev/ttydn`, where *n* is a number indicating which dial-in line it is (so that `/dev/ttyd0` is the first dial-in line), and the dial-out line corresponding to that dial-in line has a minor device number 128 greater than the minor device number of the dial-in line and is conventionally named `/dev/cuan`, where *n* is the number of the dial-in line.

The `/dev/cuan` lines are special in that they can be opened even when there is no carrier on the line. Once a `/dev/cuan` line is opened, the corresponding tty line cannot be opened until the `/dev/cuan` line is closed; a blocking open will wait until the `/dev/cuan` line is closed (which will drop Data Terminal Ready, after which Carrier Detect will usually drop as well) and carrier is detected again, and a non-blocking open will return an error. Also, if the `/dev/ttydn` line has been opened successfully (usually only when carrier is recognized on the modem) the corresponding `/dev/cuan` line cannot be opened. This allows a modem to be

attached to e.g. `/dev/ttyd0` (renamed from `/dev/ttyh0`) and used for dialin (by enabling the line for login in `/etc/ttytab`) and also used for dialout (by `tip(1C)` or `uucp(1C)`) as `/dev/cua0` when no one is logged in on the line. Note: the bit in the `flags` word in the configuration file (see above) must be zero for this line, which enables hardware carrier detection.

IOCTLS

The standard set of `termio ioctl()` calls are supported by the ALM-2.

If the `CRTSCTS` flag in the `c_cflag` is set, output will be generated only if CTS is high; if CTS is low, output will be frozen. If the `CRTSCTS` flag is clear, the state of CTS has no effect. Breaks can be generated by the `TCSBRK`, `TIOCSBRK`, and `TIOCCBRK` `ioctl()` calls. The modem control lines `TIOCM_CAR`, `TIOCM_CTS`, `TIOCM_RTS`, and `TIOCM_DTR` are provided.

The input and output line speeds may be set to any of the speeds supported by `termio`. The speeds cannot be set independently; when the output speed is set, the input speed is set to the same speed.

ERRORS

An `open()` on a `/dev/tty*` or a `/dev/cu*` device will fail if:

<code>ENXIO</code>	The unit being opened does not exist.
<code>EBUSY</code>	The dial-out device is being opened and the dial-in device is already open, or the dial-in device is being opened with a no-delay open and the dial-out device is already open.
<code>EBUSY</code>	The unit has been marked as exclusive-use by another process with a <code>TIOCEXCL</code> <code>ioctl()</code> call.
<code>EINTR</code>	The open was interrupted by the delivery of a signal.

DESCRIPTION (PRINTER PORT)

The printer port is Centronics-compatible and is suitable for most common parallel printers. Devices attached to this interface are normally handled by the line printer spooling system, and should not be accessed directly by the user.

Minor device numbers in the range 64 – 67 access the printer port, and the recommended naming is `/dev/mcpp[0-3]`.

IOCTLS

Various control flags and status bits may be fetched and set on an MCP printer port. The following flags and status bits are supported; they are defined in `<sundev/mcpcmd.h>`:

<code>MCPRIGNSLCT</code>	0x02	set if interface ignoring <code>SLCT</code> — on open
<code>MCPRDIAG</code>	0x04	set if printer is in self-test mode
<code>MCPRVMEINT</code>	0x08	set if VME bus interrupts enabled
<code>MCPRINTPE</code>	0x10	print message when out of paper
<code>MCPRINTSLCT</code>	0x20	print message when printer offline
<code>MCPRPE</code>	0x40	set if device ready, cleared if device out of paper
<code>MCPRSLCT</code>	0x80	set if device online (Centronics <code>SLCT</code> asserted)

The flags `MCPRINTSLCT`, `MCPRINTPE`, and `MCPRDIAG` may be changed; the other bits are status bits and may not be changed.

The `ioctl()` calls supported by MCP printer ports are listed below.

<code>MCPIOGPR</code>	The argument is a pointer to an unsigned char . The printer flags and status bits are stored in the unsigned char pointed to by the argument.
<code>MCPIOSPR</code>	The argument is a pointer to an unsigned char . The printer flags are set from the unsigned char pointed to by the argument.

ERRORS

Normally, the interface only reports the status of the device when attempting an `open(2V)` call. An `open()` on a `/dev/mcpp*` device will fail if:

ENXIO The unit being opened does not exist.

EIO The device is offline or out of paper.

Bit 17 of the configuration flags may be specified to say that the interface should ignore Centronics SLCT- and RDY/PE- when attempting to open the device, but this is normally useful only for configuration and troubleshooting: if the SLCT- and RDY lines are not asserted during an actual data transfer (as with a write(2V) call), no data is transferred.

FILES

/dev/mcphp[0-3]	parallel printer port
/dev/tty[h-k][0-9a-f]	hardwired tty lines
/dev/ttyd[0-9a-f]	dialin tty lines
/dev/cua[0-9a-f]	dialout tty lines

SEE ALSO

tip(1C), uucp(1C), mti(4S), termio(4), ldterm(4M), ttcompat(4M), zs(4S)

DIAGNOSTICS

Most of these diagnostics "should never happen;" their occurrence usually indicates problems elsewhere in the system as well.

mcpan: silo overflow.

More than *n* characters (*n* very large) have been received by the mcp hardware without being read by the software.

*****port *n* supports RS449 interface*****

Probably an incorrect jumper configuration. Consult the hardware manual.

mcp port *n* receive buffer error

The mcp encountered an error concerning the synchronous receive buffer.

Printer on mcphp *n* is out of paper

Printer on mcphp *n* paper ok

Printer on mcphp *n* is offline

Printer on mcphp *n* online

Assorted printer diagnostics, if enabled as discussed above.

NAME

mem, kmem, vme16d16, vme24d16, vme32d16, vme16d32, vme24d32, vme32d32, mbmem, mbio, atbus, zero, eeprom – main memory and bus I/O space

CONFIG

None; included with standard system.

DESCRIPTION

These devices are special files that map memory and bus I/O space. They may be read, written, seeked and (except for **kmem**) memory-mapped. See **read(2V)**, **write(2V)**, **mmap(2)**, and **directory(3)**,

mem is a special file that is an image of the physical memory of the computer. It may be used, for example, to examine (and even to patch) the system.

kmem is a special file that is an image of the kernel virtual memory of the system.

kmem is a special file which is a source of private zero pages.

Sun-2 and Sun-3 System

vme16d16 (also known as **vme16**) is a special file that is an image of VMEbus 16-bit addresses with 16-bit data. **vme16** address space extends from 0 to 64K.

vme24d16 (also known as **vme24**) is a special file that is an image of VMEbus 24-bit addresses with 16-bit data. **vme24** address space extends from 0 to 16 Megabytes. The VME 16-bit address space overlaps the top 64K of the 24-bit address space.

Sun-3 VMEbus

vme32d16 is a special file that is an image of VMEbus 32-bit addresses with 16-bit data.

vme16d32 is a special file that is an image of VMEbus 16-bit addresses with 32-bit data.

vme24d32 is a special file that is an image of VMEbus 24-bit addresses with 32-bit data.

vme32d32 (also known as **vme32**) is a special file that is an image of VMEbus 32-bit addresses with 32-bit data. **vme32** address space extends from 0 to 4 Giggabytes. The VME 24-bit address space overlaps the top 16 Megabytes of the 32-bit address space.

vme* type special files can only be accessed in VME based systems.

Sun-2 Multibus

mbmem is a special file that is an image of the Multibus memory of the system. Multibus memory is in the range from 0 to 16 Megabytes. **mbmem** can only be accessed in Multibus based systems.

mbio is a special file that is an image of the Multibus I/O space. Multibus I/O space extends from 0 to 64K. **mbio** can only be accessed in Multibus based systems.

When reading and writing **mbmem** and **mbio** odd counts or offsets cause byte accesses and even counts and offsets cause word accesses.

Sun386i

atbus is a special file that is an image of the AT bus space. It extends from 0 to 16 Megabytes.

eeprom is a special file that is an image of the NVRAM. It extends from 0 to 2Kb.

FILES

/dev/mem
 /dev/kmem
 /dev/mbmem
 /dev/mbio
 /dev/vme16d16
 /dev/vme16
 /dev/vme24d16
 /dev/vme24
 /dev/vme32d16

/dev/vme16d32
/dev/vme24d32
/dev/vme32d32
/dev/vme32
/dev/atbus
/dev/zero
/dev/eeprom

SEE ALSO

mmap(2), read(2V), write(2V), directory(3)

NAME

mouse – Sun mouse

CONFIG

None; included in standard system.

DESCRIPTION

The mouse indirect device provides access to the Sun Workstation mouse. When opened, it redirects operations to the standard mouse device for the workstation (attached either to a CPU serial or parallel port), and pushes the `ms(4M)` and `ttcompat(4M)` STREAMS modules on top of that device.

FILES

`/dev/mouse`

SEE ALSO

`ms(4M)`, `ttcompat(4M)`, `win(4S)`, `zs(4S)`

NAME

`ms` – Sun mouse STREAMS module

CONFIG

`pseudo-device:msn`

SYNOPSIS

```
#include <sys/stream.h>
#include <sys/stropt.h>
#include <sundev/vuid_event.h>
#include <sundev/msio.h>
ioctl(fd, I_PUSH, "ms");
```

DESCRIPTION

The `ms` STREAMS module processes byte streams generated by mice attached to a CPU serial or parallel port. When this module is pushed onto a stream, it sends a `TCSETS` `ioctl` downstream, setting the baud rate to 1200 baud and the character size to 8 bits, and enabling the receiver. All other flag words are cleared. It assumes only that the `termios(4)` functions provided by the `zs(4S)` driver are supported; no other functions need be supported.

The mouse is expected to generate a stream of bytes encoding mouse motions and changes in the state of the buttons.

Each mouse sample in the byte stream consists of three bytes: the first byte gives the button state with value `0x87|but`, where `but` is the low three bits giving the mouse buttons, where a 0 (zero) bit means that a button is pressed, and a 1 (one) bit means a button is not pressed. Thus if the left button is down the value of this sample is `0x83`, while if the right button is down the byte is `0x86`.

The next two bytes of each sample give the `x` and `y` deltas of this sample as signed bytes. The mouse uses a lower-left coordinate system, so moves to the right on the screen yield positive `x` values and moves down the screen yield negative `y` values.

The beginning of a sample is identifiable because the delta's are constrained to not have values in the range `0x80-0x87`.

A stream with `ms` pushed onto it can be used as a device that emits `firm_events` as specified by the protocol of a *Virtual User Input Device*. It understands `VIDSFORMAT`, `VIDGFORMAT`, `VIDSADDR` and `VIDGADDR` `ioctls` (see reference below).

IOCTLS

`ms` responds to the following `ioctls`, as defined in `<sundev/msio.h>` and `<sundev/vuid_event.h>`. All other `ioctls` are passed downstream. As `ms` sets the parameters of the serial port when it is opened, no `termios(4)` `ioctls` should be performed on a stream with `ms` on it, as `ms` expects the device parameters to remain as it set them.

The `MSIOGETPARMS` and `MSIOSETPARMS` calls use a structure of type `Ms_parms`, which is a structure defined in `<sundev/msio.h>`:

```
typedef struct {
    int    jitter_thresh;
    int    speed_low;
    int    speed_limit;
} Ms_parms;
```

`jitter_thresh` is the "jitter threshold" of the mouse. Motions of fewer than `jitter_thresh` units along both axes that occur in less than 1/12 second are treated as "jitter" and ignored. Thus, if the mouse moves fewer than `jitter_thresh units` and then moves back to its original position in less than 1/12 of a second, the motion is considered to be "noise" and ignored. If it moves fewer than `jitter_thresh` units and continues to move so that it has not returned to its original position after 1/12 of a second, the motion is considered to be real and is reported.

speed_limit indicates whether extremely large motions are to be ignored. If it is 1, a "speed limit" is applied to mouse motions; motions along either axis of more than *speed_limit* units are discarded.

Note: these parameters are global; if they are set for any mouse on a workstation, they apply to any other mice attached to that workstation as well.

VIDSFORMAT

VIDGFORMAT

VIDSADDR

VIDGADDR

These are standard *Virtual User Input Device ioctls*. See *SunView 1 System Programmer's Guide* for a description of their operation.

MSIOGETPARMS

The argument is a pointer to a *Ms_parms*. The current mouse parameters are stored in that structure.

MSIOSETPARMS

The argument is a pointer to a *ms_parms*. The current mouse parameters are set from the values in that structure.

SEE ALSO

mouse(4S), termios(4), win(4S), zs(4S)

SunView 1 System Programmer's Guide

NAME

mti – Systech MTI-800/1600 multi-terminal interface

CONFIG — SUN-3 SYSTEM

```
device mti0 at vme16d16 ? csr 0x620 flags 0xffff priority 4 vector mtiintr 0x88
device mti1 at vme16d16 ? csr 0x640 flags 0xffff priority 4 vector mtiintr 0x89
device mti2 at vme16d16 ? csr 0x660 flags 0xffff priority 4 vector mtiintr 0x8a
device mti3 at vme16d16 ? csr 0x680 flags 0xffff priority 4 vector mtiintr 0x8b
```

CONFIG — SUN-2 SYSTEM

```
device mti0 at mbio ? csr 0x620 flags 0xffff priority 4
device mti1 at mbio ? csr 0x640 flags 0xffff priority 4
device mti2 at mbio ? csr 0x660 flags 0xffff priority 4
device mti3 at mbio ? csr 0x680 flags 0xffff priority 4
device mti0 at vme16 ? csr 0x620 flags 0xffff priority 4 vector mtiintr 0x88
device mti1 at vme16 ? csr 0x640 flags 0xffff priority 4 vector mtiintr 0x89
device mti2 at vme16 ? csr 0x660 flags 0xffff priority 4 vector mtiintr 0x8a
device mti3 at vme16 ? csr 0x680 flags 0xffff priority 4 vector mtiintr 0x8b
```

SYNOPSIS

```
#include <fcntl.h>
#include <sys/termios.h>
open("/dev/ttyxy", mode);
open("/dev/ttydn", mode);
open("/dev/cuan", mode);
```

DESCRIPTION

The Systech MTI card provides 8 (MTI-800) or 16 (MTI-1600) serial communication lines with modem control. Each port supports those `termio(4)` device control functions specified by flags in the `c_cflag` word of the `termios` structure and by the `IGNBRK`, `IGNPAR`, `PARMRK`, or `INPCK` flags in the `c_iflag` word of the `termios` structure are performed by the `mti` driver. All other `termio(4)` functions must be performed by `STREAMS` modules pushed atop the driver; when a device is opened, the `ldterm(4M)` and `ttcompat(4M)` `STREAMS` modules are automatically pushed on top of the stream, providing the standard `termio(4)` interface.

Bit *i* of `flags` may be specified to say that a line is not properly connected, and that the line *i* should be treated as hard-wired with carrier always present. Thus specifying `flags 0x0004` in the specification of `mti0` would treat line `/dev/tty02` in this way.

Minor device numbers in the range 0 – 63 correspond directly to the normal tty lines and are named `/dev/ttyXY`, where *X* is the physical board number (0 – 3), and *Y* is the line number on the board as a single hexadecimal digit. (Thus the first line on the first board is `/dev/tty00`, and the sixteenth line on the third board is `/dev/tty2f`.)

To allow a single tty line to be connected to a modem and used for both incoming and outgoing calls, a special feature, controlled by the minor device number, has been added. Minor device numbers in the range 128 – 191 correspond to the same physical lines as those above (that is, the same line as the minor device number minus 128).

A dial-in line has a minor device in the range 0 – 63 and is conventionally renamed `/dev/ttydn`, where *n* is a number indicating which dial-in line it is (so that `/dev/ttyd0` is the first dial-in line), and the dial-out line corresponding to that dial-in line has a minor device number 128 greater than the minor device number of the dial-in line and is conventionally named `/dev/cuan`, where *n* is the number of the dial-in line.

The `/dev/cuan` lines are special in that they can be opened even when there is no carrier on the line. Once a `/dev/cuan` line is opened, the corresponding tty line can not be opened until the `/dev/cuan` line is closed; a blocking open will wait until the `/dev/cuan` line is closed (which will drop Data Terminal Ready, after which Carrier Detect will usually drop as well) and carrier is detected again, and a non-blocking open will return an error. Also, if the `/dev/ttydn` line has been opened successfully (usually only when carrier is

recognized on the modem) the corresponding `/dev/cuan` line can not be opened. This allows a modem to be attached to e.g. `/dev/ttyd0` (renamed from `/dev/tty00`) and used for dialin (by enabling the line for login in `/etc/ttytab`) and also used for dialout (by `tip(1C)` or `uucp(1C)`) as `/dev/cua0` when no one is logged in on the line. Note: the bit in the `flags` word in the configuration file (see above) must be zero for this line, which enables hardware carrier detection.

WIRING

The Systech requires the CTS modem control signal to operate. If the device does not supply CTS then RTS should be jumpered to CTS at the distribution panel (short pins 4 to 5). Also, the CD (carrier detect) line does not work properly. When connecting a modem, the modem's CD line should be wired to DSR, which the software will treat as carrier detect.

IOCTLS

The standard set of `termio ioctl()` calls are supported by `mti`.

The state of the `CRTSCTS` flag in the `c_cflag` word has no effect; no output will be generated unless CTS is high. Breaks can be generated by the `TCSBRK`, `TIOCSBRK`, and `TIOCCBRK ioctl()` calls. The modem control lines `TIOCM_CAR`, `TIOCM_CTS`, `TIOCM_RTS`, and `TIOCM_DTR` are provided; however, as described above, the DSR line is treated as CD and the CD line is ignored.

The input and output line speeds may be set to any of the speeds supported by `termio`. The speeds cannot be set independently; when the output speed is set, the input speed is set to the same speed. The baud rates `B200` and `B38400` are not supported by the hardware; `B200` selects 2000 baud, and `B38400` selects 7200 baud.

ERRORS

An `open()` will fail if:

- | | |
|--------------------|---|
| <code>ENXIO</code> | The unit being opened does not exist. |
| <code>EBUSY</code> | The dial-out device is being opened and the dial-in device is already open, or the dial-in device is being opened with a no-delay open and the dial-out device is already open. |
| <code>EBUSY</code> | The unit has been marked as exclusive-use by another process with a <code>TIOCEXCL ioctl()</code> call. |
| <code>EINTR</code> | The open was interrupted by the delivery of a signal. |

FILES

<code>/dev/tty[0-3][0-9a-f]</code>	hardwired tty lines
<code>/dev/ttyd[0-9a-f]</code>	dialin tty lines
<code>/dev/cua[0-9a-f]</code>	dialout tty lines

SEE ALSO

`tip(1C)`, `uucp(1C)`, `mcp(4S)`, `termio(4)`, `ldterm(4M)`, `ttcompat(4M)`, `zs(4S)`

DIAGNOSTICS

Most of these diagnostics "should never happen" and their occurrence usually indicates problems elsewhere in the system.

`mtin, n`: silo overflow.

More than 512 characters have been received by the `mti` hardware without being read by the software. Extremely unlikely to occur.

`mtin`: read error code `<n>`. Probable hardware fault

The `mti` returned the indicated error code. See the MTI manual.

`mtin`: DMA output error.

The `mti` encountered an error while trying to do DMA output.

`mtin`: impossible response `n`.

The `mti` returned an error it could not understand.

NAME

mtio – general magnetic tape interface

SYNOPSIS

```
#include <sys/ioctl.h>
#include <sys/mtio.h>
```

DESCRIPTION

Both 1/2" and 1/4" magnetic tape drives share the same general interface, no matter what hardware is involved. The remainder of this section discusses the common features of this interface.

The "cooked" magnetic tape device files read and write magnetic tape in 2048 byte blocks (the 2048 is actually `BLKDEV_IOSIZE` in `<sys/param.h>`). The name of such a device file might be `/dev/mt0`. The final component of the name is composed of a name that represents the type of device the file refers to, and the unit number of that device.

These files are rewound when closed; the "no-rewind" versions of these files are not. The name of "no-rewind" device files include the letter `n` at the beginning of the final component of the name; the "no-rewind" version of `/dev/mt0` would be `/dev/nmt0`. When a 1/2" tape file, open for writing or just written, is closed, two tape marks are written; if the tape is not to be rewound it is positioned with the head between the two tapemarks. When a 1/4" tape file, (due to a bug, only if) just written, is closed, only one end of file mark is written because of the inability to overwrite data on a 1/4" tape; see below.

The files discussed above are useful when you want to access the tape in a way compatible with ordinary files. This interface requires that all blocks be 2048 bytes long, and does not permit special operations (such as spacing the tape forward or backward) to be performed. When using foreign tapes, and especially when reading or writing long records, the "raw" interface is appropriate. The name of "raw" device files include the letter `r` before the device type; the "raw" version of `/dev/mt0` would be `/dev/rmt0`, and the "raw" version of `/dev/nmt0` would be `/dev/nrmt0`. Each `read(2V)` or `write(2V)` 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. 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.

1/4" tapes are not able to back up and always write fixed sized blocks. Since they cannot back up, they cannot support backward space file and backward space record. Since they always write fixed sized blocks, the size of transfers using the "raw" interface must be a multiple of the underlying blocksize, usually 512 bytes.

1/4" tapes also have an unusual tape format. They have parallel tracks, but only record information on one track at a time, switching to another track near the physical end of the medium. They erase all the tracks at once while writing the first track. Therefore, they cannot, in general, overwrite previously written data. If the old data were not on the first track, it would not be erased before being overwritten, and the result would be unreadable.

A number of additional `ioctl` operations are available on "raw" devices. The following definitions are from `<sys/mtio.h>`:

```
/*
 * Structures and definitions for mag tape I/O 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 */
#define MTRETEN    8          /* retension the tape */
#define MTERASE    9          /* erase the entire tape */
#define MTEOM     10         /* position to end of media (SCSI 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     /* vax: unibus ts-11 */
#define MT_ISHT     0x02     /* vax: massbus tu77, etc */
#define MT_ISTM     0x03     /* vax: unibus tm-11 */
#define MT_ISMT     0x04     /* vax: massbus tu78 */
#define MT_ISUT     0x05     /* vax: unibus gcr */
#define MT_ISCPC    0x06     /* sun: Multibus tapemaster */
#define MT_ISAR     0x07     /* sun: Multibus archive */
#define MT_ISSC     0x08     /* sun: SCSI archive */
#define MT_ISXY     0x09     /* sun: Xylogics 472 */
#define MT_ISSYGEN  0x0a     /* sun: SCSI Sysgen */
#define MT_ISMT02   0x0b     /* sun: SCSI Emulex MT02 */
#define MT_ISCCS    0x0c     /* sun: SCSI generic (unknown) CCS */

/* 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

```

SEE ALSO

mt(1), tar(1), read(2V), write(2V), ar(4S), tm(4S), st(4S), xt(4S)

NAME

nfs, NFS – network file system

CONFIG

options NFS

DESCRIPTION

The Network File System, or NFS, allows a client workstation to perform transparent file access over the network. Using it, a client workstation can operate on files that reside on a variety of servers, server architectures and across a variety of operating systems. Client file access calls are converted to NFS protocol requests, and are sent to the server system over the network. The server receives the request, performs the actual file system operation, and sends a response back to the client.

The Network File System operates in a stateless fashion using remote procedure (RPC) calls built on top of external data representation (XDR) protocol. These protocols are documented in *Network Programming*. The RPC protocol provides for version and authentication parameters to be exchanged for security over the network.

A server can grant access to a specific filesystem to certain clients by adding an entry for that filesystem to the server's `/etc/exports` file.

A client gains access to that filesystem with the `mount(2)` system call, which requests a file handle for the filesystem itself. Once the filesystem is mounted by the client, the server issues a file handle to the client for each file (or directory) the client accesses. If the file is somehow removed on the server side, the file handle becomes stale (dissociated with a known file).

A server may also be a client with respect to filesystems it has mounted over the network, but its clients cannot gain access to those filesystems. Instead, the client must mount a filesystem directly from the server on which it resides.

The user ID and group ID mappings must be the same between client and server. However, the server maps uid 0 (the super-user) to uid -2 before performing access checks for a client. This inhibits super-user privileges on remote filesystems.

NFS-related routines and structure definitions are described in *Network Programming*.

ERRORS

Generally physical disk I/O errors detected at the server are returned to the client for action. If the server is down or inaccessible, the client will see the console message:

NFS: file server not responding: still trying.

The client continues (forever) to resend the request until it receives an acknowledgement from the server. This means the server can crash or power down, and come back up, without any special action required by the client. It also means the client process requesting the I/O will block and remain insensitive to signals, sleeping inside the kernel at `PRIBIO`.

FILES

`/etc/exports`

SEE ALSO

`mount(2)`, `exports(5)`, `fstab(5)`, `fstab(5)`, `mount(8)`, `nfsd(8)`

Network Programming

NAME

nit – Network Interface Tap

CONFIG

```

pseudo-device  clone
pseudo-device  snit
pseudo-device  pf
pseudo-device  nbuf

```

SYNOPSIS

```

#include <sys/file.h>
#include <sys/ioctl.h>
#include <net/nit_pf.h>
#include <net/nit_buf.h>

    fd = open("/dev/nit", mode);
    ioctl(fd, I_PUSH, "pf");
    ioctl(fd, I_PUSH, "nbuf");

```

DESCRIPTION

NIT (the Network Interface Tap) is a facility composed of several STREAMS modules and drivers. These components collectively provide facilities for constructing applications that require link-level network access. Examples of such applications include `rarpd(8C)`, which is a user-level implementation of the Reverse ARP protocol, and `etherfind(8C)`, which is a network monitoring and trouble-shooting program.

NIT consists of several components that are summarized below. See their Reference Manual entries for detailed information about their specification and operation.

- nit_if(4M)** This component is a STREAMS device driver that interacts directly with the system's Ethernet drivers. After opening an instance of this device it must be bound to a specific Ethernet interface before becoming usable. Subsequently, `nit_if` transcribes packets arriving on the interface to the read side of its associated stream and delivers messages reaching it on the write side of its stream to the raw packet output code for transmission over the interface.
- nit_pf(4M)** This module provides packet-filtering services, allowing uninteresting incoming packets to be discarded with minimal loss of efficiency. It passes through unaltered all outgoing messages (those on the stream's write side).
- nit_buf(4M)** This module buffers incoming messages into larger aggregates, thereby reducing the overhead incurred by repeated `read(2V)` system calls.

NIT clients mix and match these components, based on their particular requirements. For example, the reverse ARP daemon concerns itself only with packets of a specific type and deals with low traffic volumes. Thus, it uses `nit_if` for access to the network and `nit_pf` to filter out all incoming packets except reverse ARP packets, but omits the `nit_buf` buffering module since traffic isn't high enough to justify the additional complexity of unpacking buffered packets. On the other hand, the `etherd(8C)` program, which collects Ethernet statistics for `traffic(1C)` to display, must examine every packet on the network. Therefore, it omits the `nit_if` module, since there's nothing it wishes to screen out, and includes the `nit_buf` module, since most networks have very heavy aggregate packet traffic.

EXAMPLES

The following code fragments outline how to program against parts of the NIT interface. For the sake of brevity, all error-handling code has been elided.

`initdevice` comes from `etherfind` and sets up its input stream configuration.

```

initdevice(if_flags, snaplen, chunksize)
    u_long    if_flags,
             snaplen,
             chunksize;
{
    struct strioctl    si;

```

```

struct ifreq      ifr;
struct timeval    timeout;

if_fd = open(NIT_DEV, O_RDONLY);

/* Arrange to get discrete messages from the stream. */
ioctl(if_fd, I_SRDOPT, (char *)RMSGD);

sl.ic_timeout = INFTIM;

/* Push and configure the buffering module. */
ioctl(if_fd, I_PUSH, "nbuf");

timeout.tv_sec = 1;
timeout.tv_usec = 0;
sl.ic_cmd = NIOCSTIME;
sl.ic_len = sizeof timeout;
sl.ic_dp = (char *)&timeout;
ioctl(if_fd, I_STR, (char *)&sl);

sl.ic_cmd = NIOCSCHUNK;
sl.ic_len = sizeof chunksize;
sl.ic_dp = (char *)&chunksize;
ioctl(if_fd, I_STR, (char *)&sl);

/* Configure the nit device, binding it to the proper
   underlying interface, setting the snapshot length,
   and setting nit_if-level flags. */
strncpy(ifr.ifr_name, device, sizeof ifr.ifr_name);
ifr.ifr_name[sizeof ifr.ifr_name - 1] = ' ';
sl.ic_cmd = NIOCBIND;
sl.ic_len = sizeof ifr;
sl.ic_dp = (char *)&ifr;
ioctl(if_fd, I_STR, (char *)&sl);

if (snaplen > 0) {
    sl.ic_cmd = NIOCSSNAP;
    sl.ic_len = sizeof snaplen;
    sl.ic_dp = (char *)&snaplen;
    ioctl(if_fd, I_STR, (char *)&sl);
}

if (if_flags != 0) {
    sl.ic_cmd = NIOCSFLAGS;
    sl.ic_len = sizeof if_flags;
    sl.ic_dp = (char *)&if_flags;
    ioctl(if_fd, I_STR, (char *)&sl);
}

/* Flush the read queue, to get rid of anything that accumulated
   before the device reached its final configuration. */
ioctl(if_fd, I_FLUSH, (char *)FLUSHR);
}

```

Here is the skeleton of the packet reading loop from `etherfind`. It illustrates how to cope with dismantling the headers the various NIT components glue on.

```

while ((cc = read(if_fd, buf, chunksize)) >= 0) {
    register u_char      *bp = buf,
                        *bufstop = buf + cc;

    /* Loop through each message in the chunk. */
    while (bp < bufstop) {
        register u_char      *cp = bp;
        struct nit_bufhdr    *hdrp;

```

```

struct timeval          *tvp = NULL;
u_long                 drops = 0;
u_long                 pktlen;

/* Extract information from the successive objects
   embedded in the current message. Which ones we
   have depends on how we set up the stream (and
   therefore on what command line flags were set).

   If snaplen is positive then the packet was truncated
   before the buffering module saw it, so we must
   obtain its length from the nit_if-level nit_iflen
   header. Otherwise the value in *hdrp suffices. */
hdrp = (struct nit_bufhdr *)cp;
cp += sizeof *hdrp;
if (tflag) {
    struct nit_iftime   *ntp;

    ntp = (struct nit_iftime *)cp;
    cp += sizeof *ntp;

    tvp = &ntp->nh_timestamp;
}
if (dflag) {
    struct nit_ifdrops  *ndp;

    ndp = (struct nit_ifdrops *)cp;
    cp += sizeof *ndp;

    drops = ndp->nh_drops;
}
if (snaplen > 0) {
    struct nit_iflen    *nlp;

    nlp = (struct nit_iflen *)cp;
    cp += sizeof *nlp;

    pktlen = nlp->nh_pktlen;
}
else
    pktlen = hdrp->nbb_msglen;

sp = (struct sample *)cp;
bp += hdrp->nbb_totlen;

/* Process the packet. */
}
}

```

FILES

/dev/nit clone device instance referring to nit_if

SEE ALSO

traffic(1C), read(2V), nit_if(4M), nit_pf(4M), nit_buf(4M), etherd(8C), etherfind(8C), rarpd(8C)

NAME

`nit_buf` – STREAMS NIT buffering module

CONFIG

pseudo-device `nbuf`

SYNOPSIS

```
#include <sys/ioctl.h>
#include <net/nit_buf.h>
ioctl(fd, I_PUSH, "nbuf");
```

DESCRIPTION

`nit_buf` is a STREAMS module that buffers incoming messages, thereby reducing the number of system calls and associated overhead required to read and process them. Although designed to be used in conjunction with the other components of NIT (see `nit(4P)`), `nit_buf` is a general-purpose module and can be used anywhere STREAMS input buffering is required.

Read-side Behavior

`nit_buf` collects incoming `M_DATA` and `M_PROTO` messages into *chunks*, passing each chunk upward when either the chunk becomes full or the current read timeout expires. When a message arrives, it is processed in two steps. First, the message is prepared for inclusion in a chunk, and then it is added to the current chunk. The following paragraphs discuss each step in turn.

Upon receiving a message from below, `nit_buf` immediately converts all leading `M_PROTO` blocks in the message to `M_DATA` blocks, altering only the message type field and leaving the contents alone. It then prepends a header to the converted message. This header is defined as follows.

```
struct nit_bufhdr {
    u_int   nhb_msglen;
    u_int   nhb_totlen;
};
```

The first field of this header gives the length in bytes of the converted message. The second field gives the distance in bytes from the start of the message in the current chunk (described below) to the start of the next message in the chunk; the value reflects any padding necessary to insure correct data alignment for the host machine and includes the length of the header itself.

After preparing a message, `nit_buf` attempts to add it to the end of the current chunk, using the chunk size and timeout values to govern the addition. (The chunk size and timeout values are set and inspected using the `ioctl` calls described below.) If adding the new message would make the current chunk grow larger than the chunk size, `nit_buf` closes off the current chunk, passing it up to the next module in line, and starts a new chunk, seeding it with a zero-length message. If adding the message would still make the current chunk overflow, the module passes it upward in an over-size chunk of its own. Otherwise, the module concatenates the message to the end of the current chunk.

To ensure that messages do not languish forever in an accumulating chunk, `nit_buf` maintains a read timeout. Whenever this timeout expires, the module closes off the current chunk, regardless of its length, and passes it upward; if no incoming messages have arrived, the chunk passed upward will have zero length. Whenever the module passes a chunk upward, it restarts the timeout period. These two rules insure that `nit_buf` minimizes the number of chunks it produces during periods of intense message activity and that it periodically disposes of all messages during slack intervals.

`nit_buf` handles other message types as follows. Upon receiving an `M_FLUSH` message specifying that the read queue be flushed, the module does so, clearing the currently accumulating chunk as well, and passes the message on to the module or driver above. It passes all other messages through unaltered to its upper neighbor.

Write-side Behavior

`nit_buf` intercepts `M_IOCTL` messages for the *ioctls* described below. Upon receiving an `M_FLUSH` message specifying that the write queue be flushed, the module does so and passes the message on to the module or driver below. The module passes all other messages through unaltered to its lower neighbor.

IOCTLS

nit_buf responds to the following *ioctl*s.

- NIOCSTIME** Set the read timeout value to the value referred to by the *struct timeval* pointer given as argument. Setting the timeout value to zero has the side-effect of forcing the chunk size to zero as well, so that the module will pass all incoming messages upward immediately upon arrival.
- NIOCGTIME** Return the read timeout in the *struct timeval* pointed to by the argument. If the timeout has been cleared with the **NIOCCTIME** *ioctl*, return with an ERANGE error.
- NIOCCTIME** Clear the read timeout, effectively setting its value to infinity.
- NIOCSCHUNK** Set the chunk size to the value referred to by the *u_int* pointer given as argument.
- NIOCGCHUNK** Return the chunk size in the *u_int* pointed to by the argument.

CAVEAT

The module name “nbuf” used in the system configuration file and as argument to the **I_PUSH** *ioctl* is provisional and subject to change.

SEE ALSO

nit(4P), **nit_if(4M)**, **nit_pf(4M)**

NAME

`nit_if` – STREAMS NIT device interface module

CONFIG

`pseudo-device snit`

SYNOPSIS

```
#include <sys/file.h>
open("/dev/nit", mode);
```

DESCRIPTION

`nit_if` is a STREAMS pseudo-device driver that provides STREAMS access to network interfaces. It is designed to be used in conjunction with the other components of NIT (see `nit(4P)`), but can be used by itself as a raw STREAMS network interface.

`nit_if` is an exclusive-open device that is intended to be opened indirectly through the clone device; `/dev/nit` is a suitable instance of the clone device. Before the stream resulting from opening an instance of `nit_if` may be used to read or write packets, it must first be bound to a specific network interface, using the `NIOCSBIND` `ioctl` described below.

Read-side Behavior

`nit_if` copies leading prefixes of selected packets from its associated network interface and passes them up the stream. If the `NI_PROMISC` flag is set, it passes along all packets; otherwise it passes along only packets addressed to the underlying interface.

The amount of data copied from a given packet depends on the current *snapshot length*, which is set with the `NIOCSSNAP` `ioctl` described below.

Before passing each packet prefix upward, `nit_if` optionally prepends one or more headers, as controlled by the state of the flag bits set with the `NIOCSFLAGS` `ioctl`. The driver collects headers into `M_PROTO` message blocks, with the headers guaranteed to be completely contained in a single message block, whereas the packet itself goes into one or more `M_DATA` message blocks.

Write-side Behavior

`nit_if` accepts packets from the module above it in the stream and relays them to the associated network interface for transmission. Packets must be formatted with the destination address in a leading `M_PROTO` message block, followed by the packet itself, complete with link-level header, in a sequence of `M_DATA` message blocks. The destination address must be expressed as a `'struct sockaddr'` whose `sa_family` field is `AF_UNSPEC` and whose `sa_data` field is a copy of the link-level header. (See `<sys/socket.h>` for the definition of this structure.)

`nit_if` processes `M_IOCTL` messages as described below. Upon receiving an `M_FLUSH` message specifying that the write queue be flushed, `nit_if` does so and transfers the message to the read side of the stream. It discards all other messages.

IOCTLS

`nit_if` responds to the following `ioctls`, as defined in `<net/nit_if.h>`. It generates an `M_IOCNAK` message for all others, returning this message to the invoker along the read side of the stream.

<code>SIOCGIFADDR</code>	<code>nit_if</code> passes this <code>ioctl</code> on to the underlying interface's driver and returns its response in a <code>'struct ifreq'</code> instance, as defined in <code><net/if.h></code> . (See the description of this <code>ioctl</code> in <code>if(4N)</code> for more details.)
<code>NIOCBIND</code>	This <code>ioctl</code> attaches the stream represented by its first argument to the network interface designated by its third argument, which should be a pointer to an <code>ifreq</code> structure whose <code>ifr_name</code> field names the desired interface. See <code><net/if.h></code> for the definition of this structure.
<code>NIOCSSNAP</code>	Set the current snapshot length to the value given in the <code>u_long</code> pointed to by the <code>ioctl</code> 's final argument. <code>nit_if</code> interprets a snapshot length value of zero as meaning infinity, so that it will copy all selected packets in their entirety. It constrains

positive snapshot lengths to be at least the length of an Ethernet header, so that it will pass at least the link-level header of all selected packets to its upstream neighbor.

NIOCGSNAP Returns the current snapshot length for this device instance in the *u_long* pointed to by the *ioctl*'s final argument.

NIOCSFLAGS *nit_if* recognizes the following flag bits, which must be given in the *u_long* pointed to by the *ioctl*'s final argument. This set may be augmented in future releases. All but the **NI_PROMISC** bit control the addition of headers that precede the packet body. These headers appear in the order given below, with the last-mentioned enabled header adjacent to the packet body.

NI_PROMISC Requests that the underlying interface be set into promiscuous mode and that all packets that the interface receives be passed up through the stream. *nit_if* only honors this bit for the super-user.

NI_TIMESTAMP Prepend to each selected packet a header containing the packet arrival time expressed as a 'struct timeval'.

NI_DROPS Prepend to each selected packet a header containing the cumulative number of packets that this instance of *nit_if* has dropped because of flow control requirements or resource exhaustion. The header value is expressed as a *u_long*. Note: it accounts only for events occurring within *nit_if*, and does not count packets dropped at the network interface level or by upstream modules.

NI_LEN Prepend to each selected packet a header containing the packet's original length (including link-level header), as it was before being trimmed to the snapshot length. The header value is expressed as a *u_long*.

NIOCGFLAGS Returns the current state of the flag bits for this device instance in the *u_long* pointed to by the *ioctl*'s final argument.

FILES

/dev/nit clone device instance referring to *nit_if* device
<net/nit_if.h> header file containing definitions for the *ioctls* and packet headers described above.

SEE ALSO

clone(4), *nit(4P)*, *nit_buf(4M)*, *nit_pf(4M)*

NAME

`nif_pf` – STREAMS NIT packet filtering module

CONFIG

`pseudo-device pf`

SYNOPSIS

```
#include <sys/ioctl.h>
#include <net/nit_pf.h>
    ioctl(fd, I_PUSH, "pf");
```

DESCRIPTION

`nit_pf` is a STREAMS module that subjects messages arriving on its read queue to a packet filter and passes only those messages that the filter accepts on to its upstream neighbor. Such filtering can be very useful for user-level protocol implementations and for networking monitoring programs that wish to view only specific types of events.

Read-side Behavior

`nit_pf` applies the current packet filter to all `M_DATA` and `M_PROTO` messages arriving on its read queue. The module prepares these messages for examination by first skipping over all leading `M_PROTO` message blocks to arrive at the beginning of the message's data portion. If there is no data portion, `nit_pf` accepts the message and passes it along to its upstream neighbor. Otherwise, the module ensures that the part of the message's data that the packet filter might examine lies in contiguous memory, calling the `pullupmsg` utility routine if necessary to force contiguity. (Note: this action destroys any sharing relationships that the subject message might have had with other messages.) Finally, it applies the packet filter to the message's data, passing the entire message upstream to the next module if the filter accepts, and discarding the message otherwise. See **PACKET FILTERS** below for details on how the filter works.

If there is no packet filter yet in effect, the module acts as if the filter exists but does nothing, implying that all incoming messages are accepted. **IOCTLS** below describes how to associate a packet filter with an instance of `nit_pf`.

`nit_pf` handles other message types as follows. Upon receiving an `M_FLUSH` message specifying that the read queue be flushed, the module does so, and passes the message on to its upstream neighbor. It passes all other messages through unaltered to its upper neighbor.

Write-side Behavior

`nit_pf` intercepts `M_IOCTL` messages for the `ioctl` described below. Upon receiving an `M_FLUSH` message specifying that the write queue be flushed, the module does so and passes the message on to the module or driver below. The module passes all other messages through unaltered to its lower neighbor.

IOCTLS

`nit_pf` responds to the following `ioctl`.

NIOCSETF This `ioctl` directs the module to replace its current packet filter, if any, with the filter specified by the 'struct `packetfilt`' pointer named by its final argument. This structure is defined in `<net/packetfilt.h>` as

```
struct packetfilt {
    u_char  Pf_Priority; /* priority of filter */
    u_char  Pf_FilterLen; /* # of cmds in list */
    u_short Pf_Filter[ENMAXFILTERS];
                                /* filter command list */
};
```


The *Pf_Priority* field is included only for compatibility with other packet filter implementations and is otherwise ignored. The packet filter itself is specified in the *Pf_Filter* array as a sequence of two-byte commands, with the *Pf_FilterLen* field giving the number of commands in the sequence. This implementation restricts the maximum number of commands in a filter (ENMAXFILTERS) to 40. The next section describes the available commands and their semantics.

PACKET FILTERS

A packet filter consists of the filter command list length (in units of *u_shorts*), and the filter command list itself. (The priority field mentioned above is ignored in this implementation.) Each filter command list specifies a sequence of actions that operate on an internal stack of *u_shorts* ("shortwords"). Each shortword of the command list specifies one of the actions ENF_PUSHLIT, ENF_PUSHZERO, or ENF_PUSHWORD+n, which respectively push the next shortword of the command list, zero, or shortword *n* of the subject message on the stack, and a binary operator from the set { ENF_EQ, ENF_NEQ, ENF_LT, ENF_LE, ENF_GT, ENF_GE, ENF_AND, ENF_OR, ENF_XOR } which then operates on the top two elements of the stack and replaces them with its result. When both an action and operator are specified in the same shortword, the action is performed followed by the operation.

The binary operator can also be from the set { ENF_COR, ENF_CAND, ENF_CNOR, ENF_CNAND }. These are "short-circuit" operators, in that they terminate the execution of the filter immediately if the condition they are checking for is found, and continue otherwise. All pop two elements from the stack and compare them for equality; ENF_CAND returns false if the result is false; ENF_COR returns true if the result is true; ENF_CNAND returns true if the result is false; ENF_CNOR returns false if the result is true. Unlike the other binary operators, these four do not leave a result on the stack, even if they continue.

The short-circuit operators should be used when possible, to reduce the amount of time spent evaluating filters. When they are used, you should also arrange the order of the tests so that the filter will succeed or fail as soon as possible; for example, checking the IP destination field of a UDP packet is more likely to indicate failure than the packet type field.

The special action ENF_NOPUSH and the special operator ENF_NOP can be used to only perform the binary operation or to only push a value on the stack. Since both are (conveniently) defined to be zero, indicating only an action actually specifies the action followed by ENF_NOP, and indicating only an operation actually specifies ENF_NOPUSH followed by the operation.

After executing the filter command list, a non-zero value (true) left on top of the stack (or an empty stack) causes the incoming packet to be accepted and a zero value (false) causes the packet to be rejected. (If the filter exits as the result of a short-circuit operator, the top-of-stack value is ignored.) Specifying an undefined operation or action in the command list or performing an illegal operation or action (such as pushing a shortword offset past the end of the packet or executing a binary operator with fewer than two shortwords on the stack) causes a filter to reject the packet.

EXAMPLES

The reverse ARP daemon program (rarpd(8C)) uses code similar to the following fragment to construct a filter that rejects all but RARP packets. That is, it accepts only packets whose Ethernet type field has the value ETHERTYPE_REVARP.

```

struct ether_header eh;          /* used only for offset values */
struct packetfilt pf;
register u_short *fwp = pf.Pf_Filter;
u_short offset;

/*
 * Set up filter. Offset is the displacement of the Ethernet
 * type field from the beginning of the packet in units of
 * u_shorts.
 */

```

```

        offset = ((u_int) &eh.ether_type - (u_int) &eh.ether_dhost) /
sizeof (u_short);
        *fwp++ = ENF_PUSHWORD + offset;
        *fwp++ = ENF_PUSHLIT;
        *fwp++ = htons(ETHERTYPE_REVARP);
        *fwp++ = ENF_EQ;
        pf.Pf_FilterLen = fwp - &pf.Pf_Filter[0];

```

This filter can be abbreviated by taking advantage of the ability to combine actions and operations:

```

...
        *fwp++ = ENF_PUSHWORD + offset;
        *fwp++ = ENF_PUSHLIT | ENF_EQ;
        *fwp++ = htons(ETHERTYPE_REVARP);
...

```

WARNINGS

The module name 'pf' used in the system configuration file and as argument to the `I_PUSH ioctl` is provisional and subject to change.

The `Pf_Priority` field of the `packetfilt` structure is likely to be removed.

SEE ALSO

`inet(4F)`, `nit(4P)`, `nit_buf(4M)`, `nit_if(4M)`

NAME

null – data sink

CONFIG

None; included with standard system.

SYNOPSIS

```
#include <fcntl.h>
```

```
open("/dev/null", mode);
```

DESCRIPTION

Data written on the **null** special file is discarded.

Reads from the **null** special file always return an end-of-file indication.

FILES

/dev/null

NAME

pp – Centronics-compatible parallel printer port

CONFIG

device pp0 at obio ? csr

AVAILABILITY

Sun386i systems only.

DESCRIPTION

This device driver provides an interface to the Sun386i system's on-board Centronics-compatible parallel printer port. It supports most standard PC printers with Centronics interfaces.

FILES

/dev/pp0

DIAGNOSTICS

pp*: printer not online

pp*: printer out of paper

NAME

pty – pseudo-terminal driver

CONFIG

pseudo-device *ptyn*

SYNOPSIS

```
#include <fcntl.h>
#include <sys/termios.h>
open("/dev/tty $n$ ", mode);
open("/dev/pty $n$ ", mode);
```

DESCRIPTION

The *pty* driver provides support for a pair of devices collectively known as a *pseudo-terminal*. The two devices comprising a pseudo-terminal are known as a *controller* and a *slave*. The slave device distinguishes between the **B0** baud rate and other baud rates specified in the **c_cflag** word of the **termios** structure, and the **CLOCAL** flag in that word. It does not support any of the other **termio(4)** device control functions specified by flags in the **c_cflag** word of the **termios** structure and by the **IGNBRK**, **IGNPAR**, **PARMRK**, or **INPCK** flags in the **c_iflag** word of the **termios** structure, as these functions apply only to asynchronous serial ports. All other **termio(4)** functions must be performed by **STREAMS** modules pushed atop the driver; when a slave device is opened, the **ldterm(4M)** and **ttcompat(4M)** **STREAMS** modules are automatically pushed on top of the stream, providing the standard **termio(4)** interface.

Instead of having a hardware interface and associated hardware that supports the terminal functions, the functions are implemented by another process manipulating the controller device of the pseudo-terminal.

The controller and the slave devices of the pseudo-terminal are tightly connected. Any data written on the controller device is given to the slave device as input, as though it had been received from a hardware interface. Any data written on the slave terminal can be read from the controller device (rather than being transmitted from a UART).

In configuring, if no optional "count" is given in the specification, 16 pseudo-terminal pairs are configured.

IOCTLS

The standard set of **termio** **ioctl**s are supported by the slave device. None of the bits in the **c_cflag** word have any effect on the pseudo-terminal, except that if the baud rate is set to **B0**, it will appear to the process on the controller device as if the last process on the slave device had closed the line; thus, setting the baud rate to **B0** has the effect of "hanging up" the pseudo-terminal, just as it has the effect of "hanging up" a real terminal.

There is no notion of "parity" on a pseudo-terminal, so none of the flags in the **c_iflag** word that control the processing of parity errors have any effect. Similarly, there is no notion of a "break", so none of the flags that control the processing of breaks, and none of the **ioctl**s that generate breaks, have any effect.

Input flow control is automatically performed; a process that attempts to write to the controller device will be blocked if too much unconsumed data is buffered on the slave device. The input flow control provided by the **IXOFF** flag in the **c_iflag** word is not supported.

The delays specified in the **c_oflag** word are not supported.

As there are no modems involved in a pseudo-terminal, the **ioctl**s that return or alter the state of modem control lines are silently ignored.

On Sun systems, an additional **ioctl** is provided:

TIOCCONS

The argument is ignored. All output that would normally be sent to the console (either from programs writing to **/dev/console** or from kernel printouts) is redirected so that it is written to the pseudo-terminal instead.

A few special `ioctl`s are provided on the controller devices of pseudo-terminals to provide the functionality needed by applications programs to emulate real hardware interfaces:

TIOCSTOP

The argument is ignored. Output to the pseudo-terminal is suspended, as if a STOP character had been typed.

TIOCSTART

The argument is ignored. Output to the pseudo-terminal is restarted, as if a START character had been typed.

TIOCPKT

The argument is a pointer to an `int`. If the value of the `int` is non-zero, *packet* mode is enabled; if the value of the `int` is zero, packet mode is disabled. When a pseudo-terminal is in packet mode, each subsequent `read(2V)` from the controller device will return data written on the slave device 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 using `^S`.

TIOCPKT_START

whenever output to the terminal is restarted.

TIOCPKT_DOSTOP

whenever XON/XOFF flow control is enabled after being disabled; it is considered "enabled" when the `IXON` flag in the `c_iflag` word is set, the `VSTOP` member of the `c_cc` array is `^S` and the `VSTART` member of the `c_cc` array is `^Q`.

TIOCPKT_NOSTOP

whenever XON/XOFF flow control is disabled after being enabled.

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 when interrupts occur; it can be used by other similar programs.

TIOCREMOTE

The argument is a pointer to an `int`. If the value of the `int` is non-zero, *remote* mode is enabled; if the value of the `int` is zero, remote mode is disabled. This mode can be enabled or disabled independently of packet mode. When a pseudo-terminal is in remote mode, input to the slave device of the pseudo-terminal is flow controlled and not input edited (regardless of the mode the slave side of the pseudo-terminal). Each write to the controller device produces a record boundary for the process reading the slave device. 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 EOF character. Note: this means that a process writing to a pseudo-terminal controller in *remote* mode must keep track of line boundaries, and write only one line at a time to the controller. If, for example, it were to buffer up several `NEWLINE` characters and write them to the controller with one `write()`, it would appear to a process reading from the slave as if a single line containing several `NEWLINE` characters had been typed (as if, for example, a user had typed the `LNEXT` character before typing all but the last of those `NEWLINE` characters). Remote mode can be used when doing remote line editing in a window manager, or whenever flow controlled input is required.

The `ioctl`s `TIOCGWINSZ`, `TIOCSWINSZ`, and, on Sun systems, `TIOCCONS`, can be performed on the controller device of a pseudo-terminal; they have the same effect as when performed on the slave device.

FILES

/dev/pty[p-s][0-9a-f] pseudo-terminal controller devices
/dev/tty[p-s][0-9a-f] pseudo-terminal slave devices
/dev/console

SEE ALSO

rlogin(1C), termio(4), ldterm(4M), ttcompat(4M), rlogind(8C)

BUGS

It is apparently not possible to send an EOT by writing zero bytes in TIOCREMOTE mode.

NAME

root – pseudo-driver for Sun386i root disk

CONFIG

pseudo-device rootdev

AVAILABILITY

Sun386i systems only.

DESCRIPTION

The root pseudo-driver provides indirect, device-independent access to the root disk on a diskful Sun workstation. The root disk is the disk where the mounted root partition resides - typically the disk from which the system was booted.

The intent of the root device is to allow uniform access to the partitions on the root disk, regardless of the disk's controller type or unit number. For example, the following version of */etc/fstab* will work for any disk (assuming the disk has the standard partitions and filesystems):

```

/dev/roota / 4.2 rw 1 1
/dev/rootg /usr 4.2 ro 1 2
/dev/rootb /export 4.2 rw 1 3

```

When the root device is opened, the open and all subsequent operations on that device (*read(2V)*, *write(2V)*, *ioctl(2)*, *close(2)*) are redirected to the real disk. Therefore, all device-dependent operations on a particular disk are still accessible via the root device (see *dkio(4S)*).

FILES

```

/dev/root[a-h]    block partitions
/dev/rroot[a-h]  raw partitions

```

SEE ALSO

fstab(5), *sd(4S)*, *open(2V)*, *dkio(4S)*,

NAME

routing – system supporting for local network packet routing

DESCRIPTION

The network facilities provided general packet routing, leaving routing table maintenance to applications processes.

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 (that is, 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 all references to it are removed.

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.

FILES

`/dev/kmem`

SEE ALSO

`ioctl(2)`, `route(8C)`, `routed(8C)`

NAME

sd – Disk driver for SCSI Disk Controllers

CONFIG — SUN-3 SYSTEM

controller sc0 at vme24d16 ? csr 0x200000 priority 2 vector scintr 0x40
 controller si0 at vme24d16 ? csr 0x200000 priority 2 vector siintr 0x40
 controller si0 at obio ? csr 0x140000 priority 2
 disk sd0 at sc0 drive 0 flags 0
 disk sd1 at sc0 drive 1 flags 0
 disk sd0 at si0 drive 0 flags 0
 disk sd1 at si0 drive 1 flags 0
 disk sd2 at sc0 drive 8 flags 0
 disk sd2 at si0 drive 8 flags 0

The first two controller lines above specify the first SCSI host adapter on a Sun-3/160 system. The third controller line above specifies the first and only SCSI host adapter on a Sun-3/50 system. The first four disk lines specify the first and second disk drives on the first SCSI controller in a system. The last two disk lines specify the first disk drive on the second SCSI controller in a system.

The drive value is calculated using the formula:

$$8 * target + unit$$

where *target* is the SCSI target (controller number on host adapter), and *unit* is the SCSI logical unit.

CONFIG — SUN-2 SYSTEM

controller sc0 at mbmem ? csr 0x80000 priority 2
 controller sc1 at mbmem ? csr 0x84000 priority 2
 controller sc0 at vme24 ? csr 0x200000 priority 2 vector scintr 0x40
 disk sd0 at sc0 drive 0 flags 0
 disk sd1 at sc0 drive 1 flags 0
 disk sd2 at sc1 drive 0 flags 0
 disk sd3 at sc1 drive 1 flags 0

The first two controller lines above specify the first and second SCSI host adapters on a Sun-2/120 or Sun-2/170 system. The third controller line above specifies the first host adapter on a Sun-2/160 system. The four disk lines specify the first and second disk drives on the first and second SCSI controllers in a system (where each SCSI controller is on a different host adapter).

CONFIG — Sun386i

controller wds0 at obmem ? csr 0xFB000000 dmachan 7 irq 16 priority 2
 disk sd0 at wds0 drive 0 flags 0
 disk sd1 at wds0 drive 8 flags 0
 disk sd2 at wds0 drive 16 flags 0

The drive value is calculated as described above.

DESCRIPTION

Files with minor device numbers 0 through 7 refer to various portions of drive 0. The standard device names begin with “sd” followed by the drive number and then a letter a-h for partitions 0-7 respectively. The character ? stands here for a drive number in the range 0-7.

The block-files access the disk using the system’s normal buffering mechanism and may be read and written without regard to physical disk records. There is also a “raw” interface that provides for direct transmission between the disk and the user’s read or write buffer. A single read or write call usually results in one I/O operation; therefore raw I/O is considerably more efficient when many words are transmitted. The names of the raw files conventionally begin with an extra ‘r.’

In raw I/O, requests to the SCSI disk must have an offset on a 512 byte boundary, and their length must be a multiple of 512 bytes or the driver will return an error (EINVAL). Likewise seek calls should specify a multiple of 512 bytes.

Disk Support

On Sun-2, Sun-3, Sun-4 systems, this driver handles all ST-506 and ESDI drives (assuming the correct controller is installed), by reading a label from sector 0 of the drive which describes the disk geometry and partitioning.

On Sun386i systems, this driver supports the CDC Wren III half-height, and Wren IV full-height drives, which have embedded, CCS-compatible SCSI controllers.

The `sd?a` partition is normally used for the root file system on a disk, the `sd?b` partition as a paging area, and the `sd?c`

partition for pack-pack copying (it normally maps the entire disk). The rest of the disk is normally the `sd?g` partition.

FILES

<code>/dev/sd[0-7][a-h]</code>	block files
<code>/dev/rsd[0-7][a-h]</code>	raw files

SEE ALSO

`dkio(4S)`

Adaptec ACB 4000 and 5000 Series Disk Controllers OEM Manual (Sun-2, Sun-3, Sun-4 systems only)

Emulex MD21 SCSI Disk Controller Programmer Reference Manual (Sun-2, Sun-3, Sun-4 systems only)

Product Specification for Wren III SCSI Model 94211 (Sun386i systems only)

Product Specification for Wren IV SCSI Model 94171 (Sun386i systems only)

DIAGNOSTICS

`sd%d%c: cmd how (msg) starting blk %d, blk %d (abs blk %d).`

A command such as read or write encountered a error condition (how): either it *failed*, the unit was *restored*, or an operation was *retry*'ed. The *msg* is derived from the error number given by the controller, indicating a condition such as "drive not ready" or "sector not found". The *starting blk* is the first sector of the erroneous command, relative to the beginning of the partition involved. The *blk* is the sector in error, again relative to the beginning of the partition involved. The *abs blk* is the absolute block number of the sector in error.

NAME

sockio – ioctls that operate directly on sockets

SYNOPSIS

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

DESCRIPTION

The IOCTL's listed in this manual page apply directly to sockets, independent of any underlying protocol. Note: the `setsockopt` system call (see `getsockopt(2)`) is the primary method for operating on sockets as such, rather than on the underlying protocol or network interface. `ioctls` for a specific network interface or protocol are documented in the manual page for that interface or protocol.

- SIOCSPGRP** The argument is a pointer to an `int`. Set the process-group ID that will subsequently receive `SIGIO` or `SIGURG` signals for the socket referred to by the descriptor passed to `ioctl` to the value of that `int`.
- SIOCGPGRP** The argument is a pointer to an `int`. Set the value of that `int` to the process-group ID that is receiving `SIGIO` or `SIGURG` signals for the socket referred to by the descriptor passed to `ioctl`.
- SIOCCATMARK** The argument is a pointer to an `int`. Set the value of that `int` to 1 if the read pointer for the socket referred to by the descriptor passed to `ioctl` points to a mark in the data stream for an out-of-band message, and to 0 if it does not point to a mark.

SEE ALSO

`ioctl(2)`, `getsockopt(2)`, `filio(4)`

NAME

st – Driver for Sysgen SC 4000 (Archive) and the Emulex MT-02 Tape Controller

CONFIG — SUN-3 SYSTEM

controller sc0 at vme24d16 ? csr 0x200000 priority 2 vector scintr 0x40
 controller si0 at vme24d16 ? csr 0x200000 priority 2 vector siintr 0x40
 controller si0 at obio ? csr 0x140000 priority 2
 tape st0 at sc0 drive 32 flags 1
 tape st0 at si0 drive 32 flags 1
 tape st1 at sc0 drive 40 flags 1
 tape st1 at si0 drive 40 flags 1

The first two controller lines above specify the first SCSI controller on a Sun-3/160 system. The third controller line above specifies the first and only SCSI controller on a Sun-3/50 system. The four tape lines specify the first and second tape drives on the first SCSI controller in a system.

The drive value is calculated using the formula:

$$8 * target + unit$$

where *target* is the SCSI target, and *unit* is the SCSI logical unit.

CONFIG — SUN-2 SYSTEM

controller sc0 at mbmem ? csr 0x80000 priority 2
 controller sc0 at vme24 ? csr 0x200000 priority 2 vector scintr 0x40
 controller sc1 at mbmem ? csr 0x84000 priority 2
 tape st0 at sc0 drive 32 flags 1
 tape st0 at sc1 drive 32 flags 1
 tape st1 at sc0 drive 40 flags 1
 tape st1 at sc1 drive 40 flags 1

The first two controller lines above specify the first and second SCSI controllers on a Sun-2/120 or Sun-2/170 system. The third controller line specifies the first controller on a Sun-2/160 system. The four tape lines specify the first and second tape drives on the first and second SCSI controllers in a system.

The drive value is calculated as described above.

CONFIG — Sun386i

controller wds0 at obmem ? csr 0xFB000000 dmachan 7 irq 16 priority 2
 tape st0 at wds0 drive 32 flags 1

The drive value is calculated as described above.

DESCRIPTION

The Sysgen tape controller is a SCSI bus interface to an Archive streaming tape drive. It provides a standard tape interface to the device, see `mtio(4)`, with some deficiencies listed under **BUGS** below. To utilize the QIC 24 format, access the logical device that is eight more than the default physical (QIC 11) device (that is, `rst0 = QIC 11`, `rst8 = QIC 24`). QIC 24 is the preferred format on Sun386i systems.

FILES

<code>/dev/rst[0-3]</code>	QIC 11 Format
<code>/dev/rst[8-11]</code>	QIC 24 Format
<code>/dev/nrst[0-3]</code>	non-rewinding QIC 11 Format
<code>/dev/nrst[8-11]</code>	non-rewinding QIC 24 Format

SEE ALSO

`mtio(4)`

Sysgen SC4000 Intelligent Tape Controller Product Specification

DIAGNOSTICS

st*: tape not online.

st*: no cartridge loaded.

st*: cartridge is write protected.
st*: format change failed.
st*: device not supported.

BUGS

The tape cannot reverse direction so the BSF and BSR ioctls are not supported.

The FSR ioctl is not supported.

Most disk I/O over the SCSI bus is prevented when the tape is in use. This is because the controller does not free the bus while the tape is in motion (even during rewind).

When using the raw device, the number of bytes in any given transfer must be a multiple of 512. If it is not, the device driver returns an error.

The driver will only write an end of file mark on close if the last operation was a write, without regard for the mode used when opening the file. Empty files will be deleted on a raw tape copy operation.

Some older systems may not support the QIC 24 device, and may complain (or exhibit erratic behavior) when the user attempts a QIC 24 device access.

NAME

streamio – STREAMS ioctl commands

SYNOPSIS

```
#include <stropts.h>
int ioctl (fildes, command, arg)
int fildes, command;
```

DESCRIPTION

STREAMS (see [intro\(2\)](#)) ioctl commands are a subset of [ioctl\(2\)](#) commands that perform a variety of control functions on STREAMS. The arguments *command* and *arg* are passed to the file designated by *fildes* and are interpreted by the *streamhead*. Certain combinations of these arguments may be passed to a module or driver in the stream.

fildes is an open file descriptor that refers to a stream. *command* determines the control function to be performed as described below. *arg* represents additional information that is needed by this command. The type of *arg* depends upon the command, but it is generally an integer or a pointer to a *command*-specific data structure.

Since these STREAMS commands are a subset of *ioctl*, they are subject to the errors described there. In addition to those errors, the call will fail with *errno* set to EINVAL, without processing a control function, if the stream referenced by *fildes* is linked below a multiplexor, or if *command* is not a valid value for a *stream*.

Also, as described in *ioctl*, STREAMS modules and drivers can detect errors. In this case, the module or driver sends an error message to the *stream head* containing an error value. Subsequent system calls will fail with *errno* set to this value.

IOCTLS

The following *ioctl* commands, with error values indicated, are applicable to all STREAMS files:

- I_PUSH** Pushes the module whose name is pointed to by *arg* onto the top of the current stream, just below the *streamhead*. It then calls the open routine of the newly-pushed module.
- I_PUSH will fail if one of the following occurs:
- | | |
|--------|---|
| EINVAL | The module name is invalid. |
| EFAULT | <i>arg</i> points outside the allocated address space. |
| ENXIO | The open routine of the new module failed. |
| ENXIO | A hangup is received on the stream referred to by <i>fildes</i> . |
- I_POP** Removes the module just below the *stream head* of the stream pointed to by *fildes*. *arg* should be 0 in an I_POP request.
- I_POP will fail if one of the following occurs:
- | | |
|--------|---|
| EINVAL | No module is present on <i>stream</i> . |
| ENXIO | A hangup is received on the stream referred to by <i>fildes</i> . |
- I_LOOK** Retrieves the name of the module just below the *stream head* of the stream pointed to by *fildes*, and places it in a NULL terminated character string pointed at by *arg*. The buffer pointed to by *arg* should be at least FMNAMESZ+1 bytes long. An '#include <sys/conf.h>' declaration is required.
- I_LOOK will fail if one of the following occurs:
- | | |
|--------|---|
| EFAULT | <i>arg</i> points outside the allocated address space of the process. |
| EINVAL | No module is present on <i>stream</i> . |

I_FLUSH

This request flushes all input and/or output queues, depending on the value of *arg*. Legal *arg* values are:

FLUSHR Flush read queues.
FLUSHW Flush write queues.
FLUSHRW Flush read and write queues.

I_FLUSH will fail if one of the following occurs:

EAGAIN No buffers could be allocated for the flush message.
EINVAL The value of *arg* is invalid.
ENXIO A hangup is received on the stream referred to by *fildev*.

I_SETSIG

Informs the *stream head* that the user wishes the kernel to issue the SIGPOLL signal (see `sigvec(2)`) when a particular event has occurred on the stream associated with *fildev*. **I_SETSIG** supports an asynchronous processing capability in STREAMS. The value of *arg* is a bitmask that specifies the events for which the user should be signaled. It is the bitwise-OR of any combination of the following constants:

S_INPUT A non-priority message has arrived on a *stream head* read queue, and no other messages existed on that queue before this message was placed there. This is set even if the message is of zero length.
S_HIPRI A priority message is present on the *stream head* read queue. This is set even if the message is of zero length.
S_OUTPUT The write queue just below the *stream head* is no longer full. This notifies the user that there is room on the queue for sending (or writing) data downstream.
S_MSG A STREAMS signal message that contains the SIGPOLL signal has reached the front of the *stream head* read queue.

A user process may choose to be signaled only of priority messages by setting the *arg* bitmask to the value **S_HIPRI**.

Processes that wish to receive SIGPOLL signals must explicitly register to receive them using **I_SETSIG**. If several processes register to receive this signal for the same event on the same *stream*, each process will be signaled when the event occurs.

If the value of *arg* is zero, the calling process will be unregistered and will not receive further SIGPOLL signals.

I_SETSIG will fail if one of the following occurs:

EINVAL The value of *arg* is invalid or *arg* is zero and the process is not registered to receive the SIGPOLL signal.
EAGAIN A data structure could not be allocated to store the signal request.

I_GETSIG

Returns the events for which the calling process is currently registered to be sent a SIGPOLL signal. The events are returned as a bitmask pointed to by *arg*, where the events are those specified in the description of **I_SETSIG** above.

I_GETSIG will fail if one of the following occurs:

- EINVAL** The process is not registered to receive the SIGPOLL signal.
- EFAULT** *arg* points outside the allocated address space of the process.

I_FIND

This request compares the names of all modules currently present in the stream to the name pointed to by *arg*, and returns 1 if the named module is present in the stream. It returns 0 if the named module is not present.

I_FIND will fail if one of the following occurs:

- EFAULT** *arg* points outside the allocated address space of the process.
- EINVAL** *arg* does not point to a valid module name.

I_PEEK

This request allows a user to retrieve the information in the first message on the *stream head* read queue without taking the message off the queue. *arg* points to a *strpeek* structure which contains the following members:

```

    struct strbuf  ctlbuf;
    struct strbuf  databuf;
    long          flags;

```

The *maxlen* field in the *ctlbuf* and *databuf* *strbuf* structures (see *getmsg(2)*) must be set to the number of bytes of control information and/or data information, respectively, to retrieve. If the user sets *flags* to **RS_HIPRI**, **I_PEEK** will only look for a priority message on the *stream head* read queue.

I_PEEK returns 1 if a message was retrieved, and returns 0 if no message was found on the *stream head* read queue, or if the **RS_HIPRI** flag was set in *flags* and a priority message was not present on the *stream head* read queue. It does not wait for a message to arrive. On return, *ctlbuf* specifies information in the control buffer, *databuf* specifies information in the data buffer, and *flags* contains the value 0 or **RS_HIPRI**.

I_PEEK will fail if one of the following occurs:

- EFAULT** *arg* points, or the buffer area specified in *ctlbuf* or *databuf* is, outside the allocated address space of the process.

I_SRDOPT

Sets the read mode using the value of the argument *arg*. Legal *arg* values are:

- RNORM** Byte-stream mode, the default.
- RMSGD** Message-discard mode.
- RMSGN** Message-nondiscard mode.

Read modes are described in *read(2V)*.

I_SRDOPT will fail if one of the following occurs:

- EINVAL** *arg* is not one of the above legal values.

I_GRDOPT

Returns the current read mode setting in an *int* pointed to by the argument *arg*. Read modes are described in *read(2V)*.

I_GRDOPT will fail if one of the following occurs:

- EFAULT** *arg* points outside the allocated address space of the process.

I_NREAD

Counts the number of data bytes in data blocks in the first message on the *stream head* read queue, and places this value in the location pointed to by *arg*. The return value for the command is the number of messages on the *stream head* read queue. For example, if zero is returned in *arg*, but the *ioctl* return value is greater than zero, this indicates that a zero-length message is next on the queue.

I_NREAD will fail if one of the following occurs:

EFAULT *arg* points outside the allocated address space of the process.

I_FDINSERT

creates a message from user specified buffer(s), adds information about another stream and sends the message downstream. The message contains a control part and an optional data part. The data and control parts to be sent are distinguished by placement in separate buffers, as described below.

arg points to a *strfdinsert* structure which contains the following members:

```

struct strbuf  ctlbuf;
struct strbuf  databuf;
long          flags;
int           fd;
int           offset;

```

The *len* field in the *ctlbuf strbuf* structure (see *putmsg(2)*) must be set to the size of a pointer plus the number of bytes of control information to be sent with the message. *fd* specifies the file descriptor of the other stream and *offset*, which must be word-aligned, specifies the number of bytes beyond the beginning of the control buffer where **I_FDINSERT** will store a pointer to the *fd* stream's driver read queue structure. The *len* field in the *databuf strbuf* structure must be set to the number of bytes of data information to be sent with the message or zero if no data part is to be sent.

flags specifies the type of message to be created. A non-priority message is created if *flags* is set to 0, and a priority message is created if *flags* is set to **RS_HIPRI**. For non-priority messages, **I_FDINSERT** will block if the stream write queue is full due to internal flow control conditions. For priority messages, **I_FDINSERT** does not block on this condition. For non-priority messages, **I_FDINSERT** does not block when the write queue is full and **O_NDELAY** is set. Instead, it fails and sets *errno* to **EAGAIN**.

I_FDINSERT also blocks, unless prevented by lack of internal resources, waiting for the availability of message blocks in the *stream*, regardless of priority or whether **O_NDELAY** has been specified. No partial message is sent.

I_FDINSERT will fail if one of the following occurs:

EAGAIN A non-priority message was specified, the **O_NDELAY** flag is set, and the stream write queue is full due to internal flow control conditions.

EAGAIN Buffers could not be allocated for the message that was to be created.

EFAULT *arg* points, or the buffer area specified in *ctlbuf* or *databuf* is, outside the allocated address space of the process.

EINVAL *fd* in the *strfdinsert* structure is not a valid, open stream file descriptor; the size of a pointer plus *offset* is greater than the *len* field for the buffer specified through *ctlptr*;

offset does not specify a properly-aligned location in the data buffer; an undefined value is pointed to by *flags*.

ENXIO

A hangup is received on the stream referred to by *fdes*.

ERANGE

The *len* field for the buffer specified through *databuf* does not fall within the range specified by the maximum and minimum packet sizes of the topmost stream module, or the *len* field for the buffer specified through *databuf* is larger than the maximum configured size of the data part of a message, or the *len* field for the buffer specified through *ctlbuf* is larger than the maximum configured size of the control part of a message.

I_STR

Constructs an internal STREAMS ioctl message from the data pointed to by *arg*, and sends that message downstream.

This mechanism is provided to permit a process to specify timeouts and variable-sized amounts of data when sending an ioctl request to downstream modules and drivers. It allows information to be sent with the *ioctl*, and will return to the user any information sent upstream by the downstream recipient. **I_STR** blocks until the system responds with either a positive or negative acknowledgement message, or until the request “times out” after some period of time. If the request times out, it fails with *errno* set to ETIME.

At most, one **I_STR** can be active on a stream. Further **I_STR** calls will block until the active **I_STR** completes at the *stream head*. The default timeout interval for these requests is 15 seconds. The **O_NDELAY** (see **open(2V)**) flag has no effect on this call.

To send requests downstream, *arg* must point to a *striocil* structure which contains the following members:

```

int    ic_cmd;        /* downstream command */
int    ic_timeout;    /* ACK/NAK timeout */
int    ic_len;        /* length of data arg */
char   *ic_dp;        /* ptr to data arg */

```

ic_cmd is the internal ioctl command intended for a downstream module or driver and *ic_timeout* is the number of seconds (-1 = infinite, 0 = use default, >0 = as specified) an **I_STR** request will wait for acknowledgement before timing out. *ic_len* is the number of bytes in the data argument and *ic_dp* is a pointer to the data argument. The *ic_len* field has two uses: on input, it contains the length of the data argument passed in, and on return from the command, it contains the number of bytes being returned to the user (the buffer pointed to by *ic_dp* should be large enough to contain the maximum amount of data that any module or the driver in the stream can return).

The *stream head* will convert the information pointed to by the *striocil* structure to an internal ioctl command message and send it downstream.

I_STR will fail if one of the following occurs:

EAGAIN

Buffers could not be allocated for the ioctl message.

EFAULT

arg points, or the buffer area specified by *ic_dp* and *ic_len* (separately for data sent and data returned) is, outside the allocated address space of the process.

EINVAL

ic_len is less than 0 or *ic_len* is larger than the maximum

configured size of the data part of a message or *ic_timeout* is less than -1.

ENXIO A hangup is received on the stream referred to by *fildes*.
ETIME A downstream *ioctl* timed out before acknowledgement was received.

An **I_STR** can also fail while waiting for an acknowledgement if a message indicating an error or a hangup is received at the *streamhead*. In addition, an error code can be returned in the positive or negative acknowledgement message, in the event the *ioctl* command sent downstream fails. For these cases, **I_STR** will fail with *errno* set to the value in the message.

I_SENDFD

Requests the stream associated with *fildes* to send a message, containing a file pointer, to the *stream head* at the other end of a stream pipe. The file pointer corresponds to *arg*, which must be an integer file descriptor.

I_SENDFD converts *arg* into the corresponding system file pointer. It allocates a message block and inserts the file pointer in the block. The user id and group id associated with the sending process are also inserted. This message is placed directly on the read queue (see *intro(2)*) of the *stream head* at the other end of the stream pipe to which it is connected.

I_SENDFD will fail if one of the following occurs:

EAGAIN The sending stream is unable to allocate a message block to contain the file pointer.
EAGAIN The read queue of the receiving *stream head* is full and cannot accept the message sent by **I_SENDFD**.
EBADF *arg* is not a valid, open file descriptor.
EINVAL *fildes* is not connected to a stream pipe.
ENXIO A hangup is received on the stream referred to by *fildes*.

I_RECVFD

Retrieves the file descriptor associated with the message sent by an **I_SENDFD** *ioctl* over a stream pipe. *arg* is a pointer to a data buffer large enough to hold an *strrecvfd* data structure containing the following members:

```
int fd;
unsigned short uid;
unsigned short gid;
char fill[8];
```

fd is an integer file descriptor. *uid* and *gid* are the user ID and group ID, respectively, of the sending stream.

If **O_NDELAY** is not set (see *open(2V)*), **I_RECVFD** will block until a message is present at the *streamhead*. If **O_NDELAY** is set, **I_RECVFD** will fail with *errno* set to **EAGAIN** if no message is present at the *streamhead*.

If the message at the *stream head* is a message sent by an **I_SENDFD**, a new user file descriptor is allocated for the file pointer contained in the message. The new file descriptor is placed in the *fd* field of the *strrecvfd* structure. The structure is copied into the user data buffer pointed to by *arg*.

I_RECVFD will fail if one of the following occurs:

EAGAIN A message was not present at the *stream head* read queue, and the **O_NDELAY** flag is set.

EBADMSG	The message at the <i>stream head</i> read queue was not a message containing a passed file descriptor.
EFAULT	<i>arg</i> points outside the allocated address space of the process.
EMFILE	Too many descriptors are active.
ENXIO	A hangup is received on the stream referred to by <i>fildev</i> .

The following four commands are used for connecting and disconnecting multiplexed STREAMS configurations.

I_LINK

Connects two streams, where *fildev* is the file descriptor of the stream connected to the multiplexing driver, and *arg* is the file descriptor of the stream connected to another driver. The stream designated by *arg* gets connected below the multiplexing driver. I_LINK causes the multiplexing driver to send an acknowledgement message to the *stream head* regarding the linking operation. This call returns a multiplexor ID number (an identifier used to disconnect the multiplexor, see I_UNLINK) on success, and a -1 on failure.

I_LINK will fail if one of the following occurs:

ENXIO	A hangup is received on the stream referred to by <i>fildev</i> .
ETIME	The <i>ioctl</i> timed out before an acknowledgement was received.
EAGAIN	Storage could not be allocated to perform the I_LINK.
EBADF	<i>arg</i> is not a valid, open file descriptor.
EINVAL	The stream referred to by <i>fildev</i> does not support multiplexing.
EINVAL	<i>arg</i> is not a stream, or is already linked under a multiplexor.
EINVAL	The specified link operation would cause a "cycle" in the resulting configuration; that is, if a given <i>stream head</i> is linked into a multiplexing configuration in more than one place.

An I_LINK can also fail while waiting for the multiplexing driver to acknowledge the link request, if a message indicating an error or a hangup is received at the *stream head* of *fildev*. In addition, an error code can be returned in the positive or negative acknowledgement message. For these cases, I_LINK will fail with *errno* set to the value in the message.

I_UNLINK

Disconnects the two streams specified by *fildev* and *arg*. *fildev* is the file descriptor of the stream connected to the multiplexing driver. *arg* is the multiplexor ID number that was returned by the *ioctl* I_LINK command when a stream was linked below the multiplexing driver. If *arg* is -1, then all streams which were linked to *fildev* are disconnected. As in I_LINK, this command requires the multiplexing driver to acknowledge the unlink.

I_UNLINK will fail if one of the following occurs:

ENXIO	A hangup is received on the stream referred to by <i>fildev</i> .
ETIME	The <i>ioctl</i> timed out before an acknowledgement was received.
EAGAIN	Buffers could not be allocated for the acknowledgement message.

EINVAL

The multiplexor ID number was invalid.

An **I_UNLINK** can also fail while waiting for the multiplexing driver to acknowledge the link request, if a message indicating an error or a hangup is received at the *stream head* of *files*. In addition, an error code can be returned in the positive or negative acknowledgement message. For these cases, **I_UNLINK** will fail with *errno* set to the value in the message.

SEE ALSO

close(2), **fcntl(2V)**, **getmsg(2)**, **intro(2)**, **ioctl(2)**, **open(2V)**, **poll(2)**, **putmsg(2)**, **read(2V)**, **sigvec(2)**, **write(2V)**

STREAMS Programmer's Guide

STREAMS Primer

NAME

taac – Sun applications accelerator

CONFIG

taac0 at vme32d32 ? csr 0x28000000

DESCRIPTION

The **taac** interface supports the optional TAAC-1 Applications Accelerator. This add-on device is composed of a very-long-instruction-word computation engine, coupled with an 8MB memory array. This memory area can be used either as a frame buffer, or as storage for large data sets.

Programs can be downloaded for execution on the TAAC-1 directly, they can be executed by the host processor, or the host processor and the TAAC-1 engine can be used in combination. See the *TAAC-1 User's Guide* for detailed information on accessing the TAAC-1 from the host. This manual also describes the C compiler, the programming tools, and the support libraries for the TAAC-1.

Programs on the host processor gain access to the TAAC-1 registers and memory by using **mmap(2)**.

FILES

/dev/taac0
/usr/include/taac1
/usr/lib/taac1

SEE ALSO

mmap(2)
TAAC-1 Application Accelerator: User Guide

NAME

tcp – Internet Transmission Control Protocol

SYNOPSIS

```
#include <sys/socket.h>
#include <netinet/in.h>

s = socket(AF_INET, SOCK_STREAM, 0);
```

DESCRIPTION

TCP is the virtual circuit protocol of the Internet protocol family. It provides reliable, flow-controlled, in order, two-way transmission of data. It is a byte-stream protocol used to support the `SOCK_STREAM` abstraction. TCP is layered above the Internet Protocol (IP), the Internet protocol family's unreliable inter-network datagram delivery protocol.

TCP uses IP's host-level addressing and adds its own per-host collection of "port addresses". The endpoints of a TCP connection are identified by the combination of an IP address and a TCP port number. Although other protocols, such as the User Datagram Protocol (UDP), may use the same host and port address format, the port space of these protocols is distinct. See `inet(4F)` for details on the common aspects of addressing in the Internet protocol family.

Sockets utilizing TCP are either "active" or "passive". Active sockets initiate connections to passive sockets. Both types of sockets must have their local IP address and TCP port number bound with the `bind(2)` system call after the socket is created. By default, TCP sockets are active. A passive socket is created by calling the `listen(2)` system call after binding the socket with `bind`. This establishes a queueing parameter for the passive socket. After this, connections to the passive socket can be received with the `accept(2)` system call. Active sockets use the `connect(2)` call after binding to initiate connections.

By using the special value `INADDR_ANY`, the local IP address can be left unspecified in the `bind` call by either active or passive TCP sockets. This feature is usually used if the local address is either unknown or irrelevant. If left unspecified, the local IP address will be bound at connection time to the address of the network interface used to service the connection.

Once a connection has been established, data can be exchanged using the `read(2V)` and `write(2V)` system calls.

TCP supports one socket option which is set with `setsockopt` and tested with `getsockopt(2)`. Under most circumstances, TCP sends data when it is presented. When outstanding data has not yet been acknowledged, it gathers small amounts of output to be sent in a single packet once an acknowledgement is received. For a small number of clients, such as window systems that send a stream of mouse events which receive no replies, this packetization may cause significant delays. Therefore, TCP provides a boolean option, `TCP_NODELAY` (defined in `<netinet/tcp.h>`), to defeat this algorithm. The option level for the `setsockopt` call is the protocol number for TCP, available from `getprotobyname` (see `getprotoent(3N)`).

Options at the IP level may be used with TCP; see `ip(4P)`.

TCP provides an urgent data mechanism, which may be invoked using the out-of-band provisions of `send(2)`. The caller may mark one byte as "urgent" with the `MSG_OOB` flag to `send(2)`. This causes an "urgent pointer" pointing to this byte to be set in the TCP stream. The receiver on the other side of the stream is notified of the urgent data by a `SIGURG` signal. The `SIOCATMARK` ioctl returns a value indicating whether the stream is at the urgent mark. Because the system never returns data across the urgent mark in a single `read(2V)` call, it is possible to advance to the urgent data in a simple loop which reads data, testing the socket with the `SIOCATMARK` ioctl, until it reaches the mark.

Incoming connection requests that include an IP source route option are noted, and the reverse source route is used in responding.

TCP assumes the datagram service it is layered above is unreliable. A checksum over all data helps TCP implement reliability. Using a window-based flow control mechanism that makes use of positive acknowledgements, sequence numbers, and a retransmission strategy, TCP can usually recover when datagrams are damaged, delayed, duplicated or delivered out of order by the underlying communication medium.

If the local TCP receives no acknowledgements from its peer for a period of time, as would be the case if the remote machine crashed, the connection is closed and an error is returned to the user. If the remote machine reboots or otherwise loses state information about a TCP connection, the connection is aborted and an error is returned to the user.

ERRORS

A socket operation may fail if:

EISCONN	A connect operation was attempted on a socket on which a connect operation had already been performed.
ETIMEDOUT	A connection was dropped due to excessive retransmissions.
ECONNRESET	The remote peer forced the connection to be closed (usually because the remote machine has lost state information about the connection due to a crash).
ECONNREFUSED	The remote peer actively refused connection establishment (usually because no process is listening to the port).
EADDRINUSE	A bind operation was attempted on a socket with a network address/port pair that has already been bound to another socket.
EADDRNOTAVAIL	A bind operation was attempted on a socket with a network address for which no network interface exists.
EACCES	A bind operation was attempted with a "reserved" port number and the effective user ID of the process was not super-user.
ENOBUFS	The system ran out of memory for internal data structures.

SEE ALSO

accept(2), **bind(2)**, **connect(2)**, **getsockopt(2)**, **listen(2)**, **read(2V)**, **send(2)**, **write(2V)**, **getprotoent(3N)**, **inet(4F)**, **ip(4P)**

Postel, Jon, *Transmission Control Protocol - DARPA Internet Program Protocol Specification*, RFC 793, Network Information Center, SRI International, Menlo Park, Calif., September 1981.

BUGS

SIOCSHIWAT and **SIOCGLIWAT** `ioctl`'s to set and get the high water mark for the socket queue, and so that it can be changed from 2048 bytes to be larger or smaller, have been defined (in `<sys/ioctl.h>`) but not implemented.

NAME

termio – general terminal interface

SYNOPSIS

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

DESCRIPTION

Asynchronous communications ports, pseudo-terminals, and the special interface accessed by `/dev/tty` all use the same general interface, no matter what hardware (if any) is involved. The remainder of this section discusses the common features of this interface.

Opening a Terminal Device File

When a terminal file is opened, the process normally waits until a connection is established. (In practice, users' programs seldom open these files; they are opened by `getty(8)` and become a user's standard input, output, and error files.) If the `O_NDELAY` flag was set in the second argument to `open(2V)`, the `open()` will complete immediately without waiting for a connection to be established.

Process Groups

A terminal may have a distinguished process group associated with it. This distinguished process group plays a special role in handling signal-generating input characters, as discussed below in the **Special Characters** section below.

A command interpreter, such as `cs(1)`, that supports "job control" can allocate the terminal to different *jobs*, or process groups, by placing related processes in a single process group and associating this process group with the terminal. A terminal's associated process group may be set or examined by a process with sufficient privileges. The terminal interface aids in this allocation by restricting access to the terminal by processes that are not in the current process group; see **Job Access Control** below.

The Controlling Terminal

A terminal may belong to a process as its *controlling terminal*. If a process that is a "session process group leader", and that does not have a controlling terminal, opens a terminal file not already associated with a process group, the terminal associated with that terminal file becomes the controlling terminal for that process, and the terminal's distinguished process group is set to the process group of that process. (Currently, this also happens if a process that does not have a controlling terminal and is not a member of a process group opens a terminal. In this case, if the terminal is not associated with a process group, a new process group is created with a process group ID equal to the process ID of the process in question, and the terminal is assigned to that process group. The process is made a member of the terminal's process group.)

The controlling terminal is inherited by a child process during a `fork(2)`. A process relinquishes its control terminal when it changes its process group using `setpgrp(2V)` or when it issues a `TIOCNOTTY ioctl(2)` call on a file descriptor created by opening the file `/dev/tty`.

When a session process group leader that has a controlling terminal terminates, the distinguished process group of the controlling terminal is set to zero (indicating no distinguished process group). This allows the terminal to be acquired as a controlling terminal by a new session process group leader.

Closing a Terminal Device File

When a terminal device file is closed, the process closing the file waits until all output is drained; all pending input is then flushed, and finally a disconnect is performed. If `HUPCL` is set, the existing connection is severed (by hanging up the phone line, if appropriate).

Job Access Control

If a process is in the (non-zero) distinguished process group of its controlling terminal, or if the terminal's distinguished process group is zero (if either of these are true, the process is said to be a *foreground process*), then `read(2V)` operations are allowed as described below in **Input Processing and Reading Characters**. If a process is not in the (non-zero) distinguished process group of its controlling terminal (if this is true, the process is said to be a *background process*), then any attempts to read from that terminal will send that process' process group a `SIGTTIN` signal, unless the process is ignoring `SIGTTIN`, has `SIGTTIN` blocked, or is in the middle of process creation using `vfork(2)`; in that case, the read will return `-1` and set `errno` to `EIO`, and the `SIGTTIN` signal will not be sent. The `SIGTTIN` signal will normally stop the

members of that process group.

When the TOSTOP bit is set in the `c_iflag` field, attempts by a background process to write to its controlling terminal will send that process' process group a SIGTTOU signal, unless the process is ignoring SIGTTOU, has SIGTTOU blocked, or is in the middle of process creation using `vfork()`; in that case, the process will be allowed to write to the terminal and the SIGTTOU signal will not be sent. The SIGTTOU signal will normally stop the members of that process group. Certain `ioctl()` calls that set terminal parameters are treated in this same fashion, except that TOSTOP is not checked; the effect is identical to that of terminal writes when TOSTOP is set. See IOCTLS.

Input Processing and Reading Characters

A terminal associated with one of these files 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 full, 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. If the IMAXBEL mode has not been selected, all the saved characters are thrown away without notice when the input limit is reached; if the IMAXBEL mode has been selected, the driver refuses to accept any further input, and echoes a bell (ASCII BEL).

Two general kinds of input processing are available, determined by whether the terminal device file is in canonical mode or non-canonical mode (see ICANON in the Local Modes section).

The style of input processing can also be very different when the terminal is put in non-blocking I/O mode; see `read(2V)`. In this case, reads from the terminal will never block.

It is possible to simulate terminal input using the `TIOCSTI ioctl()` call, 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 process' controlling terminal unless the process' effective user ID is super-user.

Canonical Mode Input Processing

In canonical mode input processing, terminal input is processed in units of lines. A line is delimited by a NEWLINE (ASCII LF) character, an EOF (by default, an ASCII EOT) character, or one of two user-specified end-of-line characters, EOL and EOL2. This means that a `read()` will not complete until an entire line has been typed or a signal has been received. Also, 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.

Erase and kill processing occurs during input. The ERASE character (by default, the character DEL) erases the last character typed in the current input line. The WERASE character (by default, the character CTRL-W) erases the last "word" typed in the current input line (but not any preceding SPACE or TAB characters). A "word" is defined as a sequence of non-blank characters, with TAB characters counted as blanks. Neither ERASE nor WERASE will erase beyond the beginning of the line. The KILL character (by default, the character CTRL-U) kills (deletes) the entire current input line, and optionally outputs a NEWLINE character. All these characters operate on a key-stroke basis, independently of any backspacing or tabbing that may have been done.

The REPRINT character (the character CTRL-R) prints a NEWLINE followed by all characters that have not been read. Reprinting also occurs automatically if characters that would normally be erased from the screen are fouled by program output. The characters are reprinted as if they were being echoed; as a consequence, if ECHO is not set, they are not printed.

The ERASE and KILL characters may be entered literally by preceding them with the escape character (`\`). In this case the escape character is not read. The ERASE and KILL characters may be changed.

Non-Canonical Mode Input Processing

In non-canonical mode input processing, input characters are not assembled into lines, and erase and kill processing does not occur. The MIN and TIME values are used to determine how to process the characters received.

MIN represents the minimum number of characters that should be received when the read is satisfied (when the characters are returned to the user). TIME is a timer of 0.10 second granularity that is used to timeout bursty and short term data transmissions. The four possible values for MIN and TIME and their interactions are described below.

Case A: MIN > 0, TIME > 0

In this case TIME serves as an intercharacter timer and is activated after the first character is received. Since it is an intercharacter timer, it is reset after a character is received. The interaction between MIN and TIME is as follows: as soon as one character is received, the intercharacter timer is started. If MIN characters are received before the intercharacter timer expires (remember that the timer is reset upon receipt of each character), the read is satisfied. If the timer expires before MIN characters are received, the characters received to that point are returned to the user. Note: if MIN expires at least one character will be returned because the timer would not have been enabled unless a character was received. In this case (MIN > 0, TIME > 0) the read will sleep until the MIN and TIME mechanisms are activated by the receipt of the first character.

Case B: MIN > 0, TIME = 0

In this case, since the value of TIME is zero, the timer plays no role and only MIN is significant. A pending read is not satisfied until MIN characters are received (the pending read will sleep until MIN characters are received). A program that uses this case to read record-based terminal I/O may block indefinitely in the read operation.

Case C: MIN = 0, TIME > 0

In this case, since MIN = 0, TIME no longer represents an intercharacter timer. It now serves as a read timer that is activated as soon as a read() is done. A read is satisfied as soon as a single character is received or the read timer expires. Note: in this case if the timer expires, no character will be returned. If the timer does not expire, the only way the read can be satisfied is if a character is received. In this case the read will not block indefinitely waiting for a character – if no character is received within TIME*.10 seconds after the read is initiated, the read will return with zero characters.

Case D: MIN = 0, TIME = 0

In this case return is immediate. The minimum of either the number of characters requested or the number of characters currently available will be returned without waiting for more characters to be input.

Comparison of the Different Cases of MIN, TIME Interaction

Some points to note about MIN and TIME:

1. In the following explanations one may notice that the interactions of MIN and TIME are not symmetric. For example, when MIN > 0 and TIME = 0, TIME has no effect. However, in the opposite case where MIN = 0 and TIME > 0, both MIN and TIME play a role in that MIN is satisfied with the receipt of a single character.
2. Also note that in case A (MIN > 0, TIME > 0), TIME represents an intercharacter timer while in case C (TIME = 0, TIME > 0) TIME represents a read timer.

These two points highlight the dual purpose of the MIN/TIME feature. Cases A and B, where MIN > 0, exist to handle burst mode activity (for example, file transfer programs) where a program would like to process at least MIN characters at a time. In case A, the intercharacter timer is activated by a user as a safety measure; while in case B, it is turned off.

Cases C and D exist to handle single character timed transfers. These cases are readily adaptable to screen-based applications that need to know if a character is present in the input queue before refreshing the screen. In case C the read is timed; while in case D, it is not.

Another important note is that MIN is always just a minimum. It does not denote a record length. That is, if a program does a read of 20 bytes, MIN is 10, and 25 characters are present, 20 characters will be returned to the user.

Writing Characters

When one or more characters are written, they are transmitted to the terminal as soon as previously-written characters have finished typing. Input characters are echoed as they are typed if echoing has been enabled. If 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.

Special Characters

Certain characters have special functions on input and/or output. These functions and their default character values are summarized as follows:

INTR	(CTRL-C or ASCII ETX) generates a SIGINT signal, which is sent to all processes in the distinguished process group associated with the terminal. Normally, each such process is forced to terminate, but arrangements may be made either to ignore the signal or to receive a trap to an agreed-upon location; see sigvec(2).
QUIT	(CTRL- or ASCII FS) generates a SIGQUIT signal, which is sent to all processes in the distinguished process group associated with the terminal. Its treatment is identical to the interrupt signal except that, unless a receiving process has made other arrangements, it will not only be terminated but a core image file (called core) will be created in the current working directory.
ERASE	(Rubout or ASCII DEL) erases the preceding character. It will not erase beyond the start of a line, as delimited by a NL, EOF, EOL, or EOL2 character.
WERASE	(CTRL-W or ASCII ETB) erases the preceding "word". It will not erase beyond the start of a line, as delimited by a NL, EOF, EOL, or EOL2 character.
KILL	(CTRL-U or ASCII NAK) deletes the entire line, as delimited by a NL, EOF, EOL, or EOL2 character.
REPRINT	(CTRL-R or ASCII DC2) reprints all characters that have not been read, preceded by a NEWLINE.
EOF	(CTRL-D or ASCII EOT) may be used to generate an end-of-file from a terminal. When received, all the characters waiting to be read are immediately passed to the program, without waiting for a NEWLINE, and the EOF is discarded. Thus, if there are no characters waiting, which is to say the EOF occurred at the beginning of a line, zero characters will be passed back, which is the standard end-of-file indication.
NL	(ASCII LF) is the normal line delimiter. It can not be changed; it can, however, be escaped by the LNEXT character.
EOL	
EOL2	(ASCII NUL) are additional line delimiters, like NL. They are not normally used.
SUSP	(CTRL-Z or ASCII EM) is used by the job control facility to change the current job to return to the controlling job. It generates a SIGTSTP signal, which stops all processes in the terminal's process group.
STOP	(CTRL-S or ASCII DC3) can be used to temporarily suspend output. It is useful with CRT terminals to prevent output from disappearing before it can be read. While output is suspended, STOP characters are ignored and not read.
START	(CTRL-Q or ASCII DC1) is used to resume output that has been suspended by a STOP character. While output is not suspended, START characters are ignored and not read.
DISCARD	(CTRL-O or ASCII SI) causes subsequent output to be discarded until another DISCARD character is typed, more input arrives, or the condition is cleared by a program.
LNEXT	(CTRL-V or ASCII SYN) causes the special meaning of the next character to be ignored; this works for all the special characters mentioned above. This allows characters to be input that would otherwise get interpreted by the system (for example, KILL, QUIT.)

The character values for INTR, QUIT, ERASE, WERASE, KILL, REPRINT, EOF, EOL, EOL2, SUSP, STOP, START, DISCARD, and LNEXT may be changed to suit individual tastes. If the value of a special control character is 0, the function of that special control character will be disabled. The ERASE, KILL, and EOF characters may be escaped by a preceding \ character, in which case no special function is done. Any of the special characters may be preceded by the LNEXT character, in which case no special function is done.

Modem Disconnect

If a modem disconnect is detected, and the CLOCAL flag is not set in the `c_cflag` field, a SIGHUP signal is sent to all processes in the distinguished process group associated with this terminal. Unless other arrangements have been made, this signal terminates the processes. If SIGHUP is ignored or caught, any subsequent `read()` returns with an end-of-file indication until the terminal is closed. Thus, programs that read a terminal and test for end-of-file can terminate appropriately after a disconnect. Any subsequent `write()` will return `-1` and set `errno` to `EIO` until the terminal is closed.

Terminal Parameters

The parameters that control the behavior of devices and modules providing the `termios` interface are specified by the `termios` structure, defined by `<sys/termios.h>`. Several `ioctl()` system calls that fetch or change these parameters use this structure:

```
#define NCCS      17
struct termios {
    unsigned long  c_iflag;    /* input modes */
    unsigned long  c_oflag;    /* output modes */
    unsigned long  c_cflag;    /* control modes */
    unsigned long  c_lflag;    /* local modes */
    unsigned char  c_line;     /* line discipline */
    unsigned char  c_cc[NCCS]; /* control chars */
};
```

The special control characters are defined by the array `c_cc`. The relative positions and initial values for each function are as follows:

```
0  VINTR    ETX
1  VQUIT   FS
2  VERASE  DEL
3  VKILL   NAK
4  VEOF    EOT
5  VEOL    NUL
6  VEOL2   NUL
7  VSWTCH  NUL
8  VSTART  DC1
9  VSTOP   DC3
10 VSUSP   EM
12 VREPRINT DC2
13 VDISCARD SI
14 VWERASE ETB
15 VLNEXT  SYN
```

The `MIN` value is stored in the `VMIN` element of the `c_cc` array, and the `TIME` value is stored in the `VTIME` element of the `c_cc` array. The `VMIN` element is the same element as the `VEOF` element, and the `VTIME` element is the same element as the `VEOL` element.

Input Modes

The `c_iflag` field describes the basic terminal input control:

```
IGNBRK  0000001  Ignore break condition.
BRKINT  0000002  Signal interrupt on break.
IGNPAR  0000004  Ignore characters with parity errors.
```

PARMRK	0000010	Mark parity errors.
INPCK	0000020	Enable input parity check.
ISTRIP	0000040	Strip character.
INLCR	0000100	Map NL to CR on input.
IGNCR	0000200	Ignore CR.
ICRNL	0000400	Map CR to NL on input.
IUCLC	0001000	Map upper-case to lower-case on input.
IXON	0002000	Enable start/stop output control.
IXANY	0004000	Enable any character to restart output.
IXOFF	0010000	Enable start/stop input control.
IMAXBEL	0020000	Echo BEL on input line too long.

If **IGNBRK** is set, a break condition (a character framing error with data all zeros) detected on input is ignored, that is, not put on the input queue and therefore not read by any process. Otherwise, if **BRKINT** is set, a break condition will generate a **SIGINT** and flush both the input and output queues. If neither **IGNBRK** nor **BRKINT** is set, a break condition is read as a single ASCII NUL character (`\0`).

If **IGNPAR** is set, characters with framing or parity errors (other than break) are ignored. Otherwise, if **PARMRK** is set, a character with a framing or parity error that is not ignored is read as the three-character sequence: `\377`, `\0`, `X`, where `X` is the data of the character received in error. To avoid ambiguity in this case, if **ISTRIP** is not set, a valid character of `\377` is read as `\377`, `\377`. If neither **IGNPAR** nor **PARMRK** is set, a framing or parity error (other than break) is read as a single ASCII NUL character (`\0`).

If **INPCK** is set, input parity checking is enabled. If **INPCK** is not set, input parity checking is disabled. This allows output parity generation without input parity errors.

If **ISTRIP** is set, valid input characters are first stripped to 7 bits, otherwise all 8 bits are processed.

If **INLCR** is set, a received NL character is translated into a CR character. If **IGNCR** is set, a received CR character is ignored (not read). Otherwise if **ICRNL** is set, a received CR character is translated into a NL character.

If **IUCLC** is set, a received upper-case alphabetic character is translated into the corresponding lower-case character.

If **IXON** is set, start/stop output control is enabled. A received **STOP** character will suspend output and a received **START** character will restart output. The **STOP** and **START** characters will not be read, but will merely perform flow control functions. If **IXANY** is set, any input character will restart output that has been suspended.

If **IXOFF** is set, the system will transmit a **STOP** character when the input queue is nearly full, and a **START** character when enough input has been read that the input queue is nearly empty again.

If **IMAXBEL** is set, the ASCII BEL character is echoed if the input stream overflows. Further input will not be stored, but any input already present in the input stream will not be disturbed. If **IMAXBEL** is not set, no BEL character is echoed, and all input present in the input queue is discarded if the input stream overflows.

The initial input control value is **BRKINT**, **ICRNL**, **IXON**, **ISTRIP**.

Output modes

The `c_oflag` field specifies the system treatment of output:

OPOST	0000001	Postprocess output.
OLCUC	0000002	Map lower case to upper on output.
ONLCR	0000004	Map NL to CR-NL on output.
OCRNL	0000010	Map CR to NL on output.
ONOCR	0000020	No CR output at column 0.
ONLRET	0000040	NL performs CR function.
OFILL	0000100	Use fill characters for delay.
OFDEL	0000200	Fill is DEL, else NUL.

```

NLDLY  0000400 Select new-line delays:
  NL0   0
  NL1   0000400
CRDLY  0003000 Select carriage-return delays:
  CR0   0
  CR1   0001000
  CR2   0002000
  CR3   0003000
TABDLY 0014000 Select horizontal-tab delays:
  TAB0  0           or tab expansion:
  TAB1  0004000
  TAB2  0010000
XTABS  0014000 Expand tabs to spaces.
BSDLY  0020000 Select backspace delays:
  BS0   0
  BS1   0020000
VTDLY  0040000 Select vertical-tab delays:
  VT0   0
  VT1   0040000
FFDLY  0100000 Select form-feed delays:
  FF0   0
  FF1   0100000

```

If **OPOST** is set, output characters are post-processed as indicated by the remaining flags, otherwise characters are transmitted without change.

If **OLCUC** is set, a lower-case alphabetic character is transmitted as the corresponding upper-case character. This function is often used in conjunction with **IUCLC**.

If **ONLCR** is set, the **NL** character is transmitted as the **CR-NL** character pair. If **OCRNL** is set, the **CR** character is transmitted as the **NL** character. If **ONOCR** is set, no **CR** character is transmitted when at column 0 (first position). If **ONLRET** is set, the **NL** character is assumed to do the carriage-return function; the column pointer will be set to 0 and the delays specified for **CR** will be used. Otherwise the **NL** character is assumed to do just the line-feed function; the column pointer will remain unchanged. The column pointer is also set to 0 if the **CR** character is actually transmitted.

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. If **OFILL** is set, fill characters will be transmitted for delay instead of a timed delay. This is useful for high baud rate terminals that need only a minimal delay. If **OFDEL** is set, the fill character is **DEL**, otherwise **NUL**.

If a form-feed or vertical-tab delay is specified, it lasts for about 2 seconds.

New-line delay lasts about 0.10 seconds. If **ONLRET** is set, the **RETURN** delays are used instead of the **NEWLINE** delays. If **OFILL** is set, two fill characters will be transmitted.

Carriage-return delay type 1 is dependent on the current column position, type 2 is about 0.10 seconds, and type 3 is about 0.15 seconds. If **OFILL** is set, delay type 1 transmits two fill characters, and type 2, four fill characters.

Horizontal-tab delay type 1 is dependent on the current column position. Type 2 is about 0.10 seconds. Type 3, specified by **TAB3** or **XTABS**, specifies that **TAB** characters are to be expanded into **SPACE** characters. If **OFILL** is set, two fill characters will be transmitted for any delay.

Backspace delay lasts about 0.05 seconds. If **OFILL** is set, one fill character will be transmitted.

The actual delays depend on line speed and system load.

The initial output control value is **OPOST**, **ONLCR**, **XTABS**.

The **c_cflag** field describes the hardware control of the terminal:

CBAUD	0000017	Baud rate:
B0	0	Hang up
B50	0000001	50 baud
B75	0000002	75 baud
B110	0000003	110 baud
B134	0000004	134.5 baud
B150	0000005	150 baud
B200	0000006	200 baud
B300	0000007	300 baud
B600	0000010	600 baud
B1200	0000011	1200 baud
B1800	0000012	1800 baud
B2400	0000013	2400 baud
B4800	0000014	4800 baud
B9600	0000015	9600 baud
B19200	0000016	19200 baud
B38400	0000017	38400 baud
CSIZE	0000060	Character size:
CS5	0	5 bits
CS6	0000020	6 bits
CS7	0000040	7 bits
CS8	0000060	8 bits
CSTOPB	0000100	Send two stop bits, else one.
CREAD	0000200	Enable receiver.
PARENB	0000400	Parity enable.
PARODD	0001000	Odd parity, else even.
HUPCL	0002000	Hang up on last close.
CLOCAL	0004000	Local line, else dial-up.
CRTSCTS	0010000	Enable RTS/CTS flow control.
CIBAUD	03600000	Input baud rate, if different from output rate.

The **CBAUD** bits specify the baud rate. The zero baud rate, **B0**, is used to hang up the connection. If **B0** is specified, the modem control lines will cease to be asserted. Normally, this will disconnect the line. If the **CIBAUD** bits are not zero, they specify the input baud rate, with the **CBAUD** bits specifying the output baud rate; otherwise, the output and input baud rates are both specified by the **CBAUD** bits. The values for the **CIBAUD** bits are the same as the values for the **CBAUD** bits, shifted left **IBSHIFT** bits. For any particular hardware, impossible speed changes are ignored.

The **CSIZE** bits specify the character size in bits for both transmission and reception. This size does not include the parity bit, if any. If **CSTOPB** is set, two stop bits are used, otherwise one stop bit. For example, at 110 baud, two stop bits are required.

If **PARENB** is set, parity generation and detection is enabled and a parity bit is added to each character. If parity is enabled, the **PARODD** flag specifies odd parity if set, otherwise even parity is used.

If **CREAD** is set, the receiver is enabled. Otherwise no characters will be received.

If **HUPCL** is set, the modem control lines for the port will be disconnected when the last process with the line open closes it or terminates.

If **CLOCAL** is set, a connection does not depend on the state of the modem status lines. Otherwise modem control is assumed.

If **CRTSCTS** is set, and the terminal has modem control lines associated with it, the Request To Send (RTS) modem control line will be raised, and output will occur only if the Clear To Send (CTS) modem status line is raised. If the CTS modem status line is lowered, output is suspended until CTS is raised. Some hardware may not support this function, and other hardware may not permit it to be disabled; in either of these cases, the state of the **CRTSCTS** flag is ignored.

The initial hardware control value after open is **B9600, CS7, CREAD, PARENB**.

Local Modes

The **c_lflag** field of the argument structure is used by the line discipline to control terminal functions. The basic line discipline provides the following:

ISIG	0000001	Enable signals.
ICANON	0000002	Canonical input (erase and kill processing).
XCASE	0000004	Canonical upper/lower presentation.
ECHO	0000010	Enable echo.
ECHOE	0000020	Echo erase character as BS-SP-BS.
ECHOK	0000040	Echo NL after kill character.
ECHONL	0000100	Echo NL.
NOFLSH	0000200	Disable flush after interrupt or quit.
TOSTOP	0000400	Send SIGTTOU for background output.
ECHOCTL	0001000	Echo control characters as <i>^char</i> , delete as <i>^?</i> .
ECHOPRT	0002000	Echo erase character as character erased.
ECHOKE	0004000	BS-SP-BS erase entire line on line kill.
FLUSHO	0040000	Output is being flushed.
PENDIN	0100000	Retype pending input at next read or input character.

If **ISIG** is set, each input character is checked against the special control characters **INTR**, **QUIT**, and **SUSP**. If an input character matches one of these control characters, the function associated with that character is performed. If **ISIG** is not set, no checking is done. Thus these special input functions are possible only if **ISIG** is set.

If **ICANON** is set, canonical processing is enabled. This enables the erase, word erase, kill, and reprint edit functions, and the assembly of input characters into lines delimited by **NL**, **EOF**, **EOL**, and **EOL2**. If **ICANON** is not set, read requests are satisfied directly from the input queue. A read will not be satisfied until at least **MIN** characters have been received or the timeout value **TIME** has expired between characters. This allows fast bursts of input to be read efficiently while still allowing single character input. The time value represents tenths of seconds. See the *Non-canonical Mode Input Processing* section for more details.

If **XCASE** is set, and if **ICANON** is set, an upper-case letter is accepted on input by preceding it with a **** character, and is output preceded by a **** character. In this mode, the following escape sequences are generated on output and accepted on input:

<i>for:</i>	<i>use:</i>
\	\
 	
-	-
{	{
}	}
\	\

For example, **A** is input as **\a**, **\n** as **\\n**, and **\N** as **\\N**.

If **ECHO** is set, characters are echoed as received. If **ECHO** is not set, input characters are not echoed.

If **ECHOCTL** is not set, all control characters (characters with codes between 0 and 37 octal) are echoed as themselves. If **ECHOCTL** is set, all control characters other than ASCII **TAB**, ASCII **NL**, the **START** character, and the **STOP** character, are echoed as **^X**, where **X** is the character given by adding 100 octal to the control character's code (so that the character with octal code 1 is echoed as **^A**), and the ASCII **DEL**

character, with code 177 octal, is echoed as ‘?’.

When ICANON is set, the following echo functions are possible:

1. If ECHO and ECHOE are set, and ECHOPRT is not set, the ERASE and WERASE characters are echoed as one or more ASCII BS SP BS, which will clear the last character(s) from a CRT screen.
2. If ECHO and ECHOPRT are set, the first ERASE and WERASE character in a sequence echoes as a backslash (\) followed by the characters being erased. Subsequent ERASE and WERASE characters echo the characters being erased, in reverse order. The next non-erase character types a slash (/) before it is echoed.
3. If ECHOKE is set, the kill character is echoed by erasing each character on the line from the screen (using the mechanism selected by ECHOE and ECHOPRT).
4. If ECHOK is set, and ECHOKE is not set, the NL character will be echoed after the kill character to emphasize that the line will be deleted. Note: an escape character (\) or an LNEXT character preceding the erase or kill character removes any special function.
5. If ECHONL is set, the NL character will be echoed even if ECHO is not set. This is useful for terminals set to local echo (so-called half duplex).
6. If ECHOCTL is not set, the EOF character is not echoed, unless it is escaped. Because EOT is the default EOF character, this prevents terminals that respond to EOT from hanging up. If ECHOCTL is set, the EOF character is echoed; if it is not escaped, after it is echoed, one backspace character is output if it is echoed as itself, and two backspace characters are echoed if it is echoed as ^X.

If NOFLSH is set, the normal flush of the input and output queues associated with the INTR, QUIT, and SUSP characters will not be done.

If TOSTOP is set, the signal SIGTTOU is sent to a process that tries to write to its controlling terminal if it is not in the distinguished process group for that terminal. This signal normally stops the process. Otherwise, the output generated by that process is output to the current output stream. Processes that are blocking or ignoring SIGTTOU signals are excepted and allowed to produce output.

If FLUSHO is set, data written to the terminal will be discarded. This bit is set when the FLUSH character is typed. A program can cancel the effect of typing the FLUSH character by clearing FLUSHO.

If PENDIN is set, any input that has not yet been read will be reprinted when the next character arrives as input.

The initial line-discipline control value is ISIG, ICANON, ECHO.

Minimum and Timeout

The MIN and TIME values are described above under **Non-canonical Mode Input Processing**. The initial value of MIN is 1, and the initial value of TIME is 0.

Termio Structure

The System V termio structure is used by other ioctl() calls; it is defined by <sys/termio.h> as:

```
#define NCC      8
struct termio {
    unsigned short c_iflag; /* input modes */
    unsigned short c_oflag; /* output modes */
    unsigned short c_cflag; /* control modes */
    unsigned short c_lflag; /* local modes */
    char c_line; /* line discipline */
    unsigned char c_cc[NCC]; /* control chars */
};
```

The special control characters are defined by the array `c_cc`. The relative positions for each function are as follows:

```

0  VINTR
1  VQUIT
2  VERASE
3  VKILL
4  VEOF
5  VEOL
6  VEOL2
7  reserved

```

The calls that use the `termio` structure only affect the flags and control characters that can be stored in the `termio` structure; all other flags and control characters are unaffected.

Terminal Size

The number of lines and columns on the terminal's display (or page, in the case of printing terminals) is specified in the `winsize` structure, defined by `<sys/termios.h>`. Several `ioctl()` system calls that fetch or change these parameters use this structure:

```

struct winsize {
    unsigned short  ws_row;    /* rows, in characters */
    unsigned short  ws_col;    /* columns, in characters */
    unsigned short  ws_xpixel; /* horizontal size, pixels - not used */
    unsigned short  ws_ypixel; /* vertical size, pixels - not used */
};

```

Modem Lines

On special files representing serial ports, the modem control lines supported by the hardware can be read and the modem status lines supported by the hardware can be changed. The following modem control and status lines may be supported by a device; they are defined by `<sys/termios.h>`:

<code>TIOCM_LE</code>	0001	line enable
<code>TIOCM_DTR</code>	0002	data terminal ready
<code>TIOCM_RTS</code>	0004	request to send
<code>TIOCM_ST</code>	0010	secondary transmit
<code>TIOCM_SR</code>	0020	secondary receive
<code>TIOCM_CTS</code>	0040	clear to send
<code>TIOCM_CAR</code>	0100	carrier detect
<code>TIOCM_RNG</code>	0200	ring
<code>TIOCM_DSR</code>	0400	data set ready

`TIOCM_CD` is a synonym for `TIOCM_CAR`, and `TIOCM_RI` is a synonym for `TIOCM_RNG`.

Not all of these will necessarily be supported by any particular device; check the manual page for the device in question.

IOCTLS

The `ioctl()` calls supported by devices and STREAMS modules providing the `termios` interface are listed below. Some calls may not be supported by all devices or modules.

Unless otherwise noted for a specific `ioctl()` call, these functions are restricted from use by background processes. Attempts to perform these calls will cause the process group of the process performing the call to be sent a `SIGTTOU` signal. If the process is ignoring `SIGTTOU`, has `SIGTTOU` blocked, or is in the middle of process creation using `vfork()`, the process will be allowed to perform the call and the `SIGTTOU` signal will not be sent.

TCGETS The argument is a pointer to a `termios` structure. The current terminal parameters are fetched and stored into that structure. This call is allowed from a background process; however, the information may subsequently be changed by a foreground process.

TCSETS	The argument is a pointer to a termios structure. The current terminal parameters are set from the values stored in that structure. The change is immediate.
TCSETSW	The argument is a pointer to a termios structure. The current terminal parameters are set from the values stored in that structure. The change occurs after all characters queued for output have been transmitted. This form should be used when changing parameters that will affect output.
TCSETSF	The argument is a pointer to a termios structure. The current terminal parameters are set from the values stored in that structure. The change occurs after all characters queued for output have been transmitted; all characters queued for input are discarded and then the change occurs.
TCGETA	The argument is a pointer to a termio structure. The current terminal parameters are fetched, and those parameters that can be stored in a termio structure are stored into that structure. This call is allowed from a background process; however, the information may subsequently be changed by a foreground process.
TCSETA	The argument is a pointer to a termio structure. Those terminal parameters that can be stored in a termio structure are set from the values stored in that structure. The change is immediate.
TCSETAW	The argument is a pointer to a termio structure. Those terminal parameters that can be stored in a termio structure are set from the values stored in that structure. The change occurs after all characters queued for output have been transmitted. This form should be used when changing parameters that will affect output.
TCSETAF	The argument is a pointer to a termio structure. Those terminal parameters that can be stored in a termio structure are set from the values stored in that structure. The change occurs after all characters queued for output have been transmitted; all characters queued for input are discarded and then the change occurs.
TCSBRK	The argument is an int value. Wait for the output to drain. If the argument is 0, then send a break (zero-valued bits for 0.25 seconds).
TCXONC	Start/stop control. The argument is an int value. If the argument is 0, suspend output; if 1, restart suspended output; if 2, suspend input; if 3, restart suspended input.
TCFLSH	The argument is an int value. If the argument is 0, flush the input queue; if 1, flush the output queue; if 2, flush both the input and output queues.
TIOCEXCL	The argument is ignored. Exclusive-use mode is turned on; no further opens are permitted until the file has been closed, or a TIOCNXCL is issued. The default on open of a terminal file is that exclusive use mode is off.
TIOCNXCL	The argument is ignored. Exclusive-use mode is turned off.
TIOCGPGRP	The argument is a pointer to an int . Set the value of that int to the process group ID of the distinguished process group associated with the terminal. This call is allowed from a background process; however, the information may subsequently be changed by a foreground process.
TIOCSPGRP	The argument is a pointer to an int . Associate the process group whose process group ID is specified by the value of that int with the terminal. The new process group value must be in the range of valid process group ID values, or it must be zero ("no process group"). Otherwise, the error EINVAL is returned. If any processes exist with a process ID or process group ID that is the same as the new process group value, then those processes must have the same real or saved user ID as the real or effective user ID of the calling process or be descendants of the calling process, or the effective user ID of the current process must be super-user. Otherwise, the error EPERM is returned.
TIOCOUTQ	The argument is a pointer to an int . Set the value of that int to the number of characters

in the output stream that have not yet been sent to the terminal. This call is allowed from a background process.

- TIOCSTI** The argument is a pointer to a **char**. Pretend that that character had been received as input.
- TIOCGWINSZ** The argument is a pointer to a **winsize** structure. The terminal driver's notion of the terminal size is stored into that structure. This call is allowed from a background process.
- TIOCSWINSZ** The argument is a pointer to a **winsize** structure. The terminal driver's notion of the terminal size is set from the values specified in that structure. If the new sizes are different from the old sizes, a **SIGWINCH** signal is sent to the process group of the terminal.
- TIOCMGET** The argument is a pointer to an **int**. The current state of the modem status lines is fetched and stored in the **int** pointed to by the argument. This call is allowed from a background process.
- TIOCMBIS** The argument is a pointer to an **int** whose value is a mask containing modem control lines to be turned on. The control lines whose bits are set in the argument are turned on; no other control lines are affected.
- TIOCMBIC** The argument is a pointer to an **int** whose value is a mask containing modem control lines to be turned off. The control lines whose bits are set in the argument are turned off; no other control lines are affected.
- TIOCMSET** The argument is a pointer to an **int** containing a new set of modem control lines. The modem control lines are turned on or off, depending on whether the bit for that mode is set or clear.

SEE ALSO

cs(1), **stty**(1V), **fork**(2), **ioctl**(2), **open**(2V), **read**(2V), **setpgrp**(2V), **sigvec**(2), **vfork**(2), **tty**(4), **getty**(8)

NAME

tm – tapemaster 1/2 inch tape drive

CONFIG — SUN-3 SYSTEM

controller tm0 at vme16d16 ? csr 0xa0 priority 3 vector tmintr 0x60
 controller tm1 at vme16d16 ? csr 0xa2 priority 3 vector tmintr 0x61
 tape mt0 at tm0 drive 0 flags 1
 tape mt0 at tm1 drive 0 flags 1

CONFIG — SUN-2 SYSTEM

controller tm0 at mbio ? csr 0xa0 priority 3
 controller tm0 at vme16 ? csr 0xa0 priority 3 vector tmintr 0x60
 controller tm1 at mbio ? csr 0xa2 priority 3
 controller tm1 at vme16 ? csr 0xa2 priority 3 vector tmintr 0x61
 tape mt0 at tm0 drive 0 flags 1
 tape mt0 at tm1 drive 0 flags 1

DESCRIPTION

The Tapemaster tape controller controls Pertec-interface 1/2" tape drives such as the CDC Keystone, providing a standard tape interface to the device, see mtio(4).

SEE ALSO

mt(1), tar(1), ar(4S), mtio(4)

DIAGNOSTICS

tmn : no response from ctr.
 tmn : error n during config.
 mtn : not online.
 mtn : no write ring.
 tmgo: gate wasn't open. Controller lost synch.
 tmintr: can't clear interrupts.
 tmn : stray interrupts.
 mtn : hard error bn=n er=%x.
 mtn : lost interrupt.

BUGS

The Tapemaster controller does not provide for byte-swapping and the resultant system overhead prevents streaming transports from streaming.

If a non-data error is encountered on non-raw tape, it refuses to do anything more until closed.

The system should remember which controlling terminal has the tape drive open and write error messages to that terminal rather than on the console.

NAME

ttcompat – V7 and 4BSD STREAMS compatibility module

CONFIG

None; included by default.

SYNOPSIS

```
#include <sys/stream.h>
#include <sys/stropt.h>

ioctl(fd, I_PUSH, "ttcompat");
```

DESCRIPTION

ttcompat is a STREAMS module that translates the **ioctl** calls supported by the older Version 7 and 4BSD terminal drivers into the **ioctl** calls supported by the **termio(4)** interface. All other messages pass through this module unchanged; the behavior of **read** and **write** calls is unchanged, as is the behavior of **ioctl** calls other than the ones supported by **ttcompat**.

Normally, this module is automatically pushed onto a stream when a terminal device is opened; it does not have to be explicitly pushed onto a stream. This module requires that the **termio** interface be supported by the modules and driver downstream. The **TCGETS**, **TCSETS**, and **TCSETSF** **ioctl** calls must be supported; if any information set or fetched by those **ioctl** calls is not supported by the modules and driver downstream, some of the V7/4BSD functions may not be supported. For example, if the **CBAUD** bits in the **c_cflag** field are not supported, the functions provided by the **sg_ispeed** and **sg_ospeed** fields of the **sgttyb** structure (see below) will not be supported. If the **TCFLSH** **ioctl** is not supported, the function provided by the **TIOCFLUSH** **ioctl** will not be supported. If the **TCXONC** **ioctl** is not supported, the functions provided by the **TIOCSTOP** and **TIOCSTART** **ioctl** calls will not be supported. If the **TIOCMBIS** and **TIOCMBIC** **ioctl** calls are not supported, the functions provided by the **TIOCSDTR** and **TIOCCDTR** **ioctl** calls will not be supported.

The basic **ioctl** calls use the **sgttyb** structure defined by **<sys/ioctl.h>**:

```
struct sgttyb {
    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, and reflect the values in the **c_cflag** field of the **termio** structure. The **sg_erase** and **sg_kill** fields of the argument structure specify the erase and kill characters respectively, and reflect the values in the **VERASE** and **VKILL** members of the **c_cc** field of the **termio** structure.

The **sg_flags** field of the argument structure contains several flags that determine the system's treatment of the terminal. They are mapped into flags in fields of the terminal state, represented by the **termio** structure.

Delay type 0 is always mapped into the equivalent delay type 0 in the **c_oflag** field of the **termio** structure. Other delay mappings are performed as follows:

sg_flags	c_oflag
BS1	BS1
FF1	VT1
CR1	CR2
CR2	CR3
CR3	not supported
TAB1	TAB1
TAB2	TAB2
XTABS	TAB3


```

NL1      ONLRET|CR1
NL2      NL1

```

If previous `TIOCLSET` or `TIOCLBIS` `ioctl` calls have not selected `LITOUT` or `PASS8` mode, and if `RAW` mode is not selected, the `ISTRIP` flag is set in the `c_iflag` field of the `termio` structure, and the `EVENP` and `ODDP` flags control the parity of characters sent to the terminal and accepted from the terminal:

0 Parity is not to be generated on output or checked on input; the character size is set to `CS8` and the `PARENB` flag is cleared in the `c_cflag` field of the `termio` structure.

EVENP Even parity characters are to be generated on output and accepted on input; the `INPCK` flag is set in the `c_iflag` field of the `termio` structure, the character size is set to `CS7` and the `PARENB` flag is set in the `c_cflag` field of the `termio` structure.

ODDP Odd parity characters are to be generated on output and accepted on input; the `INPCK` flag is set in the `c_iflag` field, the character size is set to `CS7` and the `PARENB` and `PARODD` flags are set in the `c_cflag` field of the `termio` structure.

EVENP|ODDP

Even parity characters are to be generated on output and characters of either parity are to be accepted on input; the `INPCK` flag is cleared in the `c_iflag` field, the character size is set to `CS7` and the `PARENB` flag is set in the `c_cflag` field of the `termio` structure.

The `RAW` flag disables all output processing (the `OPOST` flag in the `c_oflag` field, and the `XCASE` flag in the `c_iflag` field, are cleared in the `termio` structure) and input processing (all flags in the `c_iflag` field other than the `IXOFF` and `IXANY` flags are cleared in the `termio` structure). 8 bits of data, with no parity bit, are accepted on input and generated on output; the character size is set to `CS8` and the `PARENB` and `PARODD` flags are cleared in the `c_cflag` field of the `termio` structure. The signal-generating and line-editing control characters are disabled by clearing the `ISIG` and `ICANON` flags in the `c_iflag` field of the `termio` structure.

The `CRMOD` flag turn input `RETURN` characters into `NEWLINE` characters, and output and echoed `NEWLINE` characters to be output as a `RETURN` followed by a `LINEFEED`. The `ICRNL` flag in the `c_iflag` field, and the `OPOST` and `ONLCR` flags in the `c_oflag` field, are set in the `termio` structure.

The `LCASE` flag maps upper-case letters in the ASCII character set to their lower-case equivalents on input (the `IUCLC` flag is set in the `c_iflag` field), and maps lower-case letters in the ASCII character set to their upper-case equivalents on output (the `OLCUC` flag is set in the `c_oflag` field). Escape sequences are accepted on input, and generated on output, to handle certain ASCII characters not supported by older terminals (the `XCASE` flag is set in the `c_iflag` field).

Other flags are directly mapped to flags in the `termio` structure:

```

sg_flags  flags in termio structure
CBREAK    complement of ICANON in c_iflag field
ECHO      ECHO in c_iflag field
TANDEM    IXOFF in c_iflag field

```

Another structure associated with each terminal specifies characters that are special in both the old Version 7 and the newer 4BSD terminal interfaces. The following structure is defined by `<sys/ioctl.h>`:

```

struct tchars {
    char  t_intrc;      /* interrupt */
    char  t_quitc;     /* 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 characters are mapped to members of the `c_cc` field of the `termio` structure as follows:

tchars	c_cc index
<code>t_intrc</code>	VINTR
<code>t_quite</code>	VQUIT
<code>t_startc</code>	VSTART
<code>t_stopc</code>	VSTOP
<code>t_eofc</code>	VEOF
<code>t_brkc</code>	VEOL

Also associated with each terminal is a local flag word, specifying flags supported by the new 4BSD terminal interface. Most of these flags are directly mapped to flags in the `termio` structure:

local flags	flags in termio structure
LCRTBS	not supported
LPRTERA	ECHOPRT in the <code>c_lflag</code> field
LCRTERA	ECHOE in the <code>c_lflag</code> field
LTILDE	not supported
LTOSTOP	TOSTOP in the <code>c_lflag</code> field
LFLUSHO	FLUSHO in the <code>c_lflag</code> field
LNOHANG	CLOCAL in the <code>c_cflag</code> field
LCRTKIL	ECHOKE in the <code>c_lflag</code> field
LCTLECH	CTLECH in the <code>c_lflag</code> field
LPENDIN	PENDIN in the <code>c_lflag</code> field
LDECCCTQ	complement of IXANY in the <code>c_iflag</code> field
LNOFLSH	NOFLSH in the <code>c_lflag</code> field

Another structure associated with each terminal is the `ltchars` structure which defines control characters for the new 4BSD terminal interface. Its structure is:

```

struct ltchars {
    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_inxctc;    /* literal next character */
};

```

The characters are mapped to members of the `c_cc` field of the `termio` structure as follows:

ltchars	c_cc index
<code>t_suspc</code>	VSUSP
<code>t_dsuspc</code>	VDSUSP
<code>t_rprntc</code>	VREPRINT
<code>t_flushc</code>	VDISCARD
<code>t_werasc</code>	VWERASE
<code>t_inxctc</code>	VLNEXT

IOCTLS

`ttcompat` responds to the following `ioctl` calls. All others are passed to the module below.

TIOCGETP The argument is a pointer to an `sgttyb` structure. The current terminal state is fetched; the appropriate characters in the terminal state are stored in that structure, as are the input and output speeds. The values of the flags in the `sg_flags` field are derived from the flags in the terminal state and stored in the structure.

- TIOCSETP** The argument is a pointer to an `sgttyb` structure. The appropriate characters and input and output speeds in the terminal state are set from the values in that structure, and the flags in the terminal state are set to match the values of the flags in the `sg_flags` field of that structure. The state is changed with a `TCSETSf ioctl`, so that the interface delays until output is quiescent, then throws away any unread characters, before changing the modes.
- TIOCSETN** The argument is a pointer to an `sgttyb` structure. The terminal state is changed as `TIOCSETP` would change it, but a `TCSETSio ioctl` is used, so that the interface neither delays nor discards input.
- TIOCHPCL** The argument is ignored. The `HUPCL` flag is set in the `c_cflag` word of the terminal state.
- TIOCFUSH** The argument is a pointer to an `int` variable. If its value is zero, all characters waiting in input or output queues are flushed. Otherwise, the value of the `int` is treated as the logical OR of the `FREAD` and `FWRITE` flags defined by `<sys/file.h>`; if the `FREAD` bit is set, all characters waiting in input queues are flushed, and if the `FWRITE` bit is set, all characters waiting in output queues are flushed.
- TIOCSBRK** The argument is ignored. The break bit is set for the device.
- TIOCCBRK** The argument is ignored. The break bit is cleared for the device.
- TIOCSDTR** The argument is ignored. The Data Terminal Ready bit is set for the device.
- TIOCCDTR** The argument is ignored. The Data Terminal Ready bit is cleared for the device.
- TIOCSTOP** The argument is ignored. Output is stopped as if the `STOP` character had been typed.
- TIOCSTART** The argument is ignored. Output is restarted as if the `START` character had been typed.
- TIOCGETC** The argument is a pointer to an `tchars` structure. The current terminal state is fetched, and the appropriate characters in the terminal state are stored in that structure.
- TIOCSETC** The argument is a pointer to an `tchars` structure. The values of the appropriate characters in the terminal state are set from the characters in that structure.
- TIOCLGET** The argument is a pointer to an `int`. The current terminal state is fetched, and the values of the local flags are derived from the flags in the terminal state and stored in the `int` pointed to by the argument.
- TIOCLBIS** The argument is a pointer to an `int` whose value is a mask containing flags to be set in the local flags word. The current terminal state is fetched, and the values of the local flags are derived from the flags in the terminal state; the specified flags are set, and the flags in the terminal state are set to match the new value of the local flags word.
- TIOCLBIC** The argument is a pointer to an `int` whose value is a mask containing flags to be cleared in the local flags word. The current terminal state is fetched, and the values of the local flags are derived from the flags in the terminal state; the specified flags are cleared, and the flags in the terminal state are set to match the new value of the local flags word.
- TIOCLSET** The argument is a pointer to an `int` containing a new set of local flags. The flags in the terminal state are set to match the new value of the local flags word.
- TIOCGLTC** The argument is a pointer to an `ltchars` structure. The values of the appropriate characters in the terminal state are stored in that structure.
- TIOCSLTC** The argument is a pointer to an `ltchars` structure. The values of the appropriate characters in the terminal state are set from the characters in that structure.

SEE ALSO`ioctl(2)`, `termio(4)`

NAME

tty – controlling terminal interface

DESCRIPTION

The file `/dev/tty` is, in each process, a synonym for the controlling terminal of that process, if any. It is useful for programs or shell sequences 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 the name of a file for output, when typed output is desired and it is tiresome to find out what terminal is currently in use.

IOCTLS

In addition to the ioctls supported by the device that tty refers to, the following ioctl is supported:

TIOCNOTTY Detach the current process from its controlling terminal, and remove it from its current process group, without attaching it to a new process group (that is, set its process group ID to zero). This ioctl call only works on file descriptors connected to `/dev/tty`; this is used by daemon processes when they are invoked by a user at a terminal. The process attempts to open `/dev/tty`; if the open succeeds, it detaches itself from the terminal by using **TIOCNOTTY**, while if the open fails, it is obviously not attached to a terminal and does not need to detach itself.

FILES

`/dev/tty`

SEE ALSO

`termio(4)`

NAME

udp – Internet User Datagram Protocol

SYNOPSIS

```
#include <sys/socket.h>
#include <netinet/in.h>

s = socket(AF_INET, SOCK_DGRAM, 0);
```

DESCRIPTION

UDP is a simple, unreliable datagram protocol which is used to support the `SOCK_DGRAM` abstraction for the Internet protocol family. It is layered directly above the Internet Protocol (IP). UDP sockets are connectionless, and are normally used with the `sendto`, `sendmsg`, `recvfrom`, and `recvmsg` system calls (see `send(2)` and `recv(2)`). If the `connect(2)` system call is used to fix the destination for future packets, then the `recv(2)` or `read(2V)` and `send(2)` or `write(2V)` system calls may be used.

UDP address formats are identical to those used by the Transmission Control Protocol (TCP). Like TCP, UDP uses a port number along with an IP address to identify the endpoint of communication. Note: the UDP port number space is separate from the TCP port number space (that is, a UDP port may not be “connected” to a TCP port). The `bind(2)` system call can be used to set the local address and port number of a UDP socket. The local IP address may be left unspecified in the `bind` call by using the special value `INADDR_ANY`. If the `bind` call is not done, a local IP address and port number will be assigned to each packet as it is sent. Broadcast packets may be sent (assuming the underlying network supports this) by using a reserved “broadcast address”; this address is network interface dependent. Broadcasts may only be sent by the super-user.

Options at the IP level may be used with UDP; see `ip(4P)`.

There are a variety of ways that a UDP packet can be lost or discarded, including a failure of the underlying communication mechanism. UDP implements a checksum over the data portion of the packet. If the checksum of a received packet is in error, the packet will be dropped with no indication given to the user. A queue of received packets is provided for each UDP socket. This queue has a limited capacity. Arriving datagrams which will not fit within its *high-water* capacity are silently discarded.

UDP processes Internet Control Message Protocol (ICMP) error messages received in response to UDP packets it has sent. See `icmp(4P)`. ICMP “source quench” messages are ignored. ICMP “destination unreachable,” “time exceeded” and “parameter problem” messages disconnect the socket from its peer so that subsequent attempts to send packets using that socket will return an error. UDP will not guarantee that packets are delivered in the order they were sent. As well, duplicate packets may be generated in the communication process.

ERRORS

A socket operation may fail if:

EISCONN	A <code>connect</code> operation was attempted on a socket on which a <code>connect</code> operation had already been performed, and the socket could not be successfully disconnected before making the new connection.
EISCONN	A <code>sendto</code> or <code>sendmsg</code> operation specifying an address to which the message should be sent was attempted on a socket on which a <code>connect</code> operation had already been performed.
ENOTCONN	A <code>send</code> or <code>write</code> operation, or a <code>sendto</code> or <code>sendmsg</code> operation not specifying an address to which the message should be sent, was attempted on a socket on which a <code>connect</code> operation had not already been performed.
EADDRINUSE	A <code>bind</code> operation was attempted on a socket with a network address/port pair that has already been bound to another socket.
EADDRNOTAVAIL	A <code>bind</code> operation was attempted on a socket with a network address for which no network interface exists.

EINVAL A `sendmsg` operation with a non-NULL `msg_accrights` was attempted.

EACCES A `bind` operation was attempted with a “reserved” port number and the effective user ID of the process was not super-user.

ENOBUFS The system ran out of memory for internal data structures.

SEE ALSO

`bind(2)`, `connect(2)`, `read(2V)`, `recv(2)`, `send(2)`, `write(2V)`, `icmp(4P)`, `inet(4F)`, `ip(4P)`, `tcp(4P)`

Postel, Jon, *User Datagram Protocol*, RFC 768, Network Information Center, SRI International, Menlo Park, Calif., August 1980. (Sun 800-1054-01)

BUGS

`SIOCSHIWAT` and `SIOCGHIWAT` `ioctl`'s to set and get the high water mark for the socket queue, and so that it can be changed from 2048 bytes to be larger or smaller, have been defined (in `<sys/ioctl.h>`) but not implemented.

Something sensible should be done with ICMP source quench error messages if the socket is bound to a peer socket.

NAME

vp – Ikon 10071-5 Versatec parallel printer interface

CONFIG — SUN-2 SYSTEM

device vp0 at mbio ? csr

DESCRIPTION

This Sun interface to the Versatec printer/plotter is supported by the Ikon parallel interface board, a word DMA device, which is output only.

The Versatec is normally handled by the line printer spooling system and should not be accessed by the user directly.

Opening the device `/dev/vp0` may yield one of two errors: `ENXIO` indicates that the device is already in use; `EIO` indicates that the device is offline.

The printer operates in either print or plot mode. To set the printer into plot mode you should include `<vcmd.h>` and use the `ioctl(2)` call

```
ioctl(f, VSETSTATE, plotmd);
```

where `plotmd` is defined to be

```
int plotmd[] = { VPLOT, 0, 0 };
```

When going back into print mode from plot mode you normally eject paper by sending it an EOT after putting into print mode:

```
int prtmd[] = { VPRINT, 0, 0 };
```

```
...
```

```
fflush(vp);
```

```
f = fileno(vp);
```

```
ioctl(f, VSETSTATE, prtmd);
```

```
write(f, "\04", 1);
```

FILES

`/dev/vp0`

SEE ALSO

`ioctl(2)`

BUGS

If you use the standard I/O library on the Versatec, be sure to explicitly set a buffer using `setbuf`, since the library will not use buffered output by default, and will run very slowly.

Writes must start on even byte boundaries and be an even number of bytes in length.

NAME

vpc – Systech VPC-2200 Versatec printer/plotter and Centronics printer interface

CONFIG — SUN-2 SYSTEM

device vpc0 at mbio ? csr 0x480 priority 2

device vpc1 at mbio ? csr 0x500 priority 2

DESCRIPTION

This Sun interface to the Versatec printer/plotter and to Centronics printers is supported by the Systech parallel interface board, an output-only byte-wide DMA device. The device has one channel for Versatec devices and one channel for Centronics devices, with an optional long lines interface for Versatec devices.

Devices attached to this interface are normally handled by the line printer spooling system and should not be accessed by the user directly.

Opening the device `/dev/vpc0` or `/dev/lp0` may yield one of two errors: ENXIO indicates that the device is already in use; EIO indicates that the device is offline.

The Versatec printer/plotter operates in either print or plot mode. To set the printer into plot mode you should include `<vcmd.h>` and use the `ioctl(2)` call:

```
ioctl(f, VSETSTATE, plotmd);
```

where `plotmd` is defined to be

```
int plotmd[ ] = { VPLOT, 0, 0 };
```

When going back into print mode from plot mode you normally eject paper by sending it an EOT after putting into print mode:

```
int prtmd[ ] = { VPRINT, 0, 0 };
```

```
...
```

```
fflush(vpc);
```

```
f = fileno(vpc);
```

```
ioctl(f, VSETSTATE, prtmd);
```

```
write(f, "\04", 1);
```

FILES

`/dev/vpc0`

`/dev/lp0`

SEE ALSO

`ioctl(2)`

BUGS

If you use the standard I/O library on the Versatec, be sure to explicitly set a buffer using `setbuf`, since the library will not use buffered output by default, and will run very slowly.

NAME

win – Sun window system

CONFIG

pseudo-device *winnumber*
pseudo-device *dtopnumber*

DESCRIPTION

The **win** pseudo-device accesses the system drivers supporting the Sun window system. *number*, in the device description line above, indicates the maximum number of windows supported by the system. *number* is set to 128 in the GENERIC system configuration file used to generate the kernel used in Sun systems as they are shipped. The *dtop* pseudo-device line indicates the number of separate “desktops” (frame buffers) that can be actively running the Sun window system at once. In the GENERIC file, this number is set to 4.

Each window in the system is represented by a */dev/win** device. The windows are organized as a tree with windows being subwindows of their parents, and covering/covered by their siblings. Each window has a position in the tree, a position on a display screen, an input queue, and information telling what parts of it are exposed.

The window driver multiplexes keyboard and mouse input among the several windows, tracks the mouse with a cursor on the screen, provides each window access to information about what parts of it are exposed, and notifies the manager process for a window when the exposed area of the window changes so that the window may repair its display.

Full information on the window system functions is given in the *SunView 1 System Programmer's Guide*.

FILES

/dev/win[0-9]
/dev/win[0-9][0-9]

SEE ALSO

SunView 1 System Programmer's Guide

NAME

xd – Disk driver for Xylogics 7053 SMD Disk Controller

CONFIG — SUN-3 SYSTEM

controller xdc0 at vme16d32 ? csr 0xee80 priority 2 vector xdintr 0x44
 controller xdc1 at vme16d32 ? csr 0xee90 priority 2 vector xdintr 0x45
 controller xdc2 at vme16d32 ? csr 0xeea0 priority 2 vector xdintr 0x46
 controller xdc3 at vme16d32 ? csr 0xeeb0 priority 2 vector xdintr 0x47
 disk xd0 at xdc0 drive 0
 disk xd1 at xdc0 drive 1
 disk xd2 at xdc0 drive 2
 disk xd3 at xdc0 drive 3
 disk xd4 at xdc1 drive 0
 disk xd5 at xdc1 drive 1
 disk xd6 at xdc1 drive 2
 disk xd7 at xdc1 drive 3
 disk xd8 at xdc2 drive 0
 disk xd9 at xdc2 drive 1
 disk xd10 at xdc2 drive 2
 disk xd11 at xdc2 drive 3
 disk xd12 at xdc3 drive 0
 disk xd13 at xdc3 drive 1
 disk xd14 at xdc3 drive 2
 disk xd15 at xdc3 drive 3

The four controller lines given in the synopsis section above specify the first, second, third, and fourth Xylogics 7053 SMD disk controller in a Sun system.

DESCRIPTION

Files with minor device numbers 0 through 7 refer to various portions of drive 0; minor devices 8 through 15 refer to drive 1, and so on. The standard device names begin with xd followed by the drive number and then a letter a-h for partitions 0-7 respectively. The character ? stands here for a drive number in the range 0-7.

The block files access the disk using the system's normal buffering mechanism and may be read and written without regard to physical disk records. There is also a "raw" interface which provides for direct transmission between the disk and the user's read or write buffer. A single read or write call usually results in only one I/O operation; therefore raw I/O is considerably more efficient when many words are transmitted. The names of the raw files conventionally begin with an extra r.

In raw I/O counts should be a multiple of 512 bytes (a disk sector). Likewise **directory(3)** calls should specify a multiple of 512 bytes.

If flags 0x1 is specified, the overlapped seeks feature for that drive is turned off. Note: to be effective, the flag must be set on all drives for a specific controller. This action is necessary for controllers with older firmware, which have bugs preventing overlapped seeks from working properly.

DISK SUPPORT

This driver handles all SMD drives by reading a label from sector 0 of the drive which describes the disk geometry and partitioning.

The xd?a partition is normally used for the root file system on a disk, the xd?b partition as a paging area, and the xd?c partition for pack-pack copying (it normally maps the entire disk). The rest of the disk is normally the xd?g partition.

FILES

/dev/xd[0-7][a-h]	block files
/dev/rxd[0-7][a-h]	raw files

SEE ALSO

lseek(2), read(2V), write(2V), directory(3), dkio(4S)

DIAGNOSTICS**xdcn: self test error**

Self test error in controller, see the Maintenance and Reference Manual.

xdn: unable to read bad sector

The bad sector forwarding information for the disk could not be read.

xdn: initialization failed

The drive could not be successfully initialized.

xdn: unable to read label

The drive geometry/partition table information could not be read.

xdn: Corrupt label

The geometry/partition label checksum was incorrect.

xdn: offline

A drive ready status is no longer detected, so the unit has been logically removed from the system. If the drive ready status is restored, the unit will automatically come back online the next time it is accessed.

xdbc: cmd how (msg) blk #n abs blk #n

A command such as read or write encountered an error condition (*how*): either it *failed*, the controller was *reset*, the unit was *restored*, or an operation was *retry*'ed. The *msg* is derived from the error number given by the controller, indicating a condition such as "drive not ready(rq, "sector not found" or "disk write protected". The *blk #* is the sector in error relative to the beginning of the partition involved. The *abs blk #* is the absolute block number of the sector in error. Some fields of the error message may be missing since the information is not always available.

BUGS

In raw I/O **read(2V)** and **write(2V)** truncate file offsets to 512-byte block boundaries, and **write(2V)** scribbles on the tail of incomplete blocks. Thus, in programs that are likely to access raw devices, **read(2V)**, **write(2V)** and **lseek(2)** should always deal in 512-byte multiples.

Older revisions of the firmware do not properly support overlapped seeks. This will only affect systems with multiple disks on a single controller. If a large number of "zero sector count" errors appear, you should use the **flags** field to disable overlapped seeks.

NAME

xt - Xylogics 472 1/2 inch tape controller

CONFIG — SUN-3 SYSTEM

controller xtc0 at vme16d16 ? csr 0xee60 priority 3 vector xtintr 0x64
controller xtc1 at vme16d16 ? csr 0xee68 priority 3 vector xtintr 0x65
tape xt0 at xtc0 drive 0 flags 1
tape xt1 at xtc1 drive 0 flags 1

CONFIG — SUN-2 SYSTEM

controller xtc0 at mbio ? csr 0xee60 priority 3
controller xtc0 at vme16 ? csr 0xee60 priority 3 vector xtintr 0x64
controller xtc1 at mbio ? csr 0xee68 priority 3
controller xtc1 at vme16 ? csr 0xee68 priority 3 vector xtintr 0x65
tape xt0 at xtc0 drive 0 flags 1
tape xt1 at xtc1 drive 0 flags 1

DESCRIPTION

The Xylogics 472 tape controller controls Pertec-interface 1/2" tape drives such as the CDC Keystone III, providing a standard tape interface to the device, see **mtio(4)**. This controller is used to support high speed or high density drives, which are not supported effectively by the older TapeMaster controller (see **tm(4S)**).

The flags field is used to control remote density select operation: a 0 specifies no remote density selection is to be attempted, a 1 specifies that the Pertec density-select line is used to toggle between high and low density; a 2 specifies that the Pertec speed-select line is used to toggle between high and low density. The default is 1, which is appropriate for the CDC Keystone III (92185) and the Telex 9250. In no case will the controller select among more than 2 densities.

SEE ALSO

mt(1), **tar(1)**, **mtio(4)**, **tm(4S)**

NAME

xy – Disk driver for Xylogics 450 and 451 SMD Disk Controllers

CONFIG — SUN-3 SYSTEM

controller xyc0 at vme16d16 ? csr 0xee40 priority 2 vector xyintr 0x48
controller xyc1 at vme16d16 ? csr 0xee48 priority 2 vector xyintr 0x49
disk xy0 at xyc0 drive 0
disk xy1 at xyc0 drive 1
disk xy2 at xyc1 drive 0
disk xy3 at xyc1 drive 1

The two controller lines given in the synopsis sections above specify the first and second Xylogics 450 or 451 SMD disk controller in a Sun system.

CONFIG — SUN-2 SYSTEM

controller xyc0 at vme16 ? csr 0xee40 priority 2 vector xyintr 0x48
controller xyc1 at vme16 ? csr 0xee48 priority 2 vector xyintr 0x49
controller xyc0 at mbio ? csr 0xee40 priority 2
controller xyc1 at mbio ? csr 0xee48 priority 2
disk xy0 at xyc0 drive 0
disk xy1 at xyc0 drive 1
disk xy2 at xyc1 drive 0
disk xy3 at xyc1 drive 1

The first two controller lines specify the first and second Xylogics 450 or 451 SMD disk controllers in a Sun-2/160 VMEbus based system. The third and fourth controller lines specify the first and second Xylogics 450 or 451 SMD disk controllers in a Sun-2/120 or a Sun-2/170 Multibus based system.

DESCRIPTION

Files with minor device numbers 0 through 7 refer to various portions of drive 0; minor devices 8 through 15 refer to drive 1, and so on. The standard device names begin with xy followed by the drive number and then a letter a-h for partitions 0-7 respectively. The character '?' stands here for a drive number in the range 0-7.

The block files access the disk using the system's normal buffering mechanism and may be read and written without regard to physical disk records. There is also a "raw" interface which provides for direct transmission between the disk and the user's read or write buffer. A single read or write call usually results in only one I/O operation; therefore raw I/O is considerably more efficient when many words are transmitted. The names of the raw files conventionally begin with an extra r.

In raw I/O counts should be a multiple of 512 bytes (a disk sector). Likewise `directory(3)` calls should specify a multiple of 512 bytes.

If `flags 0x1` is specified, the overlapped seeks feature for that drive is turned off. Note: to be effective, the flag must be set on all drives for a specific controller. This action is necessary for controllers with older firmware, which have bugs preventing overlapped seeks from working properly.

DISK SUPPORT

This driver handles all SMD drives by reading a label from sector 0 of the drive which describes the disk geometry and partitioning.

The `xy?a` partition is normally used for the root file system on a disk, the `xy?b` partition as a paging area, and the `xy?c` partition for pack-pack copying (it normally maps the entire disk). The rest of the disk is normally the `xy?g` partition.

FILES

<code>/dev/xy[0-7][a-h]</code>	block files
<code>/dev/rxy[0-7][a-h]</code>	raw files

SEE ALSO

lseek(2), read(2V), directory(3), write(2V), dkio(4S)

DIAGNOSTICS**xycn : self test error**

Self test error in controller, see the Maintenance and Reference Manual.

xycn: WARNING: *n* bit addresses

The controller is strapped incorrectly. Sun systems use 20-bit addresses for Multibus based systems and 24-bit addresses for VMEbus based systems.

xyn : unable to read bad sector info

The bad sector forwarding information for the disk could not be read.

xyn and xyn are of same type (*n*) with different geometries.

The 450 and 451 do not support mixing the drive types found on these units on a single controller.

xyn : initialization failed

The drive could not be successfully initialized.

xyn : unable to read label

The drive geometry/partition table information could not be read.

xyn : Corrupt label

The geometry/partition label checksum was incorrect.

xyn : offline

A drive ready status is no longer detected, so the unit has been logically removed from the system. If the drive ready status is restored, the unit will automatically come back online the next time it is accessed.

xync: *cmd how (msg) blk #n abs blk #n*

A command such as read or write encountered an error condition (*how*): either it *failed*, the controller was *reset*, the unit was *restored*, or an operation was *retry*'ed. The *msg* is derived from the error number given by the controller, indicating a condition such as "drive not ready", "sector not found" or "disk write protected". The *blk #* is the sector in error relative to the beginning of the partition involved. The *abs blk #* is the absolute block number of the sector in error. Some fields of the error message may be missing since the information is not always available.

BUGS

In raw I/O **read(2V)** and **write(2V)** truncate file offsets to 512-byte block boundaries, and **write(2V)** scribbles on the tail of incomplete blocks. Thus, in programs that are likely to access raw devices, **read(2V)**, **write(2V)** and **lseek(2)** should always deal in 512-byte multiples.

Older revisions of the firmware do not properly support overlapped seeks. This will only affect systems with multiple disks on a single controller. If a large number of "zero sector count" errors appear, you should use the **flags** field to disable overlapped seeks.

NAME

zero – source of zeroes

SYNOPSIS

None; included with standard system.

DESCRIPTION

A zero special file is a source of zeroed unnamed memory.

Reads from a zero special file always return a buffer full of zeroes. The file is of infinite length.

Writes to a zero special file are always successful, but the data written is ignored.

Mapping a zero special file creates a zero-initialized unnamed memory object of a length equal to the length of the mapping and rounded up to the nearest page size as returned by `getpagesize(2)`. Multiple processes can share such a zero special file object provided a common ancestor mapped the object `MAP_SHARED`.

FILES

`/dev/zero`

SEE ALSO

`fork(2)`, `getpagesize(2)`, `mmap(2)`

NAME

zs – Zilog 8530 SCC serial communications driver

CONFIG — SUN-3 SYSTEM

device zs0 at obio ? csr 0x20000 flags 3 priority 3
device zs1 at obio ? csr 0x00000 flags 0x103 priority 3

CONFIG — SUN-2 SYSTEM

device zs0 at virtual ? csr 0xeec800 flags 3 priority 3
device zs1 at virtual ? csr 0xeec000 flags 0x103 priority 3
device zs2 at mbmem ? csr 0x80800 flags 3 priority 3
device zs3 at mbmem ? csr 0x81000 flags 3 priority 3
device zs4 at mbmem ? csr 0x84800 flags 3 priority 3
device zs5 at mbmem ? csr 0x85000 flags 3 priority 3

CONFIG — Sun386i SYSTEM

device zs0 at obmem ? csr 0xFC000000 flags 3 irq 9 priority 6
device zs1 at obmem ? csr 0xA0000020 flags 0x103 irq 9 priority 6

SYNOPSIS

```
#include <fcntl.h>
#include <sys/termios.h>
open("/dev/tty $n$ ", mode);
open("/dev/ttyd $n$ ", mode);
open("/dev/cuan", mode);
```

DESCRIPTION

The Zilog 8530 provides 2 serial communication ports with full modem control in asynchronous mode. Each port supports those `termio(4)` device control functions specified by flags in the `c_cflag` word of the `termios` structure and by the `IGNBRK`, `IGNPAR`, `PARMRK`, or `INPCK` flags in the `c_iflag` word of the `termios` structure are performed by the `zs` driver. All other `termio(4)` functions must be performed by STREAMS modules pushed atop the driver; when a device is opened, the `ldterm(4M)` and `ttcompat(4M)` STREAMS modules are automatically pushed on top of the stream, providing the standard `termio(4)` interface.

Of the synopsis lines above, the line for `zs0` specifies the serial I/O port(s) provided by the CPU board, the line for `zs1` specifies the Video Board ports (which are used for keyboard and mouse), the lines for `zs2` and `zs3` specify the first and second ports on the first SCSI board in a system, and those for `zs4` and `zs5` specify the first and second ports provided by the second SCSI board in a system, respectively.

Bit i of `flags` may be specified to say that a line is not properly connected, and that the line i should be treated as hard-wired with carrier always present. Thus specifying `flags 0x2` in the specification of `zs0` would treat line `/dev/ttyb` in this way.

Minor device numbers in the range 0 – 11 correspond directly to the normal tty lines and are named `/dev/ttya` and `/dev/ttyb` for the two serial ports on the CPU board and `/dev/tty s n` for the ports on the SCSI boards; n is 0 or 1 for the ports on the first SCSI board, and 2 or 3 for the ports on the second SCSI board.

To allow a single tty line to be connected to a modem and used for both incoming and outgoing calls, a special feature, controlled by the minor device number, has been added. Minor device numbers in the range 128 – 139 correspond to the same physical lines as those above (that is, the same line as the minor device number minus 128).

A dial-in line has a minor device in the range 0 – 11 and is conventionally renamed `/dev/ttyd n` , where n is a number indicating which dial-in line it is (so that `/dev/ttyd0` is the first dial-in line), and the dial-out line corresponding to that dial-in line has a minor device number 128 greater than the minor device number of the dial-in line and is conventionally named `/dev/cuan`, where n is the number of the dial-in line.

The `/dev/cua n` lines are special in that they can be opened even when there is no carrier on the line. Once a `/dev/cua n` line is opened, the corresponding tty line can not be opened until the `/dev/cua n` line is closed; a blocking open will wait until the `/dev/cua n` line is closed (which will drop Data Terminal Ready, after which Carrier Detect will usually drop as well) and carrier is detected again, and a non-blocking open will return an error. Also, if the `/dev/ttyd n` line has been opened successfully (usually only when carrier is recognized on the modem) the corresponding `/dev/cua n` line can not be opened. This allows a modem to be attached to e.g. `/dev/ttyd0` (renamed from `/dev/ttya`) and used for dialin (by enabling the line for login in `/etc/ttytab`) and also used for dialout (by `tip(1C)` or `uucp(1C)`) as `/dev/cua0` when no one is logged in on the line. Note: the bit in the `flags` word in the configuration file (see above) must be zero for this line, which enables hardware carrier detection.

IOCTLS

The standard set of `termio ioctl()` calls are supported by `zs`.

If the `CRTSCTS` flag in the `c_cflag` is set, output will be generated only if CTS is high; if CTS is low, output will be frozen. If the `CRTSCTS` flag is clear, the state of CTS has no effect. Breaks can be generated by the `TCSBRK`, `TIOCSBRK`, and `TIOCCBRK ioctl()` calls. The modem control lines `TIOCM_CAR`, `TIOCM_CTS`, `TIOCM_RTS`, and `TIOCM_DTR` are provided.

The input and output line speeds may be set to any of the speeds supported by `termio`. The speeds cannot be set independently; when the output speed is set, the input speed is set to the same speed.

ERRORS

An `open()` will fail if:

<code>ENXIO</code>	The unit being opened does not exist.
<code>EBUSY</code>	The dial-out device is being opened and the dial-in device is already open, or the dial-in device is being opened with a no-delay open and the dial-out device is already open.
<code>EBUSY</code>	The unit has been marked as exclusive-use by another process with a <code>TIOCEXCL ioctl()</code> call.
<code>EINTR</code>	The open was interrupted by the delivery of a signal.

FILES

<code>/dev/tty{a,b,s[0-3]}</code>	hardwired tty lines
<code>/dev/ttyd[0-9a-f]</code>	dialin tty lines
<code>/dev/cua[0-9a-f]</code>	dialout tty lines

SEE ALSO

`tip(1C)`, `uucp(1C)`, `mcp(4S)`, `mti(4S)`, `termio(4)`, `ldterm(4M)`, `ttcompat(4M)`

DIAGNOSTICS

`zsn c`: silo overflow.

The 8530 character input silo overflowed before it could be serviced.

`zsn c`: ring buffer overflow.

The driver's character input ring buffer overflowed before it could be serviced.

NAME

intro – file formats used or read by various programs

DESCRIPTION

This section describes formats of files used by various programs.

LIST OF FILE FORMATS

Name	Appears on Page	Description
a.out	a.out(5)	assembler and link editor output format
acct	acct(5)	execution accounting file
addresses	aliases(5)	addresses and aliases for sendmail(8)
aliases	aliases(5)	addresses and aliases for sendmail(8)
ar	ar(5)	archive (library) file format
audit.log	audit.log(5)	the security audit trail file
audit_control	audit_control(5)	control information for system audit daemon
audit_data	audit_data(5)	current information on audit daemon
auto.home	auto.home(5)	automount map for home directories
auto.vol	auto.vol(5)	automount map for volumes
bar	bar(5)	tape archive file format
bootparams	bootparams(5)	boot parameter data base
COFF	coff(5)	common assembler and link editor output
core	core(5)	format of memory image file
cpio	cpio(5)	format of cpio archive
crontab	crontab(5)	table of times to run periodic jobs
defaults	defaults(5)	default specifications for SunView
dir	dir(5)	format of directories
dump	dump(5)	incremental dump format
dumpdates	dump(5)	incremental dump format
environ	environ(5V)	user environment
ethers	ethers(5)	Ethernet address to hostname database or YP domain
exports	exports(5)	directories to export to NFS clients
fcntl	fcntl(5)	file control options
forward	aliases(5)	addresses and aliases for sendmail(8)
fs	fs(5)	format of a 4.2 (ufs) file system volume
fspec	fspec(5)	format specification in text files
fstab	fstab(5)	static filesystem mounting table, mounted filesystems table
ftpusers	ftpusers(5)	list of users prohibited by ftp
gettytab	gettytab(5)	terminal configuration data base
group.adjunct	group(5)	group security data file
group	group(5)	group file
help	help(5)	help file format
help_viewer	help_viewer(5)	help viewer file format
hosts.equiv	hosts(5)	list of trusted hosts
hosts	hosts(5)	host name data base
inetd.conf	inetd.conf(5)	Internet servers database
inode	fs(5)	format of a 4.2 (ufs) file system volume
internat	internat(5)	key mapping table for internationalization
ipalloc.netrange	ipalloc.netrange(5)	range of addresses to allocate
lastlog	utmp(5)	login records
magic	magic(5)	file command's magic number file
mtab	fstab(5)	static filesystem mounting table, mounted filesystems table
mtab	mtab(5)	mounted file system table
netgroup	netgroup(5)	list of network groups
netmasks	netmasks(5)	network mask data base

netrc	netrc(5)	file for ftp(1) remote login data
networks	networks(5)	network name data base
passwd.adjunct	passwd(5)	user security data file
passwd	passwd(5)	password file
phones	phones(5)	remote host phone number data base
plot	plot(5)	graphics interface
policies	policies(5)	network administration policies
printcap	printcap(5)	printer capability data base
protocols	protocols(5)	protocol name data base
publickey	publickey(5)	publickey database
queuedefs	queuedefs(5)	at/batch/cron queue description file
rasterfile	rasterfile(5)	Sun's file format for raster images
remote	remote(5)	remote host description file
rgb	rgb(5)	available colors (by name) for coloredit
rootmenu	rootmenu(5)	root menu specification for SunView
rpc	rpc(5)	rpc program number data base
scsfile	scsfile(5)	format of SCCS file
services	services(5)	Internet services and aliases
sm.bak	sm(5)	in.statd directory and file structures
sm.bak	statmon(5)	statd directories and file structures
sm.state	sm(5)	in.statd directory and file structures
sm	sm(5)	in.statd directory and file structures
sm	statmon(5)	statd directories and file structures
state	statmon(5)	statd directories and file structures
sunview	sunview(5)	initialization file for SunView
syslog.conf	syslog.conf(5)	configuration file for syslogd system log daemon
tar	tar(5)	tape archive file format
term	term(5)	terminal driving tables for nroff
term	term(5V)	format of compiled term file
termcap	termcap(5)	terminal capability data base
terminfo	terminfo(5V)	terminal capability data base
textswrc	textswrc(5)	initialization file for SunView text windows
toc	toc(5)	table of contents of optional clusters
translate	translate(5)	input and output files for system message translation
ttys	ttys(5)	terminal initialization data
ttytab	ttytab(5)	terminal initialization data
ttytype	ttytype(5)	data base of terminal types by port
types	types(5)	primitive system data types
tzfile	tzfile(5)	time zone information
updaters	updaters(5)	configuration file for YP updating
utmp	utmp(5)	login records
uuencode	uuencode(5)	format of an encoded uuencode file
vfont	vfont(5)	font formats
vgrindefs	vgrindefs(5)	vgrind's language definition data base
wtmp	utmp(5)	login records
xtab	exports(5)	directories to export to NFS clients
ypfiles	ypfiles(5)	the Yellow Pages database and directory structure

NAME

a.out – assembler and link editor output format

SYNOPSIS

```
#include <a.out.h>
#include <stab.h>
#include <nlist.h>
```

AVAILABILITY

Sun-2, Sun-3, and Sun-4 systems only. For Sun386i systems refer to `coff(5)`.

DESCRIPTION

`a.out` is the output format of the assembler `as(1)` and the link editor `ld(1)`. The link editor makes `a.out` executable files.

A file in `a.out` format consists of: a header, the program text, program data, text and data relocation information, a symbol table, and a string table (in that order). In the header, the sizes of each section are given in bytes. The last three sections may be absent if the program was loaded with the `-s` option of `ld` or if the symbols and relocation have been removed by `strip(1)`.

The machine type in the header indicates the type of hardware on which the object code can be executed. Sun-2 code runs on Sun-3 systems, but not vice versa. Program files predating release 3.0 are recognized by a machine type of '0'. Sun-4 code may not be run on Sun-2 or Sun-3, nor vice versa.

Header

The header consists of a `exec` structure. The `exec` structure has the form:

```
struct exec {
    unsigned char  a_dynamic; /* has a __DYNAMIC */
    unsigned char  a_toolversion; /* version of toolset used to create this file */
    unsigned char  a_machtype; /* machine type */
    unsigned short a_magic; /* magic number */
    unsigned long  a_text; /* size of text segment */
    unsigned long  a_data; /* size of initialized data */
    unsigned long  a_bss; /* size of uninitialized data */
    unsigned long  a_syms; /* size of symbol table */
    unsigned long  a_entry; /* entry point */
    unsigned long  a_trsize; /* size of text relocation */
    unsigned long  a_drsize; /* size of data relocation */
};
```

The members of the structure are:

a_dynamic	1 if the <code>a.out</code> file is dynamically linked or is a shared object, 0 otherwise.								
a_toolversion	The version number of the toolset (<code>as</code> , <code>ld</code> , etc.) used to create the file.								
a_machtype	One of the following: <table> <tr> <td>0</td> <td>pre-3.0 executable image</td> </tr> <tr> <td>M_68010</td> <td>executable image using only MC68010 instructions that can run on a Sun-2 or Sun-3</td> </tr> <tr> <td>M_68020</td> <td>executable image using MC68020 instructions that can run only on a Sun-3</td> </tr> <tr> <td>M_SPARC</td> <td>executable image using SPARC instructions that can run only on a Sun-4</td> </tr> </table>	0	pre-3.0 executable image	M_68010	executable image using only MC68010 instructions that can run on a Sun-2 or Sun-3	M_68020	executable image using MC68020 instructions that can run only on a Sun-3	M_SPARC	executable image using SPARC instructions that can run only on a Sun-4
0	pre-3.0 executable image								
M_68010	executable image using only MC68010 instructions that can run on a Sun-2 or Sun-3								
M_68020	executable image using MC68020 instructions that can run only on a Sun-3								
M_SPARC	executable image using SPARC instructions that can run only on a Sun-4								
a_magic	One of the following: <table> <tr> <td>OMAGIC</td> <td>An text executable image which is not to be write-protected, so the data segment is immediately contiguous with the text segment.</td> </tr> </table>	OMAGIC	An text executable image which is not to be write-protected, so the data segment is immediately contiguous with the text segment.						
OMAGIC	An text executable image which is not to be write-protected, so the data segment is immediately contiguous with the text segment.								

- NMAGIC** A write-protected text executable image. The data segment begins at the first segment boundary following the text segment, and the text segment is not writable by the program. When the image is started with `execve(2)`, the entire text and data segments will be read into memory.
- ZMAGIC** A page-aligned text executable image. The data segment begins at the first segment boundary following the text segment, and the text segment is not writable by the program. The text and data sizes are both multiples of the page size, 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 the default format produced by `ld(1)`.

The macro `N_BADMAG` takes an `exec` structure as an argument; it evaluates to 1 if the `a_magic` field of that structure is invalid, and evaluates to 0 if it is valid.

- a_text** The size of the text segment, in bytes.
- a_data** The size of the initialized portion of the data segment, in bytes.
- a_bss** The size of the "uninitialized" portion of the data segment, in bytes. This portion is actually initialized to zero. The zeroes are not stored in the `a.out` file; the data in this portion of the data segment is zeroed out when it is loaded.
- a_syms** The size of the symbol table, in bytes.
- a_entry** The virtual address of the entry point of the program; when the image is started with `execve`, the first instruction executed in the image is at this address.
- a_trsize** The size of the relocation information for the text segment.
- a_drsize** The size of the relocation information for the data segment.

The macros `N_TXTADDR`, `N_DATADDR`, and `N_BSSADDR` give the memory addresses at which the text, data, and bss segments, respectively, will be loaded.

In the `ZMAGIC` format, the size of the header is included in the size of the text section; in other formats, it is not.

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 data), and a stack. For the `ZMAGIC` format, the header is loaded with the text segment; for other formats it is not.

Program execution begins at the address given by the value of the `a_entry` field.

The stack starts at the highest possible location in the memory image, and grows downwards. The stack is automatically extended as required. The data segment is extended as requested by `brk(2)` or `sbrk`.

Text and Data Segments

The text segment begins at the start of 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 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 `N_DATOFF` macro returns the absolute file position of the beginning of the data segment when given an `exec` structure as argument.

Relocation

The relocation information appears after the text and data segments. The `N_TRELOFF` macro returns the absolute file position of the relocation information for the text segment, when given an `exec` structure as argument. The `N_DRELOFF` macro returns the absolute file position of the relocation information for the data segment, when given an `exec` structure as argument. There is no relocation information if `a_trsize+a_drsize==0`.

Relocation (Sun-2 and Sun-3 Systems)

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 is added to the bytes in the file. If a byte involves a reference to a relative location, or relocatable segment, then the value stored in the file is an offset from the associated segment.

If relocation information is present, it amounts to eight bytes per relocatable datum as in the following structure:

```
struct reloc_info_68k {
    long    r_address;          /* address which is relocated */
    unsigned int 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 */
    r_baserel:1,               /* linkage table relative */
    r_jmptable:1,              /* pc-relative to jump table */
    r_relative:1,              /* relative relocation */
    :1;
};
```

If `r_extern` is 0, then `r_symbolnum` is actually an `n_type` for the relocation (for instance, `N_TEXT` meaning relative to segment text origin.)

Relocation (Sun-4 System)

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 ignored. Unlike the Sun-2 and Sun-3 system, the offset from the associated symbol is kept with the relocation record. When the file is processed by the link editor and the external symbol becomes defined, the value of the symbol is added to this offset, and the sum is inserted into the bytes in the text or data segment.

If relocation information is present, it amounts to twelve bytes per relocatable datum as in the following structure:

```
enum reloc_type
{
    RELOC_8,          RELOC_16,          RELOC_32,          /* simplest relocs */
    RELOC_DISP8,     RELOC_DISP16,     RELOC_DISP32,     /* Disp's (pc-rel) */
    RELOC_WDISP30,   RELOC_WDISP22,   /* SR word disp's */
    RELOC_HI22,     RELOC_22,        /* SR 22-bit relocs */
    RELOC_13,       RELOC_LO10,      /* SR 13&10-bit relocs */
    RELOC_SFA_BASE, RELOC_SFA_OFF13, /* SR S.F.A. relocs */
    RELOC_BASE10,   RELOC_BASE13,    RELOC_BASE22,     /* base_relative pic */
    RELOC_PC10,     RELOC_PC22,      /* special pc-rel pic */
    RELOC_JMP_TBL,  /* jmp_tbl_rel in pic */
    RELOC_SEGOFF16, /* ShLib offset-in-seg */
    RELOC_GLOB_DAT, RELOC_JMP_SLOT, RELOC_RELATIVE,  /* rtd relocs */
};
```

```
struct reloc_info_sparc /* used when header.a_machtype == M_SPARC */
{
    unsigned long int r_address;          /* relocation addr (offset in segment) */
    unsigned int     r_index   :24;      /* segment index or symbol index */
    unsigned int     r_extern  : 1;      /* if F, r_index==SEG#; if T, SYM idx */
    int              : 2;               /* <unused> */
    enum reloc_type  r_type    : 5;      /* type of relocation to perform */
    long int         r_addend;          /* addend for relocation value */
};
```

If `r_extern` is 0, then `r_symbolnum` is actually a `n_type` for the relocation (for instance, `N_TEXT` meaning relative to segment text origin.)

Symbol Table

The `N_SYMOFF` macro returns the absolute file position of the symbol table when given an `exec` structure as argument. Within this symbol table, distinct symbols point to disjoint areas in the string table (even when two symbols have the same name). The string table immediately follows the symbol table; the `N_STROFF` macro returns the absolute file position of the string table when given an `exec` structure as argument. 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; thus, the minimum string table size is 4. Layout information as given in the include file for the Sun system is shown below.

The layout of a symbol table entry and the principal flag values that distinguish symbol types are given in the include file as follows:

```

struct nlist {
    union {
        char    *n_name;           /* for use when in-memory */
        long    n_strx;           /* index into file string table */
    } n_un;
    unsigned char n_type;        /* type flag, that is, N_TEXT etc; see below */
    char        n_other;
    short       n_desc;          /* see <stab.h> */
    unsigned    n_value;         /* value of this symbol (or add 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 */

```

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. Because of the union in the `nlist` declaration, it is impossible in C to statically initialize such a structure. If this must be done (as when using `nlist(3)`) the file `<nlist.h>` should be included, rather than `<a.out.h>`; this contains the declaration without the union.

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.

FILES

```

/usr/include/a.out.h    symbolic link to /usr/include/machine/a.out.h
/usr/include/machine    symbolic link to one of /usr/include/sun[234...]

```

/usr/include/sun2/a.out.h Sun-2 a.out header
/usr/include/sun3/a.out.h Sun-3 a.out header
/usr/include/sun4/a.out.h Sun-4 a.out header

SEE ALSO

coff(5), adb(1), as(1), cc(1V), dbx(1), ld(1), nm(1), strip(1), brk(2), nlist(3)

NAME

acct – execution accounting file

SYNOPSIS

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

DESCRIPTION

The acct(2) system call makes entries in an accounting file for each process that terminates. The accounting file is a sequence of entries whose layout, as defined by the include file is:

```
/*      @(#)acct.h 2.7 87/03/12 SMI; from UCB 7.1 6/4/86*/

/*
 * Copyright (c) 1982, 1986 Regents of the University of California.
 * All rights reserved. The Berkeley software License Agreement
 * specifies the terms and conditions for redistribution.
 */

/*
 * Accounting structures;
 * these use a comp_t type which is a 3 bits base 8
 * exponent, 13 bit fraction "floating point" number.
 * Units are 1/AHZ seconds.
 */
typedef u_short comp_t;

struct acct
{
    char      ac_comm[10]; /* Accounting command name */
    comp_t    ac_untime;   /* Accounting user time */
    comp_t    ac_sstime;   /* Accounting system time */
    comp_t    ac_etime;    /* Accounting elapsed time */
    time_t    ac_btime;    /* Beginning time */
    uid_t     ac_uid;      /* Accounting user ID */
    gid_t     ac_gid;      /* Accounting group ID */
    short     ac_mem;      /* average memory usage */
    comp_t    ac_io;       /* number of disk IO blocks */
    dev_t     ac_tty;      /* control typewriter */
    char      ac_flag;     /* Accounting flag */
};

#define AFORK      0001      /* has executed fork, but no exec */
#define ASU        0002      /* used super-user privileges */
#define ACOMPAT    0004      /* used compatibility mode */
#define ACORE      0010      /* dumped core */
#define AXSIG      0020      /* killed by a signal */

/*
 * 1/AHZ is the granularity of the data encoded in the various
 * comp_t fields. This is not necessarily equal to hz.
 */
#define AHZ 64

#ifdef KERNEL
#ifdef SYSACCT
```

```
struct acct      acctbuf;
struct vnode     *acctp;
#else
#define acct()
#endif
#endif
```

If the process does an `execve(2)`, the first 10 characters of the filename appear in `ac_comm`. The accounting flag contains bits indicating whether `execve(2)` was ever accomplished, and whether the process ever had super-user privileges.

SEE ALSO

`acct(2)`, `execve(2)`, `sa(8)`

NAME

aliases, addresses, forward – addresses and aliases for sendmail(8)

SYNOPSIS

```
/etc/passwd
/etc/aliases
/etc/aliases.dir
/etc/aliases.pag
~/forward
```

DESCRIPTION

These files contain mail addresses or aliases, recognized by sendmail(8), for the local host:

<code>/etc/passwd</code>	Mail addresses (usernames) of local users.
<code>/etc/aliases</code>	Aliases for the local host, in ASCII format. This file can be edited to add, update, or delete local mail aliases.
<code>/etc/aliases.{dir,pag}</code>	The aliasing information from <code>/etc/aliases</code> , in binary, dbm(3X) format for use by sendmail(8). The program newaliases(8), which is invoked automatically by sendmail(8), maintains these files.
<code>~/forward</code>	Addresses to which a user's mail is forwarded (see Automatic Forwarding, below).

In addition, the Yellow Pages aliases map `mail.aliases` contains addresses and aliases available for use across the network.

ADDRESSES

As distributed, sendmail(8) supports the following types of addresses:

- Local usernames. These are listed in the local host's `/etc/passwd` file.
- Local filenames. When mailed to an absolute pathname, a message can be appended to a file.
- Commands. If the first character of the address is a vertical bar, (`|`), sendmail(8) pipes the message to the standard input of the command the bar precedes.
- DARPA-standard mail addresses of the form:

name@domain

If *domain* does not contain any '.' (dots), then it is interpreted as the name of a host in the current domain. Otherwise, the message is passed to a *mailhost* that determines how to get to the specified domain. Domains are divided into subdomains separated by dots, with the top-level domain on the right. Top-level domains include:

.COM	Commerical organizations.
.EDU	Educational organizations.
.GOV	Government organizations.
.MIL	Military organizations.

For example, the full address of John Smith could be:

js@jsmachine.Podunk-U.EDU

if he uses the machine named "jsmachine" at Podunk University.

- uucp(1C) addresses of the form:

... [host!]host!username

These are sometimes mistakenly referred to as "Usenet" addresses. uucp(1C) provides links to numerous sites throughout the world for the remote copying of files.

Other site-specific forms of addressing can be added by customizing the sendmail configuration file. See the sendmail(8), and *System and Network Administration* for details. Standard addresses are recommended.

ALIASES**Local Aliases**

/etc/aliases is formatted as a series of lines of the form

name: address[, address]

name is the name of the alias or alias group, and *address* is the address of a recipient in the group. Aliases can be nested. That is, an *address* can be the name of another alias group. Because of the way `sendmail` performs mapping from upper-case to lower-case, an *address* that is the name of another alias group must not contain any upper-case letters.

Lines beginning with white space are treated as continuation lines for the preceding alias. Lines beginning with `#` are comments.

Special Aliases

An alias of the form:

owner-aliasname: address

directs error-messages resulting from mail to *alias-name* to *address*, instead of back to the person who sent the message.

An alias of the form:

aliasname: :include:pathname

with colons as shown, adds the recipients listed in the file *pathname* to the *aliasname* alias. This allows a private list to be maintained separately from the aliases file.

YP Domain Aliases

Normally, the aliases file on the master YP server is used for the *mail.aliases* YP map, which can be made available to every YP client. Thus, the */etc/aliases** files on the various hosts in a network will one day be obsolete. Domain-wide aliases should ultimately be resolved into usernames on specific hosts. For example, if the following were in the domain-wide alias file:

jsmith:js@jsmachine

then any YP client could just mail to *jsmith* and not have to remember the machine and user name for John Smith. If a `.SM` YP alias does not resolve to an address with a specific host, then the name of the YP domain is used. There should be an alias of the domain name for a host in this case. For example, the alias:

jsmith:root

sends mail on a YP client to *root@podunk-u* if the name of the YP domain is *podunk-u*.

Automatic Forwarding

When an alias (or address) is resolved to the name of a user on the local host, `sendmail` checks for a `.forward` file, owned by the intended recipient, in that user's home directory, and with universal read access. This file can contain one or more addresses or aliases as described above, each of which is sent a copy of the user's mail.

Care must be taken to avoid creating addressing loops in the `.forward` file. When forwarding mail between machines, be sure that the destination machine does not return the mail to the sender through the operation of any YP aliases. Otherwise, copies of the message may "bounce." Usually, the solution is to change the YP alias to direct mail to the proper destination.

A backslash before a username inhibits further aliasing. For instance, to invoke the `vacation(1)` program, user `js` creates a `.forward` file that contains the line:

```
\js, "/usr/ucb/vacation js"
```

so that one copy of the message is sent to the user, and another is piped into the `vacation(1)` program.

FILES

`/etc/passwd`

`/etc/aliases`

`~/forward`

SEE ALSO

`uucp(1C)`, `vacation(1)`, `dbm(3X)`, `newaliases(8)`, `sendmail(8)`,

System and Network Administration

BUGS

Because of restrictions in `dbm(3X)` a single alias cannot contain more than about 1000 characters. Nested aliases can be used to circumvent this limit.

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*(1).

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 "!<arch>\n"
#define SARMAG 8

#define ARFMAG "\n"

struct ar_hdr {
    char    ar_name[16];
    char    ar_date[12];
    char    ar_uid[6];
    char    ar_gid[6];
    char    ar_mode[8];
    char    ar_size[10];
    char    ar_fmag[2];
};
```

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 an even (0 mod 2) boundary; a NEWLINE 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.

The encoding of the header is portable across machines. If an archive contains printable files, the archive itself is printable.

Sun386i DESCRIPTION

The file produced by *ar* on Sun386i systems is identical to that described above with the following changes:

Each archive containing COFF files [see *coff*(5)] includes an archive symbol table. This symbol table is used by the link editor *ld* to determine which archive members must be loaded during the link edit process. The archive symbol table (if it exists) is always the first file in the archive (but is never listed) and is automatically created and/or updated by *ar*.

The *ar_name* field of the *ar_hdr* structure described above is blank-padded and slash (/) terminated. Common format archives can be moved from system to system as long as the portable archive command *ar* is used. Conversion tools such as *convert* exist to aid in the transportation of non-common format archives to this format.

Each archive file member begins on an even byte boundary; a NEWLINE is inserted between files if necessary. Nevertheless the size given reflects the actual size of the file exclusive of padding.

If the archive symbol table exists, the first file in the archive has a zero length name (i.e., *ar_name*[0] == '\0'). The contents of this file are as follows:

- The number of symbols. Length: 4 bytes.
- The array of offsets into the archive file. Length: 4 bytes * "the number of symbols".

- The name string table. Length: *ar_size* - (4 bytes * ("the number of symbols" + 1)).

The number of symbols and the array of offsets are managed with *sgetl* and *sputl*. The string table contains exactly as many null terminated strings as there are elements in the offsets array. Each offset from the array is associated with the corresponding name from the string table (in order). The names in the string table are all the defined global symbols found in the common object files in the archive. Each offset is the location of the archive header for the associated symbol.

SEE ALSO

ar(1V), *ld*(1), *nm*(1)

Sun386i WARNINGS

strip(1) will remove all archive symbol entries from the header. The archive symbol entries must be restored via the *ts* option of the *ar*(1V) command before the archive can be used with the link editor *ld*(1).

BUGS

Filenames lose trailing blanks. Most software dealing with archives takes even an included blank as a name terminator.

NAME

audit.log – the security audit trail file

SYNOPSIS

```
#include <sys/label.h>
#include <sys/audit.h>
#include <sys/user.h>
```

DESCRIPTION

The `audit.log` file begins with a header record consisting of an `audit_header` structure followed by the previous audit file name. When the audit daemon is started (usually only at boot time), the previous audit file name is NULL.

```
struct audit_header {
    int     ah_magic;      /* magic number */
    time_t  ah_time;      /* the time */
    short   ah_namelen;   /* length of file name */
};
typedef struct audit_header audit_header_t;
```

The file may end with a trailer record consisting of an `audit_trailer` structure followed by the name of the next audit file.

```
struct audit_trailer {
    short   at_record_size; /* size of this */
    short   at_record_type; /* its type, a trailer */
    time_t  at_time;       /* the time */
    short   at_namelen;    /* length of file name */
};
typedef struct audit_trailer audit_trailer_t;
```

The `audit.log` file contains audit records in their raw form. The records are of varying size depending on the record type. Each record has a header which is an `audit_record` structure.

```
struct audit_record {
    short   au_record_size; /* size of this */
    short   au_record_type; /* its type */
    time_t  au_time;       /* the time */
    short   au_uid;        /* real uid */
    short   au_auid;       /* audit uid */
    short   au_euid;       /* effective */
    short   au_gid;        /* real group */
    short   au_pid;        /* effective */
    int     au_errno;      /* error code */
    int     au_return;     /* a return value */
    blabel_t au_label;     /* also ... */
    short   au_param_count; /* # of parameters */
};
typedef struct audit_record audit_record_t;
```

Immediately following the header is a set of two byte integers, the number of which exist for a given record is contained in the `au_param_count` field. These numbers are the lengths of the additional data items. The additional data items follow the list of lengths, the first length describing the first data item. Interpretation of this data is left to the program accessing it.

SEE ALSO

audit(2), getauditfile(2), getuseraudit(2), audit(8),

Security Features Guide

NAME

`audit_control` – control information for system audit daemon

SYNOPSIS

`/etc/security/audit/audit_control`

DESCRIPTION

The `audit_control` file contains audit control information read by `auditd(8)`. Each line consists of a title and a string, separated by a colon. There are no restrictions on the order of lines in the file, although some lines must appear only once. A line beginning with '#' is a comment.

Directory definition lines list the directories to be used when creating audit files, in the order in which they are to be used. The format of a directory line is:

dir: *directory-name*

where *directory-name* is the name of a directory in which to create audit files, with the form:

`/etc/security/audit/server/machine`

where *server* is the name of an audit file system on the machine where this audit directory resides, and *machine* is the name of the local machine, since audit files belonging to different machines are, by convention, stored in separate subdirectories of a single audit directory. The naming convention normally has *server* be the name of a server machine, and all clients mount `/etc/security/audit/server` at the same location in their local file systems. If the same server exports several different file systems for auditing, their *server* names will, of course, be different.

The audit threshold line specifies the percentage of free space that must be present in the file system containing the current audit file. The format of the threshold line is:

minfree: *percentage*

where *percentage* indicates the amount of free space required. If free space falls below this threshold, the audit daemon `auditd(8)` invokes the shell script `/etc/security/audit/audit_warn`. If no threshold is specified, the default is 0%.

The audit flags line specifies the default system audit value. This value is combined with the user audit value read from `/etc/security/passwd.adjunct` to form the process audit state. The user audit value overrides the system audit value. The format of a flags line is:

flags: *audit-flags*

where *audit-flags* specifies which event classes are to be audited. The character string representation of *audit-flags* contains a series of flag names, each one identifying a single audit class, separated by commas. A name preceded by – means that the class should be audited for failure only; successful attempts are not audited. A name preceded by + means that the class should be audited for success only; failing attempts are not audited. Without a prefix, the name indicates that the class is to be audited for both successes and failures. The special string `all` indicates that all events should be audited; `–all` indicates that all failed attempts are to be audited, and `+all` all successful attempts. The prefixes `^`, `^–`, and `^+` turn off flags specified earlier in the string (`^–` and `^+` for failing and successful attempts, `^` for both). They are typically used to reset flags.

The following table lists the audit classes:

short name	long name	short description
dr	<code>data_read</code>	Read of data, open for reading, etc.
dw	<code>data_write</code>	Write or modification of data
dc	<code>data_create</code>	Creation or deletion of any object
da	<code>data_access_change</code>	Change in object access (modes, owner)
lo	<code>login_logout</code>	Login, logout, creation by <code>at(1)</code>
ad	<code>administrative</code>	Normal administrative operation
p0	<code>minor_privilege</code>	Privileged operation
p1	<code>major_privilege</code>	Unusual privileged operation

EXAMPLE

Here is a sample `/etc/security/audit_control` file for the machine `eggplant`:

```
dir: /etc/security/audit/jedgar/eggplant
dir: /etc/security/audit/jedgar.aux/eggplant
#
# Last-ditch audit file system when jedgar fills up.
#
dir: /etc/security/audit/global/eggplant
minfree: 20
flags: lo,p0,p1,ad,-all,^-da
```

This identifies server `jedgar` with two file systems normally used for audit data, another server `global` used only when `jedgar` fills up or breaks, and specifies that the warning script is run when the file systems are 80% filled. It also specifies that all logins, privileged and administrative operations are to be audited (whether or not they succeed), and that failures of all types except failures to access data are to be audited.

FILES

```
/etc/security/audit/audit_control
/etc/security/audit/audit_warn
/etc/security/audit/**/*
/etc/security/passwd_adjunct
```

SEE ALSO

`at(1)`, `audit(2)`, `getfaudflgs(3)`, `audit.log(5)`, `audit(8)`, `auditd(8)`

NAME

audit_data – current information on audit daemon

SYNOPSIS

/etc/security/audit/audit_data

DESCRIPTION

The **audit_data** file contains information about the audit daemon. The file contains the process ID of the audit daemon, and the pathname of the current audit log file. The format of the file is:

<pid>:<pathname>

Where *pid* is the process ID for the audit daemon, and *pathname* is the full pathname for the current audit log file.

EXAMPLE

64:/etc/security/audit/auditserv/auditclient/2df0504

FILES

/etc/security/audit/audit_data

SEE ALSO

audit(2), audit.log(5), audit(8), auditd(8)

NAME

auto.home— automount map for home directories

SYNOPSIS

/etc/auto.home

AVAILABILITY

Sun386i systems only.

DESCRIPTION

auto.home resides in the */etc* directory, and contains **automount(8)** map entries for user's home directories. On Sun386i systems, this file is used to build the **auto.home** Yellow Pages map used by **automount** at system startup and reads the **auto.master** Yellow Pages database, which contains an entry for **auto.home** and */home*. The **auto.home** map contains entries for each username in the YP **passwd** map, and the *hostname:directory* to NFS mount.

References to */home/username* are translated by the automount daemon using the **auto.home** map, and cause the directory specified in the map entry to be **nfs** mounted and that directory returned to the user's program.

User accounts created using **snap(1)** or **logintool(8)** have **passwd(5)** entries where the initial (home) directory name is, in the form */home/username*. **snap** and **logintool** also automatically create the **auto.home** entry for a user account. The format of the entry is described in **automount(8)**. An example entry is:

```
mtravis      system2:/export/home/users/mtravis
```

Thus when the user **mtravis** logs into a Sun386i systems, the automounter automatically mounts his home directory from **system2**. This allows a user to log in to any RR workstation on the network and be automatically placed in his or her home directory.

The convention for the format of home directory names used by **snap** and **logintool** is:

```
/export/home/groupname/username
```

Note that this is a different map and mechanism for home directories than the one that the automount daemon provides with the **-homes** switch. This is because the Sun386i convention for the format of home directory names differs and provides directories that can be used as mount points on a per user and per group basis.

FILES

/etc/auto.home

SEE ALSO

snap(1), **passwd(5)**, **automount(8)**, **logintool(8)**

NAME

auto.vol- automount map for volumes

SYNOPSIS

/etc/auto.vol

AVAILABILITY

Sun386i systems only.

DESCRIPTION

auto.vol resides in the /etc directory, and contains automount(8) map entries for volumes. On Sun386i systems, this file is used to build the auto.vol Yellow Pages map used by automount(8) at system startup. automount reads the auto.master Yellow Pages map, which contains an entry for auto.vol and /vol.

References to /vol/*volume_name* are translated by the automount daemon using the auto.vol map, and cause the directory specified in the map entry to be mounted.

The concept of a volume is that it is a self contained directory hierarchy that can be NFS mounted. It is referenced using a known *volume_name*. The use of an automount map is suggested so that the volume and its contents can be referenced through /vol. This is advantageous because location-transparency (i.e., which host the volume is on) and replication of read-only volumes can be provided using the automount mechanism. The format of the entry is described in automount. An example entry is:

```
archive      system4:/export/archive
```

In the above example, the archive volume is currently on line on system4. Users and programs can reference it via /vol/archive. If for some reason the volume had to be moved to another system, system2 for example, the network or system administrator simply edits the map entry for the archive volume and changes the hostname to system2 and then rebuilds the Yellow Pages maps.

```
archive      system2:/export/archive
```

Users and programs can continue to refer to the archive volume using /vol/archive, unaware that the volume was moved to another system.

FILES

/etc/auto.vol

SEE ALSO

automount(8)

NAME

bar – tape archive file format

AVAILABILITY

Sun386i systems only.

DESCRIPTION

bar(1), (the tape archive command) dumps several files into one, in a medium suitable for transportation. This format is not compatible with the format generated by tar(1).

A “bar 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 that describes the file, followed by zero or more blocks that 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 bar(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, unless the *B* keyletter is used.

The header block looks like:

```
#define TBLOCK      512

union hblock {
    char dummy[TBLOCK];
    struct header {
        char mode[8];
        char uid[8];
        char gid[8];
        char size[12];
        char mtime[12];
        char chksum[8];
        char rdev[8];
        char linkflag;
        char bar_magic[2];
        char volume_num[4];
        char compressed;
        char date[12];
        char start_of_name;
    } dbuf;
};
```

start_of_name is a null-terminated string. *date* is the date of the archive. *bar_magic* is a special number indicating that this is a bar archive. *rdev* is the device type, for files that are devices. 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*, *rdev*, and *mtime*, which do not contain the trailing null. *start_of_name* is the name of the file, as specified on the *bar* command line. Files dumped because they were in a directory that 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 that own the file. *size* is the size of the file in bytes. Links and symbolic links, and special files, 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 that 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, 1 if it is a hard link, 2 if it is a symbolic link, and 3 if it is a special file (device or FIFO). The name linked-to, if any, is in a null-terminated string, following *start_of_name*. 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.

When the H modifier is used with `bar`, an additional header block (one that does not pertain to a particular file) is written to the first block of each volume of the archive. The header ID, as specified on the command line, is copied to `start_of_name`. The size reflects the number of bytes to skip to the start of the first full file (always zero on the first volume).

The encoding of the header is designed to be portable across machines.

SEE ALSO

bar(1)

NAME

bootparams – boot parameter data base

SYNOPSIS

/etc/bootparams

DESCRIPTION

The **bootparams** file contains the list of client entries that diskless clients use for booting. For each diskless client the entry should contain the following information:

name of client

a list of keys, names of servers, and pathnames.

The first item of each entry is the name of the diskless client. The subsequent item is a list of keys, names of servers, and pathnames.

Items are separated by TAB characters.

EXAMPLE

Here is an example of the **/etc/bootparams** taken from a SunOS system.

```
myclient      root=myserver:/nfsroot/myclient \  
              swap=myserver:/nfsswap/myclient \  
              dump=myserver:/nfsdump/myclient
```

FILES

/etc/bootparams

SEE ALSO

bootparamd(8)

NAME

COFF – common assembler and link editor output

SYNOPSIS

```
#include <filehdr.h>
#include <aouthdr.h>
#include <scnhdr.h>
#include <reloc.h>
#include <linenum.h>
#include <storclass.h>
#include <syms.h>
```

AVAILABILITY

Sun386i systems only.

DESCRIPTION

The output from the link editor and the assembler (named `a.out` by default) is in COFF format (Common Object File Format) on the Sun386i system.

A common object file consists of a file header, a system header (if the file is link editor output), a table of section headers, relocation information, (optional) line numbers, a symbol table, and a string table. The order is given below.

```
File header.
UNIX system header.
Section 1 header.
...
Section n header.
Section 1 data.
...
Section n data.
Section 1 relocation.
...
Section n relocation.
Section 1 line numbers.
...
Section n line numbers.
Symbol table.
String table.
```

The last three parts of an object file (line numbers, symbol table and string table) may be missing if the program was linked with the `-s` option of `ld(1)` or if they were removed by `strip(1)`. Also note that the relocation information will be absent after linking unless the `-r` option of `ld(1)` was used. The string table exists only if the symbol table contains symbols with names longer than eight characters.

The sizes of each section (contained in the header, discussed below) are in bytes.

When an `a.out` file is loaded into memory for execution, three logical segments are set up: the text segment, the data segment (initialized data followed by uninitialized, the latter actually being initialized to all 0's), and a stack. The text segment starts at location `0x1000` by default.

The `a.out` file produced by `ld(1)` has the magic number `0413` in the first field of the system header. The headers (file header, system header, and section headers) are loaded at the beginning of the text segment and the text immediately follows the headers in the user address space. The first text address will equal `0x1000` plus the size of the headers, and will vary depending upon the number of section headers in the `a.out` file. In an `a.out` file with three sections (`.text`, `.data`, `.bss`, and `.comment`), the first text address is at `0x000010D0`. The text segment is not writable by the program; if other processes are executing the same `a.out` file, the processes will share a single text segment.

The data segment starts at the next 4K boundary past the last text address. The first data address is determined by the following: If an a.out file were split into 4K chunks, one of the chunks would contain both the end of text and the beginning of data. When the a.out file is loaded into memory for execution, that chunk will appear twice; once at the end of text and once at the beginning of data (with some unused space in between). The duplicated chunk of text that appears at the beginning of data is never executed; it is duplicated so that the operating system may bring in pieces of the file in multiples of the page size without having to realign the beginning of the data section to a page boundary. Therefore the first data address is the sum of the next segment boundary past the end of text plus the remainder of the last text address divided by 4K. If the last text address is a multiple of 4K no duplication is necessary.

On the Sun386i computer the stack begins at location 0xFBFFFFFF and grows toward lower addresses. The stack is automatically extended as required. The data segment is extended only as requested by the brk(2) system call.

For relocatable files the value of a word in the text or data portions that is not a reference to an undefined external symbol is exactly the value that will appear in memory when the file is executed. If a word in the text involves a reference to an undefined external symbol, there will be a relocation entry for the word, the storage class of the symbol-table entry for the symbol will be marked as an "external symbol", and the value and section number of the symbol-table entry will be undefined. When the file is processed by the link editor and the external symbol becomes defined, the value of the symbol will be added to the word in the file.

File Header

The format of the file header is:

```
struct filehdr
{
    unsigned shortf_magic; /* magic number */
    unsigned shortf_nscns; /* number of sections */
    long          f_timdat; /* time and date stamp */
    long          f_symptr; /* file ptr to symtab */
    long          f_nsyms; /* # symtab entries */
    unsigned shortf_opthdr; /* sizeof(opt hdr) */
    unsigned shortf_flags; /* flags */
};
```

SunOS System Header

The format of the system header is:

```
typedef struct aouthdr
{
    short  magic;          /* magic number */
    short  vstamp;        /* version stamp */
    long   tsize;         /* text size in bytes, padded */
    long   dsize;         /* initialized data (.data) */
    long   bsize;         /* uninitialized data (.bss) */
    long   entry;         /* entry point */
    long   text_start;    /* base of text used for this file */
    long   data_start;    /* base of data used for this file */
} AOUTHDR;
```

Section Header

The format of the section header is:

```

struct scnhdr
{
    char        s_name[SYMNMLEN]; /* section name */
    long        s_paddr; /* physical address */
    long        s_vaddr; /* virtual address */
    long        s_size; /* section size */
    long        s_scnptr; /* file ptr to raw data */
    long        s_relptr; /* file ptr to relocation */
    long        s_lnnoptr; /* file ptr to line numbers */
    unsigned shorts_nreloc; /* # reloc entries */
    unsigned shorts_nlnno; /* # line number entries */
    long        s_flags; /* flags */
};

```

Relocation

Object files have one relocation entry for each relocatable reference in the text or data. If relocation information is present, it will be in the following format:

```

struct reloc
{
    long        r_vaddr; /* (virtual) address of reference */
    long        r_symndx; /* index into symbol table */
    ushort     r_type; /* relocation type */
};

```

The start of the relocation information is *s_relptr* from the section header. If there is no relocation information, *s_relptr* is 0.

Line Number

The `cc(1V)` command generates an entry in the object file for each C source line on which a breakpoint is possible (when invoked with the `-g` option. Users can refer to line numbers when using the appropriate debugger, such as `dbx(1)`). The structure of these line number entries appears below.

```

struct lineno
{
    union
    {
        long    l_symndx ;
        long    l_paddr ;
    }
    l_addr ;
    unsigned short l_inno ;
};

```

Numbering starts with one at the top of the source file and increments independent of transition between functions. The initial line number entry for a function has *l_inno* equal to zero, and the symbol table index of the function's entry is in *l_symndx*. Otherwise, *l_inno* is non-zero, and *l_paddr* is the physical address of the code for the referenced line. Thus the overall structure is the following:

<i>l_addr</i>	<i>l_inno</i>
function sytab index	0
physical address	line
physical address	line
...	

```

function sytab index    0
physical address       line
physical address       line
...

```

Symbol Table

The format of each symbol in the symbol table is described by the `syment` structure, shown below. This structure is compatible with System V COFF, but has an added `_n_dbx` structure which is needed by `dbx` (1).

```

#define SYMNMLEN 8
#define FILNMLEN 14
#define DIMNUM 4

struct syment
{
    union /* all ways to get a symbol name */
    {
        char      _n_name[SYMNMLEN]; /* name of symbol */
        struct
        {
            long   _n_zeroes; /* == 0L if in string table */
            long   _n_offset; /* location in string table */
        } _n_n;
        char      * _n_nptr[2]; /* allows overlaying */
        struct
        {
            char    _n_leading_zero; /* null char */
            char    _n_dbx_type; /* stab type */
            short   _n_dbx_desc; /* value of desc field */
            long    _n_stab_ptr; /* table ptr */
        } _n_dbx;
    } _n;
    long          n_value; /* value of symbol */
    short         n_snum; /* section number */
    unsigned short n_type; /* type and derived type */
    char          n_sclass; /* storage class */
    char          n_numaux; /* number of aux entries */
};

#define n_name      _n_n_name
#define n_zeroes    _n_n_n_zeroes
#define n_offset    _n_n_n_offset
#define n_nptr      _n_n_nptr[1]

```

The storage class member (`n_sclass`) is set to one of the constants defined in `<storclass.h>`. Some symbols require more information than a single entry; they are followed by *auxiliary entries* that are the same size as a symbol entry. The format follows.

```

union auxent {
    struct {
        long    x_tagndx;
        union {
            struct {
                unsigned short x_inno;
                unsigned short x_size;
            } x_insz;
            long    x_fsize;
        } x_misc;
        union {
            struct {
                long    x_innoptr;
                long    x_endndx;
            } x_fcn;
            struct {
                unsigned short x_dimen[DIMNUM];
            } x_ary;
        } x_fcrary;
        unsigned short x_tvndx;
    } x_sym;

    struct {
        char    x_fname[FILNMLEN];
    } x_file;

    struct {
        long    x_scrlen;
        unsigned short x_nreloc;
        unsigned short x_nlinno;
    } x_scn;

    struct {
        long    x_tvfill;
        unsigned short x_tvlen;
        unsigned short x_tvrans[2];
    } x_tv;
};

```

Indexes of symbol table entries begin at *zero*. The start of the symbol table is *f_symptr* (from the file header) bytes from the beginning of the file. If the symbol table is stripped, *f_symptr* is 0. The string table (if one exists) begins at *f_symptr + (f_nsyms * SYMESZ)* bytes from the beginning of the file.

SEE ALSO

as(1), cc(1V), ld(1), brk(2), ldfcn(3)

NAME

core – format of memory image file

SYNOPSIS

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

DESCRIPTION

The operating 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). Set-user-ID and set-group-ID programs do not produce core files when they terminate as this would cause a security loophole.

The maximum size of a core file is limited by `setrlimit` (see `getrlimit(2)`). Files which would be larger than the limit are not created.

The core file consists of a `core` structure, as defined in the `<sys/core.h>` file, followed by the data pages and then the stack pages of the process image. The core structure includes the program's header, the size of the text, data, and stack segments, the name of the program and the number of the signal that terminated the process. The program's header is described by the `exec` structure defined in the `<sys/exec.h>` file, except on Sun386i systems.

```
struct core {
    int     c_magic;        /* Corefile magic number */
    int     c_len;         /* Sizeof (struct core) */
    struct  regs c_regs;    /* General purpose registers */
    struct  exec c_outhdr; /* A.out header */
    int     c_signo;       /* Killing signal, if any */
    int     c_tsize;       /* Text size (bytes) */
    int     c_dsize;       /* Data size (bytes) */
    int     c_ssize;       /* Stack size (bytes) */
    char    c_cmdname[CORE_NAMELEN + 1]; /* Command name */
    struct  fpu c_fpu;     /* external FPU state */
    int     c_ucode;       /* Exception no. from u_code */
};
```

The members of the structure are:

<code>c_magic</code>	The magic number <code>CORE_MAGIC</code> , as defined in <code><sys/core.h></code> .
<code>c_len</code>	The length of the <code>core</code> structure in the core file. This need not be equal to the current size of a <code>core</code> structure as defined in <code><sys/core.h></code> , as the core file may have been produced on a different release of the SunOS operating system.
<code>c_regs</code>	The general purpose registers at the time the core file was produced. This structure is machine-dependent.
<code>c_outhdr</code>	The executable image header of the program.
<code>c_signo</code>	The number of the signal that terminated the process; see <code>sigvec(2)</code> .
<code>c_tsize</code>	The size of the text segment of the process at the time the core file was produced.
<code>c_dsize</code>	The size of the data segment of the process at the time the core file was produced. This gives the amount of data space image in the core file.
<code>c_ssize</code>	The size of the stack segment of the process at the time the core file was produced. This gives the amount of stack space image in the core file.
<code>c_cmdname</code>	The first <code>CORE_NAMELEN</code> characters of the last component of the path name of the program.

c_fpu The status of the floating point hardware at the time the core file was produced. This member is not present on Sun-2s.

c_ucose The signal code of the signal that terminated the process, if any. See sigvec(2).

SEE ALSO

adb(1), dbx(1), getrlimit(2), sigvec(2)

NAME

cpio – format of cpio archive

DESCRIPTION

The old format *header* structure, when the `-c` option of `cpio` is not used, is:

```
struct {
    short  h_magic,
           h_dev;
    ushort h_ino,
           h_mode,
           h_uid,
           h_gid;
    short  h_nlink,
           h_rdev,
           h_mtime[2],
           h_namesize,
           h_filesize[2];
    char   h_name[h_namesize rounded to a word];
} Hdr;
```

The byte order here is that of the machine on which the tape was written. If the tape is being read on a machine with a different byte order, you have to use `swab(3)` after reading the header. You can determine what byte order the tape was written with by examining the `h_magic` field; if it is equal to 0143561 (octal), which is the standard magic number 070707 (octal) with the bytes swapped, the tape was written in a byte order opposite to that of the machine on which it is being read. If you are producing a tape to be read on a machine with the opposite byte order to that of the machine on which it is being produced, you can use `swap` before writing the header.

When the `-c` option is used, the *header* information is described by the statement below:

```
sscanf(Chdr, "%6o%6o%6o%6o%6o%6o%6o%6o%6o%6o%11lo%6o%11lo%s",
       &Hdr.h_magic, &Hdr.h_dev, &Hdr.h_ino, &Hdr.h_mode,
       &Hdr.h_uid, &Hdr.h_gid, &Hdr.h_nlink, &Hdr.h_rdev,
       &Hdr.h_mtime, &Hdr.h_namesize, &Hdr.h_filesize, &Hdr.h_name);
```

Longtime and *Longfile* are equivalent to `Hdr.h_mtime` and `Hdr.h_filesize`, respectively. The contents of each file is recorded in an element of the array of varying length structures, *archive*, together with other items describing the file. Every instance of `h_magic` contains the constant 070707 (octal). The items `h_dev` through `h_mtime` have meanings explained in `stat(2)`. The length of the NULL-terminated path name `h_name`, including the NULL byte, is given by `h_namesize`.

The last record of the *archive* always contains the name TRAILER!!!. Special files, directories, and the trailer, are recorded with `h_filesize` equal to zero. Symbolic links are recorded similarly to regular files, with the “contents” of the file being the name of the file the symbolic link points to.

SEE ALSO

`cpio(1)`, `find(1)`, `stat(2)`, `swab(3)`

NAME

crontab – table of times to run periodic jobs

SYNOPSIS

`/var/spool/cron/crontabs/*`

DESCRIPTION

The `cron` utility is a permanent process, started by `/etc/rc.local`. `cron` consults the files in the directory `/var/spool/cron/crontabs` to find out what tasks are to be done, and at what time.

Each line in a `crontab` file consists of six fields, separated by spaces or tabs, as follows:

1. Minutes field, which can have values in the range 0 through 59.
2. Hours field, which can have values in the range 0 through 23.
3. Day of the month, in the range 1 through 31.
4. Month of the year, in the range 1 through 12.
5. Day of the week, in the range 0 through 6. Sunday is day 0 in this scheme of things. For backward compatibility with older systems, Sunday may also be specified as day 7.
6. (The remainder of the line) is the command to be run. A percent character in this field (unless escaped by `\`) is translated to a NEWLINE 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.

Any of fields 1 through 5 can be a list of values separated by commas. A value can either be a number, or a pair of numbers separated by a hyphen, indicating that the job is to be done for all the times in the specified range. If a field is an asterisk character (`*`) it means that the job is done for all possible values of the field.

Note: the specification of days may be made by two fields (day of the month and day of the week). If both are specified as a list of elements, both are adhered to. For example,

```
0 0 1,15 * 1
```

would run a command on the first and fifteenth of each month, as well as on every Monday. To specify days by only one field, the other field should be set to `*` (for example,

```
0 0 * * 1
```

would run a command only on Mondays).

The command is run from your home directory with an `arg0` of `sh`. Users who desire to have their `.profile` executed must explicitly do so in the command. `cron` supplies a default environment for every shell, defining `HOME`, `LOGNAME`, `USER`, `SHELL(=/bin/sh)`, and `PATH(=/usr/ucb:/bin:/usr/bin)`.

NOTE: Users should remember to redirect the standard output and standard error of their commands! If this is not done, any generated output or errors will be mailed to the user.

EXAMPLE

```
0 0 * * * calendar -
15 0 * * * /usr/etc/sa -s >/dev/null
15 4 * * * find /etc/preserve -mtime +7 -a -exec rm -f {} ;
40 4 * * * find / -name '#*' -atime +3 -exec rm -f {} ;
0 0 1-5 * * /usr/local/weekdays
0 0 0,6 * * /usr/local/weekends
```

The `calendar` command runs at minute 0 of hour 0 (midnight) of every day. The `/usr/etc/sa` command runs at 15 minutes after midnight every day. The two `find` commands run at 15 minutes past four and at 40 minutes past four, respectively, every day of the year. The `/usr/local/weekdays` command is run at midnight on weekdays. Finally, the `/usr/local/weekends` command is run at midnight on weekends.

FILES

/var/spool/cron/crontabs/*
tables of times to run periodic jobs
/etc/rc.local
.profile

SEE ALSO

cron(8), rc(8)

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 and directories must be read using the `getdirentries(2)` system call or the `directory(3)` library routines. The fact that a file is a directory is indicated by a bit in the flag word of its inode entry; see `fs(5)`.

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 (512 bytes on most machines):

```
#ifdef KERNEL
#define DIRBLKSIZ DEV_BSIZE
#else
#define DIRBLKSIZ 512
#endif
```

```
#define MAXNAMLEN 255
```

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 that 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.

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 */
};
```

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 (`'/'`), for which `'..'` has the same meaning as `'.'`.

SEE ALSO

getdirentries(2), directory(3), fs(5)

NAME

dump, dumpdates – incremental dump format

SYNOPSIS

```
#include <sys/types.h>
#include <sys/inode.h>
#include <dumprestor.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 `<dumprestor.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_CLRI6
#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 default number of 1024 byte records in a physical tape block, changeable by the **b** option to **dump**. 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:

c_type	The type of the header.
c_date	The date the dump was taken.
c_ddate	The date the file system was dumped from.
c_volume	The current volume number of the dump.
c_tapea	The current number of this (1024-byte) record.
c_inumber	The number of the inode being dumped if this is of type TS_INODE .
c_magic	This contains the value MAGIC above, truncated as needed.
c_checksum	This contains whatever value is needed to make the record sum to CHECKSUM .
c_dinode	This is a copy of the inode as it appears on the file system; see fs(5) .
c_count	The count of characters in <i>c_addr</i> .
c_addr	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:

id_name	The dumped filesystem is <i>/dev/id_nam</i> .
id_incno	The level number of the dump tape; see dump(8) .
id_ddate	The date of the incremental dump in system format see types(5) .

FILES

/etc/dumpdates
/dev/id_nam

SEE ALSO

fs(5), **types(5)**, **dump(8)**, **restore(8)**

BUGS

Should more explicitly describe format of *dumpdates* file.

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 <code>sh(1)</code> , <code>time(1V)</code> , <code>nice(1)</code> , etc., apply in searching for a file known by an incomplete path name. The prefixes are separated by ‘:’. The <code>login(1)</code> process sets <code>PATH=:/usr/ucb:/bin:/usr/bin</code> .
HOME	The name of the user’s login directory, set by <code>login(1)</code> from the password file <code>/etc/passwd</code> (see <code>passwd(5)</code>).
TERM	The type of terminal on which the user is logged in. This information is used by commands, such as <code>nroff(1)</code> or <code>plot(1G)</code> , which may exploit special terminal capabilities. See <code>/etc/termcap</code> (<code>termcap(5)</code>) for a list of terminal types.
SHELL	The path name of the user’s login shell.
TERMCAP	The string describing the terminal in <code>TERM</code> , or the name of the <code>termcap</code> file, see <code>termcap(3X)</code> , <code>termcap(5)</code> .
EXINIT	A startup list of commands read by <code>ex(1)</code> , <code>edit</code> , and <code>vi(1)</code> .
USER	
LOGNAME	The user’s login name.
TZ	The name of the time zone that the user is located in. If <code>TZ</code> is not present in the environment, the system’s default time zone, normally the time zone that the computer is located in, is used.

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`.

SYSTEM V DESCRIPTION

The description of the variable `TERMCAP` does not apply to programs built in the System V environment.

FILES

`/etc/passwd`
`etc/termcap`

SEE ALSO

`csh(1)`, `ex(1)`, `login(1)`, `nice(1)`, `nroff(1)`, `plot(1G)`, `sh(1)`, `time(1V)`, `vi(1)`, `execve(2)`, `getenv(3)`, `system(3)`, `termcap(3X)`, `passwd(5)`, `termcap(5)`

NAME

ethers – Ethernet address to hostname database or YP domain

DESCRIPTION

The **ethers** file contains information regarding the known (48 bit) Ethernet addresses of hosts on the Internet. For each host on an Ethernet, a single line should be present with the following information:

Ethernet address
official host name

Items are separated by any number of blanks and/or TAB characters. A '#' indicates the beginning of a comment extending to the end of line.

The standard form for Ethernet addresses is "x:x:x:x:x:x" where *x* is a hexadecimal number between 0 and ff, representing one byte. The address bytes are always in network order. Host names may contain any printable character other than a SPACE, TAB, NEWLINE, or comment character. It is intended that host names in the **ethers** file correspond to the host names in the **hosts(5)** file.

The **ether_line()** routine from the Ethernet address manipulation library, **ethers(3N)** may be used to scan lines of the **ethers** file.

FILES

/etc/ethers

SEE ALSO

ethers(3N), **hosts(5)**

NAME

exports, xtab – directories to export to NFS clients

SYNOPSIS

/etc/exports

/etc/xtab

DESCRIPTION

The */etc/exports* file contains entries for directories that can be exported to NFS clients. This file is read automatically by the `exportfs(8)` command. If you change this file, you must run `exportfs(8)` for the changes to affect the daemon's operation.

Only when this file is present at boot time does the `rc.local` script execute `exportfs(8)` and start the NFS file-system daemon, `nfsd(8)`.

The */etc/xtab* file contains entries for directories that are *currently* exported. This file should only be accessed by programs using `getexportent` (see `exportent(3)`). (Use the `-u` option of `exportfs` to remove entries from this file).

An entry for a directory consists of a line of the following form:

directory `-option[,option]...`

directory is the pathname of a directory (or file).

option is one of

ro Export the directory read-only. If not specified, the directory is exported read-write.

rw=hostnames[:hostname]...

Export the directory read-mostly. Read-mostly means read-only to most machines, but read-write to those specified. If not specified, the directory is exported read-write to all.

anon=uid

If a request comes from an unknown user, use *uid* as the effective user ID. Note: root users (uid 0) are always considered "unknown" by the NFS server, unless they are included in the "root" option below. The default value for this option is `-2`. Setting "anon" to `-1` disables anonymous access. Note: by default secure NFS will accept insecure requests as anonymous, and those wishing for extra security can disable this feature by setting "anon" to `-1`.

root=hostnames[:hostname]...

Give root access only to the root users from a specified *hostname*. The default is for no hosts to be granted root access.

access=client[:client]...

Give mount access to each *client* listed. A *client* can either be a host-name, or a netgroup (see `netgroup(5)`). Each *client* in the list is first checked for in the netgroup database, and then the hosts database. The default value allows any machine to mount the given directory.

secure Require clients to use a more secure protocol when accessing the directory.

A '#' (pound-sign) anywhere in the file indicates a comment that extends to the end of the line.

EXAMPLE

/usr	-access=clients	# export to my clients
/usr/local	# export to the world	
/usr2	-access=hermes:zip:tutorial	# export to only these machines
/usr/sun	-root=hermes:zip	# give root access only to these
/usr/new	-anon=0	# give all machines root access
/usr/bin	-ro	# export read-only to everyone
/usr/stuff	-access=zip,anon=-3,ro	# several options on one line

FILES

/etc/exports
/etc/xtab
/etc/hosts
/etc/netgroup
rc.local

SEE ALSO

exportent(3), hosts(5), netgroup(5), exportfs(8), nfsd(8)

WARNINGS

You cannot export either a parent directory or a subdirectory of an exported directory that is *within the same filesystem*. It would be illegal, for instance, to export both **/usr** and **/usr/local** if both directories resided on the same disk partition.

NAME

fcntl – file control options

SYNOPSIS

```
#include <fcntl.h>
```

DESCRIPTION

The fcntl(2V) function provides for control over open files. This include file describes *requests* and *arguments* to fcntl and open(2V) as shown below:

```
/*          @ (#)fcntl.h 1.2 83/12/08 SMI; from UCB 4.2 83/09/25*/
/*
 * Flag values accessible to open(2V) and fcntl(2)
 * (The first three can only be set by open)
 */
#define      O_RDONLY          0
#define      O_WRONLY          1
#define      O_RDWR            2
#define      O_NDELAY          FNDELAY      /* Non-blocking I/O */
#define      O_APPEND          FAPPEND      /* append (writes guaranteed at the end) */
#ifndef     F_DUPFD
/* fcntl(2) requests */
#define      F_DUPFD           0           /* Duplicate files */
#define      F_GETFD           1           /* Get fildes flags */
#define      F_SETFD           2           /* Set fildes flags */
#define      F_GETFL           3           /* Get file flags */
#define      F_SETFL           4           /* Set file flags */
#define      F_GETOWN          5           /* Get owner */
#define      F_SETOWN          6           /* Set owner */
/* flags for F_GETFL, F_SETFL— copied from <sys/file.h> */
#define      FNDELAY           00004/* non-blocking reads */
#define      FAPPEND           00010/* append on each write */
#define      FASYNC            00100/* signal pgrp when data ready */
#endif
```

SEE ALSO

fcntl(2V), open(2V)

NAME

fs, inode – format of a 4.2 (ufs) file system volume

SYNOPSIS

```
#include <sys/types.h>
#include <ufs/fs.h>
#include <ufs/inode.h>
```

DESCRIPTION

Standard 4.2 (ufs) file system storage volumes have 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 contain primary and secondary bootstrapping programs.

The actual file system begins at sector 16 with the *super-block*. The layout of the super block is defined by the include file `<ufs/fs.h>`

Each disk drive contains some number of file systems. A file system consists of a number of cylinder groups. Each cylinder group contains 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 cannot 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(8)`.

`fs_minfree` gives the minimum acceptable percentage of file system blocks which may be free. If the free-list 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).

Note: 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 INBSIZE. 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: 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.

Note: 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 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 <ufs/inode.h>.

SEE ALSO

mkfs(8)

NAME

fspec – format specification in text files

DESCRIPTION

It is sometimes convenient to maintain text files on the operating system with non-standard tab stop settings, (that is, tab stops that are not set at every eighth column). Such files must generally be converted to a standard format, frequently by replacing all TAB characters with the appropriate number of SPACE characters, before they can be processed by operating system commands. A format specification occurring in the first line of a text file specifies how TAB characters are to be expanded in the remainder of the file.

A format specification consists of a sequence of parameters separated by blanks and surrounded by the brackets <: and >:. Each parameter consists of a keyletter, possibly followed immediately by a value. The following parameters are recognized:

t tabs The *t* parameter specifies the tab stop settings for the file. The value of *tabs* must be one of the following:

1. A list of column numbers separated by commas, indicating tab stops set at the specified columns;
2. A '-' followed immediately by an integer *n*, indicating tab stops set at intervals of *n* columns, that is, at $1+n$, $1+2*n$, and so on;
3. A '-' followed by the name of a "canned" tab stop specification.

Up to 40 numbers are allowed in a comma-separated list of tab stop settings. If any number (except the first one) is preceded by a plus sign, it is taken as an increment to be added to the previous value. Thus, the formats *t1, 10, 20, 30* and *t1, 10, +10, +10* are considered identical.

Standard tab stops are specified by *t-8*, or equivalently, *t1, 9, 17, 25*, etc. This is the tab stop setting that most operating system utilities assume, and is the most likely setting to be found at a terminal. The specification *t-0* specifies no tab stops at all.

The canned tab stops specifications that are recognized are as follows:

- | | |
|-----------|--|
| a | 1, 10, 16, 36, 72
Assembler, IBM S/370, first format |
| a2 | 1, 10, 16, 40, 72
Assembler, IBM S/370, second format |
| c | 1, 8, 12, 16, 20, 55
COBOL, normal format |
| c2 | 1, 6, 10, 14, 49
COBOL compact format (columns 1-6 omitted). Using this code, the first typed character corresponds to card column 7, one space gets you to column 8, and a TAB reaches column 12. Files using this tab stop setup should include a format specification as follows:
<:t-c2 m6 s66 d:> |
| c3 | 1, 6, 10, 14, 18, 22, 26, 30, 34, 38, 42, 46, 50, 54, 58, 62, 67
COBOL compact format (columns 1-6 omitted), with more tab stops than c2 . This is the recommended format for COBOL. The appropriate format specification is:
<:t-c3 m6 s66 d:> |
| f | 1, 7, 11, 15, 19, 23
FORTRAN |
| p | 1, 5, 9, 13, 17, 21, 25, 29, 33, 37, 41, 45, 49, 53, 57, 61
PL/I |
| s | 1, 10, 55 |

SNOBOL

- u** 1, 12, 20, 44
UNIVAC 1100 Assembler
- s size** The *s* parameter specifies a maximum line size. The value of *size* must be an integer. Size checking is performed after TAB characters have been expanded, but before the margin is prepended.
- m margin**
The *m* parameter specifies a number of SPACE characters to be prepended to each line. The value of *margin* must be an integer.
- d** The *d* parameter takes no value. Its presence indicates that the line containing the format specification is to be deleted from the converted file.
- e** The *e* parameter takes no value. Its presence indicates that the current format is to prevail only until another format specification is encountered in the file.

Default values, which are assumed for parameters not supplied, are *t*-8 and *m*0. If the *s* parameter is not specified, no size checking is performed. If the first line of a file does not contain a format specification, the above defaults are assumed for the entire file. The following is an example of a line containing a format specification:

```
* <:t5,10,15 s72:> *
```

If a format specification can be disguised as a comment, it is not necessary to code the *d* parameter.

SEE ALSO

ed(1), tabs(1)

NAME

fstab, mtab – static filesystem mounting table, mounted filesystems table

SYNOPSIS

/etc/fstab

/etc/mtab

DESCRIPTION

The **/etc/fstab** file contains entries for filesystems and disk partitions to mount using the **mount(8)** command, which is normally invoked by the **rc.boot** script at boot time. This file is used by various utilities that mount, unmount, check the consistency of, dump, and restore file systems. It is also used by the system itself when locating the swap partition.

The **/etc/mtab** file contains entries for filesystems *currently* mounted, and is read by programs using the routines described in **getmntent(3)**. **umount** (see **mount(8)**) removes entries from this file.

Each entry consists of a line of the form:

filesystem directory type options freq pass

filesystem is the pathname of a block-special device, or the name of a remote filesystem in *host:pathname* form.

directory is the pathname of the directory on which to mount the filesystem.

type is the filesystem type, which can be one of:

4.2 to mount a block-special device

nfs to mount an exported NFS filesystem

swap to indicate a swap partition

ignore to have the **mount** command ignore the current entry (good for noting disk partitions that are not being used)

options contains a comma-separated list (no spaces) of mounting options, some of which can be applied to all types of filesystems, and others which only apply to specific types.

4.2 options:

quota | noquota

disk quotas are enforced or not enforced

nfs options:

bg | fg If the first attempt fails, retry in the background, or, in the foreground

retry=*n* The number of times to retry the mount operation.

rsize=*n* Set the read buffer size to *n* bytes.

wsize=*n* Set the write buffer size to *n* bytes.

timeo=*n* Set the NFS timeout to *n* tenths of a second.

retrans=*n* The number of NFS retransmissions.

port=*n* The server IP port number.

soft | hard Return an error if the server does not respond, or continue the retry request until the server responds.

intr Allow keyboard interrupts on hard mounts.

secure Use a more secure protocol for NFS transactions.

acregmin=*n* Hold cached attributes for at least *n* seconds after file modification.

acregmax=*n* Hold cached attributes for no more than *n* seconds after file modification.

acdirmin=*n* Hold cached attributes for at least *n* seconds after directory update.

acdirmax=*n* Hold cached attributes for no more than *n* seconds after directory update.

actimeo=*n* Set *min* and *max* times for regular files and directories to *n* seconds.

Common options:

ro|rw mount either read-only or read-write
suid|nosuid setuid execution allowed or disallowed
grpuid Create files with BSD semantics for propagation of the group ID. With this option, files inherit the group ID of the directory in which they are created, regardless of the directory's setgid bit.
noauto Do not mount this file system automatically (using **mount -a**).

freq is the interval (in days) between dumps.

pass is the **fsck(8)** pass in which to check the partition. Filesystems with the same pass number are checked simultaneously. Filesystems with *pass* equal to 0 are not checked.

A pound-sign (#) as the first non-white character indicates a comment line which is ignored by routines that read this file. The order of records in **/etc/fstab** is important because **fsck**, **mount**, and **umount** process the file sequentially; an entry for a file system must appear *after* the entry for any file system it is to be mounted on top of.

EXAMPLES

In this example, the **/home/user** directory is hard mounted read-write over the NFS, along with additional swap space in the form of a mounted swap file (see *System and Network Administration* for details on adding swap space):

```
/dev/xy0a / 4.2 rw,noquota 1 1
/dev/xy0b /usr 4.2 rw,noquota 1 1
example:/home/user /home/user nfs rw,hard,fg 0 0
/export/swap/myswap swap swap rw 0 0
```

FILES

/etc/fstab
/etc/mstab

SEE ALSO

getmntent(3), **fsck(8)**, **mount(8)**, **quotacheck(8)**, **quotaon(8)**,

NAME

ftusers – list of users prohibited by ftp

SYNOPSIS

/usr/etc/ftusers

DESCRIPTION

ftusers contains a list of users who cannot access this system using the **ftp(1C)** program. **ftusers** contains one user name per line.

SEE ALSO

ftp(1C), **ftpd(8C)**

NAME

gettytab – terminal configuration data base

SYNOPSIS

/etc/gettytab

DESCRIPTION

gettytab is a simplified version of the termcap(5) data base used to describe terminal lines. The initial terminal login process getty(8) accesses the gettytab file each time it starts, allowing simpler reconfiguration of terminal characteristics. Each entry in the data base is used to describe one class of terminals.

There is a default terminal class, **default**, that is used to set global defaults for all other classes. (That is, the default entry is read, then the entry for the class required is used to override particular settings.)

CAPABILITIES

Refer to termcap(5) for a description of the file layout. The *Default* column below lists defaults obtained if there is no entry in the table obtained, nor one in the special default table.

<i>Name</i>	<i>Type</i>	<i>Default</i>	<i>Description</i>
ab	bool	false	read a \r first and guess the baud rate from it
ap	bool	false	terminal uses 7 bits, any parity
bd	num	0	backspace delay
bk	str	0377	alternate end of line character (input break)
cb	bool	false	use crt backspace mode
cd	num	0	carriage-return delay
ce	bool	false	use crt erase algorithm
ck	bool	false	use crt kill algorithm
cl	str	NULL	screen clear sequence
co	bool	false	console - add after login prompt
dx	bool	false	set DECCTLQ
ds	str	^Y	delayed suspend character
ec	bool	false	leave echo OFF
ep	bool	false	terminal uses 7 bits, even parity
er	str	^?	erase character
et	str	^D	end of text (EOF) character
ev	str	NULL	initial environment
f0	num	unused	tty mode flags to write messages
f1	num	unused	tty mode flags to read login name
f2	num	unused	tty mode flags to leave terminal as
fd	num	0	form-feed (vertical motion) delay
fl	str	^O	output flush character
hc	bool	false	do NOT hangup line on last close
he	str	NULL	hostname editing string
hn	str	hostname	hostname
ht	bool	false	terminal has real tabs
ig	bool	false	ignore garbage characters in login name
im	str	NULL	initial (banner) message
in	str	^C	interrupt character
is	num	unused	input speed
kl	str	^U	kill character
lc	bool	false	terminal has lower case
lm	str	login:	login prompt
ln	str	^V	“literal next” character
lo	str	/usr/bin/login	program to exec when name obtained
nd	num	0	newline (line-feed) delay
nl	bool	false	terminal has (or might have) a newline character

nx	str	default	next table (for auto speed selection)
op	bool	false	terminal uses 7 bits, odd parity
os	num	unused	output speed
p8	bool	false	terminal uses 8 bits, no parity
pc	str		pad character
pe	bool	false	use printer (hard copy) erase algorithm
pf	num	0	delay between first prompt and following flush (seconds)
ps	bool	false	line connected to a MICOM port selector
qu	str	^	quit character
rp	str	^R	line retype character
rw	bool	false	do NOT use RAW for input, use CBREAK
sp	num	0	line speed (input and output)
su	str	^Z	suspend character
tc	str	none	table continuation
td	num	0	tab delay
to	num	0	timeout (seconds)
tt	str	NULL	terminal type (for environment)
ub	bool	false	do unbuffered output (of prompts etc)
uc	bool	false	terminal is known upper case only
we	str	^W	word erase character
xc	bool	false	do NOT echo control chars as ^X
xf	str	^S	XOFF (stop output) character
xn	str	^Q	XON (start output) character

If no line speed is specified, speed will not be altered from that which prevails when **getty** is entered. Specifying an input or output speed overrides line speed for stated direction only. If **ab** is specified, **getty** will initially read a character from the tty, assumed to be a carriage return, and will attempt to figure out the baud rate based on what the character appears as. It will then look for a table entry for that baud rate; if the line appears to be a 300 baud line, it will look for an entry **300-baud**, if it appears to be a 1200 baud line, it will look for an entry **1200-baud**, etc..

Terminal modes to be used for the output of the message, for input of the login name, and to leave the terminal set as upon completion, are derived from the Boolean flags specified. If the derivation should prove inadequate, any (or all) of these three may be overridden with one of the **f0**, **f1**, or **f2** numeric specifications, which can be used to specify (usually in octal, with a leading '0') the exact values of the flags. Local (new tty) flags are set in the top 16 bits of this (32 bit) value.

Should **getty** receive a NULL character (presumed to indicate a line break) it will restart using the table indicated by the **nx** entry. If there is none, it will re-use its original table.

Delays are specified in milliseconds, the nearest possible delay available in the tty driver will be used. Should greater certainty be desired, delays with values 0, 1, 2, and 3 are interpreted as choosing that particular delay algorithm from the driver.

The **cl** screen clear string may be preceded by a (decimal) number of milliseconds of delay required (as with **termcap(5)**). This delay is simulated by repeated use of the pad character **pc**.

The initial message, and login message, **im** and **lm** may include the character sequence **%h** or **%t** to obtain the hostname or tty name respectively. (**%%** obtains a single '%' character.) The hostname is normally obtained from the system, but may be set by the **hn** table entry. In either case it may be edited with **he**. The **he** string is a sequence of characters, each character that is neither '@' nor '#' is copied into the final hostname. A '@' in the **he** string, copies one character from the real hostname to the final hostname. A '#' in the **he** string, skips the next character of the real hostname. Surplus '@' and '#' characters are ignored.

When **getty** execs the login process, given in the **lo** string (usually **/usr/bin/login**), it will have set the environment to include the terminal type, as indicated by the **tt** string (if it exists). The **ev** string, can be used to enter additional data into the environment. It is a list of comma separated strings, each of which will presumably be of the form *name=value*.

If a non-zero timeout is specified, with *to*, then *getty* will exit within the indicated number of seconds, either having received a login name and passed control to *login*, or having received an alarm signal, and exited. This may be useful to hangup dial in lines.

Output from *getty* is even parity unless *op* or *p8* is specified. *op* may be specified with *ap* to allow any parity on input, but generate odd parity output. Note: this only applies while *getty* is being run, terminal driver limitations prevent a more complete implementation. *getty* does not check parity of input characters in RAW mode.

FILES

/etc/gettytab

SEE ALSO

termcap(5), *getty(8)*

NAME

group – group file

SYNOPSIS*/etc/group***DESCRIPTION**

The **group** file contains a one-line entry for each group recognized by the system, of the form:

```
groupname:password:gid:user-list
```

where:

<i>groupname</i>	is the name of the group.
<i>gid</i>	is the group's numerical ID within the system; it must be unique.
<i>user-list</i>	is a comma-separated list of users allowed in the group.

If the password field is empty, no password is demanded. The **group** file is an ASCII file. Because of the encrypted passwords, the **group** file can and does have general read permission, and can be used as a mapping of numerical group IDs to user names.

A group entry beginning with a '+' (plus sign), means to incorporate an entry or entries from the Yellow Pages. A '+' on a line by itself means to insert the entire contents of the Yellow Pages group file at that point in the file. An entry of the form: '+*groupname*' means to insert the entry (if any) for *groupname*. If a '+' entry has a non-empty *password* or *user-list* field, the contents of that field override the corresponding field from the Yellow Pages. The *gid* field cannot be overridden in this way.

An entry of the form: -*groupname* indicates that the group is disallowed. All subsequent entries for the indicated *groupname*, whether originating from the Yellow Pages, or the local **group** file, are ignored.

Malformed entries cause routines that read this file to halt, in which case group assignments specified further along are never made. To prevent this from happening, use **grpck(8)** to check the */etc/group* database from time to time.

The Sun386i system uses the following group IDs as program privileges:

operator – 5	Privilege to do backup as root.
accounts – 11	Privilege to update user accounts.
networks – 12	Privilege to change network configuration.
devices – 13	Privilege to modify printer, terminal, or modem configurations.

On all Sun systems, SunOS uses group ID 0 as privilege to run **su(1)**.

EXAMPLE

Here is a sample group file when the **group.adjunct** file does not exist:

```
primary:q.mJzTnu8icF.:10:fred,mary
+myproject:::bill,steve
+:
```

Here is a sample group file when the **group.adjunct** file does exist:

```
primary:#$primary:10:fred,mary
+myproject:::bill,steve
+:
```

If these entries appear at the end of a group file, then the group *primary* will have members **fred** and **mary**, and a group ID of **10**. The group *myproject* will have members **bill** and **steve**, and the password and group ID of the Yellow Pages entry for the group **myproject**. All groups listed in the Yellow Pages are pulled in and placed after the entry for **myproject**.

FILES

/etc/group

SEE ALSO

passwd(1), su(1), getgroups(2), initgroups(3), crypt(3), group.adjunct(5), passwd(5), grpck(8)

BUGS

The **passwd(1)** command will not change group passwords.

NAME

group.adjunct – group security data file

SYNOPSIS

/etc/security/group.adjunct

DESCRIPTION

The **group.adjunct** file contains the following information for each group:

groupname This is the group's name in the system; it must be unique.

password The encrypted password, formerly field two of the */etc/group* file.

The **group.adjunct** file is in ASCII. Fields are separated by a colon, and each group is separated from the next by a NEWLINE.

A **group.adjunct** file can have a line beginning with a '+' (plus sign), which means to incorporate entries from the Yellow Pages. There are two styles of '+' entries: all by itself, '+' means to insert the entire contents of the **group.adjunct** Yellow Pages file at that point; *+name* means to insert the entry (if any) for *name* from the Yellow Pages at that point. If a '+' entry has a non-NULL password, the contents of that field will override what is contained in the Yellow Pages.

FILES

/etc/group

SEE ALSO

crypt(3), getgraent(3), getgrent(3), group(5)

NAME

help – help file format

SYNOPSIS

/usr/lib/help/*

AVAILABILITY

Sun386i systems only.

DESCRIPTION

Each SunView application using the **help** feature has a simple ASCII file in **/usr/lib/help** with the name *application-name.info*.

This file contains the text of help messages for each SunView object within that program. Each help message is separated in the file by a line beginning with a colon and identified by a keyword which matches the **HELP_DATA** attribute of the SunView object.

The first character of each line in the file may be:

#	comment line
:	keyword line
any other	1-32 help text lines

If the line is a keyword line, it has the following structure—

:keyword[s]:datastring [pagenumber]<cr>

keyword is a 1-65 character keyword
 --any displayable characters may be used
 --several keywords may be present
 --keywords are separated by 1-or-more blanks

datastring is 1-256 ASCII bytes, and describes the path of the data files for **help_viewer**, relative to **/usr/lib/help**.

pagenumber is an optional page number within the **help_viewer** data file.

The help text which follows the **:keyword** line will be displayed in an Alert Box when help is requested for one of the keywords by pressing the help key.

The **datastring** will be sent (by RPC) to the **help_viewer** procedure when the user selects the More Help box in the Alert Box window.

EXAMPLE

Here is part of a typical help file, called **mailtool.info**.

:abort

Abort button

o Quits the Mail application (click left on button). Tentative message deletions do not become permanent.

o Provides a menu of Abort options (click right on button).

:cancel:mailtool/Writing_and_Sending_Mail 1

Cancel button

- o Closes the message composition window without sending message (click left on button).**
- o Provides a menu of Cancel options (click right on button).**

Pressing the help key while in the cancel or abort buttons triggers the display of the corresponding text. The words *cancel* and *abort* in this file are the keywords. In the case of abort, there is no More Help available. For cancel, More Help is available and it is stored in the first page of the *Writing_and_Sending_Mail* file in the mailtool directory.

FILES

*/usr/lib/help/** files for the pop-up help facility

SEE ALSO

help_viewer(1), *help_viewer(5)*

Sun386i Developer's Guide

NAME

help_viewer – help viewer file format

SYNOPSIS

/usr/lib/help/**

AVAILABILITY

Sun386i systems only.

DESCRIPTION

The **help_viewer** reads files of various types. The Top Level list of applications documented is **/usr/lib/help/Top_Level**. The Master Index shown at the top level is **/usr/lib/help/Master_Index**. These files are FrameMaker files. To add or remove a heading from this list, use FrameMaker (1.1 or later).

Each directory within **/usr/lib/help** that corresponds to a SunView application name contains detailed information about that application. These are also FrameMaker files. The ***.rf** files are rasterfiles, of standard image format created by FrameMaker. These are the pictures that are interleaved into the text.

The **Frame/** subdirectory of **/usr/lib/help** contains topic, contents, and index templates which can be used to create new Help Viewer handbooks. The **Interleaf/** subdirectory contains Interleaf templates, fonts, and initialization files.

FILES

/usr/lib/help/**

SEE ALSO

help(5), help_viewer(1)

NAME

hosts – host name data base

SYNOPSIS

/etc/hosts

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:

Internet address
official host name
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.

EXAMPLE

Here is a typical line from the /etc/hosts file:

```
192.9.1.20    gaia          # John Smith
```

FILES

/etc/hosts

SEE ALSO

`gethostent(3N)`, `inet(3N)`

NAME

hosts.equiv, rhosts – trusted hosts by system and by user

DESCRIPTION

The `/etc/hosts.equiv` file contains a list of trusted hosts. When an `rlogin(1C)` or `rsh(1C)` request is received from a host listed in this file, and when the user making the request is listed in the `/etc/passwd` file, then the remote login is allowed with no further checking. In this case, `rlogin` does not prompt for a password, and commands submitted through `rsh` are executed. Thus, a remote user with a local user ID is said to have “equivalent” access from a remote host named in this file.

The format of the `hosts.equiv` file consists of a one-line entry for each host, of the form:

hostname [*username*]

The *hostname* field normally contains the name of a trusted host from which a remote login can be made. However, an entry consisting of a single ‘+’ indicates that all known hosts are to be trusted. A hostname must be the “official” name as listed in the `hosts(5)` database. This is the first name given in the hosts database entry; hostname aliases are not recognized. Remote login access can also be given or denied for all hosts within a specific network group. An entry of the form:

`+@group`

means that all hosts in the named network group are trusted. An entry of the form:

`-@group`

means that all hosts in the group are not trusted; remote login access is denied to hosts in that group, except when an entry for a specific host appears ahead of the “minus” group entry.

The *username* field can be used to specify a user who is allowed to log in under any valid user ID. Careful thought about security should be given before providing this privilege to a user. You can also specify a network group in the *username* field with an entry of the form:

`+@group1 +@group2`

in which case any user in *group2* logging in from a host in *group1* may log in as anyone. Again, security is an important consideration here.

The User's .rhosts File

Whenever a remote login is attempted, the remote login daemon checks for a `.rhosts` file in the home directory of the user attempting to log in. A user's `.rhosts` file has the same format as the `hosts.equiv` file, and is used to give or deny access only for the *specific user* attempting to log in from a given host. While an entry in the `hosts.equiv` file allows remote login access to *any* user from the indicated host, an entry in a user's `.rhosts` file only allows access from a named host to the user in whose home directory the `.rhosts` file appears. (When this file is used, permissions in the user's home directory should allow read and search access by anyone, so it may be located and read.) When a user attempts a remote login, his `.rhosts` file is, in effect, prepended to the `hosts.equiv` file for permission checking. Thus, if a host is specified in the user's `.rhosts` file, login access is allowed, even if it would otherwise be excluded by a minus group entry in `/etc/hosts.equiv`.

The Root .rhosts File

When the user attempting a remote login is root, only the `./rhosts` file is checked, not `/etc/hosts.equiv`.

FILES

`/etc/hosts.equiv`
`/etc/passwd`
`~/.rhosts`
`/etc`

SEE ALSO

`rlogin(1C)`, `rsh(1C)`, `hosts(5)`, `netgroup(5)`, `passwd(5)`

NAME

indent.pro – default options for indent

DESCRIPTION

The **.indent.pro** file in either the current or home directory contains default command line options for the **indent(1)** program. It is a text file that contains space-separated command line options. For a description of these options, see **indent(1)**.

Explicit command line options override options taken from **.indent.pro**.

Here is a sample **.indent.pro** file:

```
-bap -nbad -nbbb -bc -br -cdb -nce  
-fc1 -ip -lp -npcs -psl -sc -nsob -cli0  
-di12 -l79 -i4 -d0 -c33
```

FILES

./indent.pro
~/indent.pro

SEE ALSO

indent(1)

NAME

inetd.conf – Internet servers database

DESCRIPTION

The *inetd.conf* file contains the list of servers that *inetd*(8C) invokes when it receives an Internet request over a socket. Each server entry is composed of a single line of the form:

service-name socket-type protocol wait-status uid server-program server-arguments

Fields can be separated by either spaces or TAB characters. A '#' (pound-sign) indicates the beginning of a comment; characters up to the end of the line are not interpreted by routines that search this file.

service-name is the name of a valid service listed in the file */etc/services*. For RPC services, the value of the *service-name* field consists of the RPC service name, followed by a slash and either a version number or a range of version numbers (for example, *mountd/1*).

socket-type can be one of:

stream	for a stream socket,
dgram	for a datagram socket,
raw	for a raw socket,
rdm	for a "reliably delivered message" socket, or
seqpacket	for a sequenced packet socket.

protocol must be a recognized protocol listed in the file */etc/protocols*. For RPC services, the field consists of the string "rpc" followed by a slash and the name of the protocol (for example, *rpc/udp* for an RPC service using the UDP protocol as a transport mechanism).

wait-status is **nowait** for all but "single-threaded" datagram servers — servers which do not release the socket until a timeout occurs (such as *comsat*(8C) and *talkd*(8C)). These must have the status **wait**. Although *tftpd*(8C) establishes separate "pseudo-connections", its forking behavior can lead to a race condition unless it is also given the status **wait**.

uid is the user ID under which the server should run. This allows servers to run with access privileges other than those for root.

server-program is either the pathname of a server program to be invoked by *inetd* to perform the requested service, or the value **internal** if *inetd* itself provides the service.

server-arguments If a server must be invoked with command-line arguments, the entire command line (including argument 0) must appear in this field (which consists of all remaining words in the entry). If the server expects *inetd* to pass it the address of its peer (for compatibility with 4.2BSD executable daemons), then the first argument to the command should be specified as '%A'.

FILES

/etc/inetd.conf
/etc/services
/etc/protocols

SEE ALSO

services(5), *comsat*(8C), *inetd*(8C), *talkd*(8C), *tftpd*(8C)

NAME

internat – key mapping table for internationalization

AVAILABILITY

Sun386i systems only.

DESCRIPTION

This file format is used for the file specified by the `-f` flag of `oldsetkeys(1)`.

The file has three columns. First column is keytable identifier, one of: `BASE`, `CTRL`, `SHIFT`, `CAPS`, `UP`, `BASE_ISO`, `SHIFT_ISO` or `ALTG`. The second column is a decimal keystation number. The third column is hexadecimal keytable entry value. The file must end with line of "END, 0, 0". As usual, comment lines start with #.

EXAMPLES

This is the file for mapping keys to Canadian standards:

```
# /usr/lib/.setkeys: Key remapping, used by "setkeys remap"
#
# First column is keytable identifier:
#           BASE, CTRL, SHIFT, CAPS, UP, BASE_ISO, SHIFT_ISO or ALTG
# Second column is decimal keystation number
# Third column is hexadecimal keytable entry value
# File must end with line of "END, 0, 0"
# Comment lines must start with #
#
#
# --- Keymaps for Canadian keyboard ---
# > Define Alt Graph key (SHIFTKEYS+ALTGRAPH=86)
BASE 119 86
CTRL 119 86
SHIFT 119 86
CAPS 119 86
UP 119 86
# > Define Caps key (SHIFTKEYS+CAPSLOCK=80)
BASE 13 80
CTRL 13 80
SHIFT 13 80
CAPS 13 80
# > Define Floating Accent keys
#           FA_UMLAUT = A9
#           FA_CFLEX = AA
#           FA_TILDE = AB
#           FA_CEDILLA = AC
#           FA_ACUTE = AD
#           FA_GRAVE = AE
BASE 64 AA
SHIFT 64 A9
CAPS 64 A9
BASE 65 AC
SHIFT 65 AB
CAPS 65 AB
BASE 87 AE
SHIFT 87 AD
CAPS 87 AD
# > Define ASCII values
BASE 88 5B
```

SHIFT 88 7B
CAPS 88 7B
BASE 15 5D
SHIFT 15 7D
CAPS 15 7D
SHIFT 31 22
SHIFT 32 2F
SHIFT 35 3F
SHIFT 107 27
CAPS 107 27
SHIFT 108 60
CAPS 108 60
BASE 124 3C
SHIFT 124 3E
CAPS 124 3E
> Define ISO values
BASE_ISO 109 E9
SHIFT_ISO 109 C9
> Define Alternate Graph ISO values
ALTG 88 AB
ALTG 15 BB
ALTG 30 B1
ALTG 31 B2
ALTG 32 B3
ALTG 33 A2
ALTG 34 A4
ALTG 35 5E
ALTG 36 40
ALTG 37 A3
ALTG 38 5C
ALTG 40 AC
ALTG 41 23
ALTG 63 B6
ALTG 64 BC
ALTG 65 BD
ALTG 42 BE
ALTG 106 B5
ALTG 105 BA
> End of file
END 0 0

SEE ALSO

oldsetkeys(1)

The *Sun386i Developer's Guide* for keystation number diagrams.

NAME

ipalloc.netrange – range of addresses to allocate

SYNOPSIS

/etc/ipalloc.netrange

AVAILABILITY

Sun386i systems only.

DESCRIPTION

This file, if it exists on the YP master of the `hosts.byaddr` YP map, specifies the ranges of IP addresses that can be allocated by the `ipallocald(8C)` daemon. This allows multiple address assignment authorities, probably in multiple administrative domains, to coexist on the same IP network by preallocating ranges of addresses. If the file does not exist, the daemon assumes that all addresses not listed in the hosts map may be freely allocated.

This file can contain blank lines. Comments begin with a `#` character and extend to the end of the current line. Ranges of free addresses are specified on one line per-network or subnetwork.

The first token on the line is the IP address, in four part "dot" notation as also used in the `hosts` file, of the network or subnetwork described. It is separated from the second token by white space. The second token is a comma-separated list of local host number ranges on that network. These ranges take two forms: a single number specifies just that local host number, and two numbers separated by a dash specify all local host numbers starting at the first number and ending at the second. (In the case of a subnet, host numbers not in that subnet are excluded.)

For example, the following file would specify that a subset of the addresses on the class C network 192.9.200.0 may be allocated, and only some of the addresses on two particular subnets of the class B network 128.255.0.0 may be allocated. In any case, only non-broadcast addresses not listed in the hosts map are subject to allocation:

```
# We have three network cables administered using automatic # IP address allocation.
```

```
192.9.200.0          50-100,200-254      128.255.210.0      3,5,7,9,100-110
128.255.211.0       1-254
```

SEE ALSO

`hosts(5)`, `netmasks(5)`, `ipallocald(8C)`

BUGS

There is a silent limit of twenty ranges per network.

NAME

link – link editor interfaces

SYNOPSIS

```
#include <link.h>
```

DESCRIPTION

Dynamically linked executables created by `ld(1)` contain a number of data structures that are used by the dynamic link editor to finish link-editing the program during program execution. These data structures are described with a `link_dynamic` structure, as defined in the `<link.h>` file. `ld` always identifies the location of this structure in the executable file with the symbol `__DYNAMIC`. This symbol is `ld`-defined and if referenced in an executable that does not require dynamic linking will have the value zero.

The program stub linked with “main” programs by compiler drivers such as `cc(1)` (called `crt0`) tests the definition of `__DYNAMIC` to determine whether or not the dynamic link editor should be invoked. Programs supplying a substitute for `crt0` must either duplicate this functionality or else require that the programs with which they are linked be linked *statically*. Otherwise, such replacement `crt0`'s must open and map in the executable `/usr/lib/ld.so` using `mmap(2)`. Care should be taken to ensure that the expected mapping relationship between the “text” and “data” segments of the executable is maintained in the same manner that the `execve(2)` system call does. The first location following the `a.out` header of this executable is the entry point to a function that begins the dynamic link-editing process. This function must be called and supplied with two arguments. The first argument is an integer representing the revision level of the argument list, and should have the value “1”. The second should be a pointer to an argument list structure of the form:

```
struct {
    int     crt_ba;           /* base address of ld.so */
    int     crt_dzfd;        /* open fd to /dev/zero */
    int     crt_ldfd;        /* open fd to ld.so */
    struct  link_dynamic *crt_dp; /* pointer to program's __DYNAMIC */
    char    **crt_ep;        /* environment strings */
    caddr_t crt_bp;         /* debugger hook */
}
```

The members of the structure are:

<code>crt_ba</code>	The address at which <code>/usr/lib/ld.so</code> has been mapped.
<code>crt_dzfd</code>	An open file descriptor for <code>/dev/zero</code> . <code>ld.so</code> will close this file descriptor before returning.
<code>crt_ldfd</code>	The file descriptor used to map <code>/usr/lib/ld.so</code> . <code>ld.so</code> will close this file descriptor before returning.
<code>crt_dp</code>	A pointer to the label <code>__DYNAMIC</code> in the executable which is calling <code>ld.so</code> .
<code>crt_ep</code>	A pointer to the environment strings provided to the program.
<code>crt_bp</code>	A location in the executable which contains an instruction that will be executed after the call to <code>ld.so</code> returns. This location is used as a breakpoint in programs that are being executed under the control of a debugger such as <code>adb(1)</code> .

SEE ALSO

`ld(1)`, `mmap(2)`, `a.out(5)`

BUGS

These interfaces are under development and are subject to rapid change.

NAME

magic – file command's magic number file

DESCRIPTION

The **file(1)** command identifies the type of a file using, among other tests, a test for whether the file begins with a certain *magic number*. The file **/etc/magic** specifies what magic numbers are to be tested for, what message to print if a particular magic number is found, and additional information to extract from the file.

Each line of the file specifies a test to be performed. A test compares the data starting at a particular offset in the file with a 1-byte, 2-byte, or 4-byte numeric value or a string. If the test succeeds, a message is printed. The line consists of the following fields:

offset A number specifying the offset, in bytes, into the file of the data which is to be tested.

type The type of the data to be tested. The possible values are:

byte A one-byte value.

short A two-byte value.

long A four-byte value.

string A string of bytes.

The types **byte**, **short**, and **long** may optionally be followed by a mask specifier of the form **&number**. If a mask specifier is given, the value is AND'ed with the *number* before any comparisons are done. The *number* is specified in C form; for instance, 13 is decimal, 013 is octal, and 0x13 is hexadecimal.

test The value to be compared with the value from the file. If the type is numeric, this value is specified in C form; if it is a string, it is specified as a C string with the usual escapes permitted (for instance, **\n** for NEWLINE).

Numeric values may be preceded by a character indicating the operation to be performed. It may be '=', to specify that the value from the file must equal the specified value, '<', to specify that the value from the file must be less than the specified value, '>', to specify that the value from the file must be greater than the specified value, '&', to specify that all the bits in the specified value must be set in the value from the file, '^', to specify that at least one of the bits in the specified value must not be set in the value from the file, or x to specify that any value will match. If the character is omitted, it is assumed to be '='.

For string values, the byte string from the file must match the specified byte string; the byte string from the file which is matched is the same length as the specified byte string.

message The message to be printed if the comparison succeeds. If the string contains a **printf(3S)** format specification, the value from the file (with any specified masking performed) is printed using the message as the format string.

Some file formats contain additional information which is to be printed along with the file type. A line which begins with the character '>' indicates additional tests and messages to be printed. If the test on the line preceding the first line with a '>' succeeds, the tests specified in all the subsequent lines beginning with '>' are performed, and the messages printed if the tests succeed. The next line which does not begin with a '>' terminates this.

FILES

/etc/magic

SEE ALSO

file(1), **printf(3S)**

BUGS

There should be more than one level of subtests, with the level indicated by the number of '>' at the beginning of the line.

NAME

mtab – mounted file system table

SYNOPSIS

/etc/mtab

#include <mntent.h>

DESCRIPTION

mtab resides in the */etc* directory, and contains a table of filesystems currently mounted by the **mount(8)** command. **umount** removes entries from this file.

The file contains a line of information for each mounted filesystem, structurally identical to the contents of */etc/fstab*, described in **fstab(5)**. There are a number of lines of the form:

fsname dir type opts freq passno

for example:

/dev/xy0a / 4.2 rw,noquota 1 2

The file is accessed by programs using **getmntent(3)**, and by the system administrator using a text editor.

FILES

/etc/mtab

/etc/fstab

SEE ALSO

getmntent(3), **fstab(5)**, **mount(8)**

NAME

netgroup – list of network groups

DESCRIPTION

netgroup defines network wide groups, used for permission checking when doing remote mounts, remote logins, and remote shells. For remote mounts, the information in **netgroup** is used to classify machines; for remote logins and remote shells, it is used to classify users. Each line of the **netgroup** file defines a group and has the format

groupname member1 member2

where *memberi* is either another group name, or a triple:

(hostname, username, domainname)

Any of these three fields can be empty, in which case it signifies a wild card. Thus

universal (,,)

defines a group to which everyone belongs.

A gateway machine should be listed under all possible hostnames by which it may be recognized:

wan (gateway,,) (gateway-ebb,,)

Field names that begin with something other than a letter, digit or underscore (such as '-') work in precisely the opposite fashion. For example, consider the following entries:

justmachines (analytica,-,sun)

justpeople (-,babbage,sun)

The machine **analytica** belongs to the group **justmachines** in the domain **sun**, but no users belong to it. Similarly, the user **babbage** belongs to the group **justpeople** in the domain **sun**, but no machines belong to it.

The *domainname* field refers to the domain *n* which the triple is valid, not the name containing the trusted host.

FILES

/etc/netgroup

SEE ALSO

getnetgrent(3N), **exports(5)**, **makedbm(8)**, **ypserv(8)**

NAME

netmasks – network mask data base

DESCRIPTION

The **netmasks** file contains network masks used to implement IP standard subnetting. For each network that is subnetted, a single line should exist in this file with the network number, any number of SPACE or TAB characters, and the network mask to use on that network. Network numbers and masks may be specified in the conventional IP '.' notation (like IP host addresses, but with zeroes for the host part). For example,

128.32.0.0 255.255.255.0

can be used to specify the the Class B network 128.32.0.0 should have eight bits of subnet field and eight bits of host field, in addition to the standard sixteen bits in the network field. When running Yellow Pages, this file on the master is used for the **netmasks.byaddr** map.

FILES

/etc/netmasks

SEE ALSO

ifconfig(8C)

Postel, Jon, and Mogul, Jeff, *Internet Standard Subnetting Procedure*, RFC 950, Network Information Center, SRI International, Menlo Park, Calif., August 1985.

NAME

netrc – file for ftp(1C) remote login data

DESCRIPTION

The **.netrc** file contains data for logging in to a remote host over the network for file transfers by **ftp(1C)**. This file resides in the user's home directory on the machine initiating the file transfer. Its permissions should be set to disallow read access by group and others (see **chmod(1V)**).

The following tokens are recognized; they may be separated by SPACE, TAB, or NEWLINE characters:

machinename

Identify a remote machine name. The auto-login process searches the **.netrc** file for a **machine** token that matches the remote machine specified on the **ftp** command line or as an **open** command argument. Once a match is made, the subsequent **.netrc** tokens are processed, stopping when the EOF is reached or another **machine** token is encountered.

loginname

Identify a user on the remote machine. If this token is present, the auto-login process will initiate a login using the specified name.

passwordstring

Supply a password. If this token is present, the auto-login process will supply the specified string if the remote server requires a password as part of the login process. Note: if this token is present in the **.netrc** file, **ftp** will abort the auto-login process if the **.netrc** is readable by anyone besides the user.

accountstring

Supply an additional account password. If this token is present, the auto-login process will supply the specified string if the remote server requires an additional account password, or the auto-login process will initiate an **ACCT** command if it does not.

macdefname

Define a macro. This token functions as the **ftp macdef** command functions. A macro is defined with the specified name; its contents begin with the next **.netrc** line and continue until a NULL line (consecutive NEWLINE characters) is encountered. If a macro named **init** is defined, it is automatically executed as the last step in the auto-login process.

EXAMPLE

The command:

```
machine ray login demo password mypassword
```

allows an autologin to the machine **ray** using the login name **demo** with password **mypassword**.

FILES

~/netrc

SEE ALSO

chmod(1V), **ftp(1C)**, **ftpd(8C)**

NAME

networks – network name data base

DESCRIPTION

The **networks** file contains information regarding the known networks which comprise the DARPA Internet. For each network a single line should be present with the following information:

- official network name
- network number
- aliases

Items are separated by any number of blanks and/or TAB characters. A '#' indicates the beginning of a comment; characters up to the end of the line are not interpreted by routines which search the file. This file is normally created from the official network data base maintained at the Network Information Control Center (NIC), though local changes may be required to bring it up to date regarding unofficial aliases and/or unknown networks.

Network number may be specified in the conventional '.' notation using the `inet_network()` routine from the Internet address manipulation library, `inet(3N)`. Network names may contain any printable character other than a field delimiter, NEWLINE, or comment character.

FILES

`/etc/networks`

SEE ALSO

`getnetent(3N)`, `inet(3N)`

BUGS

A name server should be used instead of a static file. A binary indexed file format should be available for fast access.

NAME

passwd – password file

SYNOPSIS

/etc/passwd

DESCRIPTION

The `passwd` file contains basic information about each user's account. This file contains a one-line entry for each authorized user, of the form:

```
username:password:uid:gid:gcoss-field:home-dir:login-shell
```

where

<i>username</i>	is the user's login name. This field contains no uppercase characters, and must not be more than eight characters in length.
<i>password</i>	is the user's encrypted password, or a string of the form: <code>##name</code> if the encrypted password is in the <code>/etc/security/passwd.adjunct</code> file (see <code>passwd.adjunct(5)</code>). If this field is empty, <code>login(1)</code> does not request a password before logging the user in.
<i>uid</i>	is the user's numerical ID for the system, which must be unique. <i>uid</i> is generally a value between 0 and 32767.
<i>gid</i>	is the numerical ID of the group that the user belongs to. <i>gid</i> is generally a value between 0 and 32767.
<i>gcoss-field</i>	is the user's real name, along with information to pass along in a mail-message heading. It is called the <i>gcoss-field</i> for historical reasons. A <code>&</code> in this field stands for the login name (in cases where the login name appears in a user's real name).
<i>home-dir</i>	is the pathname to the directory in which the user is initially positioned upon logging in.
<i>login-shell</i>	is the user's initial shell program. If this field is empty, the default shell is <code>/usr/bin/sh</code> .

The `passwd` file can also have lines beginning with a '+' (plus sign) which means to incorporate entries from the Yellow Pages. There are three styles of + entries in this file: by itself, + means to insert the entire contents of the Yellow Pages password file at that point; `+name` means to insert the entry (if any) for *name* from the Yellow Pages at that point; `+@netgroup` means to insert the entries for all members of the network group *netgroup* at that point. If a `+name` entry has a non-NULL *password*, *gcoss*, *home-dir*, or *login-shell* field, the value of that field overrides what is contained in the Yellow Pages. The *uid* and *gid* fields cannot be overridden.

The `passwd` file can also have lines beginning with a '-' (minus sign) which means to disallow entries from the Yellow Pages. There are two styles of '-' entries in this file: `-name` means to disallow any subsequent entries (if any) for *name* (in this file or in the Yellow Pages); `-@netgroup` means to disallow any subsequent entries for all members of the network group *netgroup*.

The password file is an ASCII file that resides in the `/etc` directory. Because the encrypted passwords on a secure system are kept in the `passwd.adjunct` file, `/etc/passwd` has general read permission on all systems, and can be used by routines that map numerical user IDs to names.

Appropriate precautions must be taken to lock the `/etc/passwd` file against simultaneous changes if it is to be edited with a text editor; `vipw(8)` does the necessary locking.

EXAMPLE

Here is a sample `passwd` file when `passwd.adjunct` does not exist:

```
root:q.mJzTnu8icF.:0:10:God:/:bin/csh
fred:6k/7KCFRPNVXg:508:10:% Fredericks:/usr2/fred:/bin/csh
+john:
+@documentation:no-login:
+::::Guest
```

Here is a sample `passwd` file when `passwd.adjunct` does exist:

```
root:##root:0:10:God:/:bin/csh
fred:##fred:508:10:& Fredericks:/usr2/fred:/bin/csh
+john:
+@documentation:no-login:
+::::Guest
```

In this example, there are specific entries for users `root` and `fred`, to assure that they can log in even when the system is running standalone. The user `john` will have his password entry in the Yellow Pages incorporated without change; anyone in the netgroup `documentation` will have their password field disabled, and anyone else will be able to log in with their usual password, shell, and home directory, but with a `gcos`-field of `Guest`.

FILES

```
/etc/passwd
/etc/security/passwd.adjunct
```

SEE ALSO

`login(1)`, `mail(1)`, `passwd(1)`, `crypt(3)`, `getpwent(3)`, `group(5)`, `passwd.adjunct(5)`, `adduser(8)`, `sendmail(8)`, `vipw(8)`

BUGS

`mail(1)` and `sendmail(8)` use the `gcos`-field to compose the `From:` line for addressing mail messages, but these programs get confused by nested parentheses when composing replies. This problem can be avoided by using different types of brackets within the `gcos`-field; for example:

```
(& Fredricks [Podunk U <EE/CIS>] {818}-555-5555)
```

NAME

passwd.adjunct – user security data file

SYNOPSIS

/etc/security/passwd.adjunct

DESCRIPTION

The **passwd.adjunct** file contains the following information for each user:

<i>name</i>	This is the user's login name in the system and it must be unique.
<i>password</i>	The encrypted password.
<i>minimum label</i>	The lowest security level at which this user is allowed to login (not used at C2 level).
<i>maximum label</i>	The highest security level at which this user is allowed to login (not used at C2 level).
<i>default label</i>	The security level at which this user will run unless a label is specified at login.
<i>always audit flags</i>	Flags specifying events always to be audited for this user's processes; see audit_control(5) .
<i>never audit flags</i>	Flags specifying events never to be audited for this user's processes; see audit_control(5) .

Fields are separated by a colon, and each user from the next by a NEWLINE.

The **passwd.adjunct** file can also have lines beginning with a '+' (plus sign), which means to incorporate entries from the Yellow Pages. There are three styles of '+' entries: all by itself, '+' means to insert the entire contents of the Yellow Pages **passwd.adjunct** file at that point; *+name* means to insert the entry (if any) for *name* from the Yellow Pages at that point; *+@name* means to insert the entries for all members of the network group *name* at that point. If a '+' entry has a non-NULL password, it will override what is contained in the Yellow Pages.

EXAMPLE

Here is a sample **/etc/security/passwd.adjunct** file:

```
root:q.mJzTnu8icF:::::::::
ignatz:7KsI8CFRPNVXg::b,ap,bp,gp,dp,ic,r,d,l::+dc,+da:-dr:
rex:7HU8UUGRPNVXg:b,ap:b,ap,bp:b,ap::+ad:
+fred:9x.FFUw6xcJBa:::::::::
+:
```

The user **root** is the super-user, who has no special label constraints nor audit interest. The user **ignatz** may have any label from the lowest to the level **b** and any of a large number of categories. **ignatz** will run at system low unless he specifies otherwise. He is being audited on the system default event classes as well as data creations and access changes, but never for failed data reads. The user **rex** can function only at the level **b** and only in the categories **ap** or **ap** and **bp**. By default, he will run at '**b,bp**'. He is audited with the system defaults, except that successful administrative operations are not audited. The user **fred** will have the labels and audit flags that are specified in the Yellow Pages **passwd.adjunct** file. Any other users specified in the Yellow Pages will be able to log in on this system.

The user security data file resides in the **/etc/security** directory. Because it contains encrypted passwords, it does not have general read permission.

FILES

/etc/security/passwd.adjunct
/etc/security

SEE ALSO

login(1), **passwd(1)**, **crypt(3)**, **getpwaent(3)**, **getpwent(3)**, **audit_control(5)**, **passwd(5)**, **adduser(8)**

NAME

phones – remote host phone number data base

SYNOPSIS

/etc/phones

DESCRIPTION

The file */etc/phones* contains the system-wide private phone numbers for the *tip(1C)* program. */etc/phones* is normally unreadable, and so may contain privileged information. The format of */etc/phones* is a series of lines of the form:

<system-name>[\t]<phone-number>.*

The system name is one of those defined in the *remote(5)* file and the phone number is constructed from *[0123456789-=*%]*. The '=' and '*' characters are indicators to the auto call units to pause and wait for a second dial tone (when going through an exchange). The '=' is required by the DF02-AC and the '*' is required by the BIZCOMP 1030.

Comment lines are lines containing a '#' sign in the first column of the line.

Only one phone number per line is permitted. However, if more than one line in the file contains the same system name *tip(1C)* will attempt to dial each one in turn, until it establishes a connection.

FILES

/etc/phones

SEE ALSO

tip(1C), *remote(5)*

NAME

plot – graphics interface

DESCRIPTION

Files of this format are produced by routines described in `plot(3X)`, and are interpreted for various devices by commands described in `plot(1G)`. A graphics file is a stream of plotting instructions. Each instruction consists of an ASCII letter usually followed by bytes of binary information. The instructions are executed in order. A point is designated by four bytes representing the x and y values; each value is a signed integer. The last designated point in an l , m , n , or p instruction becomes the “current point” for the next instruction.

Each of the following descriptions begins with the name of the corresponding routine in `plot(3X)`.

- m** Move: the next four bytes give a new current point.
 - n** Cont: draw a line from the current point to the point given by the next four bytes. See `plot(1G)`.
 - p** Point: plot the point given by the next four bytes.
 - l** Line: draw a line from the point given by the next four bytes to the point given by the following four bytes.
 - t** Label: place the following ASCII string so that its first character falls on the current point. The string is terminated by a NEWLINE.
 - a** Arc: the first four bytes give the center, the next four give the starting point, and the last four give the end point of a circular arc. The least significant coordinate of the end point is used only to determine the quadrant. The arc is drawn counter-clockwise.
 - c** Circle: the first four bytes give the center of the circle, the next two the radius.
 - e** Erase: start another frame of output.
 - f** Linemod: take the following string, up to a NEWLINE, as the style for drawing further lines. The styles are “dotted,” “solid,” “longdashed,” “shortdashed,” and “dotdashed.” Effective only in `plot 4014` and `plot ver`.
 - s** Space: the next four bytes give the lower left corner of the plotting area; the following four give the upper right corner. The plot will be magnified or reduced to fit the device as closely as possible.
- Space settings that exactly fill the plotting area with unity scaling appear below for devices supported by the filters of `plot(1G)`. The upper limit is just outside the plotting area. In every case the plotting area is taken to be square; points outside may be displayable on devices whose face is not square.

```
4014    space(0, 0, 3120, 3120);
ver     space(0, 0, 2048, 2048);
300, 300s space(0, 0, 4096, 4096);
450     space(0, 0, 4096, 4096);
```

SEE ALSO

`graph(1G)`, `plot(1G)`, `plot(3X)`

NAME

policies – network administration policies

AVAILABILITY

Sun386i systems only.

DESCRIPTION

The **policies** file contains information relevant to domain-wide administration policies. Each line contains two tokens, separated by white space; the first token is the name of an administrative policy, and the second is the value of that policy.

FILES

/etc/policies

/var/yp/domainname/policies.{dir,pag}

SEE ALSO

pnpd(8C), **rarpd(8C)**, **logintool(8)**

NAME

printcap – printer capability data base

SYNOPSIS

/etc/printcap

DESCRIPTION

printcap is a simplified version of the **termcap(5)** data base for describing printers. The spooling system accesses the **printcap** file every time it is used, allowing dynamic addition and deletion of printers. Each entry in the data base describes one printer. This data base may not be substituted for, as is possible for **termcap**, because it may allow accounting to be bypassed.

The default printer is normally **lp**, though the environment variable **PRINTER** may be used to override this. Each spooling utility supports a **-Pprinter** option to explicitly name a destination printer.

Refer to *System and Network Administration* for a discussion of how to set up the database for a given printer. On Sun386i systems, refer to **snap(1)** for information on setting up printers with the system and network administration program.

Each entry in the **printcap** file describes a printer, and is a line consisting of a number of fields separated by ':' characters. The first entry for each printer gives the names which are known for the printer, separated by '|' characters. The first name is conventionally a number. The second name given is the most common abbreviation for the printer, and the last name given should be a long name fully identifying the printer. The second name should contain no blanks; the last name may well contain blanks for readability. Entries may continue onto multiple lines by giving a '\' as the last character of a line, and empty fields may be included for readability.

Capabilities in **printcap** are all introduced by two-character codes, and are of three types:

Boolean Capabilities that indicate that the printer has some particular feature. Boolean capabilities are simply written between the ':' characters, and are indicated by the word **'bool'** in the **type** column of the capabilities table below.

Numeric Capabilities that supply information such as baud-rates, number of lines per page, and so on. Numeric capabilities are indicated by the word **num** in the **type** column of the capabilities table below. Numeric capabilities are given by the two-character capability code followed by the '#' character, followed by the numeric value. For example:

:br#1200:

is a numeric entry stating that this printer should run at 1200 baud.

String Capabilities that give a sequence which can be used to perform particular printer operations such as cursor motion. String valued capabilities are indicated by the word **str** in the **type** column of the capabilities table below. String valued capabilities are given by the two-character capability code followed by an '=' sign and then a string ending at the next following ':'. For example,

:rp=spinwriter:

is a sample entry stating that the remote printer is named **spinwriter**.

Sun386i DESCRIPTION

On Sun386i systems, **lpr(1)** and related printing commands use the Yellow Pages name service to obtain the **printcap** entry for a named printer if the entry does not exist in the local **/etc/printcap** file. For example, when a user issues the command

lpr -Pnewprinter foo

lpr searches **/etc/printcap** on the local system for an entry for **newprinter**. If no local entry for **newprinter** exists, then **lpr** searches the YP map called **printcap**. The search is invisible to the user.

lpr creates the spooling directory for the printer automatically if no spooling directory exists.

System administrators can make a printer available to the entire YP domain by placing an entry for that printer in the YP **printcap** map, typically using **snap**. Otherwise, the system administrator must edit the **/etc/printcap** file on the YP master and then rebuild the YP map.

CAPABILITIES

<i>Name</i>	<i>Type</i>	<i>Default</i>	<i>Description</i>
af	str	NULL	name of accounting file
br	num	none	if lp is a tty, set the baud rate (ioctl call)
cf	str	NULL	cifplot data filter
df	str	NULL	TeX data filter (DVI format)
du	str	0	User ID of user 'daemon'.
fc	num	0	if lp is a tty, clear flag bits
ff	str	"\f"	string to send for a form feed
fo	bool	false	print a form feed when device is opened
fs	num	0	like 'fc' but set bits
gf	str	NULL	graph data filter (plot(3X) format)
hl	bool	false	print the burst header page last
ic	bool	false	driver supports (non standard) ioctl to indent printout
if	str	NULL	name of text filter which does accounting
lf	str	"/dev/console"	error logging file name
lo	str	"lock"	name of lock file
lp	str	"/dev/lp"	device name to open for output
mc	num	0	maximum number of copies
ms	str	NULL	list of terminal modes to set or clear
mx	num	1000	maximum file size (in BUFSIZ blocks), zero = unlimited
nd	str	NULL	next directory for list of queues (unimplemented)
nf	str	NULL	ditroff data filter (device independent troff)
of	str	NULL	name of output filtering program
pc	num	200	price per foot or page in hundredths of cents
pl	num	66	page length (in lines)
pw	num	132	page width (in characters)
px	num	0	page width in pixels (horizontal)
py	num	0	page length in pixels (vertical)
rf	str	NULL	filter for printing FORTRAN style text files
rg	str	NULL	restricted group. Only members of group allowed access
rm	str	NULL	machine name for remote printer
rp	str	"lp"	remote printer name argument
rs	bool	false	restrict remote users to those with local accounts
rw	bool	false	open printer device read/write instead of read-only
sb	bool	false	short banner (one line only)
sc	bool	false	suppress multiple copies
sd	str	"/var/spool/lpd"	spool directory
sf	bool	false	suppress form feeds
sh	bool	false	suppress printing of burst page header
st	str	"status"	status file name
tc	str	NULL	name of similar printer; must be last
tf	str	NULL	troff data filter (C/A/T phototypesetter)
tr	str	NULL	trailer string to print when queue empties
vf	str	NULL	raster image filter
xc	num	0	if lp is a tty, clear local mode bits
xs	num	0	like 'xc' but set bits

If the local line printer driver supports indentation, the daemon must understand how to invoke it.

Note: the **fs**, **fc**, **xs**, and **xc** fields are flag *masks* rather than flag *values*. Certain default device flags are set when the device is opened by the line printer daemon if the device is connected to a terminal port. The flags indicated in the **fc** field are then cleared; the flags in the **fs** field are then set (or vice-versa, depending on the order of **fc#nnnn** and **fs#nnnn** in the **/etc/printcap** file). The bits cleared by the **fc** field and set by the **fs** field are those in the **sg_flags** field of the **sgtty** structure, as set by the **TIOCSETP** **ioctl** call, and the bits cleared by the **xc** field and set by the **xs** field are those in the "local flags" word, as set by the **TIOCLSET** **ioctl** call. See **ttcompat(4M)** for a description of these flags. For example, to set exactly the flags 06300 in the **fs** field, which specifies that the EVENP, ODDP, and XTABS modes are to be set, and all other flags are to be cleared, do:

```
:fc#0177777:fs#06300:
```

The same process applies to the **xc** and **xs** fields. Alternatively, the **ms** field can be used to specify modes to be set and cleared. These modes are specified as **stty(1V)** modes; any mode supported by **stty** may be specified, except for the baud rate which must be specified with the **br** field. This permits modes not supported by the older terminal interface described in **ttcompat(4M)** to be set or cleared. Thus, to set the terminal port to which the printer is attached to even parity, tab expansion, no newline to carriage-return/line-feed translation, and RTS/CTS flow control enabled, do:

```
:ms=evenp,-tabs,nl,crtscts:
```

On Sun386i systems, the **tc** field, as in the **termcap(5)** file, must appear last in the list of capabilities. It is recommended that each type of printer have a general entry describing common capabilities; then an individual printer can be defined with its particular capabilities plus a **tc** field that points to the general entry for that type of printer.

FILES

/etc/printcap

SEE ALSO

lpq(1), **lpr(1)**, **lprm(1)**, **snap(1)**, **stty(1V)**, **plot(3X)**, **ttcompat(4M)**, **termcap(5)**, **lpc(8)**, **lpd(8)**, **pac(8)**

System and Network Administration

NAME

protocols – protocol name data base

SYNOPSIS

/etc/protocols

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.

EXAMPLE

The following example is taken from SunOS.

```
#
# Internet (IP) protocols
#
ip          0          IP          # internet protocol, pseudo protocol number
icmp       1          ICMP        # internet control message protocol
ggp        3          GGP         # gateway-gateway protocol
tcp        6          TCP         # transmission control protocol
pup        12         PUP         # PARC universal packet protocol
udp        17         UDP         # user datagram protocol
```

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

publickey – public key database

SYNOPSIS

/etc/publickey

DESCRIPTION

/etc/publickey is the public key database used for secure networking. Each entry in the database consists of a network user name (which may either refer to a user or a hostname), followed by the user's public key (in hex notation), a colon, and then the user's secret key encrypted with its login password (also in hex notation).

This file is altered either by the user through the **chkey(1)** command or by the system administrator through the **newkey(8)** command. The file **/etc/publickey** should only contain data on the Yellow Pages master machine, where it is converted into the YP database **publickey.byname**.

SEE ALSO

chkey(1), **publickey(3R)**, **newkey(8)**, **ypupdated(8C)**

NAME

queuedefs – queue description file for at, batch, and cron

SYNOPSIS

`/var/spool/cron/queuedefs`

DESCRIPTION

The `queuedefs` file describes the characteristics of the queues managed by `cron(8)`. Each non-comment line in this file describes one queue. The format of the lines are as follows:

`q.[njob][nicen][nwaitw]`

The fields in this line are:

- `q` The name of the queue. `a` is the default queue for jobs started by `at(1)`; `b` is the default queue for jobs started by `batch` (see `at(1)`); `c` is the default queue for jobs run from a `crontab(5)` file.
- `njob` The maximum number of jobs that can be run simultaneously in that queue; if more than `njob` jobs are ready to run, only the first `njob` jobs will be run, and the others will be run as jobs that are currently running terminate. The default value is 100.
- `nice` The `nice(1)` value to give to all jobs in that queue that are not run with a user ID of super-user. The default value is 2.
- `nwait` The number of seconds to wait before rescheduling a job that was deferred because more than `njob` jobs were running in that job's queue, or because more than 25 jobs were running in all the queues. The default value is 60.

Lines beginning with `#` are comments, and are ignored.

EXAMPLE

```
#
# @(#)queuedefs 1.1 87/02/18 SMI; from S5R3
#
a.4j1n
b.2j2n90w
```

This file specifies that the `a` queue, for `at` jobs, can have up to 4 jobs running simultaneously; those jobs will be run with a `nice` value of 1. As no `nwait` value was given, if a job cannot be run because too many other jobs are running `cron` will wait 60 seconds before trying again to run it. The `b` queue, for `batch` jobs, can have up to 2 jobs running simultaneously; those jobs will be run with a `nice` value of 2. If a job cannot be run because too many other jobs are running, `cron` will wait 90 seconds before trying again to run it. All other queues can have up to 100 jobs running simultaneously; they will be run with a `nice` value of 2, and if a job cannot be run because too many other jobs are running `cron` will wait 60 seconds before trying again to run it.

FILES

`/var/spool/cron/queuedefs`

SEE ALSO

`at(1)`, `nice(1)`, `crontab(5)`, `cron(8)`

NAME

rasterfile – Sun's file format for raster images

SYNOPSIS

```
#include <rasterfile.h>
```

DESCRIPTION

A rasterfile is composed of three parts: first, a header containing 8 integers; second, a (possibly empty) set of colormap values; and third, the pixel image, stored a line at a time, in increasing y order. The image is laid out in the file as in a memory pixrect. Each line of the image is rounded up to the nearest 16 bits.

The header is defined by the following structure:

```
struct rasterfile {
    int    ras_magic;
    int    ras_width;
    int    ras_height;
    int    ras_depth;
    int    ras_length;
    int    ras_type;
    int    ras_maptype;
    int    ras_maplength;
};
```

The *ras_magic* field always contains the following constant:

```
#define RAS_MAGIC    0x59a66a95
```

The *ras_width*, *ras_height*, and *ras_depth* fields contain the image's width and height in pixels, and its depth in bits per pixel, respectively. The depth is either 1 or 8, corresponding to standard frame buffer depths. The *ras_length* field contains the length in bytes of the image data. For an unencoded image, this number is computable from the *ras_width*, *ras_height*, and *ras_depth* fields, but for an encoded image it must be explicitly stored in order to be available without decoding the image itself. Note: the length of the header and of the (possibly empty) colormap values are not included in the value of the *ras_length* field; it is only the image data length. For historical reasons, files of type RT_OLD will usually have a 0 in the *ras_length* field, and software expecting to encounter such files should be prepared to compute the actual image data length if needed. The *ras_maptype* and *ras_maplength* fields contain the type and length in bytes of the colormap values, respectively. If *ras_maptype* is not RMT_NONE and the *ras_maplength* is not 0, then the colormap values are the *ras_maplength* bytes immediately after the header. These values are either uninterpreted bytes (usually with the *ras_maptype* set to RMT_RAW) or the equal length red, green and blue vectors, in that order (when the *ras_maptype* is RMT_EQUAL_RGB). In the latter case, the *ras_maplength* must be three times the size in bytes of any one of the vectors.

FILES

/usr/include/rasterfile.h

SEE ALSO

SunView 1 Programmer's Guide

NAME

remote – remote host description file

SYNOPSIS

/etc/remote

DESCRIPTION

The systems known by **tip**(1C) and their attributes are stored in an ASCII file which is structured somewhat like the **termcap**(5) file. Each line in the file provides a description for a single *system*. Fields are separated by a colon ':'. Lines ending in a '\ character with an immediately following NEWLINE are continued on the next line.

The first entry is the name(s) of the host system. If there is more than one name for a system, the names are separated by vertical bars. After the name of the system comes the fields of the description. A field name followed by an '=' sign indicates a string value follows. A field name followed by a '#' sign indicates a following numeric value.

Entries named **tip*** and **cu*** are used as default entries by **tip**, and the **cu** interface to **tip**, as follows. When **tip** is invoked with only a phone number, it looks for an entry of the form **tip300**, where 300 is the baud rate with which the connection is to be made. When the **cu** interface is used, entries of the form **cu300** are used.

CAPABILITIES

Capabilities are either strings (**str**), numbers (**num**), or boolean flags (**bool**). A string capability is specified by *capability=value*; for example, '**dv=/dev/harris**'. A numeric capability is specified by *capability#value*; for example, '**xa#99**'. A boolean capability is specified by simply listing the capability.

- at** (str) Auto call unit type.
- br** (num) The baud rate used in establishing a connection to the remote host. This is a decimal number. The default baud rate is 300 baud.
- cm** (str) An initial connection message to be sent to the remote host. For example, if a host is reached through port selector, this might be set to the appropriate sequence required to switch to the host.
- cu** (str) Call unit if making a phone call. Default is the same as the *dv* field.
- di** (str) Disconnect message sent to the host when a disconnect is requested by the user.
- du** (bool) This host is on a dial-up line.
- dv** (str) *device(s)toopen*toestablish If this file refers to a terminal line, **tip**(1C) attempts to perform an exclusive open on the device to insure only one user at a time has access to the port.
- el** (str) Characters marking an end-of-line. The default is NULL. **tip** only recognizes '"' escapes after one of the characters in **el**, or after a carriage-return.
- fs** (str) Frame size for transfers. The default frame size is equal to BUFSIZ.
- hd** (bool) The host uses half-duplex communication, local echo should be performed.
- ie** (str) Input EOF marks. The default is NULL.
- oe** (str) Output EOF string. The default is NULL. When **tip** is transferring a file, this string is sent at EOF.
- pa** (str) The type of parity to use when sending data to the host. This may be one of "even", "odd", "none", "zero" (always set bit 8 to zero), "one" (always set bit 8 to 1). The default is "none".
- pn** (str) Telephone number(s) for this host. If the telephone number field contains an '@' sign, **tip** searches the */etc/phones* file for a list of telephone numbers — see **phones**(5). A '%' sign in the telephone number indicates a 5-second delay for the Ventel Modem.
- tc** (str) Indicates that the list of capabilities is continued in the named description. This is used primarily to share common capability information.

Here is a short example showing the use of the capability continuation feature:

```
UNIX-1200:\
    :dv=/dev/cua0:el=^D^U^C^S^Q^O@:du:at=ventel:ie=#$:oe=^D:br#1200:
arpavax|ax:\
    :pn=7654321%:tc=UNIX-1200
```

FILES

/etc/remote
/etc/phones

SEE ALSO

tip(1C), phones(5), termcap(5)

NAME

`resolve.conf` – configuration file for name server routines

DESCRIPTION

The resolver configuration file contains information that is read by the resolver routines the first time they are invoked in a process. The file is designed to be human readable and contains a list of name-value pairs that provide various types of resolver information.

The different configuration options are:

nameserver *address* The Internet address (in dot notation) of a name server that the resolver should query. At least one name server should be listed. Up to MAXNS (currently 3) name servers may be listed, in that case the resolver library queries tries them in the order listed. (The algorithm used is to try a name server, and if the query times out, try the next, until out of name servers, then repeat trying all the name servers until a maximum number of retries are made).

domain *name* The default domain to append to names that do not have a dot in them. This defaults to the domain set by the `domainname(1)` command.

address *address* An Internet address (in dot notation) of any preferred networks. The list of addresses returned by the resolver will be sorted to put any addresses on this network before any others.

The name value pair must appear on a single line, and the keyword (for instance, `nameserver`) must start the line. The value follows the keyword, separated by white space.

FILES

`/etc/resolve.conf`

SEE ALSO

`domainname(1)`, `gethostent(3N)`, `resolver(3)`, `named(8C)`

NAME

rgb – available colors (by name) for coloredit

SYNOPSIS

.rgb

\$HOME/.rgb

/usr/lib/.rgb

AVAILABILITY

Sun386i systems only.

DESCRIPTION

.rgb is an ASCII file containing consecutive lines terminated by newlines. Each line starts with three integers, each in the range 0-255. These integers are the RGB equivalent for the color named on the same line. At least one tab character delimits the last integer from the name field. The coloreditor searches for this file, first in the current directory; next, in the users home directory; and finally, in /usr/lib. The user can add to or delete from the .rgb file that he or she has access to, thus changing the available color table for subsequent invocations of coloredit.

EXAMPLE

0 0 0	Black
0 0 255	Blue
95 159 159	Cadet Blue
66 66 111	Cornflower Blue
107 35 142	Dark Slate Blue

SEE ALSO

coloredit(1)

NAME

rpc – rpc program number data base

SYNOPSIS

/etc/rpc

DESCRIPTION

The *rpc* file contains user readable names that can be used in place of rpc program numbers. Each line has the following information:

name of server for the rpc program
 rpc program 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.

Here is an example of the */etc/rpc* file from the SunOS System.

```

#
#          rpc 1.10 87/04/10
#
portmapper 100000      portmap sunrpc
rstatd      100001      rstat rup perfmeter
rusersd     100002      rusers
nfs         100003      nfsprog
ypserv      100004      ypprog
mountd      100005      mount showmount
ypbind      100007
walld       100008      rwall shutdown
yppasswdd   100009      yppasswd
etherstatd  100010      etherstat
rquotad     100011      rquotaprog quota rquota
sprayd      100012      spray
3270_mapper 100013
rje_mapper  100014
selection_svc 100015      selnsvc
database_svc 100016
rex         100017      rex
alis        100018
sched       100019
llockmgr    100020
nlockmgr    100021
x25.inr     100022
statmon     100023
status      100024
bootparam   100026
ypupdated   100028      ypupdate
keyserv     100029      keyserver

```

FILES

/etc/rpc

SEE ALSO

getrpcent(3N)

NAME

sccsfile – format of SCCS file

DESCRIPTION

An SCCS file is an ASCII file. It consists of six logical parts: the *checksum*, the *delta table* (contains information about each delta), *user names* (contains login names and/or numerical group IDs of users who may add deltas), *flags* (contains definitions of internal keywords), *comments* (contains arbitrary descriptive information about the file), and the *body* (contains the actual text lines intermixed with control lines).

Throughout an SCCS file there are lines which begin with the ASCII SOH (start of heading) character (octal 001). This character is hereafter referred to as *the control character* and will be represented graphically as '@'. Any line described below which is not depicted as beginning with the control character is prevented from beginning with the control character.

Entries of the form DDDDD represent a five digit string (a number between 00000 and 99999).

Each logical part of an SCCS file is described in detail below.

Checksum

The checksum is the first line of an SCCS file. The form of the line is:

@hDDDDD

The value of the checksum is the sum of all characters, except those of the first line. The @h provides a *magic number* of (octal) 064001.

Delta table

The delta table consists of a variable number of entries of the form:

```
@s DDDDD/DDDDD/DDDDD
@d <type> <SCCS ID> yr/mo/da hr:mi:se <pgmr> DDDDD DDDDD
@i DDDDD ...
@x DDDDD ...
@g DDDDD ...
@m <MR number>
.
.
.
@c <comments> ...
.
.
.
@e
```

The first line (@s) contains the number of lines inserted/deleted/unchanged respectively. The second line (@d) contains the type of the delta (currently, normal: D< and removed: R), the SCCS ID of the delta, the date and time of creation of the delta, the login name corresponding to the real user ID at the time the delta was created, and the serial numbers of the delta and its predecessor, respectively.

The @i, @x, and @g lines contain the serial numbers of deltas included, excluded, and ignored, respectively. These lines are optional.

The @m lines (optional) each contain one MR number associated with the delta; the @c lines contain comments associated with the delta.

The @e line ends the delta table entry.

User names

The list of login names and/or numerical group IDs of users who may add deltas to the file, separated by NEWLINE characters. The lines containing these login names and/or numerical group IDs are surrounded by the bracketing lines @u and @U. An empty list allows anyone to

make a delta.

Flags

Keywords used internally (see `admin(1)` for more information on their use). Each flag line takes the form:

```
@f <flag> <optional text>
```

The following flags are defined:

```
@f t      <type of program>
@f v      <program name>
@f i
@f b
@f m      <module name>
@f f      <floor>
@f c      <ceiling>
@f d      <default-sid>
@f n
@f j
@f l      <lock-releases>
@f q      <user defined>
@f e      <0|1>
```

The **t** flag defines the replacement for the **identification keyword**. The **v** flag controls prompting for MR numbers in addition to comments; if the optional text is present it defines an MR number validity checking program. The **i** flag controls the warning/error aspect of the 'No id keywords' message. When the **i** flag is not present, this message is only a warning; when the **i** flag is present, this message will cause a "fatal" error (the file will not be gotten, or the delta will not be made). When the **b** flag is present the **-b** keyletter may be used on the `get` command to cause a branch in the delta tree. The **m** flag defines the first choice for the replacement text of the `scsfile.5` identification keyword. The **f** flag defines the "floor" release; the release below which no deltas may be added. The **c** flag defines the "ceiling" release; the release above which no deltas may be added. The **d** flag defines the default SID to be used when none is specified on a `get` command. The **n** flag causes `delta` to insert a "null" delta (a delta that applies *no* changes) in those releases that are skipped when a delta is made in a *new* release (for example, when delta 5.1 is made after delta 2.7, releases 3 and 4 are skipped). The absence of the **n** flag causes skipped releases to be completely empty. The **j** flag causes `get` to allow concurrent edits of the same base SID. The **l** flag defines a *list* of releases that are *locked* against editing (`get(1)` with the **-e** keyletter). The **q** flag defines the replacement for the **identification keyword**. The **e** flag indicates whether a file is encoded or not. A **1** indicates that the file is encoded. Files need to be encoded when they contain control characters, or when they do not end with a NEWLINE. The **e** flag allows for any type of file to be checked in.

Comments

Arbitrary text surrounded by the bracketing lines `@t` and `@T`. The comments section typically will contain a description of the file's purpose.

Body

The body consists of text lines and control lines. Text lines do not begin with the control character, control lines do. There are three kinds of control lines: *insert*, *delete*, and *end*, represented by:

```
@I DDDDD
@D DDDDD
@E DDDDD
```

respectively. The digit string is the serial number corresponding to the delta for the control line.

SEE ALSO

admin(1), delta(1), get(1), prs(1)

NAME

services – Internet services and aliases

DESCRIPTION

The services file contains an entry for each service available through the DARPA Internet. Each entry consists of a line of the form:

service-name port/protocol aliases

service-name This is the official Internet service name.

port/protocol This field is composed of the port number and protocol through which the service is provided (for instance, 512/tcp).

aliases This is a list of alternate names by which the service might be requested.

Fields can be separated by any number of spaces or TAB's. A '#' (pound-sign) 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), inetd.conf(5)

BUGS

A name server should be used instead of a static file.

NAME

sm, sm.bak, sm.state – in.statd directory and file structures

SYNOPSIS

/etc/sm, /etc/sm.bak, /etc/sm.state

DESCRIPTION

/etc/sm and /etc/sm.bak are directories generated by in.statd. Each entry in /etc/sm represents the name of the machine to be monitored by the in.statd daemon. Each entry in /etc/sm.bak represents the name of the machine to be notified by the in.statd daemon upon its recovery.

/etc/sm.state is a file generated by rpc.statd to record the its version number. This version number is incremented each time a crash or recovery takes place.

FILES

/etc/sm
/etc/sm.bak
/etc/sm.state

SEE ALSO

lockd(8C), statd(8C)

NAME

sm, sm.bak, state – statd directories and file structures

SYNOPSIS

/etc/sm /etc/sm.bak /etc/state

DESCRIPTION

/etc/sm and */etc/sm.bak* are directories generated by *statd*. Each entry in */etc/sm* represents the name of the machine to be monitored by the *statd* daemon. Each entry in */etc/sm.bak* represents the name of the machine to be notified by the *statd* daemon upon its recovery.

/etc/state is a file generated by *statd* to record its version number. This version number is incremented each time a crash or recovery takes place.

FILES

/etc/sm
/etc/sm.bak
/etc/state

SEE ALSO

lockd(8C), *statd(8C)*

NAME

syslog.conf – configuration file for syslogd system log daemon

SYNOPSIS

/etc/syslog.conf

DESCRIPTION

The file */etc/syslog.conf* contains information used by the system log daemon, *syslogd(8)*, to forward a system message to appropriate log files and/or users. *syslog* preprocesses this file through *m4(1V)* to obtain the correct information for certain log files.

A configuration entry is composed of two TAB-separated fields:

selector *action*

The *selector* field contains a semicolon-separated list of priority specifications of the form:

facility.level[;facility.level]

where *facility* is a system facility, or comma-separated list of facilities, and *level* is an indication of the severity of the condition being logged. Recognized values for *facility* include:

user	Messages generated by user processes. This is the default priority for messages from programs or facilities not listed in this file.
kern	Messages generated by the kernel.
mail	The mail system.
daemon	System daemons, such as <i>ftpd(8C)</i> , <i>routed(8C)</i> , etc.
auth	The authorization system: <i>login(1)</i> , <i>su(1)</i> , <i>getty(8)</i> , etc.
lpr	The line printer spooling system: <i>lpr(1)</i> , <i>lpc(8)</i> , <i>lpd(8)</i> , etc.
news	Reserved for the USENET network news system.
uucp	Reserved for the UUCP system; it does not currently use the <i>syslog</i> mechanism.
cron	The <i>cron/at</i> facility; <i>crontab(1)</i> , <i>at(1)</i> , <i>cron(8)</i> , etc.
local0-7	Reserved for local use.
mark	For timestamp messages produced internally by <i>syslogd</i> .
*	An asterisk indicates all facilities except for the <i>mark</i> facility.

Recognized values for *level* are (in descending order of severity):

emerg	For panic conditions that would normally be broadcast to all users.
alert	For conditions that should be corrected immediately, such as a corrupted system database.
crit	For warnings about critical conditions, such as hard device errors.
err	For other errors.
warning	For warning messages.
notice	For conditions that are not error conditions, but may require special handling.
info	Informational messages.
debug	For messages that are normally used only when debugging a program.
none	Do not send messages from the indicated <i>facility</i> to the selected file. For example, a <i>selector</i> of

***.debug;mail.none**

will send all messages *except* mail messages to the selected file.

The *action* field indicates where to forward the message. Values for this field can have one of four forms:

- A filename, beginning with a leading slash, which indicates that messages specified by the *selector* are to be written to the specified file. The file will be opened in append mode.
- The name of a remote host, prefixed with an @, as with: @*server*, which indicates that messages specified by the *selector* are to be forwarded to the syslogd on the named host.
- A comma-separated list of usernames, which indicates that messages specified by the *selector* are to be written to the named users if they are logged in.
- An asterisk, which indicates that messages specified by the *selector* are to be written to all logged-in users.

Blank lines are ignored. Lines for which the first nonwhite character is a '#' are treated as comments.

Sun386i DESCRIPTION

The file is as described above, except that there is an additional valid entry type, for translation. A line containing the keyword "translate," if present, specifies how system error messages are translated, suppressed, or forwarded to appropriate log files and/or users.

A translation entry in the file is composed of five TAB-separated fields:

<i>translate</i>	<i>source</i>	<i>facility</i>	<i>input</i>	<i>output</i>
------------------	---------------	-----------------	--------------	---------------

The *translate* field consists of the word **translate** and is used to indicate that this is a translation entry.

The *source* field contains a comma separated list of source names. Recognized sources are:

klog	Messages placed in /dev/klog by the kernel.
log	Messages placed in /dev/log file by local programs.
syslog	Messages placed in the internet socket by programs on other systems.
*	An asterisk indicates all three sources (klog, log and syslog).

The *facility* field contains a comma-separated list of facilities.

The *input* field is the name of the file used to map error messages (in printf format strings) to numbers. This number is used to locate a new string in the file specified in the output field. The format of both files is described in `translate(5)`.

The output file specified by the output field translates the numbers from the input file into the desired error messages, and also specifies the format to be used to output each message.

EXAMPLE

With the following configuration file:

*.notice;mail.info	/var/log/notice
*.crit	/var/log/critical
kern,mark.debug	/dev/console
kern.err	@server
*.emerg	*
*.alert	root,operator
*.alert;auth.warning	/var/log/auth

syslogd will log all mail system messages except debug messages and all notice (or higher) messages into a file named /var/log/notice. It logs all critical messages into /var/log/critical, and all kernel messages and 20-minute marks onto the system console.

Kernel messages of **err** (error) severity or higher are forwarded to the machine named *server*. Emergency messages are forwarded to all users. The users "root" and "operator" are informed of any alert messages. All messages from the authorization system of **warning** level or higher are logged in the file */var/log/auth*.

FILES

/etc/syslog.conf
/var/log/notice
/var/log/critical
/var/log/auth

SEE ALSO

at(1), **crontab(1)**, **logger(1)**, **login(1)**, **lpr(1)**, **m4(1V)**, **su(1)**, **syslog(3)**, **translate(5)**, **cron(8)**, **ftpd(8C)**, **getty(8)**, **lpc(8)**, **lpd(8)**, **routed(8C)**, **syslogd(8)**

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 EOF 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, unless the **B** keyletter is used.

The header block looks like:

```
#define TBLOCK512
#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 *lfilename* 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 inode 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

term – terminal driving tables for nroff

SYNOPSIS

/usr/lib/term/tabname

DESCRIPTION

nroff(1) uses driving tables to customize its output for various types of output devices, such as terminals, line printers, daisy-wheel printers, or special output filter programs. These driving tables are written as C programs, compiled, and installed in the directory */usr/lib/term*. The *name* of the output device is specified with the **-T** option of **nroff**. The structure of the terminal table is as follows:

```
#define    INCH    240

struct {
    int bset;
    int breset;
    int Hor;
    int Vert;
    int Newline;
    int Char;
    int Em;
    int Halfline;
    int Adj;
    char *twinit;
    char *twrest;
    char *twnl;
    char *hhr;
    char *hlf;
    char *flr;
    char *bdon;
    char *bdoff;
    char *ploton;
    char *plotoff;
    char *up;
    char *down;
    char *right;
    char *left;
    char *codetab[256-32];
    char *zzz;
} t;
```

The meanings of the various fields are as follows:

bset	Bits to set in the sg_flags field of the sgtty structure before output; see ttcompat(4M) .
breset	Bits to reset in the sg_flags field of the sgtty structure after output; see ttcompat(4M) .
Hor	Horizontal resolution in fractions of an inch.
Vert	Vertical resolution in fractions of an inch.
Newline	Space moved by a NEWLINE (LINEFEED) character in fractions of an inch.
Char	Quantum of character sizes, in fractions of an inch. (that is, a character is a multiple of Char units wide)
Em	Size of an em in fractions of an inch.
Halfline	Space moved by a half- LINEFEED (or half-reverse- LINEFEED) character in fractions of an inch.

Adj	Quantum of white space, in fractions of an inch. (that is, white spaces are a multiple of Adj units wide) Note: if this is less than the size of the SPACE character (in units of Char ; see below for how the sizes of characters are defined), nroff will output fractional SPACE characters using plot mode. Also, if the -e switch to nroff is used, Adj is set equal to Hor by nroff .										
twinit	Set of characters used to initialize the terminal in a mode suitable for nroff .										
twrest	Set of characters used to restore the terminal to normal mode.										
twnl	Set of characters used to move down one line.										
hhr	Set of characters used to move up one-half line.										
hlf	Set of characters used to move down one-half line.										
flr	Set of characters used to move up one line.										
bdon	Set of characters used to turn on hardware boldface mode, if any.										
bdoff	Set of characters used to turn off hardware boldface mode, if any.										
ploton	Set of characters used to turn on hardware plot mode (for Diablo type mechanisms), if any.										
plotoff	Set of characters used to turn off hardware plot mode (for Diablo type mechanisms), if any.										
up	Set of characters used to move up one resolution unit (Vert) in plot mode, if any.										
down	Set of characters used to move down one resolution unit (Vert) in plot mode, if any.										
right	Set of characters used to move right one resolution unit (Hor) in plot mode, if any.										
left	Set of characters used to move left one resolution unit (Hor) in plot mode, if any.										
codetab	Definition of characters needed to print an nroff character on the terminal. The first byte is the number of character units (Char) needed to hold the character; that is, \001 is one unit wide, \002 is two units wide, etc. The high-order bit (0200) is on if the character is to be underlined in underline mode (.ul). The rest of the bytes are the characters used to produce the character in question. If the character has the sign (0200) bit on, it is a code to move the terminal in plot mode. It is encoded as: <table> <tr> <td>0100 bit on</td> <td>vertical motion.</td> </tr> <tr> <td>0100 bit off</td> <td>horizontal motion.</td> </tr> <tr> <td>040 bit on</td> <td>negative (up or left) motion.</td> </tr> <tr> <td>040 bit off</td> <td>positive (down or right) motion.</td> </tr> <tr> <td>037 bits</td> <td>number of such motions to make.</td> </tr> </table>	0100 bit on	vertical motion.	0100 bit off	horizontal motion.	040 bit on	negative (up or left) motion.	040 bit off	positive (down or right) motion.	037 bits	number of such motions to make.
0100 bit on	vertical motion.										
0100 bit off	horizontal motion.										
040 bit on	negative (up or left) motion.										
040 bit off	positive (down or right) motion.										
037 bits	number of such motions to make.										

zzz A zero terminator at the end.

All quantities which are in units of fractions of an inch should be expressed as '**INCH**num*/*denom***', where *num* and *denom* are respectively the numerator and denominator of the fraction; that is, 1/48 of an inch would be written as '**INCH/48**'.

If any sequence of characters does not pertain to the output device, that sequence should be given as a null string.

The following is a sample **codetab** encoding.

"\001 " ,	/*space*/
"\001!" ,	/*!*/
"\001\"" ,	/*\"*/
"\001#" ,	/*#*/
"\001\$" ,	/*\$*/

"\001%",	/*%*/
"\001&",	/*&*/
"\001'",	/*'*/
"\001(",	/*(*/
"\001)",	/*)*/
"\001*",	/****/
"\001+",	/*+*/
"\001,",	/*,*/*
"\001-",	/*-*/
"\001.",	/*.*/*
"\001/",	/*/*/*
"\2010",	/*0*/
"\2011",	/*1*/
"\2012",	/*2*/
"\2013",	/*3*/
"\2014",	/*4*/
"\2015",	/*5*/
"\2016",	/*6*/
"\2017",	/*7*/
"\2018",	/*8*/
"\2019",	/*9*/
"\001:",	/*:*/
"\001;",	/*;*/
"\001<",	/*<*/
"\001=",	/*=*/
"\001>",	/*>*/
"\001?",	/*?*/
"\001@",	/*@*/
"\201A",	/*A*/
"\201B",	/*B*/
"\201C",	/*C*/
"\201D",	/*D*/
"\201E",	/*E*/
"\201F",	/*F*/
"\201G",	/*G*/
"\201H",	/*H*/
"\201I",	/*I*/
"\201J",	/*J*/
"\201K",	/*K*/
"\201L",	/*L*/
"\201M",	/*M*/
"\201N",	/*N*/
"\201O",	/*O*/
"\201P",	/*P*/
"\201Q",	/*Q*/
"\201R",	/*R*/
"\201S",	/*S*/
"\201T",	/*T*/
"\201U",	/*U*/
"\201V",	/*V*/
"\201W",	/*W*/
"\201X",	/*X*/
"\201Y",	/*Y*/

"\201Z",	/*Z*/
"\001[",	/*[*/
"\001\\",	/**/
"\001]",	/*]*/
"\001^",	/*^*/
"\001_",	/*_*/
"\001'",	/*'*/
"\201a",	/*a*/
"\201b",	/*b*/
"\201c",	/*c*/
"\201d",	/*d*/
"\201e",	/*e*/
"\201f",	/*f*/
"\201g",	/*g*/
"\201h",	/*h*/
"\201i",	/*i*/
"\201j",	/*j*/
"\201k",	/*k*/
"\201l",	/*l*/
"\201m",	/*m*/
"\201n",	/*n*/
"\201o",	/*o*/
"\201p",	/*p*/
"\201q",	/*q*/
"\201r",	/*r*/
"\201s",	/*s*/
"\201t",	/*t*/
"\201u",	/*u*/
"\201v",	/*v*/
"\201w",	/*w*/
"\201x",	/*x*/
"\201y",	/*y*/
"\201z",	/*z*/
"\001{",	/*{*/
"\001 ",	/* */
"\001}",	/*}*/
"\001~",	/*~*/
"\000\0",	/*narrow sp*/
"\001-",	/*hyphen*/
"\001\016Z\017",	/*bullet*/
"\002[",	/*square*/
"\002--",	/*3/4 em dash*/
"\001_",	/*rule*/
"\0031/4",	/*1/4*/
"\0031/2",	/*1/2*/
"\0033/4",	/*3/4*/
"\001-",	/*minus*/
"\202fi",	/*fi*/
"\202ff",	/*ff*/
"\202fff",	/*fff*/
"\203ffi",	/*ffi*/
"\203fffi",	/*fffi*/
"\001\016p\017",	/*degree*/

"\001\b\342-\302",	/*dagger*/
"\001\301s\343s\302",	/*section*/
"\001'",	/*foot mark*/
"\001\033Z",	/*acute accent*/
"\001'",	/*grave accent*/
"\001_",	/*underrule*/
"\001/",	/*long slash*/
"\000\0",	/*half narrow space*/
"\001 ",	/*unpaddable space*/
"\001\016A\017",	/*alpha*/
"\001\016B\017",	/*beta*/
"\001\016C\017",	/*gamma*/
"\001\016D\017",	/*delta*/
"\001\016E\017",	/*epsilon*/
"\001\016F\017",	/*zeta*/
"\001\016G\017",	/*eta*/
"\001\016H\017",	/*theta*/
"\001\016I\017",	/*iota*/
"\001\016J\017",	/*kappa*/
"\001\016K\017",	/*lambda*/
"\001\016L\017",	/*mu*/
"\001\016M\017",	/*nu*/
"\001\016N\017",	/*xi*/
"\001\016O\017",	/*omicron*/
"\001\016P\017",	/*pi*/
"\001\016Q\017",	/*rho*/
"\001\016R\017",	/*sigma*/
"\001\016S\017",	/*tau*/
"\001\016T\017",	/*upsilon*/
"\001\016U\017",	/*phi*/
"\001\016V\017",	/*chi*/
"\001\016W\017",	/*psi*/
"\001\016X\017",	/*omega*/
"\001\016#\017",	/*Gamma*/
"\001\016\$\017",	/*Delta*/
"\001\016(\017",	/*Theta*/
"\001\016+\017",	/*Lambda*/
"\001\016.\017",	/*Xi*/
"\001\0160\017",	/*Pi*/
"\001\0169\017",	/*Sigma*/
"\000",	/**/
"\001\0164\017",	/*Upsilon*/
"\001\0165\017",	/*Phi*/
"\001\0167\017",	/*Psi*/
"\001\0168\017",	/*Omega*/
"\001\016[\017",	/*square root*/
"\001\016Y\017",	/*(ts yields script-l*/
"\001\016k\017",	/*root en*/
"\001>\b_",	/*>=*/
"\001<\b_",	/*<=*/
"\001=\b_",	/*identically equal*/
"\001-",	/*equation minus*/
"\001\016o\017",	/*approx =*/

"\001\016n\017",	/*approximates*/
"\001=\b/",	/*not equal*/
"\002-\242-\202>",	/*right arrow*/
"\002<\b\202-\242\200-",	/*left arrow*/
"\001\b^",	/*up arrow*/
"\001\b\302v\342",	/*down arrow*/
"\001=",	/*equation equal*/
"\001\016\017",	/*multiply*/
"\001\016)\017",	/*divide*/
"\001\016j\017",	/*plus-minus*/
"\001\243\203\203\243",	/*cup (union)*/
"\001\243\203\351\311\203\243",	/*cap (intersection)*/
"\001\243\203\302-\345-\303",	/*subset of*/
"\001\302-\345-\303\203\243",	/*superset of*/
"\001\b\243\203\302-\345-\303",	/*improper subset*/
"\001\b\302-\345-\303\203\243",	/*improper superset*/
"\001\016^\017",	/*infinity*/
"\001\200o\201\301^\241\341^\241\341^\201\301",	/*partial derivative*/
"\001\016:\017",	/*gradient*/
"\001\200-\202\341,\301\242",	/*not*/
"\001\016?\017",	/*integral sign*/
"\002o\242c\202",	/*proportional to*/
"\001O\b/",	/*empty set*/
"\001<\b\341-\302",	/*member of*/
"\001+",	/*equation plus*/
"\003(R)",	/*registered*/
"\003(C)",	/*copyright*/
"\001 ",	/*box rule */
"\001\033Y",	/*cent sign*/
"\001\b\342=\302",	/*double dagger*/
"\002=>",	/*right hand*/
"\002<=",	/*left hand*/
"\001*",	/*math * */
"\001\0162\017",	/*\ (bs yields small sigma)*/
"\001 ",	/*or (was star)*/
"\001O",	/*circle*/
"\001 ",	/*left top of big brace*/
"\001 ",	/*left bot of big brace*/
"\001 ",	/*right top of big brace*/
"\001 ",	/*right bot of big brace*/
"\001\016]\017",	/*left center of big brace*/
"\001\016\\017",	/*right center of big brace*/
"\001 ",	/*bold vertical*/
"\001 ",	/*left floor (lb of big bracket)*/
"\001 ",	/*right floor (rb of big bracket)*/
"\001 ",	/*left ceiling (lt of big bracket)*/
"\001 ",	/*right ceiling (rt of big bracket)*/

FILES

/usr/lib/term/tabname
/usr/lib/term/README

driving tables
list of terminals supported by nroff(1)

SEE ALSO

nroff(1), ttcompat(4M)

NAME

term – format of compiled term file

SYNOPSIS

term

DESCRIPTION

Compiled terminfo descriptions are placed under the directory `/usr/share/lib/terminfo`. In order to avoid a linear search of a huge system directory, a two-level scheme is used: `/usr/share/lib/terminfo/c/name` where *name* is the name of the terminal, and *c* is the first character of *name*. Thus, *act4* can be found in the file `/usr/share/lib/terminfo/a/act4`. Synonyms for the same terminal are implemented by multiple links to the same compiled file.

The format has been chosen so that it will be the same on all hardware. An 8 or more bit byte is assumed, but no assumptions about byte ordering or sign extension are made.

The compiled file is created with the `tic(8V)` program, and read by the routine `setupterm` (see `curses(3V)`). Both of these pieces of software are part of `curses(3V)`. The file is divided into six parts:

- the header,
- terminal names,
- boolean flags,
- numbers,
- strings,
- and
- string table.

The header section begins the file. This section contains six short integers in the format described below. These integers are:

- (1) the magic number (octal 0432);
- (2) the size, in bytes, of the names section;
- (3) the number of bytes in the boolean section;
- (4) the number of short integers in the numbers section;
- (5) the number of offsets (short integers) in the strings section;
- (6) the size, in bytes, of the string table.

Short integers are stored in two 8-bit bytes. The first byte contains the least significant 8 bits of the value, and the second byte contains the most significant 8 bits. (Thus, the value represented is $256 \times \text{second} + \text{first}$.) The value `-1` is represented by `0377, 0377`, other negative values are illegal. The `-1` generally means that a capability is missing from this terminal. Note: this format corresponds to the hardware of the VAX and PDP-11. Machines where this does not correspond to the hardware read the integers as two bytes and compute the result.

The terminal names section comes next. It contains the first line of the terminfo description, listing the various names for the terminal, separated by the `'|'` character. The section is terminated with an ASCII NUL character.

The boolean flags have one byte for each flag. This byte is either 0 or 1 as the flag is present or absent. The capabilities are in the same order as the file `<term.h>`.

Between the boolean section and the number section, a NULL byte will be inserted, if necessary, to ensure that the number section begins on an even byte. All short integers are aligned on a short word boundary.

The numbers section is similar to the flags section. Each capability takes up two bytes, and is stored as a short integer. If the value represented is `-1`, the capability is taken to be missing.

The strings section is also similar. Each capability is stored as a short integer, in the format above. A value of `-1` means the capability is missing. Otherwise, the value is taken as an offset from the beginning of the string table. Special characters in `^X` or `\c` notation are stored in their interpreted form, not the printing representation. Padding information `$<nn>` and parameter information `%x` are stored intact in uninterpreted form.

The final section is the string table. It contains all the values of string capabilities referenced in the string section. Each string is NULL terminated.

Note: it is possible for `setupterm` to expect a different set of capabilities than are actually present in the file. Either the database may have been updated since `setupterm` has been recompiled (resulting in extra unrecognized entries in the file) or the program may have been recompiled more recently than the database was updated (resulting in missing entries). The routine `setupterm` must be prepared for both possibilities — this is why the numbers and sizes are included. Also, new capabilities must always be added at the end of the lists of boolean, number, and string capabilities.

As an example, an octal dump of the description for the Microterm ACT 4 is included:

```
microterm|act4|microterm act iv,
cr=^M, cud1=^J, ind=^J, bel=^G, am, cub1=^H,
ed=^_, el=^^, clear=^L, cup=^T%p1%c%p2%c,
cols#80, lines#24, cuf1=^X, cuu1=^Z, home=^],

000 032 001      \0 025 \0 \b \0 212 \0 " \0 m i c r
020 o t e r m | a c t 4 | m i c r o
040 t e r m      a c t      i v \0 \0 001 \0 \0
060 \0 \0 \0 \0 \0 \0 \0 \0 \0 \0 \0 \0 \0 \0 \0
100 \0 \0 P \0 377 377 030 \0 377 377 377 377 377 377 377
120 377 377 377 377 \0 \0 002 \0 377 377 377 377 004 \0 006 \0
140 \b \0 377 377 377 377 \n \0 026 \0 030 \0 377 377 032 \0
160 377 377 377 377 034 \0 377 377 036 \0 377 377 377 377 377
200 377 377 377 377 377 377 377 377 377 377 377 377 377 377
*
520 377 377 377 377      \0 377 377 377 377 377 377 377 377 377
540 377 377 377 377 377 377 007 \0 \r \0 \f \0 036 \0 037 \0
560 024 % p 1 % c % p 2 % c \0 \n \0 035 \0
600 \b \0 030 \0 032 \0 \n \0
```

Some limitations: total compiled entries cannot exceed 4096 bytes. The name field cannot exceed 128 bytes.

FILES

`/usr/share/lib/terminfo/**`
 compiled terminal capability data base

SEE ALSO

`curses(3V)`, `terminfo(5V)`, `tic(8V)`

NAME

termcap – terminal capability data base

DESCRIPTION

termcap is a data base describing the capabilities of terminals. Terminals are described in **termcap** source descriptions by giving a set of capabilities which they have, by describing how operations are performed, by describing padding requirements, and by specifying initialization sequences. This database is used by applications programs such as **vi(1)**, and libraries such as **curses(3X)**, so they can work with a variety of terminals without changes to the programs.

Each **termcap** entry consist of a number of colon-separated (:) fields. The first field for each terminal lists the various names by which it is known, separated by bar (|) characters. The first name is always two characters long, and is used by older (version 6) systems (which store the terminal type in a 16-bit word in a system-wide database). The second name given is the most common abbreviation for the terminal (this is the one to which the environment variable **TERM** would normally be set). The last name should fully identify the terminal's make and model. All other names are taken as synonyms for the initial terminal name. All names but the first and last should be in lower case and contain no blanks; the last name may well contain upper case and blanks for added readability.

Terminal names (except for the last, verbose entry) should be chosen using the following conventions:

- The particular piece of hardware making up the terminal should have a root name chosen; for example, for the Hewlett-Packard 2621, **hp2621**. This name should not contain hyphens.
- Modes that the hardware can be in or user preferences should be indicated by appending a hyphen and an indicator of the mode. Thus, a **vt100** in 132-column mode would be given as: **vt100-w**. The following suffixes should be used where possible:

<i>Suffix</i>	<i>Meaning</i>	<i>Example</i>
-w	wide mode (more than 80 columns)	vt100-w
-am	with automatic margins (usually default)	vt100-am
-nam	without automatic margins	vt100-nam
-n	number of lines on the screen	aaa-60
-na	no arrow keys (leave them in local)	concept100-na
-np	number of pages of memory	concept100-4p
-rv	reverse video	concept100-rv

Terminal entries may continue onto multiple lines by giving a \ as the last character of a line, and empty fields may be included for readability (here between the last field on a line and the first field on the next). Comments may be included on lines beginning with #.

Types of Capabilities

Terminal capabilities each have a two-letter code, and are of three types:

<i>boolean</i>	These indicate particular features of the terminal. For instance, an entry for a terminal that has automatic margins (an automatic RETURN and LINEFEED when the end of a line is reached) would contain a field with the boolean capability am .
<i>numeric</i>	These give the size of the display of some other attribute. Numeric capabilities are followed by the character '#', and a number. An entry for a terminal with an 80-column display would have a field containing co#80 .
<i>string</i>	These indicate the character sequences used to perform particular terminal operations. String-valued capabilities, such as ce (clear-to-end-of-line sequence) are given by the two-letter code, followed by the character '=', and a string (which ends at the following : field delimiter).

A delay factor, in milliseconds may appear after the '='. Padding characters are supplied by **tputs** after the remainder of the string is sent. The delay can be either a number, or a number followed by the character '*', which indicates that the proportional padding is required, in which case the number given is the

amount of padding for each line affected by an operation using that capability. (In the case of an insert-character operation, the factor is still the number of *lines* affected; this is always 1 unless the terminal has *in* and the software uses it.)

When a *** is specified, it is sometimes useful to give a delay of the form 3.5 to specify a delay per line to tenths of milliseconds. (Only one decimal place is allowed.)

Comments

To comment-out a capability field, insert a '.' (period) as the first character in that field (following the :).

Escape Sequence Codes

A number of escape sequences are provided in the string-valued capabilities for easy encoding of characters there:

<code>\E</code>	maps to ESC
<code>^X</code>	maps to CTRL- <i>X</i> for any appropriate character <i>X</i>
<code>\n</code>	maps to LINEFEED
<code>\r</code>	maps to RETURN
<code>\t</code>	maps to TAB
<code>\b</code>	maps to BACKSPACE
<code>\f</code>	maps to FORMFEED

Finally, characters may be given as three octal digits after a backslash (for example, `\123`), and the characters ^ (caret) and \ (backslash) may be given as `\^` and `\\` respectively.

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 NUL character in a string capability it must be encoded as `\200`. (The routines that 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.)

Parameterized Strings

Cursor addressing and other strings requiring parameters are described by a parameterized string capability, with `printf(3S)`-like escapes (`%x`) in it; other characters are passed through unchanged. For example, to address the cursor, the `cm` capability is given, using two parameters: the row and column to move to. (Rows and columns are numbered from zero and refer to the physical screen visible to the user, not to any unseen memory. If the terminal has memory-relative cursor addressing, that can be indicated by an analogous `CM` capability.)

The `%` escapes have the following meanings:

<code>%%</code>	produce the character <code>%</code>
<code>%d</code>	output <i>value</i> as in <code>printf %d</code>
<code>%2</code>	output <i>value</i> as in <code>printf %2d</code>
<code>%3</code>	output <i>value</i> as in <code>printf %3d</code>
<code>%.<i>x</i></code>	output <i>value</i> as in <code>printf %c</code>
<code>%+<i>x</i></code>	add <i>x</i> to <i>value</i> , then do ' <code>%.</code> '
<code>%><i>xy</i></code>	if <i>value</i> > <i>x</i> then add <i>y</i> , no output
<code>%r</code>	reverse order of two parameters, no output
<code>%i</code>	increment by one, no output
<code>%n</code>	exclusive-or all parameters with 0140 (Datamedia 2500)
<code>%B</code>	BCD ($16 * (\text{value} / 10) + (\text{value} \% 10)$), no output
<code>%D</code>	Reverse coding ($\text{value} - 2 * (\text{value} \% 16)$), no output (Delta Data)

Consider the Hewlett-Packard 2645, which, to get to row 3 and column 12, needs to be sent `\E&a12c03Y` padded for 6 milliseconds. Note: the order of the row and column coordinates is reversed here and that the row and column are sent as two-digit integers. Thus its `cm` capability is '`:cm=6\E&%r%2c%2Y:`'. Terminals that use '`%.`' need to be able to backspace the cursor (`\e`) and to move the cursor up one line on the screen (`\u`). This is necessary because it is not always safe to transmit `\n`, `^D`, and `\r`, as the system may change or discard them. (Programs using `termcap` must set terminal modes so that TAB characters are not

expanded, making `\t` safe to send. This turns out to be essential for the Ann Arbor 4080.)

A final example is the Lear Siegler ADM-3a, which offsets row and column by a blank character, thus it requires `:cm=\E=%+ %+:'`.

Row or column absolute cursor addressing can be given as single-parameter capabilities `ch` (horizontal position absolute) and `cv` (vertical position absolute). Sometimes these are shorter than the more general two-parameter sequence (as with the Hewlett-Packard 2645) and can be used in preference to `cm`. If there are parameterized local motions (for example, move n positions to the right) these can be given as `DO`, `LE`, `RI`, and `UP` with a single parameter indicating how many positions to move. These are primarily useful if the terminal does not have `cm`, such as the Tektronix 4025.

Delays

Certain capabilities control padding in the terminal driver. These are primarily needed by hardcopy terminals and are used by the `tset (1)` program to set terminal driver modes appropriately. Delays embedded in the capabilities `cr`, `sf`, `le`, `ff`, and `ta` will set the appropriate delay bits in the terminal driver. If `pb` (padding baud rate) is given, these values can be ignored at baud rates below the value of `pb`. For 4.2BSD `tset`, the delays are given as numeric capabilities `dC`, `dN`, `dB`, `dF`, and `dT` instead.

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 entries must not exceed 1024. The capabilities given before `tc` override those in the terminal type invoked by `tc`. A capability can be canceled by placing `xx@` to the left of the `tc` invocation, where `xx` is the capability. For example, the entry

```
hn|2621-nl:ks@:ke@:tc=2621:
```

defines a `2621-nl` that does not have the `ks` or `ke` capabilities, 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.

CAPABILITIES

The characters in the *Notes* field in the next table have the following meanings (more than one may apply to a capability):

- N** indicates numeric parameter(s)
- P** indicates that padding may be specified
- *** indicates that padding may be based on the number of lines affected
- o** indicates capability is obsolete

Obsolete capabilities have no `terminfo` equivalents, since they were considered useless, or are subsumed by other capabilities. New software should not rely on them.

<i>Name</i>	<i>Type</i>	<i>Notes</i>	<i>Description</i>
<code>!1</code>	<i>str</i>		sent by shifted save key
<code>!2</code>	<i>str</i>		sent by shifted suspend key
<code>!3</code>	<i>str</i>		sent by shifted undo key
<code>#1</code>	<i>str</i>		sent by shifted help key
<code>#2</code>	<i>str</i>		sent by shifted home key
<code>#3</code>	<i>str</i>		sent by shifted input key
<code>#4</code>	<i>str</i>		sent by shifted left-arrow key
<code>%0</code>	<i>str</i>		sent by redo key
<code>%1</code>	<i>str</i>		sent by help key
<code>%2</code>	<i>str</i>		sent by mark key
<code>%3</code>	<i>str</i>		sent by message key
<code>%4</code>	<i>str</i>		sent by move key
<code>%5</code>	<i>str</i>		sent by next-object key
<code>%6</code>	<i>str</i>		sent by open key
<code>%7</code>	<i>str</i>		sent by options key

%8	<i>str</i>		sent by previous-object key
%9	<i>str</i>		sent by print or copy key
%a	<i>str</i>		sent by shifted message key
%b	<i>str</i>		sent by shifted move key
%c	<i>str</i>		sent by shifted next-object key
%d	<i>str</i>		sent by shifted options key
%e	<i>str</i>		sent by shifted previous-object key
%f	<i>str</i>		sent by shifted print or copy key
%g	<i>str</i>		sent by shifted redo key
%h	<i>str</i>		sent by shifted replace key
%i	<i>str</i>		sent by shifted right-arrow key
%j	<i>str</i>		sent by shifted resume key
&0	<i>str</i>		sent by shifted cancel key
&1	<i>str</i>		sent by ref(erence) key
&2	<i>str</i>		sent by refresh key
&3	<i>str</i>		sent by replace key
&4	<i>str</i>		sent by restart key
&5	<i>str</i>		sent by resume key
&6	<i>str</i>		sent by save key
&7	<i>str</i>		sent by suspend key
&8	<i>str</i>		sent by undo key
&9	<i>str</i>		sent by shifted beg(inning) key
+0	<i>str</i>		sent by shifted find key
+1	<i>str</i>		sent by shifted cmd (command) key
+2	<i>str</i>		sent by shifted copy key
+3	<i>str</i>		sent by shifted create key
+4	<i>str</i>		sent by shifted delete-char key
+5	<i>str</i>		sent by shifted delete-line key
+6	<i>str</i>		sent by select key
+7	<i>str</i>		sent by shifted end key
+8	<i>str</i>		sent by shifted clear-line key
+9	<i>str</i>		sent by shifted exit key
5l	<i>bool</i>		printer will not echo on screen
@0	<i>str</i>		sent by find key
@1	<i>str</i>		sent by beg(inning) key
@2	<i>str</i>		sent by cancel key
@3	<i>str</i>		sent by close key
@4	<i>str</i>		sent by cmd (command) key
@5	<i>str</i>		sent by copy key
@6	<i>str</i>		sent by create key
@7	<i>str</i>		sent by end key
@8	<i>str</i>		sent by enter/send key (unreliable)
@9	<i>str</i>		sent by exit key
AL	<i>str</i>	(NP*)	add <i>n</i> new blank lines
CC	<i>str</i>		terminal settable command character in prototype
CM	<i>str</i>	(NP)	memory-relative cursor motion to row <i>m</i> , column <i>n</i>
DC	<i>str</i>	(NP*)	delete <i>n</i> characters
DL	<i>str</i>	(NP*)	delete <i>n</i> lines
DO	<i>str</i>	(NP*)	move cursor down <i>n</i> lines
EP	<i>bool</i>	(o)	even parity
F1-F9	<i>str</i>		sent by function keys 11-19
FA-FZ	<i>str</i>		sent by function keys 20-45
Fa-Fr	<i>str</i>		sent by function keys 46-63
HC	<i>bool</i>		cursor is hard to see
HD	<i>bool</i>	(o)	half-duplex
IC	<i>str</i>	(NP*)	insert <i>n</i> blank characters
K1	<i>str</i>		sent by keypad upper left
K2	<i>str</i>		sent by keypad center

K3	<i>str</i>		sent by keypad upper right
K4	<i>str</i>		sent by keypad lower left
K5	<i>str</i>		sent by keypad lower right
LC	<i>bool</i>	(<i>o</i>)	lower-case only
LE	<i>str</i>	(<i>NP</i>)	move cursor left <i>n</i> positions
LF	<i>str</i>	(<i>P</i>)	turn off soft labels
LO	<i>str</i>	(<i>P</i>)	turn on soft labels
MC	<i>str</i>	(<i>P</i>)	clear left and right soft margins
ML	<i>str</i>	(<i>P</i>)	set soft left margin
MR	<i>str</i>	(<i>P</i>)	set soft right margin
NL	<i>bool</i>	(<i>o</i>)	\n is NEWLINE, not LINEFEED
NP	<i>bool</i>		pad character does not exist
NR	<i>bool</i>		ti does not reverse te
NI	<i>num</i>		number of labels on screen (start at 1)
OP	<i>bool</i>	(<i>o</i>)	odd parity
RA	<i>str</i>	(<i>P</i>)	turn off automatic margins
RF	<i>str</i>		send next input character (for pty)
RI	<i>str</i>	(<i>NP</i>)	move cursor right <i>n</i> positions
RX	<i>str</i>	(<i>P</i>)	turn off xoff/xon handshaking
SA	<i>str</i>	(<i>P</i>)	turn on automatic margins
SF	<i>str</i>	(<i>NP*</i>)	scroll forward <i>n</i> lines
SR	<i>str</i>	(<i>NP*</i>)	scroll backward <i>n</i> lines
SX	<i>str</i>	(<i>P</i>)	turn on xoff/xon handshaking
UC	<i>bool</i>	(<i>o</i>)	upper-case only
UP	<i>str</i>	(<i>NP*</i>)	move cursor up <i>n</i> lines
XF	<i>str</i>		x-off character (default DC3)
XN	<i>str</i>		x-on character (default DC1)
ac	<i>str</i>		graphic character set pairs aAbBcC – def=VT100
ae	<i>str</i>	(<i>P</i>)	end alternate character set
al	<i>str</i>	(<i>P*</i>)	add new blank line
am	<i>bool</i>		terminal has automatic margins
as	<i>str</i>	(<i>P</i>)	start alternate character set
bc	<i>str</i>	(<i>o</i>)	backspace if not ^H
bl	<i>str</i>	(<i>P</i>)	audible signal (bell)
bs	<i>bool</i>	(<i>o</i>)	terminal can backspace with ^H
bt	<i>str</i>	(<i>P</i>)	back-tab
bw	<i>bool</i>		le (backspace) wraps from column 0 to last column
cb	<i>str</i>	(<i>P</i>)	clear to beginning of line, inclusive
cd	<i>str</i>	(<i>P*</i>)	clear to end of display
ce	<i>str</i>	(<i>P</i>)	clear to end of line
ch	<i>str</i>	(<i>NP</i>)	set cursor column (horizontal position)
cl	<i>str</i>	(<i>P*</i>)	clear screen and home cursor
cm	<i>str</i>	(<i>NP</i>)	screen-relative cursor motion to row <i>m</i> , column <i>n</i>
co	<i>num</i>		number of columns in a line
cr	<i>str</i>	(<i>P*</i>)	RETURN
cs	<i>str</i>	(<i>NP</i>)	change scrolling region to lines <i>m</i> through <i>n</i> (VT100)
ct	<i>str</i>	(<i>P</i>)	clear all tab stops
cv	<i>str</i>	(<i>NP</i>)	set cursor row (vertical position)
dB	<i>num</i>	(<i>o</i>)	milliseconds of bs delay needed (default 0)
dC	<i>num</i>	(<i>o</i>)	milliseconds of cr delay needed (default 0)
dF	<i>num</i>	(<i>o</i>)	milliseconds of ff delay needed (default 0)
dN	<i>num</i>	(<i>o</i>)	milliseconds of nl delay needed (default 0)
dT	<i>num</i>	(<i>o</i>)	milliseconds of horizontal tab delay needed (default 0)
dV	<i>num</i>	(<i>o</i>)	milliseconds of vertical tab delay needed (default 0)
da	<i>bool</i>		display may be retained above the screen
db	<i>bool</i>		display may be retained below the screen
dc	<i>str</i>	(<i>P*</i>)	delete character
dl	<i>str</i>	(<i>P*</i>)	delete line

dm	<i>str</i>		enter delete mode
do	<i>str</i>		down one line
ds	<i>str</i>		disable status line
eA	<i>str</i>	(P)	enable graphic character set
ec	<i>str</i>	(NP)	erase <i>n</i> characters
ed	<i>str</i>		end delete mode
ei	<i>str</i>		end insert mode
eo	<i>bool</i>		can erase overstrikes with a blank
es	<i>bool</i>		escape can be used on the status line
ff	<i>str</i>	(P*)	hardcopy terminal page eject
fs	<i>str</i>		return from status line
gn	<i>bool</i>		generic line type (for example dialup, switch)
hc	<i>bool</i>		hardcopy terminal
hd	<i>str</i>		half-line down (forward 1/2 linefeed)
ho	<i>str</i>	(P)	home cursor
hs	<i>bool</i>		has extra "status line"
hu	<i>str</i>		half-line up (reverse 1/2 linefeed)
hz	<i>bool</i>		cannot print '~'s (Hazeltine)
iI	<i>str</i>		terminal initialization string (terminfo only)
i3	<i>str</i>		terminal initialization string (terminfo only)
iP	<i>str</i>		pathname of program for initialization (terminfo only)
ic	<i>str</i>	(P*)	insert character
if	<i>str</i>		name of file containing initialization string
im	<i>str</i>		enter insert mode
in	<i>bool</i>		insert mode distinguishes nulls
ip	<i>str</i>	(P*)	insert pad after character inserted
is	<i>str</i>		terminal initialization string
it	<i>num</i>		tab stops initially every <i>n</i> positions
k0-k9	<i>str</i>		sent by function keys 0-9
k;	<i>str</i>		sent by function key 10
kA	<i>str</i>		sent by insert-line key
kB	<i>str</i>		sent by back-tab key
kC	<i>str</i>		sent by clear-screen or erase key
kD	<i>str</i>		sent by delete-character key
kE	<i>str</i>		sent by clear-to-end-of-line key
kF	<i>str</i>		sent by scroll-forward/down key
kH	<i>str</i>		sent by home-down key
kI	<i>str</i>		sent by insert-character or enter-insert-mode key
kL	<i>str</i>		sent by delete-line key
kM	<i>str</i>		sent by insert key while in insert mode
kN	<i>str</i>		sent by next-page key
kP	<i>str</i>		sent by previous-page key
kR	<i>str</i>		sent by scroll-backward/up key
kS	<i>str</i>		sent by clear-to-end-of-screen key
kT	<i>str</i>		sent by set-tab key
ka	<i>str</i>		sent by clear-all-tabs key
kb	<i>str</i>		sent by backspace key
kd	<i>str</i>		sent by down-arrow key
ke	<i>str</i>		out of "keypad transmit" mode
kh	<i>str</i>		sent by home key
kl	<i>str</i>		sent by left-arrow key
km	<i>bool</i>		has a "meta" key (shift, sets parity bit)
kn	<i>num</i>	(o)	number of function (k0-k9) keys (default 0)
ko	<i>str</i>	(o)	termcap entries for other non-function keys
kr	<i>str</i>		sent by right-arrow key
ks	<i>str</i>		put terminal in "keypad transmit" mode
kt	<i>str</i>		sent by clear-tab key
ku	<i>str</i>		sent by up-arrow key

10-19	<i>str</i>		labels on function keys 0-9 if not f0-f9
la	<i>str</i>		label on function key 10 if not f10
le	<i>str</i>	(P)	move cursor left one position
lh	<i>num</i>		number of rows in each label
li	<i>num</i>		number of lines on screen or page
ll	<i>str</i>		last line, first column
lm	<i>num</i>		lines of memory if > ll (0 means varies)
lw	<i>num</i>		number of columns in each label
ma	<i>str</i>	(o)	arrow key map (used by vi version 2 only)
mb	<i>str</i>		turn on blinking attribute
md	<i>str</i>		turn on bold (extra bright) attribute
me	<i>str</i>		turn off all attributes
mh	<i>str</i>		turn on half-bright attribute
mi	<i>bool</i>		safe to move while in insert mode
mk	<i>str</i>		turn on blank attribute (characters invisible)
ml	<i>str</i>	(o)	memory lock on above cursor
mm	<i>str</i>		turn on "meta mode" (8th bit)
mo	<i>str</i>		turn off "meta mode"
mp	<i>str</i>		turn on protected attribute
mr	<i>str</i>		turn on reverse-video attribute
ms	<i>bool</i>		safe to move in standout modes
mu	<i>str</i>	(o)	memory unlock (turn off memory lock)
nc	<i>bool</i>	(o)	no correctly-working cr (Datamedia 2500, Hazeltine 2000)
nd	<i>str</i>		non-destructive space (cursor right)
nl	<i>str</i>	(o)	NEWLINE character if not
ns	<i>bool</i>	(o)	terminal is a CRT but does not scroll
nw	<i>str</i>	(P)	NEWLINE (behaves like cr followed by do)
nx	<i>bool</i>		padding will not work, xoff/xon required
os	<i>bool</i>		terminal overstrikes
pO	<i>str</i>	(N)	turn on the printer for <i>n</i> bytes
pb	<i>num</i>		lowest baud where delays are required
pc	<i>str</i>		pad character (default NUL)
pf	<i>str</i>		turn off the printer
pk	<i>str</i>		program function key <i>n</i> to type string <i>s</i> (terminfo only)
pl	<i>str</i>		program function key <i>n</i> to execute string <i>s</i> (terminfo only)
pn	<i>str</i>	(NP)	program label <i>n</i> to show string <i>s</i> (terminfo only)
po	<i>str</i>		turn on the printer
ps	<i>str</i>		print contents of the screen
pt	<i>bool</i>	(o)	has hardware tab stops (may need to be set with is)
px	<i>str</i>		program function key <i>n</i> to transmit string <i>s</i> (terminfo only)
r1	<i>str</i>		reset terminal completely to sane modes (terminfo only)
r2	<i>str</i>		reset terminal completely to sane modes (terminfo only)
r3	<i>str</i>		reset terminal completely to sane modes (terminfo only)
rP	<i>str</i>	(P)	like lp but when in replace mode
rc	<i>str</i>	(P)	restore cursor to position of last se
rf	<i>str</i>		name of file containing reset string
ri	?		unknown at present
rp	<i>str</i>	(NP*)	repeat character <i>c</i> <i>n</i> times
rs	<i>str</i>		reset terminal completely to sane modes
sa	<i>str</i>	(NP)	define the video attributes (9 parameters)
sc	<i>str</i>	(P)	save cursor position
se	<i>str</i>		end standout mode
sf	<i>str</i>	(P)	scroll text up
sg	<i>num</i>		number of garbage chars left by so or se (default 0)
so	<i>str</i>		begin standout mode
sr	<i>str</i>	(P)	scroll text down
st	<i>str</i>		set a tab stop in all rows, current column
ta	<i>str</i>	(P)	move cursor to next 8-position hardware tab stop

tc	<i>str</i>		entry of similar terminal – must be last
te	<i>str</i>		string to end programs that use termcap
ti	<i>str</i>		string to begin programs that use termcap
ts	<i>str</i>	(<i>N</i>)	go to status line, column <i>n</i>
uc	<i>str</i>		underscore one character and move past it
ue	<i>str</i>		end underscore mode
ug	<i>num</i>		number of garbage chars left by us or ue (default 0)
ul	<i>bool</i>		underline character overstrikes
up	<i>str</i>		upline (cursor up)
us	<i>str</i>		start underscore mode
vb	<i>str</i>		visible bell (must not move cursor)
ve	<i>str</i>		make cursor appear normal (undo vs/vi)
vi	<i>str</i>		make cursor invisible
vs	<i>str</i>		make cursor very visible
vt	<i>num</i>		virtual terminal number (not supported on all systems)
wi	<i>str</i>	(<i>N</i>)	set current window to lines <i>i</i> through <i>j</i> , columns <i>m</i> through <i>n</i>
ws	<i>num</i>		number of columns in status line
xb	<i>bool</i>		Beehive (f1=ESC, f2=^C)
xn	<i>bool</i>		NEWLINE ignored after 80 cols (Concept)
xo	<i>bool</i>		terminal uses xoff/xon handshaking
xr	<i>bool</i>	(<i>o</i>)	RETURN acts like ce cr nl (Delta Data)
xs	<i>bool</i>		standout not erased by overwriting (Hewlett-Packard)
xt	<i>bool</i>		TAB characters destructive, magic so char (Telera 1061)
xx	<i>bool</i>	(<i>o</i>)	Tektronix 4025 insert-line

ENVIRONMENT

If the environment variable **TERMCAP** contains an absolute pathname, programs look to that file for terminal descriptions, rather than **/usr/share/lib/termcap**. If the value of this variable is in the form of a **termcap** entry, programs use that value for the terminal description.

FILES

/usr/share/lib/termcap file containing terminal descriptions

SEE ALSO

ex(1), **more(1)**, **tset(1)**, **ul(1)**, **vi(1)**, **curses(3X)**, **printf(3S)**, **termcap(3X)**, **term(5V)**, **terminfo(5V)**

System and Network Administration

CAVEATS AND BUGS

UNIX System V uses **terminfo(5V)** rather than **termcap**. SunOS supports either **termcap** or **terminfo(5V)** terminal databases, depending on whether you link with the **termcap(3X)** or **curses(3V)** libraries. Transitions between the two should be relatively painless if capabilities flagged as “obsolete” are avoided.

vi 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 NEWLINE characters) may not exceed 1024.

Not all programs support all entries.

NAME

terminfo – terminal capability data base

SYNOPSIS

`/usr/share/lib/terminfo/?/*`

AVAILABILITY

This database is available with the *System V* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

terminfo is a compiled database (see **tic(8V)**) describing the capabilities of terminals. Terminals are described in **terminfo** source descriptions by giving a set of capabilities which they have, by describing how operations are performed, by describing padding requirements, and by specifying initialization sequences. This database is used by applications programs, and by libraries such as **curses(3V)**, so they can work with a variety of terminals without changes to the programs. To obtain the source description for a terminal, use the **-I** option of **infocmp(8V)**.

Entries in **terminfo** source files consist of a number of comma-separated fields. White space after each comma is ignored. The first line of each terminal description in the **terminfo** database gives the name by which **terminfo** knows the terminal, separated by pipe (|) characters. The first name given is the most common abbreviation for the terminal (this is the one to which the environment variable **TERM** would normally be set), the last name given should be a long name fully identifying the terminal, and all others are understood as synonyms for the terminal name. All names but the last should contain no blanks; the last name may contain blanks for readability.

Terminal names (except for the last, verbose entry) should be chosen using the following conventions:

- The particular piece of hardware making up the terminal should have a root name chosen; for example, for the Hewlett-Packard 2621, **hp2621**. This name should not contain hyphens.
- Modes that the hardware can be in or user preferences should be indicated by appending a hyphen and an indicator of the mode. Thus, a **vt100** in 132-column mode would be given as: **vt100-w**. The following suffixes should be used where possible:

<i>Suffix</i>	<i>Meaning</i>	<i>Example</i>
-w	wide mode (more than 80 columns)	vt100-w
-am	with automatic margins (usually default)	vt100-am
-nam	without automatic margins	vt100-nam
-n	number of lines on the screen	aaa-60
-na	no arrow keys (leave them in local)	concept100-na
-np	number of pages of memory	concept100-4p
-rv	reverse video	concept100-rv

CAPABILITIES

In the table below, the **Variable** is the name by which the C programmer (at the **terminfo** level) accesses the capability. The **capname** is the short name for this variable used in the text of the database. It is used by a person updating the database and by the **tput(1V)** command when asking what the value of the capability is for a particular terminal. The **Termcap Code** is a two-letter code that corresponds to the old **termcap** capability name.

Capability names have no hard length limit, but an informal limit of 5 characters has been adopted to keep them short. Whenever possible, names are chosen to be the same as or similar to the ANSI X3.64-1979 standard. Semantics are also intended to match those of the specification.

All string capabilities listed below may have padding specified, with the exception of those used for input. Input capabilities, listed under the Strings section in the table below, have names beginning with 'key_'. The following indicators may appear at the end of the Description for a variable.

- (G) indicates that the string is passed through `tparam()` with parameters (parms) as given (#).
- (*) indicates that padding may be based on the number of lines affected.
- (#_i) indicates the *i*th parameter.

<i>Variable</i>	<i>Capname</i>	<i>Termcap</i>	<i>Description</i>
<i>Boolean</i>			
<code>auto_left_margin</code>	<code>bw</code>	<code>bw</code>	<code>cub1</code> wraps from column 0 to last column
<code>auto_right_margin</code>	<code>am</code>	<code>am</code>	Terminal has automatic margins
<code>no_esc_ctlc</code>	<code>xsb</code>	<code>xb</code>	Beehive (f1=ESC, f2=^C)
<code>ceol_standout_glitch</code>	<code>xhp</code>	<code>xs</code>	Standout not erased by overwriting (Hewlett-Packard)
<code>eat_newline_glitch</code>	<code>xenl</code>	<code>xn</code>	NEWLINE ignored after 80 cols (Concept)
<code>erase_overstrike</code>	<code>eo</code>	<code>eo</code>	Can erase overstrikes with a blank
<code>generic_type</code>	<code>gn</code>	<code>gn</code>	Generic line type (for example, dialup, switch).
<code>hard_copy</code>	<code>hc</code>	<code>hc</code>	Hardcopy terminal
<code>hard_cursor</code>	<code>chts</code>	<code>HC</code>	Cursor is hard to see.
<code>has_meta_key</code>	<code>km</code>	<code>km</code>	Has a meta key (shift, sets parity bit)
<code>has_status_line</code>	<code>hs</code>	<code>hs</code>	Has extra "status line"
<code>insert_null_glitch</code>	<code>in</code>	<code>in</code>	Insert mode distinguishes nulls
<code>memory_above</code>	<code>da</code>	<code>da</code>	Display may be retained above the screen
<code>memory_below</code>	<code>db</code>	<code>db</code>	Display may be retained below the screen
<code>move_insert_mode</code>	<code>mir</code>	<code>mi</code>	Safe to move while in insert mode
<code>move_standout_mode</code>	<code>msgr</code>	<code>ms</code>	Safe to move in standout modes
<code>needs_xon_xoff</code>	<code>nxon</code>	<code>nx</code>	Padding will not work, xon/xoff required
<code>non_rev_rmcup</code>	<code>nrrmc</code>	<code>NR</code>	<code>smcup</code> does not reverse <code>rmcup</code>
<code>no_pad_char</code>	<code>npc</code>	<code>NP</code>	Pad character does not exist
<code>over_strike</code>	<code>os</code>	<code>os</code>	Terminal overstrikes on hard-copy terminal
<code>prtr_silent</code>	<code>mc5l</code>	<code>5l</code>	Printer will not echo on screen.
<code>status_line_esc_ok</code>	<code>eslok</code>	<code>es</code>	Escape can be used on the status line
<code>dest_tabs_magic_sms0</code>	<code>xt</code>	<code>xt</code>	Destructive TAB characters, magic <code>sms0</code> char (Telera 1061)
<code>tlide_glitch</code>	<code>hz</code>	<code>hz</code>	Hazeltine; cannot print tildes(^)
<code>transparent_underline</code>	<code>ul</code>	<code>ul</code>	Underline character overstrikes
<code>xon_xoff</code>	<code>xon</code>	<code>xo</code>	Terminal uses xon/xoff handshaking
<i>Number</i>			
<code>columns</code>	<code>cols</code>	<code>co</code>	Number of columns in a line
<code>init_tabs</code>	<code>it</code>	<code>it</code>	tab stops initially every # spaces.
<code>label_height</code>	<code>lh</code>	<code>lh</code>	Number of rows in each label
<code>label_width</code>	<code>lw</code>	<code>lw</code>	Number of cols in each label
<code>lines</code>	<code>lines</code>	<code>ll</code>	Number of lines on screen or page
<code>lines_of_memory</code>	<code>lm</code>	<code>lm</code>	Lines of memory if > lines; 0 means varies
<code>magic_cookie_glitch</code>	<code>xmc</code>	<code>sg</code>	Number blank chars left by <code>sms0</code> or <code>rmso</code>
<code>num_labels</code>	<code>nlab</code>	<code>Nl</code>	Number of labels on screen (start at 1)
<code>padding_baud_rate</code>	<code>pb</code>	<code>pb</code>	Lowest baud rate where padding needed
<code>virtual terminal</code>	<code>vt</code>	<code>vt</code>	Virtual terminal number (not supported on all systems)
<code>width_status_line</code>	<code>wsl</code>	<code>ws</code>	Number of columns in status line
<i>String</i>			
<code>acs_chars</code>	<code>acsc</code>	<code>ac</code>	Graphic charset pairs aAbBcC - def=VT100
<code>back_tab</code>	<code>cbt</code>	<code>bt</code>	Back tab
<code>bell</code>	<code>bel</code>	<code>bl</code>	Audible signal (bell)
<code>carriage_return</code>	<code>cr</code>	<code>cr</code>	RETURN (*)
<code>change_scroll_region</code>	<code>csr</code>	<code>cs</code>	Change to lines #1 through #2 (VT100) (G)

char_padding	rmp	rP	Like ip but when in replace mode
clear_all_tabs	tbc	ct	Clear all tab stops
clear_margins	mgc	MC	Clear left and right soft margins
clear_screen	clear	cl	Clear screen and home cursor (*)
clr_bol	el1	cb	Clear to beginning of line, inclusive
clr_eol	el	ce	Clear to end of line
clr_eos	ed	cd	Clear to end of display (*)
column_address	hpa	ch	Horizontal position absolute (G)
command_character	cmdch	CC	Terminal settable command char in prototype
cursor_address	cup	cm	Cursor motion to row #1 col #2 (G)
cursor_down	cuD1	do	Down one line
cursor_home	home	ho	Home cursor (if no cup)
cursor_invisible	civis	vi	Make cursor invisible
cursor_left	cub1	le	Move cursor left one SPACE.
cursor_mem_address	mrcup	CM	Memory relative cursor addressing (G)
cursor_normal	cnorm	ve	Make cursor appear normal (undo cvvis/civis)
cursor_right	cuf1	nd	Non-destructive space (cursor right)
cursor_to_ll	ll	ll	Last line, first column (if no cup)
cursor_up	cuu1	up	Upline (cursor up)
cursor_visible	cvvis	vs	Make cursor very visible
delete_character	dch1	dc	Delete character (*)
delete_line	dll	dl	Delete line (*)
dis_status_line	dsl	ds	Disable status line
down_half_line	hd	hd	Half-line down (forward 1/2 LINEFEED)
ena_acs	enacs	eA	Enable alternate char set
enter_alt_charset_mode	smacs	as	Start alternate character set
enter_am_mode	smam	SA	Turn on automatic margins
enter_blink_mode	blink	mb	Turn on blinking
enter_bold_mode	bold	md	Turn on bold (extra bright) mode
enter_ca_mode	smcup	tl	String to begin programs that use cup
enter_delete_mode	smdc	dm	Delete mode (enter)
enter_dim_mode	dim	mh	Turn on half-bright mode
enter_insert_mode	smir	im	Insert mode (enter);
enter_protected_mode	prot	mp	Turn on protected mode
enter_reverse_mode	rev	mr	Turn on reverse video mode
enter_secure_mode	invis	mk	Turn on blank mode (chars invisible)
enter_standout_mode	smso	so	Begin standout mode
enter_underline_mode	smul	us	Start underscore mode
enter_xon_mode	smxon	SX	Turn on xon/xoff handshaking
erase_chars	ech	ec	Erase #1 characters (G)
exit_alt_charset_mode	rmacs	ae	End alternate character set
exit_am_mode	rmam	RA	Turn off automatic margins
exit_attribute_mode	sgr0	me	Turn off all attributes
exit_ca_mode	rncup	te	String to end programs that use cup
exit_delete_mode	rmdc	ed	End delete mode
exit_insert_mode	rmir	ei	End insert mode;
exit_standout_mode	rmso	se	End standout mode
exit_underline_mode	rmul	ue	End underscore mode
exit_xon_mode	rmxon	RX	Turn off xon/xoff handshaking
flash_screen	flash	vb	Visible bell (must not move cursor)
form_feed	ff	ff	Hardcopy terminal page eject (*)
from_status_line	fsl	fs	Return from status line
init_1string	is1	i1	Terminal initialization string
init_2string	is2	is	Terminal initialization string
init_3string	is3	i3	Terminal initialization string
init_file	if	if	Name of initialization file containing is
init_prog	iprog	iP	Path name of program for init.
insert_character	ich1	ic	Insert character

insert_line	il1	al	Add new blank line (*)
insert_padding	ip	ip	Insert pad after character inserted (*)
key_a1	ka1	K1	KEY_A1, 0534, Upper left of keypad
key_a3	ka3	K3	KEY_A3, 0535, Upper right of keypad
key_b2	kb2	K2	KEY_B2, 0536, Center of keypad
key_backspace	kbs	kb	KEY_BACKSPACE, 0407, Sent by BACKSPACE key
key_beg	kbeg	@1	KEY_BEG, 0542, Sent by beg(inning) key
key_btab	kcbt	kB	KEY_BTAB, 0541, Sent by back-tab key
key_c1	kc1	K4	KEY_C1, 0537, Lower left of keypad
key_c3	kc3	K5	KEY_C3, 0540, Lower right of keypad
key_cancel	kcan	@2	KEY_CANCEL, 0543, Sent by cancel key
key_catab	ktbc	ka	KEY_CATAB, 0526, Sent by clear-all-tabs key
key_clear	kclr	kC	KEY_CLEAR, 0515, Sent by clear-screen or erase key
key_close	kclo	@3	KEY_CLOSE, 0544, Sent by close key
key_command	kcnd	@4	KEY_COMMAND, 0545, Sent by cmd (command) key
key_copy	kcpy	@5	KEY_COPY, 0546, Sent by copy key
key_create	kcrt	@6	KEY_CREATE, 0547, Sent by create key
key_ctab	kctab	kt	KEY_CTAB, 0525, Sent by clear-tab key
key_dc	kdch1	kD	KEY_DC, 0512, Sent by delete-character key
key_dl	kd11	kL	KEY_DL, 0510, Sent by delete-line key
key_down	kcud1	kd	KEY_DOWN, 0402, Sent by terminal down-arrow key
key_eic	krmir	kM	KEY_EIC, 0514, Sent by rmir or smir in insert mode
key_end	kend	@7	KEY_END, 0550, Sent by end key
key_enter	kent	@8	KEY_ENTER, 0527, Sent by enter/send key
key_eol	kel	kE	KEY_EOL, 0517, Sent by clear-to-end-of-line key
key_eos	ked	kS	KEY_EOS, 0516, Sent by clear-to-end-of-screen key
key_exit	kext	@9	KEY_EXIT, 0551, Sent by exit key
key_f0	kf0	k0	KEY_F(0), 0410, Sent by function key f0
key_f1	kf1	k1	KEY_F(1), 0411, Sent by function key f1
key_f2	kf2	k2	KEY_F(2), 0412, Sent by function key f2
key_f3	kf3	k3	KEY_F(3), 0413, Sent by function key f3
key_f4	kf4	k4	KEY_F(4), 0414, Sent by function key f4
key_f5	kf5	k5	KEY_F(5), 0415, Sent by function key f5
key_f6	kf6	k6	KEY_F(6), 0416, Sent by function key f6
key_f7	kf7	k7	KEY_F(7), 0417, Sent by function key f7
key_f8	kf8	k8	KEY_F(8), 0420, Sent by function key f8
key_f9	kf9	k9	KEY_F(9), 0421, Sent by function key f9
key_f10	kf10	k;	KEY_F(10), 0422, Sent by function key f10
key_f11	kf11	F1	KEY_F(11), 0423, Sent by function key f11
key_f12	kf12	F2	KEY_F(12), 0424, Sent by function key f12
key_f13	kf13	F3	KEY_F(13), 0425, Sent by function key f13
key_f14	kf14	F4	KEY_F(14), 0426, Sent by function key f14
key_f15	kf15	F5	KEY_F(15), 0427, Sent by function key f15
key_f16	kf16	F6	KEY_F(16), 0430, Sent by function key f16
key_f17	kf17	F7	KEY_F(17), 0431, Sent by function key f17
key_f18	kf18	F8	KEY_F(18), 0432, Sent by function key f18
key_f19	kf19	F9	KEY_F(19), 0433, Sent by function key f19
key_f20	kf20	FA	KEY_F(20), 0434, Sent by function key f20
key_f21	kf21	FB	KEY_F(21), 0435, Sent by function key f21
key_f22	kf22	FC	KEY_F(22), 0436, Sent by function key f22
key_f23	kf23	FD	KEY_F(23), 0437, Sent by function key f23
key_f24	kf24	FE	KEY_F(24), 0440, Sent by function key f24
key_f25	kf25	FF	KEY_F(25), 0441, Sent by function key f25
key_f26	kf26	FG	KEY_F(26), 0442, Sent by function key f26
key_f27	kf27	FH	KEY_F(27), 0443, Sent by function key f27
key_f28	kf28	FI	KEY_F(28), 0444, Sent by function key f28
key_f29	kf29	FJ	KEY_F(29), 0445, Sent by function key f29
key_f30	kf30	FK	KEY_F(30), 0446, Sent by function key f30

key_f31	kf31	FL	KEY_F(31), 0447, Sent by function key f31
key_f32	kf32	FM	KEY_F(32), 0450, Sent by function key f32
key_f33	kf33	FN	KEY_F(13), 0451, Sent by function key f13
key_f34	kf34	FO	KEY_F(34), 0452, Sent by function key f34
key_f35	kf35	FP	KEY_F(35), 0453, Sent by function key f35
key_f36	kf36	FQ	KEY_F(36), 0454, Sent by function key f36
key_f37	kf37	FR	KEY_F(37), 0455, Sent by function key f37
key_f38	kf38	FS	KEY_F(38), 0456, Sent by function key f38
key_f39	kf39	FT	KEY_F(39), 0457, Sent by function key f39
key_f40	kf40	FU	KEY_F(40), 0460, Sent by function key f40
key_f41	kf41	FV	KEY_F(41), 0461, Sent by function key f41
key_f42	kf42	FW	KEY_F(42), 0462, Sent by function key f42
key_f43	kf43	FX	KEY_F(43), 0463, Sent by function key f43
key_f44	kf44	FY	KEY_F(44), 0464, Sent by function key f44
key_f45	kf45	FZ	KEY_F(45), 0465, Sent by function key f45
key_f46	kf46	Fa	KEY_F(46), 0466, Sent by function key f46
key_f47	kf47	Fb	KEY_F(47), 0467, Sent by function key f47
key_f48	kf48	Fc	KEY_F(48), 0470, Sent by function key f48
key_f49	kf49	Fd	KEY_F(49), 0471, Sent by function key f49
key_f50	kf50	Fe	KEY_F(50), 0472, Sent by function key f50
key_f51	kf51	Ff	KEY_F(51), 0473, Sent by function key f51
key_f52	kf52	Fg	KEY_F(52), 0474, Sent by function key f52
key_f53	kf53	Fh	KEY_F(53), 0475, Sent by function key f53
key_f54	kf54	Fi	KEY_F(54), 0476, Sent by function key f54
key_f55	kf55	Fj	KEY_F(55), 0477, Sent by function key f55
key_f56	kf56	Fk	KEY_F(56), 0500, Sent by function key f56
key_f57	kf57	Fl	KEY_F(57), 0501, Sent by function key f57
key_f58	kf58	Fm	KEY_F(58), 0502, Sent by function key f58
key_f59	kf59	Fn	KEY_F(59), 0503, Sent by function key f59
key_f60	kf60	Fo	KEY_F(60), 0504, Sent by function key f60
key_f61	kf61	Fp	KEY_F(61), 0505, Sent by function key f61
key_f62	kf62	Fq	KEY_F(62), 0506, Sent by function key f62
key_f63	kf63	Fr	KEY_F(63), 0507, Sent by function key f63
key_find	kfnd	@0	KEY_FIND, 0552, Sent by find key
key_help	khlp	%1	KEY_HELP, 0553, Sent by help key
key_home	khome	kh	KEY_HOME, 0406, Sent by home key
key_ic	kich1	kI	KEY_IC, 0513, Sent by ins-char/enter ins-mode key
key_il	kill	kA	KEY_IL, 0511, Sent by insert-line key
key_left	kcub1	kl	KEY_LEFT, 0404, Sent by terminal left-arrow key
key_ll	kll	kH	KEY_LL, 0533, Sent by home-down key
key_mark	kmrk	%2	KEY_MARK, 0554, Sent by mark key
key_message	kmsg	%3	KEY_MESSAGE, 0555, Sent by message key
key_move	kmov	%4	KEY_MOVE, 0556, Sent by move key
key_next	knxt	%5	KEY_NEXT, 0557, Sent by next-object key
key_npage	knp	kN	KEY_NPAGE, 0522, Sent by next-page key
key_open	kopn	%6	KEY_OPEN, 0560, Sent by open key
key_options	kopt	%7	KEY_OPTIONS, 0561, Sent by options key
key_ppage	kpp	kP	KEY_PPAGE, 0523, Sent by previous-page key
key_previous	kprv	%8	KEY_PREVIOUS, 0562, Sent by previous-object key
key_print	kpnt	%9	KEY_PRINT, 0532, Sent by print or copy key
key_redo	krdo	%0	KEY_REDO, 0563, Sent by redo key
key_reference	kref	&1	KEY_REFERENCE, 0564, Sent by ref(erence) key
key_refresh	krfr	&2	KEY_REFRESH, 0565, Sent by refresh key
key_replace	krpl	&3	KEY_REPLACE, 0566, Sent by replace key
key_restart	krst	&4	KEY_RESTART, 0567, Sent by restart key
key_resume	kres	&5	KEY_RESUME, 0570, Sent by resume key
key_right	kcuf1	kr	KEY_RIGHT, 0405, Sent by terminal right-arrow key
key_save	ksav	&6	KEY_SAVE, 0571, Sent by save key

key_sbeg	kBEG	&9	KEY_SBEG, 0572, Sent by shifted beginning key
key_scancel	kCAN	&0	KEY_SCANCEL, 0573, Sent by shifted cancel key
key_scommand	kCMD	*1	KEY_SCOMMAND, 0574, Sent by shifted command key
key_scopy	kCPY	*2	KEY_SCOPY, 0575, Sent by shifted copy key
key_screate	kCRT	*3	KEY_SCREATE, 0576, Sent by shifted create key
key_sdc	kDC	*4	KEY_SDC, 0577, Sent by shifted delete-char key
key_sdl	kDL	*5	KEY_SDL, 0600, Sent by shifted delete-line key
key_select	kslt	*6	KEY_SELECT, 0601, Sent by select key
key_send	kEND	*7	KEY_SEND, 0602, Sent by shifted end key
key_seol	KEOL	*8	KEY_SEOL, 0603, Sent by shifted clear-line key
key_sexit	kEXT	*9	KEY_SEXIT, 0604, Sent by shifted exit key
key_sf	kind	kF	KEY_SF, 0520, Sent by scroll-forward/down key
key_sfnd	kFND	*0	KEY_SFIND, 0605, Sent by shifted find key
key_shelp	kHLP	#1	KEY_SHELP, 0606, Sent by shifted help key
key_shome	kHOM	#2	KEY_SHOME, 0607, Sent by shifted home key
key_sic	kIC	#3	KEY_SIC, 0610, Sent by shifted input key
key_sleft	kLFT	#4	KEY_SLEFT, 0611, Sent by shifted left-arrow key
key_smessage	kMSG	%a	KEY_SMESSAGE, 0612, Sent by shifted message key
key_smove	kMOV	%b	KEY_SMOVE, 0613, Sent by shifted move key
key_snext	kNXT	%c	KEY_SNEXT, 0614, Sent by shifted next key
key_soptions	kOPT	%d	KEY_SOPTIONS, 0615, Sent by shifted options key
key_sprevious	kPRV	%e	KEY_SPREVIOUS, 0616, Sent by shifted prev key
key_sprint	kPRT	%f	KEY_SPRINT, 0617, Sent by shifted print key
key_sr	kri	kR	KEY_SR, 0521, Sent by scroll-backward/up key
key_sredo	kRDO	%g	KEY_SREDO, 0620, Sent by shifted redo key
key_sreplace	kRPL	%h	KEY_SREPLACE, 0621, Sent by shifted replace key
key_sright	kRIT	%i	KEY_SRIGHT, 0622, Sent by shifted right-arrow key
key_sresume	kRES	%j	KEY_SRSUME, 0623, Sent by shifted resume key
key_ssave	kSAV	!1	KEY_SSAVE, 0624, Sent by shifted save key
key_ssuspend	kSPD	!2	KEY_SSUSPEND, 0625, Sent by shifted suspend key
key_stab	khts	kT	KEY_STAB, 0524, Sent by set-tab key
key_sundo	kUND	!3	KEY_SUNDO, 0626, Sent by shifted undo key
key_suspend	kspd	&7	KEY_SUSPEND, 0627, Sent by suspend key
key_undo	kund	&8	KEY_UNDO, 0630, Sent by undo key
key_up	kcuu1	ku	KEY_UP, 0403, Sent by terminal up-arrow key
keypad_local	rmkx	ke	Out of "keypad-transmit" mode
keypad_xmit	smkx	ks	Put terminal in "keypad-transmit" mode
lab_f0	lf0	l0	Labels on function key f0 if not f0
lab_f1	lf1	l1	Labels on function key f1 if not f1
lab_f2	lf2	l2	Labels on function key f2 if not f2
lab_f3	lf3	l3	Labels on function key f3 if not f3
lab_f4	lf4	l4	Labels on function key f4 if not f4
lab_f5	lf5	l5	Labels on function key f5 if not f5
lab_f6	lf6	l6	Labels on function key f6 if not f6
lab_f7	lf7	l7	Labels on function key f7 if not f7
lab_f8	lf8	l8	Labels on function key f8 if not f8
lab_f9	lf9	l9	Labels on function key f9 if not f9
lab_f10	lf10	la	Labels on function key f10 if not f10
label_off	rmln	LF	Turn off soft labels
label_on	smln	LO	Turn on soft labels
meta_off	rmm	mo	Turn off "meta mode"
meta_on	smm	mm	Turn on "meta mode" (8th bit)
newline	nel	nw	NEWLINE (behaves like cr followed by lf)
pad_char	pad	pc	Pad character (rather than null)
parm_dch	dch	DC	Delete #1 chars (G*)
parm_delete_line	dl	DL	Delete #1 lines (G*)
parm_down_cursor	cud	DO	Move cursor down #1 lines. (G*)
parm_ich	ich	IC	Insert #1 blank chars (G*)

parm_index	indn	SF	Scroll forward #1 lines. (G)
parm_insert_line	il	AL	Add #1 new blank lines (G*)
parm_left_cursor	cub	LE	Move cursor left #1 spaces (G)
parm_right_cursor	cuf	RI	Move cursor right #1 spaces. (G*)
parm_rindex	rin	SR	Scroll backward #1 lines. (G)
parm_up_cursor	cuu	UP	Move cursor up #1 lines. (G*)
pkey_key	pfkey	pk	Prog funct key #1 to type string #2
pkey_local	pfloc	pl	Prog funct key #1 to execute string #2
pkey_xmit	px	px	Prog funct key #1 to xmit string #2
plab_norm	pln	pn	Prog label #1 to show string #2
print_screen	mc0	ps	Print contents of the screen
prtr_non	mc5p	pO	Turn on the printer for #1 bytes
prtr_off	mc4	pf	Turn off the printer
prtr_on	mc5	po	Turn on the printer
repeat_char	rep	rp	Repeat char #1 #2 times (G*)
req_for_input	rfi	RF	Send next input char (for pty)
reset_1string	rs1	r1	Reset terminal completely to sane modes
reset_2string	rs2	r2	Reset terminal completely to sane modes
reset_3string	rs3	r3	Reset terminal completely to sane modes
reset_file	rf	rf	Name of file containing reset string
restore_cursor	rc	rc	Restore cursor to position of last sc
row_address	vpa	cv	Vertical position absolute (G)
save_cursor	sc	sc	Save cursor position.
scroll_forward	ind	sf	Scroll text up
scroll_reverse	ri	sr	Scroll text down
set_attributes	sgr	sa	Define the video attributes #1-#9 (G)
set_left_margin	smgl	ML	Set soft left margin
set_right_margin	smgr	MR	Set soft right margin
set_tab	hts	st	Set a tab stop in all rows, current column.
set_window	wind	wi	Current window is lines #1-#2 cols #3-#4 (G)
tab	ht	ta	Move the cursor to the next 8 space hardware tab stop.
to_status_line	tsl	ts	Go to status line, col #1 (G)
underline_char	uc	uc	Underscore one char and move past it
up_half_line	hu	hu	Half-line up (reverse 1/2 line-feed)
xoff_character	xoffc	XF	X-off character
xon_character	xonc	XN	X-on character

SAMPLE ENTRY

The following entry, which describes the Concept 100 terminal, is among the more complex entries in the terminfo file as of this writing.

```
concept100|c100|concept|c104|c100-4p|concept 100,
am,db,eo,in,mir,ul,xenl,cols#80,lines#24,pb#9600,vt#8,
bel='G,blank=\EH,blink=\EC,clear='L$<2*>,cnorm=\Ew,cr='M$<9>,
cub1='H,cud1='J,cuf1=\E=,cup=\Ea%p1%' '%+%c%p2%' '%+%c,cuu1=\E;,
cvvis=\EW,dch1=\E^A$<16*>,dim=\EE,dll=\E^B$<3*>,
ed=\E^C$<16*>,el=\E^U$<16>,flash=\Ek$<20>\EK,ht=\t$<8>,
ill=\E^R$<3*>,ind='J,.ind='J$<9>,ip=$<16*>,
is2=\EU\EAE7\E5\E8\E1\ENH\EK\E\0\Eo&\0\Eo\47\E,
kbs='h,kcub1=\E>,kcud1=\E<,kcuf1=\E=,kcuu1=\E; ,kf1=\E5,
kf2=\E6,kf3=\E7,khome=\E?,prot=\EI,
rep=\Er%p1%c%p2%' '%+%c$<2*>,rev=\ED,
rmcup=\Ev\s\s\s\s$<6>\Ep\r\n,rmir=\E\0,rmkx=\Ex,
rmso=\Ed\Ee,rmul=\Eg,rmul=\Eg,sgr0=\EN\0,
smcup=\EU\Ev\s\s8p\Ep\r,smir=\E^P,smkx=\EX,sms0=\EE\ED,
smul=\EG,
```

Entries may continue onto multiple lines by placing white space at the beginning of each line except the first. Lines beginning with # are taken as comment lines. Capabilities in `terminfo` are of three types: boolean capabilities which indicate that the terminal has some particular feature, numeric capabilities giving the size of the terminal or particular features, and string capabilities, which give a sequence which can be used to perform particular terminal operations.

Types of Capabilities

All capabilities have names. For instance, the fact that the Concept has *automatic margins* (that is, an automatic RETURN and LINEFEED when the end of a line is reached) is indicated by the capability `am`. Hence the description of the Concept includes `am`. Numeric capabilities are followed by the character # and then the value. Thus `cols`, which indicates the number of columns the terminal has, gives the value 80 for the Concept. The value may be specified in decimal, octal or hexadecimal using normal C conventions.

Finally, string-valued capabilities, such as `el` (clear to end of line sequence) are given by the two- to five-character capname, an '=', and then a string ending at the next following comma. A delay in milliseconds may appear anywhere in such a capability, enclosed in `$<. .>` brackets, as in `'el=\EK$<3>'`, and padding characters are supplied by `tputs()` (see `curses(3V)`) to provide this delay. The delay can be either a number, for example, 20, or a number followed by an * (for example, 3*), a / (for example, 5/), or both (for example, 10*/). A * indicates that the padding required is proportional to the number of lines affected by the operation, and the amount given is the per-affected-unit padding required. (In the case of insert character, the factor is still the number of lines affected. This is always one unless the terminal has `in` and the software uses it.) When a * is specified, it is sometimes useful to give a delay of the form 3.5 to specify a delay per unit to tenths of milliseconds. (Only one decimal place is allowed.) A / indicates that the padding is mandatory. Otherwise, if the terminal has `xon` defined, the padding information is advisory and will only be used for cost estimates or when the terminal is in raw mode. Mandatory padding will be transmitted regardless of the setting of `xon`.

A number of escape sequences are provided in the string-valued capabilities for easy encoding of characters there:

<code>\E, \e</code>	map to ESC
<code>^X</code>	maps to CTRL- <i>X</i> for any appropriate character <i>X</i>
<code>\n</code>	maps to NEWLINE
<code>\l</code>	maps to LINEFEED
<code>\r</code>	maps to RETURN
<code>\t</code>	maps to TAB
<code>\b</code>	maps to BACKSPACE
<code>\f</code>	maps to FORMFEED
<code>\s</code>	maps to SPACE
<code>\0</code>	maps to NUL

(`\0` will actually produce `\200`, which does not terminate a string but behaves as a null character on most terminals.) Finally, characters may be given as three octal digits after a backslash (for example, `\123`), and the characters ^ (caret), \ (backslash), : (colon), and , (comma) may be given as `\^`, `\\`, `\:`, and `\,` respectively.

Sometimes individual capabilities must be commented out. To do this, put a period before the capability name. For example, see the second `ind` in the example above. Note: capabilities are defined in a left-to-right order and, therefore, a prior definition will override a later definition.

Preparing Descriptions

The most effective way to prepare a terminal description is by imitating the description of a similar terminal in `terminfo` and to build up a description gradually, using partial descriptions with some *curses*-based application to check that they are correct. Be aware that a very unusual terminal may expose deficiencies in the ability of the `terminfo` file to describe it or bugs in the application. To test a new terminal description, set the environment variable `TERMINFO` to a pathname of a directory containing the compiled description you are working on and programs will look there rather than in `/usr/share/lib/terminfo`. To get the padding for insert-line correct (if the terminal manufacturer did not document it) a severe test is to insert 16 lines into the middle of a full screen at 9600 baud. If the display is corrupted, more padding is usu-

ally needed. A similar test can be used for insert-character.

Basic Capabilities

The number of columns on each line for the terminal is given by the `cols` numeric capability. If the terminal has a screen, then the number of lines on the screen is given by the `lines` capability. If the terminal wraps around to the beginning of the next line when it reaches the right margin, then it should have the `am` capability. If the terminal can clear its screen, leaving the cursor in the home position, then this is given by the `clear` string capability. If the terminal overstrikes (rather than clearing a position when a character is struck over) then it should have the `os` capability. If the terminal is a printing terminal, with no soft copy unit, give it both `hc` and `os`. (`os` applies to storage scope terminals, such as Tektronix 4010 series, as well as hard-copy and APL terminals.) If there is a code to move the cursor to the left edge of the current row, give this as `cr`. (Normally this will be RETURN, CTRL-M.) If there is a code to produce an audible signal (bell, beep, etc) give this as `bel`. If the terminal uses the xon-xoff flow-control protocol, like most terminals, specify `xon`.

If there is a code to move the cursor one position to the left (such as backspace) that capability should be given as `cub1`. Similarly, codes to move to the right, up, and down should be given as `cuf1`, `cuu1`, and `cud1`. These local cursor motions should not alter the text they pass over; for example, you would not normally use `cuf1=s` because the SPACE would erase the character moved over.

A very important point here is that the local cursor motions encoded in `terminfo` are undefined at the left and top edges of a screen terminal. Programs should never attempt to backspace around the left edge, unless `bw` is given, and should never attempt to go up locally off the top. In order to scroll text up, a program will go to the bottom left corner of the screen and send the `ind` (index) string.

To scroll text down, a program goes to the top left corner of the screen and sends the `ri` (reverse index) string. The strings `ind` and `ri` are undefined when not on their respective corners of the screen.

Parameterized versions of the scrolling sequences are `indn` and `rin` which have the same semantics as `ind` and `ri` except that they take one parameter, and `scroll` that many lines. They are also undefined except at the appropriate edge of the screen.

The `am` capability tells whether the cursor sticks at the right edge of the screen when text is output, but this does not necessarily apply to a `cuf1` from the last column. The only local motion which is defined from the left edge is if `bw` is given, then a `cub1` from the left edge will move to the right edge of the previous row. If `bw` is not given, the effect is undefined. This is useful for drawing a box around the edge of the screen, for example. If the terminal has switch selectable automatic margins, the `terminfo` file usually assumes that this is on; that is, `am`. If the terminal has a command which moves to the first column of the next line, that command can be given as `nel` (NEWLINE). It does not matter if the command clears the remainder of the current line, so if the terminal has no `cr` and if it may still be possible to craft a working `nel` out of one or both of them.

These capabilities suffice to describe hardcopy and screen terminals. Thus the model 33 teletype is described as

```
33|tty33|tty|model 33 teletype,
    bel=^G, cols#72, cr=^M, cud1=^J, hc, ind=^J, os,
```

while the Lear Siegler ADM-3 is described as

```
adm3|lsi adm3,
    am, bel=^G, clear=^Z, cols#80, cr=^M, cub1=^H, cud1=^J,
    ind=^J, lines#24,
```

Parameterized Strings

Cursor addressing and other strings requiring parameters in the terminal are described by a parameterized string capability, with `printf(3S)`-like escapes (`%x`) in it. For example, to address the cursor, the `cup` capability is given, using two parameters: the row and column to address to. (Rows and columns are numbered from zero and refer to the physical screen visible to the user, not to any unseen memory.) If the terminal has memory relative cursor addressing, that can be indicated by `mrcup`.

The parameter mechanism uses a stack and special % codes to manipulate it in the manner of a Reverse Polish Notation (postfix) calculator. Typically a sequence will push one of the parameters onto the stack and then print it in some format. Often more complex operations are necessary. Binary operations are in postfix form with the operands in the usual order. That is, to get x-5 one would use '%gx%{5}%-'.

The % encodings have the following meanings:

%%	outputs %
%[[:]flags][width[.precision]][doxXs]	as in printf(3S), flags are [-+#] and SPACE
%c	print pop() gives %c
%p[1-9]	push i^{th} parm
%P[a-z]	set variable [a-z] to pop()
%g[a-z]	get variable [a-z] and push it
%'c'	push char constant c
%{nn}	push decimal constant nn
%l	push strlen(pop())
%+ %- %* %/ %m	arithmetic (%m is mod): push(pop() op pop())
%& % %^	bit operations: push(pop() op pop())
%= %> %<	logical operations: push(pop() op pop())
%A %O	logical operations: and, or
%! %~	unary operations: push(op pop())
%i	(for ANSI terminals) add 1 to first parm, if one parm present, or first two parms, if more than one parm present
%?expr %tthenpart %eelsepart %;	if-then-else, '%eelsepart' is optional; else-if's are possible in Algol 68: %? c ₁ %t b ₁ %e c ₂ %t b ₂ %e c ₃ %t b ₃ %e c ₄ %t b ₄ %e b ₅ %; c _i are conditions, b _i are bodies.

If the '-' flag is used with '%[doxXs]', then a colon (:) must be placed between the '%' and the '-' to differentiate the flag from the binary '%-' operator, for example, '%:-16.16s'.

Consider the Hewlett-Packard 2645, which, to get to row 3 and column 12, needs to be sent \E&a12c03Y padded for 6 milliseconds. Note: the order of the rows and columns is inverted here, and that the row and column are zero-padded as two digits. Thus its cup capability is:

```
cup=\E&a%p2%2.2dc%p1%2.2dY$<6>
```

The Micro-Term ACT-IV needs the current row and column sent preceded by a ^T, with the row and column simply encoded in binary, 'cup=^T%p1%c%p2%c'. Terminals which use %c need to be able to backspace the cursor (cub1), and to move the cursor up one line on the screen (cuu1). This is necessary because it is not always safe to transmit \n, ^D, and \r, as the system may change or discard them. (The library routines dealing with terminfo set tty modes so that TAB characters are never expanded, so \t is safe to send. This turns out to be essential for the Ann Arbor 4080.)

A final example is the LSI ADM-3a, which uses row and column offset by a blank character, thus 'cup=\E=%p1%'s'%'+%c%p2%'s'%'+%c'. After sending '\E=', this pushes the first parameter, pushes the ASCII value for a space (32), adds them (pushing the sum on the stack in place of the two previous values), and outputs that value as a character. Then the same is done for the second parameter. More complex arithmetic is possible using the stack.

Cursor Motions

If the terminal has a fast way to home the cursor (to very upper left corner of screen) then this can be given as home; similarly a fast way of getting to the lower left-hand corner can be given as ll; this may involve going up with cuu1 from the home position, but a program should never do this itself (unless ll does) because it can make no assumption about the effect of moving up from the home position. Note: the home

position is the same as addressing to (0,0): to the top left corner of the screen, not of memory. (Thus, the \EH sequence on Hewlett-Packard terminals cannot be used for home without losing some of the other features on the terminal.)

If the terminal has row or column absolute-cursor addressing, these can be given as single parameter capabilities **hpa** (horizontal position absolute) and **vpa** (vertical position absolute). Sometimes these are shorter than the more general two-parameter sequence (as with the Hewlett-Packard 2645) and can be used in preference to **cup**. If there are parameterized local motions (for example, move *n* spaces to the right) these can be given as **cud**, **cub**, **cuf**, and **cuu** with a single parameter indicating how many spaces to move. These are primarily useful if the terminal does not have **cup**, such as the Tektronix 4025.

Area Clears

If the terminal can clear from the current position to the end of the line, leaving the cursor where it is, this should be given as **el**. If the terminal can clear from the beginning of the line to the current position inclusive, leaving the cursor where it is, this should be given as **el1**. If the terminal can clear from the current position to the end of the display, then this should be given as **ed**. **ed** is only defined from the first column of a line. (Thus, it can be simulated by a request to delete a large number of lines, if a true **ed** is not available.)

Insert/Delete Line

If the terminal can open a new blank line before the line where the cursor is, this should be given as **il1**; 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 **dl1**; this is done only from the first position on the line to be deleted. Versions of **il1** and **dl1** which take a single parameter and insert or delete that many lines can be given as **il** and **dl**.

If the terminal has a settable destructive scrolling region (like the VT100) the command to set this can be described with the **csr** capability, which takes two parameters: the top and bottom lines of the scrolling region. The cursor position is, alas, undefined after using this command. It is possible to get the effect of insert or delete line using this command — the **sc** and **rc** (save and restore cursor) commands are also useful. Inserting lines at the top or bottom of the screen can also be done using **ri** or **ind** on many terminals without a true insert/delete line, and is often faster even on terminals with those features.

To determine whether a terminal has destructive scrolling regions or non-destructive scrolling regions, create a scrolling region in the middle of the screen, place data on the bottom line of the scrolling region, move the cursor to the top line of the scrolling region, and do a reverse index (**ri**) followed by a delete line (**dl1**) or index (**ind**). If the data that was originally on the bottom line of the scrolling region was restored into the scrolling region by the **dl1** or **ind**, then the terminal has non-destructive scrolling regions. Otherwise, it has destructive scrolling regions. Do not specify **csr** if the terminal has non-destructive scrolling regions, unless **ind**, **ri**, **indn**, **rin**, **dl**, and **dl1** all simulate destructive scrolling.

If the terminal has the ability to define a window as part of memory, which all commands affect, it should be given as the parameterized string **wind**. The four parameters are the starting and ending lines in memory and the starting and ending columns in memory, in that order.

If the terminal can retain display memory above, then the **da** capability should be given; if display memory can be retained below, then **db** should be given. These indicate that deleting a line or scrolling a full screen may bring non-blank lines up from below or that scrolling back with **ri** may bring down non-blank lines.

Insert/Delete Character

There are two basic kinds of intelligent terminals with respect to insert/delete character operations which can be described using **terminfo**. The most common insert/delete character operations affect only the characters on the current line and shift characters off the end of the line rigidly. Other terminals, such as the Concept 100 and the Perkin Elmer Owl, make a distinction between typed and untyped blanks on the screen, shifting upon an insert or delete only to an untyped blank on the screen which is either eliminated, or expanded to two untyped blanks. You can determine the kind of terminal you have by clearing the screen and then typing text separated by cursor motions. Type **'abc def'** using local cursor motions (not SPACE characters) between the **abc** and the **def**. Then position the cursor before the **abc** and put the termi-

nal 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". While these are two logically separate attributes (one line versus multiline insert mode, and special treatment of untyped blanks) we have seen no terminals whose insert mode cannot be described with the single attribute.

`terminfo` can describe both terminals which have an insert mode and terminals which send a simple sequence to open a blank position on the current line. Give as `smir` the sequence to get into insert mode. Give as `rmir` the sequence to leave insert mode. Now give as `ich1` any sequence needed to be sent just before sending the character to be inserted. Most terminals with a true insert mode will not give `ich1`; terminals which send a sequence to open a screen position should give it here. (If your terminal has both, insert mode is usually preferable to `ich1`. Do not give both unless the terminal actually requires both to be used in combination.) If post-insert padding is needed, give this as a number of milliseconds padding 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`. If your terminal needs both to be placed into an "insert mode" and a special code to precede each inserted character, then both `smir/rmir` and `ich1` can be given, and both will be used. The `ich` capability, with one parameter, `n`, will repeat the effects of `ich1` `n` times.

If padding is necessary between characters typed while not in insert mode, give this as a number of milliseconds padding in `rmp`.

It is occasionally necessary to move around while in insert mode to delete characters on the same line (for example, if there is a TAB character after the insertion position). If your terminal allows motion while in insert mode you can give the capability `mir` to speed up inserting in this case. Omitting `mir` will affect only speed. Some terminals (notably Datamedia's) must not have `mir` because of the way their insert mode works.

Finally, you can specify `dch1` to delete a single character, `dch` with one parameter, `n`, to delete `n` characters, and delete mode by giving `smdc` and `rmdc` to enter and exit delete mode (any mode the terminal needs to be placed in for `dch1` to work).

A command to erase `n` characters (equivalent to outputting `n` blanks without moving the cursor) can be given as `ech` with one parameter.

Highlighting, Underlining, and Visible Bells

If your terminal has one or more kinds of display attributes, these can be represented in a number of different ways. You should choose one display form as *standout mode* (see `curses(3V)`), representing a good, high contrast, easy-on-the-eyes, format for highlighting error messages and other attention getters. (If you have a choice, reverse-video plus half-bright is good, or reverse-video alone; however, different users have different preferences on different terminals.) The sequences to enter and exit standout mode are given as `sms0` and `rms0`, respectively. If the code to change into or out of standout mode leaves one or even two blanks on the screen, as the TVI 912 and Teleray 1061 do, then `xmc` should be given to tell how many blanks are left.

Codes to begin underlining and end underlining can be given as `smul` and `rmul` respectively. If the terminal has a code to underline the current character and move the cursor one position to the right, such as the Micro-Term MIME, this can be given as `uc`.

Other capabilities to enter various highlighting modes include `blink` (blinking), `bold` (bold or extra-bright), `dim` (dim or half-bright), `invis` (blanking or invisible text), `prot` (protected), `rev` (reverse-video), `sgr0` (turn off all attribute modes), `smacs` (enter alternate-character-set mode), and `rmacs` (exit alternate-character-set mode). Turning on any of these modes singly may or may not turn off other modes. If a command is necessary before alternate character set mode is entered, give the sequence in `enacs` (enable alternate-character-set mode).

If there is a sequence to set arbitrary combinations of modes, this should be given as `sgr` (set attributes), taking nine parameters. Each parameter is either 0 or non-zero, as the corresponding attribute is on or off. The nine parameters are, in order: standout, underline, reverse, blink, dim, bold, blank, protect, alternate character set. Not all modes need be supported by `sgr`, only those for which corresponding separate attribute commands exist. (See the example at the end of this section.)

Terminals with the “magic cookie” glitch (`xmc`) deposit special “cookies” when they receive mode-setting sequences, which affect the display algorithm rather than having extra bits for each character. Some terminals, such as the Hewlett-Packard 2621, automatically leave standout mode when they move to a new line or the cursor is addressed. Programs using standout mode should exit standout mode before moving the cursor or sending a newline, unless the `msg` capability, asserting that it is safe to move in standout mode, is present.

If the terminal has a way of flashing the screen to indicate an error quietly (a bell replacement), then this can be given as `flash`; it must not move the cursor. A good flash can be done by changing the screen into reverse video, pad for 200 ms, then return the screen to normal video.

If the cursor needs to be made more visible than normal when it is not on the bottom line (to make, for example, a non-blinking underline into an easier to find block or blinking underline) give this sequence as `cvvis`. The boolean `chts` should also be given. If there is a way to make the cursor completely invisible, give that as `civis`. The capability `cnorm` should be given which undoes the effects of either of these modes.

If the terminal needs to be in a special mode when running a program that uses these capabilities, the codes to enter and exit this mode can be given as `smcup` and `rmcup`. This arises, for example, from terminals like the Concept with more than one page of memory. If the terminal has only memory relative cursor addressing and not screen relative cursor addressing, a one screen-sized window must be fixed into the terminal for cursor addressing to work properly. This is also used for the Tektronix 4025, where `smcup` sets the command character to be the one used by `terminfo`. If the `smcup` sequence will not restore the screen after an `rmcup` sequence is output (to the state prior to outputting `rmcup`), specify `nrrmc`.

If your terminal generates underlined characters by using the underline character (with no special codes needed) even though it does not otherwise overstrike characters, then you should give the capability `ul`. For terminals where a character overstriking another leaves both characters on the screen, give the capability `os`. If overstrikes are erasable with a blank, then this should be indicated by giving `eo`.

Example of highlighting: assume that the terminal under question needs the following escape sequences to turn on various modes.

tparm parameter	attribute	escape sequence
	none	\E[0m
p1	standout	\E[0;4;7m
p2	underline	\E[0;3m
p3	reverse	\E[0;4m
p4	blink	\E[0;5m
p5	dim	\E[0;7m
p6	bold	\E[0;3;4m
p7	invis	\E[0;8m
p8	protect	not available
p9	altcharset	^O (off) ^N(on)

Note: each escape sequence requires a 0 to turn off other modes before turning on its own mode. Also note that, as suggested above, *standout* is set up to be the combination of *reverse* and *dim*. Also, since this terminal has no *bold* mode, *bold* is set up as the combination of *reverse* and *underline*. In addition, to allow combinations, such as *underline+blink*, the sequence to use would be '\E[0;3;5m'. The terminal does not have *protect* mode, either, but that cannot be simulated in any way, so `p8` is ignored. The *altcharset* mode is different in that it is either ^O or ^N depending on whether it is off or on. If all modes were to be turned on, the sequence would be '\E[0;3;4;5;7;8m^N'.

Now look at when different sequences are output. For example, ‘;3’ is output when either ‘p2’ or ‘p6’ is true, that is, if either *underline* or *bold* modes are turned on. Writing out the above sequences, along with their dependencies, gives the following:

sequence	when to output	terminfo translation
\E[0	always	\E[0
;3	if p2 or p6	;%?%p2%p6% %;3%;
;4	if p1 or p3 or p6	;%?%p1%p3% p6% %;4%;
;5	if p4	;%?%p4%;5%;
;7	if p1 or p5	;%?%p1%p5% %;7%;
;8	if p7	;%?%p7%;8%;
m	always	m
^N or ^O	if p9 ^N, else ^O	;%?%p9%t^N%e^O%;

Putting this all together into the sgr sequence gives:

```
sgr=\E[0;%?%p2%p6%|%;3%;;%?%p1%p3%|p6%|%;4%;;%?%p5%;5%;;%?%p1%p5%|%;7%;;%?%p7%;8%;m;%?%p9%t^N%e^O%;
```

Keypad

If the terminal has a keypad that transmits codes when the keys are pressed, this information can be given. Note: it is not possible to handle terminals where the keypad only works in local (this applies, for example, to the unshifted Hewlett-Packard 2621 keys). If the keypad can be set to transmit or not transmit, give these codes as **smkx** and **rmkx**. Otherwise the keypad is assumed to always transmit.

The codes sent by the left arrow, right arrow, up arrow, down arrow, and home keys can be given as **kcub1**, **kcufl**, **kcuu1**, **kcud1**, and **khome** respectively. If there are function keys such as f0, f1, ..., f63, the codes they send can be given as **kf0**, **kf1**, ..., **kf63**. If the first 11 keys have labels other than the default f0 through f10, the labels can be given as **lf0**, **lf1**, ..., **lf10**. The codes transmitted by certain other special keys can be given: **kl** (home down), **kbs** (BACKSPACE), **ktbc** (clear all tab stops), **kctab** (clear the tab stop in this column), **kclr** (clear screen or erase key), **kdch1** (delete character), **kdll** (delete line), **krmir** (exit insert mode), **kel** (clear to end of line), **ked** (clear to end of screen), **kich1** (insert character or enter insert mode), **kill** (insert line), **kn** (next page), **kpp** (previous page), **kind** (scroll forward/down), **kri** (scroll backward/up), **khts** (set a tab stop in this column). In addition, if the keypad has a 3 by 3 array of keys including the four arrow keys, the other five keys can be given as **ka1**, **ka3**, **kb2**, **kc1**, and **kc3**. These keys are useful when the effects of a 3 by 3 directional pad are needed. Further keys are defined above in the capabilities list.

Strings to program function keys can be given as **pfkey**, **pfloc**, and **px**. A string to program their soft-screen labels can be given as **pln**. Each of these strings takes two parameters: the function key number to program (from 0 to 10) and the string to program it with. Function key numbers out of this range may program undefined keys in a terminal-dependent manner. The difference between the capabilities is that **pfkey** causes pressing the given key to be the same as the user typing the given string; **pfloc** executes the string by the terminal in local mode; and **px** transmits the string to the computer. The capabilities **nlab**, **lw** and **lh** define how many soft labels there are and their width and height. If there are commands to turn the labels on and off, give them in **smln** and **rmln**. **smln** is normally output after one or more **pln** sequences to make sure that the change becomes visible.

Tabs and Initialization

If the terminal has hardware tab stops, the command to advance to the next tab stop can be given as **ht** (usually CTRL-I). A “backtab” command which moves leftward to the next tab stop can be given as **cbt**. By convention, if the teletype modes indicate that TAB characters are being expanded by the computer rather than being sent to the terminal, programs should not use **ht** or **cbt** even if they are present, since the user may not have the tab stops properly set. If the terminal has hardware tab stops which are initially set every *n* spaces when the terminal is powered up, the numeric parameter it is given, showing the number of spaces the tab stops are set to. This is normally used by ‘tput init’ (see **tput(1V)**) to determine whether to set the mode for hardware TAB expansion and whether to set the tab stops. If the terminal has tab stops

that can be saved in nonvolatile memory, the **terminfo** description can assume that they are properly set. If there are commands to set and clear tab stops, they can be given as **tbc** (clear all tab stops) and **hts** (set a tab stop in the current column of every row).

Other capabilities include: **is1**, **is2**, and **is3**, initialization strings for the terminal; **iprogram**, the path name of a program to be run to initialize the terminal; and **if**, the name of a file containing long initialization strings. These strings are expected to set the terminal into modes consistent with the rest of the **terminfo** description. They must be sent to the terminal each time the user logs in and be output in the following order: run the program **iprogram**; output **is1**; output **is2**; set the margins using **mgc**, **smgl** and **smgr**; set the tab stops using **tbc** and **hts**; print the file **if**; and finally output **is3**. This is usually done using the **init** option of **tput(1V)**.

Most initialization is done with **is2**. Special terminal modes can be set up without duplicating strings by putting the common sequences in **is2** and special cases in **is1** and **is3**. Sequences that do a harder reset from a totally unknown state can be given as **rs1**, **rs2**, **rf**, and **rs3**, analogous to **is1**, **is2**, **is3**, and **if**. (The method using files, **if** and **rf**, is used for a few terminals, from **/usr/share/lib/tabset/***; however, the recommended method is to use the initialization and reset strings.) These strings are output by **'tput reset'**, which is used when the terminal gets into a wedged state. Commands are normally placed in **rs1**, **rs2**, **rs3**, and **rf** only if they produce annoying effects on the screen and are not necessary when logging in. For example, the command to set a terminal into 80-column mode would normally be part of **is2**, but on some terminals it causes an annoying glitch on the screen and is not normally needed since the terminal is usually already in 80-column mode.

If a more complex sequence is needed to set the tab stops than can be described by using **tbc** and **hts**, the sequence can be placed in **is2** or **if**.

If there are commands to set and clear margins, they can be given as **mgc** (clear all margins), **smgl** (set left margin), and **smgr** (set right margin).

Delays

Certain capabilities control padding in the terminal driver. These are primarily needed by hard-copy terminals, and are used by **'tput init'** to set tty modes appropriately. Delays embedded in the capabilities **cr**, **ind**, **cub1**, **ff**, and **tab** can be used to set the appropriate delay bits to be set in the tty driver. If **pb** (padding baud rate) is given, these values can be ignored at baud rates below the value of **pb**.

Status Lines

If the terminal has an extra "status line" that is not normally used by software, this fact can be indicated. If the status line is viewed as an extra line below the bottom line, into which one can cursor address normally (such as the Heathkit H19's 25th line, or the 24th line of a VT100 which is set to a 23-line scrolling region), the capability **hs** should be given. Special strings that go to a given column of the status line and return from the status line can be given as **tsl** and **fsl**. (**fsl** must leave the cursor position in the same place it was before **tsl**. If necessary, the **sc** and **rc** strings can be included in **tsl** and **fsl** to get this effect.) The capability **tsl** takes one parameter, which is the column number of the status line the cursor is to be moved to.

If escape sequences and other special commands, such as **TAB**, work while in the status line, the flag **eslok** can be given. A string which turns off the status line (or otherwise erases its contents) should be given as **dsl**. If the terminal has commands to save and restore the position of the cursor, give them as **sc** and **rc**. The status line is normally assumed to be the same width as the rest of the screen, for example, **cols**. If the status line is a different width (possibly because the terminal does not allow an entire line to be loaded) the width, in columns, can be indicated with the numeric parameter **wsl**.

Line Graphics

If the terminal has a line drawing alternate character set, the mapping of glyph to character would be given in `acsc`. The definition of this string is based on the alternate character set used in the DEC VT100 terminal, extended slightly with some characters from the AT&T 4410v1 terminal.

glyph name	VT100+ character
arrow pointing right	+
arrow pointing left	,
arrow pointing down	.
solid square block	0
lantern symbol	I
arrow pointing up	-
diamond	'
checker board (stipple)	a
degree symbol	f
plus/minus	g
board of squares	h
lower right corner	j
upper right corner	k
upper left corner	l
lower left corner	m
plus	n
scan line 1	o
horizontal line	q
scan line 9	s
left tee (├)	t
right tee (┤)	u
bottom tee (┴)	v
top tee (┬)	w
vertical line	x
bullet	-

The best way to describe a new terminal's line graphics set is to add a third column to the above table with the characters for the new terminal that produce the appropriate glyph when the terminal is in the alternate character set mode. For example,

glyph name	VT100+ char	new tty char
upper left corner	l	R
lower left corner	m	F
upper right corner	k	T
lower right corner	j	G
horizontal line	q	,
vertical line	x	.

Now write down the characters left to right, as in `'acsc=IRmFkTjGq\,x.'`

Miscellaneous

If the terminal requires other than a null (zero) character as a pad, then this can be given as `pad`. Only the first character of the pad string is used. If the terminal does not have a pad character, specify `npc`.

If the terminal can move up or down half a line, this can be indicated with `hu` (half-line up) and `hd` (half-line down). This is primarily useful for superscripts and subscripts on hardcopy terminals. If a hardcopy terminal can eject to the next page (form feed), give this as `ff` (usually CTRL-L).

If there is a command to repeat a given character a given number of times (to save time transmitting a large number of identical characters) this can be indicated with the parameterized string **rep**. The first parameter is the character to be repeated and the second is the number of times to repeat it. Thus, **'tparm(repeat_char, 'x', 10)** is the same as **'xxxxxxxxxx'**.

If the terminal has a settable command character, such as the Tektronix 4025, this can be indicated with **cmdch**. A prototype command character is chosen which is used in all capabilities. This character is given in the **cmdch** capability to identify it. On some UNIX systems, when the environment variable **CC** is set to a single-character value, all occurrences of the prototype character are replaced with that character.

Terminal descriptions that do not represent a specific kind of known terminal, such as **switch**, **dialup**, **patch**, and **network**, should include the **gn** (generic) capability so that programs can complain that they do not know how to talk to the terminal. (This capability does not apply to **virtual** terminal descriptions for which the escape sequences are known.) If the terminal is one of those supported by the UNIX system virtual terminal protocol, the terminal number can be given as **vt**. A line-turn-around sequence to be transmitted before doing reads should be specified in **rft**.

If the terminal uses **xon/xoff** handshaking for flow control, give **xon**. Padding information should still be included so that routines can make better decisions about costs, but actual pad characters will not be transmitted. Sequences to turn on and off **xon/xoff** handshaking may be given in **smxon** and **rmxon**. If the characters used for handshaking are not **^S** and **^Q** (CTRL-S and CTRL-Q, respectively), they may be specified with **xonc** and **xoffc**.

If the terminal has a "meta key" which acts as a shift key, setting the 8th bit of any character transmitted, this fact can be indicated with **km**. Otherwise, software will assume that the 8th bit is parity and it will usually be cleared. If strings exist to turn this "meta mode" on and off, they can be given as **smm** and **rmm**.

If the terminal has more lines of memory than will fit on the screen at once, the number of lines of memory can be indicated with **lm**. A value of **lm#0** indicates that the number of lines is not fixed, but that there is still more memory than fits on the screen.

Media copy strings which control an auxiliary printer connected to the terminal can be given as **mc0**: print the contents of the screen, **mc4**: turn off the printer, and **mc5**: turn on the printer. When the printer is on, all text sent to the terminal will be sent to the printer. A variation, **mc5p**, takes one parameter, and leaves the printer on for as many characters as the value of the parameter, then turns the printer off. The parameter should not exceed 255. If the text is not displayed on the terminal screen when the printer is on, specify **mc5i** (silent printer). All text, including **mc4**, is transparently passed to the printer while an **mc5p** is in effect.

Special Cases

The working model used by **terminfo** fits most terminals reasonably well. However, some terminals do not completely match that model, requiring special support by **terminfo**. These are not meant to be construed as deficiencies in the terminals; they are just differences between the working model and the actual hardware. They may be unusual devices or, for some reason, do not have all the features of the **terminfo** model implemented.

Terminals which can not display tilde (~) characters, such as certain Hazeltine terminals, should indicate **hz**.

Terminals which ignore a **LINEFEED** immediately after an **am** wrap, such as the Concept 100, should indicate **xenl**. Those terminals whose cursor remains on the right-most column until another character has been received, rather than wrapping immediately upon receiving the right-most character, such as the VT100, should also indicate **xenl**.

If **el** is required to get rid of standout (instead of writing normal text on top of it), **xhp** should be given.

Those Teleray terminals whose tabs turn all characters moved over to blanks, should indicate **xt** (destructive TAB characters). This capability is also taken to mean that it is not possible to position the cursor on top of a "magic cookie" therefore, to erase standout mode, it is instead necessary to use delete and insert

line.

Those Beehive Superbee terminals which do not transmit the escape or CTRL-C characters, should specify `xsib`, indicating that the f1 key is to be used for escape and the f2 key for CTRL-C.

Similar Terminals

If there are two very similar terminals, one can be defined as being just like the other with certain exceptions. The string capability `use` can be given with the name of the similar terminal. The capabilities given before `use` override those in the terminal type invoked by `use`. A capability can be canceled by placing `xx@` to the left of the capability definition, where `xx` is the capability. For example, the entry

```
att4424-2|Teletype 4424 in display function group ii,
    rev@, sgr@, smul@, use=att4424,
```

defines an AT&T 4424 terminal that does not have the `rev`, `sgr`, and `smul` capabilities, and hence cannot do highlighting. This is useful for different modes for a terminal, or for different user preferences. More than one `use` capability may be given.

FILES

```
/usr/share/lib/terminfo/?/* compiled terminal description database
/usr/share/lib/tabset/* tab stop settings for some terminals, in a format appropriate to be output to the terminal (escape sequences that set margins and tab stops)
```

SEE ALSO

`tput(1V)`, `curses(3V)`, `printf(3S)`, `term(5V)`, `captainfo(8V)`, `infocmp(8V)`, `tic(8V)`

WARNING

As described in the **Tabs and Initialization** section above, a terminal's initialization strings, `is1`, `is2`, and `is3`, if defined, must be output before a `curses(3V)` program is run. An available mechanism for outputting such strings is `tput init` (see `tput(1V)`).

Tampering with entries in `/usr/share/lib/terminfo/?/*` (for example, changing or removing an entry) can affect programs that expect the entry to be present and correct. In particular, removing the description for the "dumb" terminal will cause unexpected problems.

NAME

toc – table of contents of optional clusters in Application SunOS and Developer's Toolkit

SYNOPSIS

`/usr/lib/load/toc`

AVAILABILITY

Sun386i systems only.

DESCRIPTION

The **toc** file contains information specifying the organization of the optional clusters in Application SunOS and Developer's Toolkit on the Sun386i distribution media. For each cluster, a single line should be present with the following information:

cluster name
 set containing the cluster (Application SunOS or Developer's Toolkit)
 size of the cluster (in kilobytes)
 diskette volume of the cluster in the set (for loading from 3.5" diskette)
 tape and file number of the cluster (for loading from 1/4" tape)

Items are separated by a ':'.

Cluster names can contain any printable character other than a ':', space, tab, or newline character. The set containing the cluster is specified by an 'A' for Application SunOS or 'D' for Developer's Toolkit. The diskette volume is the number of the diskette within the diskette set on which the cluster begins. The tape and file number specifies the tape and file position of the cluster on the tape.

EXAMPLE

The following is an example to the **toc** file.

```
accounting:A:55:14:1@12
advanced_admin:A:628:14:1@4
audit:A:144:14:1@8
comm:A:312:13:1@9
disk_quotas:A:56:14:1@11
doc_prep:A:790:13:1@10
extended_commands:A:276:13:1@5
games:A:2351:19:1@17
mail_plus:A:135:14:1@7
man_pages:A:5586:16:1@14
name_server:A:339:14:1@13
networking_plus:A:610:13:1@6
old:A:131:14:1@16
plot:A:227:14:1@14
spellcheck:A:455:13:1@2
sysV_commands:A:2505:14:1@3
base_devel:D:5389:1:2@2
plot_devel:D:247:5:2@3
sccs:D:328:5:2@4
sunview_devel:D:1768:5:2@5
sysV_devel:D:4287:3:2@6
proffibs:D:4755:4:2@7
config:D:3065:6:2@8
```

The first line specifies that the accounting cluster is part of Application SunOS and requires 55 kilobytes of disk storage. In the diskette distribution, it begins on diskette 14 of Application SunOS optional clusters. In the tape distribution, it can be found on file 12 of tape 1. The last line specifies that the *config* cluster is part of Developer's Toolkit and requires 3065 kilobytes of disk storage. In the diskette distribution, it begins on diskette 6 of Developer's Toolkit. In the tape distribution, it can be found on file 8 of tape 2.

FILES

`/usr/lib/load/toc`

SEE ALSO

`cluster(1)` `load(1)` `unload(1)`

NAME

translate – input and output files for system message translation

AVAILABILITY

Sun386i systems only.

DESCRIPTION

These files are used by `syslogd(8)` to translate systems messages. The input file is used to map system messages (in `printf(3S)` format strings) to numbers. This number is then used to locate a new string in the output file.

An initial part of each line in the input file may specify that the message should be suppressed. Recognized suppression specifications are:

- (NONE) Suppress the message always.
- (n) Allow only one message every n seconds. ((10) for example).
- () Do not suppress the message. This can be used in a message that begins with a '('.

Note that the message suppression specification is optional. If not present, the message is not suppressed.

Each line in the output file translates the numbers from the input file into the desired error messages, and also specifies the format to be used to output each message. The order of parameters passed from the input message can be changed, by replacing the % of a format phrase with a %num\$ where num is a digit string. For example, if num is 2, the second parameter on the input file line will be used. The value of num can be from 1 to the number of parameters in the input message.

If a string is translated to a number that is not found in the output file, the message is suppressed.

EXAMPLES

An example input file:

```
$quote "
1      "(NONE)(1) logopen test code: %s\n"
2      "(10)(2) logopen test code: %s\n"
3      "() (3) logopen test code: %s\n"
4      "() (4) logopen test code: %s\n"
5      "(10)(5) logopen testcode: %s * 100\n"
6      "(10)(6) logopen testcode: %s * 100\n"
7      "(10)(7) logopen testcode: %s * 100\n"
8      "(10)%s: %s\n"
9      "(10)\n%s: write failed, file system is full\n"
10     "(10)NFS server %s not responding still trying\n"
11     "(10)NFS %s failed for server %s: %s\n"
12     "(10)NFS server %s ok\n"
13     "(NONE)\n%s: write failed, file system is full\n"
14     "(10)NFS server %s not responding still trying\n"
15     "(10)NFS %s failed for server %s: %s\n"
```

An example output file:

```
$quote "  
1 "TRANSLATION:(1) logopen test code: %s\n"  
2 "TRANSLATION: (2) logopen test code: %s IS REALLY\n"  
3 "TRANSLATION: (3) logopen test code: %s\n"  
4 "TRANSLATION: (4) logopen test code: %s\n"  
5 "TRANSLATION: (5) logopen testcode: %s * 100\n"  
6 "TRANSLATION: (6) logopen testcode: %s * 100\n"  
7 "TRANSLATION: (7) logopen testcode: %s * 100\n"  
8 "TRANSLATION: %s: %s\n"  
9 "TRANSLATION: \n%s: write failed, file system is full\n"  
10 "TRANSLATION: NFS server %s not responding still trying\n"  
11 "TRANSLATION: NFS %s failed for server %s: %s\n"  
12 "TRANSLATION: NFS server %s ok\n"  
13 "Out of disk on file system %s\n"  
14 "Network file server %s not ok. Check your cable\n"  
15 "Network file server %2$s down (%1$s, %3$s)\n"
```

SEE ALSO

syslogd(8)

NAME

ttytab, ttys – terminal initialization data

DESCRIPTION

The `/etc/ttytab` file contains information that is used by various routines to initialize and control the use of terminal special files. This information is read with the `gettyent(3)` library routines. There is one line in `/etc/ttytab` file per special file.

The `/etc/ttys` file should not be edited; it is derived from `/etc/ttytab` by `init(8)` at boot time, and is only included for backward compatibility with programs that may still require it.

Fields are separated by TAB and/or SPACE characters. Some fields may contain more than one word and should be enclosed in double quotes. Blank lines and comments can appear anywhere in the file; comments are delimited by '#' and NEWLINE. Unspecified fields default to NULL. The first field is the terminal's entry in the device directory, `/dev`. The second field of the file is the command to execute for the line, typically `getty(8)`, which performs such tasks as baud-rate recognition, reading the login name, and calling `login(1)`. It can be, however, any desired command, for example the start up for a window system terminal emulator or some other daemon process, and can contain multiple words if quoted. The third field is the type of terminal normally connected to that tty line, as found in the `termcap(5)` data base file. The remaining fields set flags in the `ty_status` entry (see `gettyent(3)`) or specify a window system process that `init(8)` will maintain for the terminal line.

As flag values, the strings `on` and `off` specify whether `init` should execute the command given in the second field, while `secure` in addition to `on` allows "root" to login on this line. If the console is not marked "secure," the system prompts for the root password before coming up in single-user mode. These flag fields should not be quoted. The string `window=` is followed by a quoted command string which `init` will execute before starting `getty`. If the line ends in a comment, the comment is included in the `ty_comment` field of the `ttvent` structure.

EXAMPLE

```
console "/usr/etc/getty std.1200" vt100      on secure
ttyd0  "/usr/etc/getty d1200"    dialup   on      # 555-1234
ttyh0  "/usr/etc/getty std.9600" hp2621-nl on      # 254MC
ttyh1  "/usr/etc/getty std.9600" plugboard on      # John's office
ttyp0  none                      network
ttyp1  none                      network  off
ttyv0  "/usr/new/xterm -L :0"    vs100   on window="/usr/new/Xvs100 0"
```

The first line permits "root" login on the console at 1200 baud, and indicates that the console is secure for single-user operation. The second example allows dialup at 1200 baud without "root" login, and the third and fourth examples allow login at 9600 baud with terminal types of `hp2621-nl` and `plugboard`, respectively. The fifth and sixth lines are examples of network pseudo-ttys, for which `getty` should not be enabled. The last line shows a terminal emulator and window-system startup entry.

FILES

`/dev`
`/etc/ttytab`

SEE ALSO

`login(1)`, `gettyent(3)`, `gettytab(5)`, `termcap(5)`, `getty(8)`, `init(8)`

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.

FILES

/dev/

SEE ALSO

login(1), **tset(1)**, **termcap(5)**

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 the system code; some data of these types are accessible to user code:

```
/*
 * Copyright (c) 1982, 1986 Regents of the University of California.
 * All rights reserved. The Berkeley software License Agreement
 * specifies the terms and conditions for redistribution.
 */
```

```
#ifndef _TYPES_
#define _TYPES_
```

```
/*
 * Basic system types.
 */
```

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

```
typedef unsigned char  u_char;
typedef unsigned short u_short;
typedef unsigned int   u_int;
typedef unsigned long  u_long;
typedef unsigned short ushort; /* System V compatibility */
typedef unsigned int   uint; /* System V compatibility */
```

```
#ifdef vax
typedef struct  _physadr { int r[1]; } *physadr;
typedef struct label_t {
    int    val[14];
} label_t;
#endif
#ifdef mc68000
typedef struct  _physadr { short r[1]; } *physadr;
typedef struct label_t {
    int    val[13];
} label_t;
#endif
#ifdef sparc
typedef struct _physadr { int r[1]; } *physadr;
typedef struct label_t {
    int    val[2];
} label_t;
#endif
#ifdef i386
typedef struct  _physadr { short r[1]; } *physadr;
typedef struct label_t {
    int    val[8];
} 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 long      time_t;
typedef short     dev_t;
typedef long      off_t;
typedef u_short   uid_t;
typedef u_short   gid_t;
typedef long      key_t;

#define NBBY      8      /* number of bits in a byte */
/*
 * Select uses bit masks of file descriptors in longs.
 * These macros manipulate such bit fields (the filesystem macros use chars).
 * FD_SETSIZE may be defined by the user, but the default here
 * should be >= NOFILE (param.h).
 */
#ifndef FD_SETSIZE
#define FD_SETSIZE 256
#endif

typedef long      fd_mask;
#define NFDBITS   (sizeof(fd_mask) * NBBY)/* bits per mask */
#ifndef howmany
#ifdef sun386
#define howmany(x, y) (((u_int)(x))+(((u_int)(y))-1))/((u_int)(y))
#else
#define howmany(x, y) (((x)+((y)-1))/(y))
#endif
#endif
#endif

typedef struct fd_set {
    fd_mask fds_bits[howmany(FD_SETSIZE, NFDBITS)];
} fd_set;

typedef char *    addr_t;

#define FD_SET(n, p) ((p)->fds_bits[(n)/NFDBITS] |= (1 << ((n) % NFDBITS)))
#define FD_CLR(n, p) ((p)->fds_bits[(n)/NFDBITS] &= ~(1 << ((n) % NFDBITS)))
#define FD_ISSET(n, p) ((p)->fds_bits[(n)/NFDBITS] & (1 << ((n) % NFDBITS)))
#define FD_ZERO(p) bzero((char *) (p), sizeof(*(p)))

#ifdef sparc
/*
 * routines that call setjmp have strange control flow graphs,
 * since a call to a routine that calls resume/longjmp will eventually
 * return at the setjmp site, not the original call site. This
 * utterly wrecks control flow analysis.
 */

```

```
extern int setjmp();  
#pragma unknown_control_flow(setjmp)  
#endif sparc
```

```
#endif _TYPES_
```

The form *daddr_t* is used for disk addresses, 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

adb(1), lseek(2), time(3C), fs(5)

NAME

tzfile – time zone information

SYNOPSIS

#include <tzfile.h>

DESCRIPTION

The time zone information files used by `tzset(3V)` begin with bytes reserved for future use, followed by three four-byte values of type `long`, written in a “standard” byte order (the high-order byte of the value is written first). These values are, in order:

<code>tzh_timecnt</code>	The number of “transition times” for which data is stored in the file.
<code>tzh_typecnt</code>	The number of “local time types” for which data is stored in the file (must not be zero).
<code>tzh_charcnt</code>	The number of characters of “time zone abbreviation strings” stored in the file.

The above header is followed by `tzh_timecnt` four-byte values of type `long`, sorted in ascending order. These values are written in “standard” byte order. Each is used as a transition time (as returned by `gettimeofday(2)`) at which the rules for computing local time change. Next come `tzh_timecnt` one-byte values of type `unsigned char`; each one tells which of the different types of “local time” types described in the file is associated with the same-indexed transition time. These values serve as indices into an array of `ttinfo` structures that appears next in the file; these structures are defined as follows:

```
struct ttinfo {
    long      tt_gmtoff;
    int       tt_isdst;
    unsigned int tt_abbrind;
};
```

Each structure is written as a four-byte value for `tt_gmtoff` of type `long`, in a standard byte order, followed by a one-byte value for `tt_isdst` and a one-byte value for `tt_abbrind`. In each structure, `tt_gmtoff` gives the number of seconds to be added to GMT, `tt_isdst` tells whether `tm_isdst` should be set by `localtime` (see `ctime(3)`) and `tt_abbrind` serves as an index into the array of time zone abbreviation characters that follow the `ttinfo` structure(s) in the file.

`localtime` uses the first standard-time `ttinfo` structure in the file (or simply the first `ttinfo` structure in the absence of a standard-time structure) if either `tzh_timecnt` is zero or the time argument is less than the first transition time recorded in the file.

SEE ALSO

gettimeofday(2), ctime(3), localtime(3), tzset(3V)

NAME

updaters – configuration file for YP updating

SYNOPSIS

/var/yp/updaters

DESCRIPTION

The file */var/yp/updaters* is a makefile (see **make(1)**) which is used for updating YP databases. Databases can only be updated in a secure network, that is, one that has a **publickey(5)** database. Each entry in the file is a make target for a particular YP database. For example, if there is a YP database named **passwd.byname** that can be updated, there should be a make target named **passwd.byname** in the *updaters* file with the command to update the file.

The information necessary to make the update is passed to the update command through standard input. The information passed is described below (all items are followed by a NEWLINE, except for 4 and 6)

- Network name of client wishing to make the update (a string)
- Kind of update (an integer)
- Number of bytes in key (an integer)
- Actual bytes of key
- Number of bytes in data (an integer)
- Actual bytes of data

After getting this information through standard input, the command to update the particular database should decide whether the user is allowed to make the change. If not, it should exit with the status **YPERR_ACCESS**. If the user is allowed to make the change, the command should make the change and exit with a status of zero. If there are any errors that may prevent the updater from making the change, it should exit with the status that matches a valid YP error code described in *<rpcsvc/ypclnt.h>*.

FILES

/var/yp/updaters

SEE ALSO

make(1), **ypupdate(3N)**, **publickey(5)**, **ypupdated(8C)**

NAME

utmp, wtmp, lastlog – login records

SYNOPSIS

```
#include <utmp.h>
#include <lastlog.h>
```

DESCRIPTION**utmp file**

The **utmp** file records information about who is currently using the system. The file is a sequence of **utmp** structure entries. That structure is defined in **<utmp.h>**, and contains the following members:

ut_line	Character array containing the name of the terminal on which the user logged in.
ut_name	Character array containing the name of the user who logged in.
ut_host	Character array containing the name of the host from which the user remotely logged in, if they logged in from another host; otherwise, a null string.
ut_time	long containing the time at which the user logged in, in seconds since 00:00 GMT, January 1, 1970.

Whenever a user logs in, **login(1)** fills in the entry in **/etc/utmp** for the terminal on which the user logged in. When they log out, **init(8)** clears that entry by setting **ut_name** and **ut_host** to null strings and **ut_time** to the time at which the user logged out.

Some window systems will make entries in **utmp** for terminal emulation windows running shells, so that library routines such as **getlogin** will work correctly in that window. These entries do not directly represent logged-in users; they are associated with a user who has already logged into the system on another terminal. These entries generally have a **ut_line** field that refers to a pseudo-terminal, and a **ut_host** field that is a null string. The macro **nonuser**, defined in **<utmp.h>**, takes a pointer to a **utmp** structure as an argument and, if the entry has a **ut_line** field that refers to a pseudo-terminal, and a **ut_host** field that is a null string, will return 1; otherwise, it will return 0. This can be used by programs that print information about logged-in users if they should not list entries made for logged-in users' additional windows.

wtmp file

The **wtmp** file records all logins and logouts. It also consists of a sequence of **utmp** entries.

Whenever a user logs in, **login** appends a record identical to the record it placed in **utmp** to the end of **/var/adm/wtmp**. Whenever a user logs out, **init** appends a record with **ut_line** equal to the terminal that the user was logged in on, **ut_name** and **ut_host** null, and **ut_time** equal to the time at which the user logged out.

When the system is shut down, **init** appends a record with a **ut_line** of **~**, a **ut_name** of **shutdown**, a null **ut_host**, and a **ut_time** equal to the time at which the shutdown occurred. When the system is rebooted, **init** appends a record with a **ut_line** of **~**, a **ut_name** of **reboot**, a null **ut_host**, and a **ut_time** equal to the time at which **init** wrote the record.

When the **date** command is used to change the system-maintained time, **date** appends a record with a **ut_line** of **|**, **ut_name** and **ut_host** null, and **ut_time** equal to the system time before the change, and then appends a record with a **ut_line** of **{**, **ut_name** and **ut_host** null, and **ut_time** equal to the system time after the change.

None of the programs that maintain **wtmp** create the file, so that if record-keeping is to be enabled, it must be created by hand as a zero-length file, and if it is removed, record-keeping is turned off. It is summarized by **ac(8)**.

As **wtmp** is appended to whenever a user logs in or out, it should be truncated periodically so that it does not consume all the disk space on its file system.

lastlog file

The **lastlog** file records the most recent login-date for every user logged in. The file is a sequence of **lastlog** structure entries. That structure is defined in **<lastlog.h>**, and contains the following members:

ll_time long containing the time at which the user logged in, in seconds since 00:00 GMT, January 1, 1970.

ll_line Character array containing the name of the terminal on which the user logged in.

ll_host Character array containing the name of the host from which the user remotely logged in, if they logged in from another host; otherwise, a null string.

When reporting (and updating) the most recent login date, **login** performs an **lseek(2)** to a byte-offset in **/var/adm/lastlog** corresponding to the **userid**. Because the count of **userids** may be high, whereas the number actual users may be small within a network environment, the bulk of this file may never be allocated by the file system even though an offset may appear to be quite large. Although **ls(1V)** may show it to be large, chances are that this file need not truncated. **du(1V)** will report the correct (smaller) amount of space actually allocated to it.

FILES

/etc/utmp
/var/adm/wtmp
/var/adm/lastlog

SEE ALSO

login(1), **who(1)**, **ac(8)**, **init(8)**

NAME

uuencode – format of an encoded uuencode file

DESCRIPTION

Files output by `uuencode(1C)` consist of a header line, followed by a number of body lines, and a trailer line. `uudecode` (see `uuencode(1C)`) will ignore any lines preceding the header or following the trailer. Lines preceding a header must not, of course, look like a header.

The header line is distinguished by having the first 6 characters 'begin'. The word `begin` is followed by a mode (in octal), and a string which names the remote file. Spaces separate the three items in the header line.

The body consists of a number of lines, each at most 62 characters long (including the trailing `NEWLINE`). These consist of a character count, followed by encoded characters, followed by a `NEWLINE`. The character count is a single printing character, and represents an integer, the number of bytes the rest of the line represents. Such integers are always in the range from 0 to 63 and can be determined by subtracting the character space (octal 40) from the character.

Groups of 3 bytes are stored in 4 characters, 6 bits per character. All are offset by a `SPACE` to make the characters printing. The last line may be shorter than the normal 45 bytes. If the size is not a multiple of 3, this fact can be determined by the value of the count on the last line. Extra garbage will be included to make the character count a multiple of 4. The body is terminated by a line with a count of zero. This line consists of one ASCII `SPACE`.

The trailer line consists of `end` on a line by itself.

SEE ALSO

`mail(1)`, `uuencode(1C)`, `uucp(1C)`, `uusend(1C)`

NAME

vfont – font formats

SYNOPSIS

```
#include <vfont.h>
```

DESCRIPTION

The fonts used by the window system and printer/plotters have the following format. Each font is in a file, which contains a header, an array of character description structures, and an array of bytes containing the bit maps for the characters. The header has the following format:

```
struct header {
    short      magic;           /* Magic number VFONT_MAGIC */
    unsigned shortsize;       /* Total # bytes of bitmaps */
    short      maxx;          /* Maximum horizontal glyph size */
    short      maxy;          /* Maximum vertical glyph size */
    short      xtend;         /* (unused) */
};
#define VFONT_MAGIC           0436
```

maxx and *maxy* are intended to be the maximum horizontal and vertical size of any glyph in the font, in raster lines. (A glyph is just a printed representation of a character, in a particular size and font.) The size is the total size of the bit maps for the characters in bytes. The *xtend* field is not currently used.

After the header is an array of NUM_DISPATCH structures, one for each of the possible characters in the font. Each element of the array has the form:

```
struct dispatch {
    unsigned shortaddr;       /* &(glyph) - &(start of bitmaps) */
    short      nbytes;        /* # bytes of glyphs (0 if no glyph) */
    char      up, down, left, right; /* Widths from baseline point */
    short      width;         /* Logical width, used by troff */
};
#define NUM_DISPATCH         256
```

The *nbytes* field is nonzero for characters which actually exist. For such characters, the *addr* field is an offset into the bit maps to where the character's bit map begins. The *up*, *down*, *left*, and *right* fields are offsets from the base point of the glyph to the edges of the rectangle which the bit map represents. (The imaginary "base point" is a point which is vertically on the "base line" of the glyph (the bottom line of a glyph which does not have a descender) and horizontally near the left edge of the glyph; often 3 or so pixels past the left edge.) The bit map contains *up+down* rows of data for the character, each of which has *left+right* columns (bits). Each row is rounded up to a number of bytes. The *width* field represents the logical width of the glyph in bits, and shows the horizontal displacement to the base point of the next glyph.

FILES

```
/usr/lib/vfont/*
/usr/lib/fonts/fixedwidthfonts/*
```

SEE ALSO

```
troff(1), vfontinfo(1), vswap(1)
```

BUGS

A machine-independent font format should be defined. The shorts in the above structures contain different bit patterns depending whether the font file is for use on a VAX or a Sun. The vswap program must be used to convert one to the other.

NAME

vgrindefs – vgrind's language definition data base

SYNOPSIS

/usr/lib/vgrindefs

DESCRIPTION

vgrindefs contains all language definitions for vgrind. The data base is very similar to termcap(5). Capabilities in vgrindefs are of two types: Boolean capabilities which indicate that the language has some particular feature and string capabilities which give a regular expression or keyword list. Entries may continue onto multiple lines by giving a \ as the last character of a line. Lines starting with # are comments.

Capabilities

The following table names and describes each capability.

Name	Type	Description
ab	str	Regular expression for the start of an alternate form comment
ae	str	Regular expression for the end of an alternate form comment
bb	str	Regular expression for the start of a block
be	str	Regular expression for the end of a lexical block
cb	str	Regular expression for the start of a comment
ce	str	Regular expression for the end of a comment
id	str	String giving characters other than letters and digits that may legally occur in identifiers (default '_')
kw	str	A list of keywords separated by spaces
lb	str	Regular expression for the start of a character constant
le	str	Regular expression for the end of a character constant
oc	bool	Present means upper and lower case are equivalent
pb	str	Regular expression for start of a procedure
pl	bool	Procedure definitions are constrained to the lexical level matched by the 'px' capability
px	str	A match for this regular expression indicates that procedure definitions may occur at the next lexical level. Useful for lisp-like languages in which procedure definitions occur as subexpressions of defuns.
sb	str	Regular expression for the start of a string
se	str	Regular expression for the end of a string
tc	str	Use the named entry as a continuation of this one
tl	bool	Present means procedures are only defined at the top lexical level

Regular Expressions

vgrindefs uses regular expressions similar to those of ex(1) and lex(1). The characters '^', '\$', ':', and '\ are reserved characters and must be 'quoted' with a preceding \ if they are to be included as normal characters. The metasympols and their meanings are:

\$	The end of a line
^	The beginning of a line
\d	A delimiter (space, tab, newline, start of line)
\a	Matches any string of symbols (like '.' in lex)
\p	Matches any identifier. In a procedure definition (the 'pb' capability) the string that matches this symbol is used as the procedure name.
()	Grouping
	Alternation
?	Last item is optional
\e	Preceding any string means that the string will not match an input string if the input string is preceded by an escape character (\). This is typically used for languages (like C) that can include the string delimiter in a string by escaping it.

Unlike other regular expressions in the system, these match words and not characters. Hence something like '(tramp|steamer)flies?' would match 'tramp', 'steamer', 'trampflies', or 'steamerflies'. Contrary to some forms of regular expressions, vgrindef alternation binds very tightly. Grouping parentheses are likely to be necessary in expressions involving alternation.

Keyword List

The keyword list is just a list of keywords in the language separated by spaces. If the 'oc' boolean is specified, indicating that upper and lower case are equivalent, then all the keywords should be specified in lower case.

EXAMPLE

The following entry, which describes the C language, is typical of a language entry.

```
C|c|the C programming language:\
:pb=`d?*?d?\p\d??):bb={:be=}:cb=/*:ce=/:sb=":se=\e":\
:lb=':le=\e':tl:\
:kw=asm auto break case char continue default do double else enum\
extern float for fortran goto if int long register return short\
sizeof static struct switch typedef union unsigned while #define\
#else #endif #if #ifdef #ifndef #include #undef # define else endif\
if ifdef ifndef include undef:
```

Note that the first field is just the language name (and any variants of it). Thus the C language could be specified to vgrind(1) as 'c' or 'C'.

FILES

/usr/lib/vgrindefs file containing terminal descriptions

SEE ALSO

vgrind(1), troff(1)

NAME

ypfiles – the Yellow Pages database and directory structure

DESCRIPTION

The Yellow Pages (YP) network lookup service uses a distributed, replicated database of **dbm** files contained in the `/var/yp` directory hierarchy on each YP server. A **dbm** database consists of two files, created by calls to the **ndbm(3)** library package. One has the filename extension `.pag` and the other has the filename extension `.dir`. For instance, the database named `hosts.byname`, is implemented by the pair of files `hosts.byname.pag` and `hosts.byname.dir`.

A **dbm** database served by the YP is called a YP *map*. A YP *domain* is a subdirectory of `/var/yp` containing a set of YP maps. Any number of YP domains can exist. Each may contain any number of maps.

No maps are required by the YP lookup service itself, although they may be required for the normal operation of other parts of the system. There is no list of maps which YP serves — if the map exists in a given domain, and a client asks about it, the YP will serve it. For a map to be accessible consistently, it must exist on all YP servers that serve the domain. To provide data consistency between the replicated maps, an entry to run **ypxfr** periodically should be made in the super-user's `crontab` file on each server. More information on this topic is in **ypxfr(8)**.

YP maps should contain two distinguished key-value pairs. The first is the key `YP_LAST_MODIFIED`, having as a value a ten-character ASCII order number. The order number should be the system time in seconds when the map was built. The second key is `YP_MASTER_NAME`, with the name of the YP master server as a value. **makedbm(8)** generates both key-value pairs automatically. A map that does not contain both key-value pairs can be served by the YP, but the **ypserv** process will not be able to return values for “Get order number” or “Get master name” requests. See **ypserv(8)**. In addition, values of these two keys are used by **ypxfr** when it transfers a map from a master YP server to a slave. If **ypxfr** cannot figure out where to get the map, or if it is unable to determine whether the local copy is more recent than the copy at the master, you must set extra command line switches when you run it.

YP maps must be generated and modified only at the master server. They are copied to the slaves using **ypxfr(8)** to avoid potential byte-ordering problems among YP servers running on machines with different architectures, and to minimize the amount of disk space required for the **dbm** files. The YP database can be initially set up for both masters and slaves by using **ypinit(8)**.

After the server databases are set up, it is probable that the contents of some maps will change. In general, some ASCII source version of the database exists on the master, and it is changed with a standard text editor. The update is incorporated into the YP map and is propagated from the master to the slaves by running `/var/yp/Makefile`. All Sun-supplied maps have entries in `/var/yp/Makefile`; if you add a YP map, edit this file to support the new map. The makefile uses **makedbm(8)** to generate the YP map on the master, and **yppush(8)** to propagate the changed map to the slaves. **yppush** is a client of the map **ypservers**, which lists all the YP servers. For more information on this topic, see **yppush(8)**.

FILES

`/var/yp`
`/var/yp/Makefile`

SEE ALSO

dbm(3X), **makedbm(8)**, **rpcinfo(8C)**, **ypinit(8)**, **ypmake(8)**, **yppoll(8)**, **yppush(8)**, **ypserv(8)**, **ypxfr(8)**,



NAME

intro – introduction to games and demos

DESCRIPTION

This section describes available games and demos.

LIST OF GAMES AND DEMOS

Name	Appears on Page	Description
adventure	adventure(6)	an exploration game
arithmetic	arithmetic(6)	provide drill in number facts
backgammon	backgammon(6)	the game of backgammon
banner	banner(6)	print large banner on printer
battlestar	battlestar(6)	a tropical adventure game
bcd	bcd(6)	convert to antique media
bj	bj(6)	the game of black jack
boggle	boggle(6)	play the game of boggle
boggletool	boggletool(6)	play a game of boggle
bouncedemo	graphics_demos(6)	graphics demonstration programs
canfield	canfield(6)	Canfield solitaire card game
canfieldtool	canfield(6)	Canfield solitaire card game
canvas_demo	sunview_demos(6)	Window-System demonstration programs
cfcores	canfield(6)	Canfield solitaire card game
chase	chase(6)	try to escape to killer robots
chess	chess(6)	the game of chess
chesstool	chesstool(6)	window-based front-end to chess program
ching	ching(6)	the book of changes and other cookies
craps	craps(6)	the game of craps
cribbage	cribbage(6)	the card game cribbage
cursor_demo	sunview_demos(6)	Window-System demonstration programs
factor	factor(6)	factor a number, generate large primes
fish	fish(6)	play "Go Fish"
fortune	fortune(6)	print a random, hopefully interesting, adage
framedemo	graphics_demos(6)	graphics demonstration programs
gammontool	gammontool(6)	play a game of backgammon
graphics_demos	graphics_demos(6)	graphics demonstration programs
hack	hack(6)	replacement for rogue
hangman	hangman(6)	computer version of the game hangman
hunt	hunt(6)	a multiplayer multiterminal game
jumpdemo	graphics_demos(6)	graphics demonstration programs
life	life(6)	John Conway's game of life
mille	mille(6)	play Mille Bornes
monop	monop(6)	Monopoly game
moo	moo(6)	guessing game
number	number(6)	convert Arabic numerals to English
ppt	bcd(6)	convert to antique media
primes	factor(6)	factor a number, generate large primes
primes	primes(6)	print all primes larger than some given number
quiz	quiz(6)	test your knowledge
rain	rain(6)	animated raindrops display
random	random(6)	select lines randomly from a file
robots	robots(6)	fight off villainous robots
snake	snake(6)	display chase game
snscore	snake(6)	display chase game
spheresdemo	graphics_demos(6)	graphics demonstration programs

sunview_demos
trek
worm
worms
wump

sunview_demos(6)
trek(6)
worm(6)
worms(6)
wump(6)

Window-System demonstration programs
trekkie game
play the growing worm game
animate worms on a display terminal
the game of hunt-the-wumpus

NAME

adventure – an exploration game

SYNOPSIS

/usr/games/adventure

DESCRIPTION

The object of the game is to locate and explore Colossal Cave, find the treasures hidden there, and bring them back to the building with you. The program is self-describing to a point, but part of the game is to discover its rules.

To terminate a game, type **quit**; to save a game for later resumption, type **suspend**.

BUGS

Saving a game creates a large executable file instead of just the information needed to resume the game.

NAME

arithmetic – provide drill in number facts

SYNOPSIS

/usr/games/arithmetic [+-x/] [range]

DESCRIPTION

arithmetic types out simple arithmetic problems, and waits for an answer to be typed in. If the answer is correct, it types back "Right!", and a new problem. If the answer is wrong, it replies "What?", and waits for another answer. Every twenty problems, it publishes statistics on correctness and the time required to answer.

To quit the program, type an interrupt (DELETE character).

The first optional argument determines the kind of problem to be generated; '+', '-', 'x', '/' respectively cause addition, subtraction, multiplication, and division problems to be generated. One or more characters can be given; if more than one is given, the different types of problems will be mixed in random order; default is +-.

range is a decimal number; all addends, subtrahends, differences, multiplicands, divisors, and quotients will be less than or equal to the value of *range*. Default *range* is 10.

At the start, all numbers less than or equal to *range* are equally likely to appear. If the respondent makes a mistake, the numbers in the problem which was missed become more likely to reappear.

As a matter of educational philosophy, the program will not give correct answers, since the learner should, in principle, be able to calculate them. Thus the program is intended to provide drill for someone just past the first learning stage, not to teach number facts *de novo*. For almost all users, the relevant statistic should be time per problem, not percent correct.

NAME

backgammon – the game of backgammon

SYNOPSIS

backgammon [-] [n r w b pr pw pb tterm *sfilename*]

DESCRIPTION

backgammon lets you play backgammon against the computer or against a 'friend'. All commands only are one letter, so you don't need to type a carriage return, except at the end of a move. **backgammon** is mostly self documenting, so that a q ? (question mark) will usually get some help. If you answer y when **backgammon** asks if you want the rules, you will get text explaining the rules of the game, some hints on strategy, instruction on how to use **backgammon**, and a tutorial consisting of a practice game against the computer. A description of how to use **backgammon** can be obtained by answering y when it asks if you want instructions. The possible arguments for **backgammon** (most are unnecessary but some are very convenient) consist of:

n	don't ask for rules or instructions
r	player is red (implies n)
w	player is white (implies n)
b	two players, red and white (implies n)
pr	print the board before red's turn
pw	print the board before white's turn
pb	print the board before both player's turn
tterm	terminal is type <i>term</i> , uses <i>letc/termcap</i> , otherwise uses the TERM environment variable.
sfile	recover previously saved game from <i>file</i> . This can also be done by executing the saved file, that is, typing its name in as a command.

Arguments may be optionally preceded by a - sign. Several arguments may be concatenated together, but not after s or t arguments, since they can be followed by an arbitrary string. Any unrecognized arguments are ignored. An argument of a lone - gets a description of possible arguments.

If **term** has capabilities for direct cursor movement. **backgammon** 'fixes' the board after each move, so the board does not need to be reprinted, unless the screen suffers some horrendous malady. Also, any 'p' option will be ignored.

QUICK REFERENCE

When **backgammon** prompts by typing only your color, type a space or carriage return to roll, or

d	to double
p	to print the board
q	to quit
s	to save the game for later

When **backgammon** prompts with 'Move:', type

p	to print the board
q	to quit
s	to save the game

or a *move*, which is a sequence of

s-f	move from s to f
s/r	move one man on s the roll r separated by commas or spaces and ending with a newline. Available abbreviations are

s-f1-f2 means **s-f1,f1-f2**

s/r1r2 means **s/r1,s/r2**

Use **b** for bar and **h** for home, or **0** or **25** as appropriate.

FILES

/usr/games/teachgammon	rules and tutorial
/etc/termcap	terminal capabilities

BUGS

backgammon's strategy needs much work.

NAME

banner – print large banner on printer

SYNOPSIS

/usr/games/banner [**-wn**] message ...

DESCRIPTION

banner prints a large, high quality banner on the standard output. If the message is omitted, it prompts for and reads one line of its standard input. If **-w** is given, the output is reduced from a width of 132 to *n*, suitable for a narrow terminal. If *n* is omitted, it defaults to 80.

The output should be printed on a hard-copy device, up to 132 columns wide, with no breaks between the pages. The volume is enough that you want a printer or a fast hardcopy terminal, but if you are patient, a decwriter or other 300 baud terminal will do.

BUGS

Several ASCII characters are not defined, notably '<', '>', '[', ']', '\', '^', '_', '{', '}', '|', and '~'.

Also, the characters '"', "'", and '&' are funny looking (but in a useful way.)

The **-w** option is implemented by skipping some rows and columns. The smaller it gets, the grainier the output. Sometimes it runs letters together.

NAME

battlestar – a tropical adventure game

SYNOPSIS

battlestar [-r]

DESCRIPTION

battlestar is an adventure game in the classic style. However, it is slightly less of a puzzle and more a game of exploration. There are a few magical words in the game, but on the whole, simple English should suffice to make one's desires understandable to the parser.

OPTIONS

-r Recover a saved game.

THE SETTING

In the days before the darkness came, when battlestars ruled the heavens...

Three He made and gave them to His daughters,
 Beautiful nymphs, the goddesses of the waters.
 One to bring good luck and simple feats of wonder,
 Two to wash the lands and churn the waves asunder,
 Three to rule the world and purge the skies with thunder.

In those times great wizards were known and their powers were beyond belief. They could take any object from thin air, and, uttering the word 'su', could disappear.

In those times men were known for their lust of gold and desire to wear fine weapons. Swords and coats of mail were fashioned that could withstand a laser blast.

But when the darkness fell, the rightful reigns were toppled. Swords and helms and heads of state went rolling across the grass. The entire fleet of battlestars was reduced to a single ship.

USAGE**Sample Commands**

take	---	take an object
drop	---	drop an object
wear	---	wear an object you are holding
draw	---	carry an object you are wearing
puton	---	take an object and wear it
take off	---	draw an object and drop it
throw	<object> <direction>	
!	<shell esc>	

Implied Objects

>: take watermelon
 watermelon:
 Taken.
 >: eat
 watermelon:
 Eaten.
 >: take knife and sword and apple, drop all
 knife:
 Taken.
 broadsword:
 Taken.
 apple:
 Taken.
 knife:
 Dropped.

broadsword:
Dropped.
apple:
Dropped.
>: get
knife:
Taken.

Notice that the "shadow" of the next word stays around if you want to take advantage of it. That is, saying 'take knife' and then 'drop' will drop the knife you just took.

Score and Inven

The two commands `score` and `inven` will print out your current status in the game.

Saving a Game

The command `save` will save your game in a file called `Bstar`. You can recover a saved game by using the `-r` option when you start up the game.

Directions

The compass directions N, S, E, and W can be used if you have a compass. If you do not have a compass, you will have to say R, L, A, or B, which stand for Right, Left, Ahead, and Back. Directions printed in room descriptions are always printed in R, L, A, & B relative directions.

BUGS

Countless.

NAME

bcd, ppt – convert to antique media

SYNOPSIS

/usr/games/bcd *text*

/usr/games/ppt

DESCRIPTION

bcd converts the literal *text* into a form familiar to old-timers.

ppt converts the standard input into yet another form.

SEE ALSO

dd(1)

NAME

bj – the game of black jack

SYNOPSIS

/usr/games/bj

DESCRIPTION

bj is a serious attempt at simulating the dealer in the game of black jack (or twenty-one) as might be found in Reno. The following rules apply:

The bet is \$2 every hand.

A player “natural” (black jack) pays \$3. A dealer natural loses \$2. Both dealer and player naturals is a “push” (no money exchange).

If the dealer has an ace up, the player is allowed to make an “insurance” bet against the chance of a dealer natural. If this bet is not taken, play resumes as normal. If the bet is taken, it is a side bet where the player wins \$2 if the dealer has a natural and loses \$1 if the dealer does not.

If the player is dealt two cards of the same value, he is allowed to “double”. He is allowed to play two hands, each with one of these cards. (The bet is doubled also; \$2 on each hand.)

If a dealt hand has a total of ten or eleven, the player may “double down”. He may double the bet (\$2 to \$4) and receive exactly one more card on that hand.

Under normal play, the player may “hit” (draw a card) as long as his total is not over twenty-one. If the player “busts” (goes over twenty-one), the dealer wins the bet.

When the player “stands” (decides not to hit), the dealer hits until he attains a total of seventeen or more. If the dealer busts, the player wins the bet.

If both player and dealer stand, the one with the largest total wins. A tie is a push.

The machine deals and keeps score. The following questions will be asked at appropriate times. Each question is answered by y followed by a new-line for “yes”, or just new-line for “no”.

? (this means, “do you want a hit?”)

Insurance?

Double down?

Every time the deck is shuffled, the dealer so states and the “action” (total bet) and “standing” (total won or lost) is printed. To exit, hit the interrupt key (CTRL-C) and the action and standing will be printed.

NAME

boggle – play the game of boggle

SYNOPSIS

`/usr/games/boggle [+] [++]`

AVAILABILITY

This game is available with the *Games* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

This program is intended for people wishing to sharpen their skills at Boggle (TM Parker Bros.). If you invoke the program with 4 arguments of 4 letters each, (e.g. “`boggle appl epie moth erhd`”) the program forms the obvious Boggle grid and lists all the words from `/usr/dict/words` found therein. If you invoke the program without arguments, it will generate a board for you, let you enter words for 3 minutes, and then tell you how well you did relative to `/usr/dict/words`.

The object of Boggle is to find, within 3 minutes, as many words as possible in a 4 by 4 grid of letters. Words may be formed from any sequence of 3 or more adjacent letters in the grid. The letters may join horizontally, vertically, or diagonally. However, no position in the grid may be used more than once within any one word. In competitive play amongst humans, each player is given credit for those of his words which no other player has found.

In interactive play, enter your words separated by spaces, tabs, or newlines. A bell will ring when there is 2:00, 1:00, 0:10, 0:02, 0:01, and 0:00 time left. You may complete any word started before the expiration of time. You can surrender before time is up by hitting ‘break’. While entering words, your erase character is only effective within the current word and your line kill character is ignored.

Advanced players may wish to invoke the program with 1 or 2 +’s as the first argument. The first + removes the restriction that positions can only be used once in each word. The second + causes a position to be considered adjacent to itself as well as its (up to) 8 neighbors.

NAME

boggletool – play a game of boggle

SYNOPSIS

`/usr/games/chesstool [number] [+[+]] [16-character string]`

AVAILABILITY

This game is available with the *Games* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

boggletool allows you to play the game of Boggle (TM Parker Bros.) against the computer. The *number* argument specifies the time limit in minutes (the default is 3 minutes). If a 16 character long string is placed on the command line, it is interpreted as a Boggle board: the first four letters form the top row, the next four letters the second row, etc. If no letters are specified, a board is randomly rolled by the computer from a set of Boggle cubes. The +[+] argument is explained below under **Advanced Play**.

PLAYING THE GAME**Rules of the Game**

The object of Boggle is to find as many words as possible in a 4 by 4 grid of letters within a certain time limit. Words may be formed from any sequence of 3 or more adjacent letters in the grid. The letters may join horizontally, vertically, or diagonally. Normally, no letter in the grid may be used more than once in a word (see **Advanced Play** for exceptions).

Playing the Game

When invoked, boggletool displays a grid of letters and an hourglass. To enter words, simply type in lower case letters to spell the word you want. Use any whitespace (SPACE, TAB, or NEWLINE) to finish a word. To correct mistakes you make, use BACKSPACE or DEL to delete the last character, or use CTRL-U to delete an entire word. boggletool verifies that words you enter are both in the grid and are valid English words. If you type in a character which would form a word which is not in the grid, the display will flash and the character you typed will not be echoed. When you type any whitespace to end the current word, boggletool will verify that the word is three or more letters long and that it appears in the dictionary. If the word you typed is illegal for either reason, the display will flash and you will have to either erase the word or change it. If you try to enter a valid word which you have already entered, the display will flash and the previous occurrence of the word will be highlighted. Again, you will have to erase the word before continuing. As you enter words, the "sand" in the hourglass will fall. At the end of the time limit, the display will flash and you will no longer be allowed to enter words. After a moment, the computer will display two lists of words: the words you found, and other words which also appear in the grid. To play another game, just type any capital letter (or use the pop-up menu).

Using the Menu

The pop-up menu is invoked by pressing the RIGHT mouse button. There are four items in it, and they work as follows.

Restart Game

Create a new boggletool a new board, reset the timer, and allow you to start from scratch.

Restart Timer

Allows you to cheat by resetting the hourglass timer to zero.

Give Up

End the game and print the results immediately.

Quit Allows you to quit running the boggletool program. A prompt appears asking you to confirm the quit; when it does, click the LEFT mouse button to quit or the RIGHT mouse button to abort the quit.

Advanced Play

There are two options for advanced players. If a single + appears on the command line, letters in the grid may be reused. If two +'s are on the command line, letters may also be considered adjacent to themselves

as well as to their neighbors. Although it is far easier to find words with these two options, there are also many more possible words in the grid and it is therefore difficult to find them all.

FILES

/usr/games/boggledict dictionary file for computer's words

NAME

canfield, canfieldtool, cfscores – Canfield solitaire card game

SYNOPSIS

`/usr/games/canfield` [`-ac`]

`/usr/games/canfieldtool` [`-ac`]

`/usr/games/cfscores` [`-ac`] [*username*]

AVAILABILITY

These games are available with the *Games* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

canfield can be played on any terminal. **canfieldtool** is the SunView version with attractive graphics.

If you have never played solitaire before, it is recommended that you consult a solitaire instruction book. In **canfield**, tableau cards may be built on each other downward in alternate colors. An entire pile must be moved as a unit in building. Top cards of the piles are available to be able to be played on foundations, but never into empty spaces.

Spaces must be filled from the stock. The top card of the stock also is available to be played on foundations or built on tableau piles. After the stock is exhausted, tableau spaces may be filled from the talon and the player may keep them open until he wishes to use them.

Cards are dealt from the hand to the talon by threes and this repeats until there are no more cards in the hand or the player quits. To have cards dealt onto the talon the player types **ht** for his move. Foundation base cards are also automatically moved to the foundation when they become available.

Canfieldtool

Once you understand the rules, **canfieldtool** is self-explanatory.

Canfield

The rules for betting are somewhat less strict than those used in the official version of the game. The initial deal costs \$13. You may quit at this point or inspect the game. Inspection costs \$13 and allows you to make as many moves as is possible without moving any cards from your hand to the talon. (The initial deal places three cards on the talon; if all these cards are used, three more are made available.) Finally, if the game seems interesting, you must pay the final installment of \$26. At this point you are credited at the rate of \$5 for each card on the foundation; as the game progresses you are credited with \$5 for each card that is moved to the foundation. Each run through the hand after the first costs \$5. The card counting feature costs \$1 for each unknown card that is identified. If the information is toggled on, you are only charged for cards that became visible since it was last turned on. Thus the maximum cost of information is \$34. Playing time is charged at a rate of \$1 per minute. If the `-a` flag is specified, it prints out the canfield accounts for all users that have played the game since the database was set up.

OPTIONS

- a** Print out canfield accounts for all users that have played the game since the database was set up.
- c** Maintain card counting statistics on the bottom of the screen. When properly used this can greatly increase the chances of winning.

With no arguments, **cfscores** prints out the current status of your canfield account. If *username* is specified, it prints out the status of their account.

FILES

`/usr/games/canfield` the game itself
`/usr/games/lib/cfscores` the database of scores

BUGS

It is impossible to cheat.

NAME

chase – try to escape to killer robots

SYNOPSIS

`/usr/games/chase [nrobots] [nfences]`

DESCRIPTION

The object of the game `chase` is to move around inside of the box on the screen without getting eaten by the robots chasing and without running into anything.

If a robot runs into another robot while chasing you, they crash and leave a junk heap. If a robot runs into a fence, it is destroyed.

If you can survive until all the robots are destroyed, you have won!

If you do not specify either *nrobots* or *nfences*, `chase` will prompt you for them.

NAME

chess – the game of chess

SYNOPSIS

`/usr/games/chess`

AVAILABILITY

This game is available for Sun-2, Sun-3 and Sun-4 systems with the *Games* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

chess is a computer program that plays class D chess. Moves may be given either in standard (descriptive) notation or in algebraic notation. The symbol '+' is used to specify check; 'o-o' and 'o-o-o' specify castling. To play black, type 'first'; to print the board, type an empty line.

Each move is echoed in the appropriate notation followed by the program's reply.

DIAGNOSTICS

The most cryptic diagnostic is 'eh?' which means that the input was syntactically incorrect.

FILES

`/usr/games/lib/chess.book`
book of opening moves

BUGS

Pawns may be promoted only to queens.

NAME

chesstool – window-based front-end to chess program

SYNOPSIS

`/usr/games/chesstool [chess_program]`

AVAILABILITY

This game is available for Sun-2, Sun-3 and Sun-4 systems, with the *Games* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

chesstool is a window-based front-end to the **chess(6)** program. Used without options, **chesstool** uses `/usr/games/chess`; you can designate any alternate program which uses the same command syntax as **chess(6)** with the *chess_program* argument.

When **chesstool** starts up, it displays a large window with three subwindows. The first subwindow displays messages 'Illegal move', for example. The second subwindow is an options subwindow; options are described below. The final subwindow is a chessboard display with white and black pieces and two (advisory only) timekeeping clocks.

Make your moves with the mouse: select a piece by positioning the arrow cursor over the piece and pressing the left mouse button down, then drag the piece to the destination square, and release the button. The cursor will then turn to an hourglass icon while the system plays.

Items in the subwindow may be selected with either the left or middle mouse buttons. These options are:

- | | |
|----------------------|--|
| Last Play | Show the last play made. |
| Undo | Undo your last move and the machine's response.
Once the game is over, it is not possible to restart it, so undo will update the board, but the game cannot be continued from that position. |
| Flash | Flash when the machine has completed its move.
When this command is selected, a check mark will appear next to the word Flash . In flash mode, if chesstool is open, the piece moved by the system on its play will flash until you make your move. If chesstool is iconic, the entire icon will flash when the machine has made its move. Thus you can "Close" chesstool and be alerted when it's your turn to move. To turn flash mode off, select flash again. |
| Machine White | Start a new game with the machine playing white. |
| Human White | Start a new game with the machine playing black. |
| Quit | Exit from chesstool . |

There are two moves which are special: castling and capturing a pawn *enpassant*. To castle, move the king only. The position of the rook will automatically be updated. Since the king moves two squares when castling, the move is unambiguous. To capture *enpassant*, move the pawn to the square occupied by the opposing pawn which will be captured.

SEE ALSO

chess(6)

NAME

ching – the book of changes and other cookies

SYNOPSIS

/usr/games/ching [*hexagram*]

DESCRIPTION

The *I Ching* or *Book of Changes* is an ancient Chinese oracle that has been in use for centuries as a source of wisdom and advice.

The text of the *oracle* (as it is sometimes known) consists of sixty-four *hexagrams*, each symbolized by a particular arrangement of six straight (—) and broken (--) lines. These lines have values ranging from six through nine, with the even values indicating the broken lines.

Each *hexagram* consists of two major sections. The **Judgement** relates specifically to the matter at hand (For instance, "It furthers one to have somewhere to go.") while the **Image** describes the general attributes of the *hexagram* and how they apply to one's own life ("Thus the superior man makes himself strong and untiring.").

When any of the lines has the value six or nine, it is a moving line; for any such line there is an appended judgement which becomes significant. Furthermore, the moving lines are inherently unstable and change into their opposites; a second *hexagram* (and thus an additional judgement) is formed.

Normally, one consults the oracle by fixing the desired question firmly in mind and then casting a set of changes (lines) using yarrow-stalks or tossed coins. The resulting *hexagram* will be the answer to the question.

Using an algorithm suggested by S. C. Johnson, this oracle simply reads a question from the standard input (up to an EOF) and hashes the individual characters in combination with the time of day, process ID and any other magic numbers which happen to be lying around the system. The resulting value is used as the seed of a random number generator which drives a simulated coin-toss divination. The answer is then piped through `nroff` for formatting and will appear on the standard output.

For those who wish to remain steadfast in the old traditions, the oracle will also accept the results of a personal divination using, for example, coins. To do this, cast the change and then type the resulting line values as an argument.

The impatient modern may prefer to settle for Chinese cookies; try `fortune(6)`.

SEE ALSO

It furthers one to see the great man.

DIAGNOSTICS

The great prince issues commands,
Founds states, vests families with fiefs.
Inferior people should not be employed.

BUGS

Waiting in the mud
Brings about the arrival of the enemy.
If one is not extremely careful,
Somebody may come up from behind and strike him.
Misfortune.

NAME

craps – the game of craps

SYNOPSIS

`/usr/games/craps`

DESCRIPTION

craps is a form of the game of craps that is played in Las Vegas. The program simulates the *roller*, while the user (the *player*) places bets. The player may choose, at any time, to bet with the roller or with the *House*. A bet of a negative amount is taken as a bet with the House, any other bet is a bet with the roller.

The player starts off with a "bankroll" of \$2,000.

The program prompts with:

bet?

The bet can be all or part of the player's bankroll. Any bet over the total bankroll is rejected and the program prompts with **bet?** until a proper bet is made.

Once the bet is accepted, the roller throws the dice. The following rules apply (the player wins or loses depending on whether the bet is placed with the roller or with the House; the odds are even). The *first* roll is the roll immediately following a bet:

1. On the first roll:

7 or 11	wins for the roller;
2, 3, or 12	wins for the House;
any other number	is the <i>point</i> , roll again (Rule 2 applies).

2. On subsequent rolls:

<i>point</i>	roller wins;
7	House wins;
any other number	roll again.

If a player loses the entire bankroll, the House will offer to lend the player an additional \$2,000. The program will prompt:

marker?

A **yes** (or **y**) consummates the loan. Any other reply terminates the game.

If a player owes the House money, the House reminds the player, before a bet is placed, how many markers are outstanding.

If, at any time, the bankroll of a player who has outstanding markers exceeds \$2,000, the House asks:

Repay marker?

A reply of **yes** (or **y**) indicates the player's willingness to repay the loan. If only 1 marker is outstanding, it is immediately repaid. However, if more than 1 marker are outstanding, the House asks:

How many?

markers the player would like to repay. If an invalid number is entered (or just a carriage return), an appropriate message is printed and the program will prompt with **How many?** until a valid number is entered.

If a player accumulates 10 markers (a total of \$20,000 borrowed from the House), the program informs the player of the situation and exits.

Should the bankroll of a player who has outstanding markers exceed \$50,000, the *total* amount of money borrowed will be *automatically* repaid to the House.

Any player who accumulates \$100,000 or more breaks the bank. The program then prompts:

New game?

to give the House a chance to win back its money.

Any reply other than yes is considered to be a no (except in the case of **bet?** or **How many?**). To exit, send an interrupt (break), DELETE character or CTRL-D The program will indicate whether the player won, lost, or broke even.

MISCELLANEOUS

The random number generator for the die numbers uses the seconds from the time of day. Depending on system usage, these numbers, at times, may seem strange but occurrences of this type in a real dice situation are not uncommon.

NAME

cribbage – the card game cribbage

SYNOPSIS

/usr/games/cribbage [-eqr] name ...

DESCRIPTION

cribbage plays the card game cribbage, with **cribbage** playing one hand and the user the other. **cribbage** initially asks the user if the rules of the game are needed – if so, **cribbage** displays the appropriate section from *According to Hoyle* with **more**(1).

OPTIONS

- e Provide an explanation of the correct score when the player makes mistakes scoring his hand or crib. This is especially useful for beginning players.
- q Print a shorter form of all messages – this is only recommended for users who have played the game without specifying this option.
- r Instead of asking the player to cut the deck, **cribbage** will randomly cut the deck.

PLAYING CRIBBAGE

cribbage first asks the player whether he wishes to play a short game (“once around”, to 61) or a long game (“twice around”, to 121). A response of ‘s’ results in a short game, any other response plays a long game.

At the start of the first game, **cribbage** asks the player to cut the deck to determine who gets the first crib. The user should respond with a number between 0 and 51, indicating how many cards down the deck is to be cut. The player who cuts the lower ranked card gets the first crib. If more than one game is played, the loser of the previous game gets the first crib in the current game.

For each hand, **cribbage** first prints the player’s hand, whose crib it is, and then asks the player to discard two cards into the crib. The cards are prompted for one per line, and are typed as explained below.

After discarding, **cribbage** cuts the deck (if it is the player’s crib) or asks the player to cut the deck (if it’s its crib); in the latter case, the appropriate response is a number from 0 to 39 indicating how far down the remaining 40 cards are to be cut.

After cutting the deck, play starts with the non-dealer (the person who doesn’t have the crib) leading the first card. Play continues, as per cribbage, until all cards are exhausted. **cribbage** keeps track of the scoring of all points and the total of the cards on the table.

After play, the hands are scored. **cribbage** requests the player to score his hand (and the crib, if it is his) by printing out the appropriate cards (and the cut card enclosed in brackets). Play continues until one player reaches the game limit (61 or 121).

A carriage return when a numeric input is expected is equivalent to typing the lowest legal value; when cutting the deck this is equivalent to choosing the top card.

SPECIFYING CARDS

Cards are specified as *rank* followed by *suit*. The *ranks* may be specified as one of a, 2, 3, 4, 5, 6, 7, 8, 9, t, j, q, and k, or alternatively, one of ace, two, three, four, five, six, seven, eight, nine, ten, jack, queen, and king. *Suits* may be specified as s, h, d, and c, or alternatively as spades, hearts, diamonds, and clubs. A card may be specified as *rank suit*, or *rank of suit*. If the single letter *rank* and *suit* designations are used, the space separating the *suit* and *rank* may be left out. Also, if only one card of the desired *rank* is playable, typing the *rank* is sufficient. For example, if your hand was 2h, 4d, 5c, 6h, jc, kd and you wanted to discard the king of diamonds, you could type any of k, king, kd, k d, k of d, king d, king of d, k diamonds, k of diamonds, king diamonds, or king of diamonds,

FILES

/usr/games/cribbage

CRIBBAGE(6)

GAMES AND DEMOS

CRIBBAGE(6)

SEE ALSO

more(1)

NAME

factor, primes – factor a number, generate large primes

SYNOPSIS

/usr/games/factor [*number*]

/usr/games/primes [*number*]

DESCRIPTION

factor reads lines from its standard input. If it reads a positive number, **factor** will factor the number and print its prime factors, printing each one the proper number of times. **factor** exits when it reads zero, a negative number, or something other than a number. If a *number* is given, **factor** will factor the number, print its prime factors, and exit.

primes reads a number from the standard input and prints all primes larger than the given number and smaller than 2^{32} (about 4.3×10^9). If a *number* is given, **primes** will use that number rather than reading one from the standard input.

DIAGNOSTICS

Ouch. Input out of range or for garbage input.

NAME

fish – play “Go Fish”

SYNOPSIS

/usr/games/fish

DESCRIPTION

fish plays the game of “Go Fish”, a children’s card game. The object is to accumulate “books” of 4 cards with the same face value. The players alternate turns; each turn begins with one player selecting a card from his hand, and asking the other player for all cards of that face value. If the other player has one or more cards of that face value in his hand, he gives them to the first player, and the first player makes another request. Eventually, the first player asks for a card which is not in the second player’s hand: he replies ‘GO FISH!’ The first player then draws a card from the “pool” of undealt cards. If this is the card he had last requested, he draws again. When a book is made, either through drawing or requesting, the cards are laid down and no further action takes place with that face value.

To play the computer, simply make guesses by typing a, 2, 3, 4, 5, 6, 7, 8, 9, 10, j, q, or k k when asked. Hitting a RETURN character gives you information about the size of my hand and the pool, and tells you about my books. Saying ‘p’ as a first guess puts you into “pro” level; the default is pretty dumb.

NAME

fortune – print a random, hopefully interesting, adage

SYNOPSIS

`/usr/games/fortune [-] [-alsw] [filename]`

DESCRIPTION

fortune with no arguments prints out a random adage. The flags mean:

- `-a` Choose from either list of adages.
- `-l` Long messages only.
- `-s` Short messages only.
- `-w` Waits before termination for an amount of time calculated from the number of characters in the message. This is useful if it is executed as part of the logout procedure to guarantee that the message can be read before the screen is cleared.

FILES

`/usr/games/lib/fortunes.dat`

NAME

gammontool – play a game of backgammon

SYNOPSIS

/usr/games/gammontool [*path*]

AVAILABILITY

This game is available with the *Games* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

gammontool paints a backgammon board on the screen, and then lets you play against the computer. It must be run in SunWindows. The optional *path* argument specifies an alternate move-generating program, which must be specially designed to run with **gammontool**.

The game has three subwindows: an option window on top, a message window in the middle, and a large board on the bottom. The buttons in the option window are used to restart, double, etc. The message window has two lines: the first tells whose turn it is, and the second displays any errors that occur.

The Initial Roll

To start the game, roll the dice to determine who goes first. Move the mouse arrow onto the board and click the left button. One die appears on each side of the board: the die on the left is yours, and the die on the right is the computer's. If your roll is greater, then you move; if not, the computer makes a move.

Making Your Move

When it is your turn, 'Yourmove' appears in the message window. Place the mouse over any piece of your color, and click the left button. While holding down the button, move the mouse to drag the piece; the piece follows the mouse until you release the button. The tool checks each move and does not allow illegal moves. When you have made as many moves as you can, the computer takes its turn; after it finishes, you may either roll again, or double.

Doubling

To double, click the *Double* button in the option window and wait for the computer's response. If the computer doubles you, a message is displayed and you must answer with the **Accept Double** or **Refuse Double** buttons. The **Forfeit** button can also be used to refuse a double. If the game is doubled, a doubling cube with the proper value is displayed on the bar strip. If the number is facing up, then you may double next. If the number is upside down, it is the computer's turn to double.

Other Buttons

If you want to change your move before you have finished it, use the **Redo Move** or **Redo Entire Move** buttons in the option window. **Redo Entire Move** replaces all of the pieces you have moved so that you can redo them all. **Redo Move** only replaces the last piece you moved, so it is useful when you roll doubles and want to redo only the last piece you moved. Note that once you have made all of the moves your roll permits, play passes immediately to the computer, so you cannot redo the very last move. The **Show Last Move** button allows you to see the last move again.

Leaving the Game

If you want to quit playing backgammon, use the **Quit** button. If you want to forfeit the game, use the **Forfeit** button. The computer penalizes you by taking a certain number of points, but the program does not terminate.

To play another game after winning, losing, or forfeiting, click the **New Game** button. To change the color of your pieces, click the mouse button while pointing at either the **White** or **Black** checkboxes. You may change colors at any time, even in the middle of a game. Changing colors in the middle of a game does not mean that you trade places with the computer; your pieces stay where they are, but they are repainted with the new color. Your pieces always move from the top right to the bottom right of the board, regardless of your color. Your pieces always move from the top right to the bottom right of the board, regardless of your color. As an additional cue as to your color, your dice are always displayed on the left half of the board.

Log File

If there is a **gammonlog** file in your home directory, **gammontool** keeps a log of the games played. Each move and double gets recorded, along with the winners and accumulated scores.

FILES

~/gammonlog log of games played
/usr/games/lib/gammonscores
 log of wins and losses

BUGS

The default strategy used by the computer is very poor.

If a single move uses more than one die (for instance if you roll 5, 6 and move 11 spaces without touching down in the middle) it is unpredictable where the program will make the piece touch down. This may be important if there is a blot on one of these middle points. The program will always make the move if possible, but if two midpoints would work and there is a blot on one of them, it is much better to explicitly hit the blot and then move the piece the rest of the way.

NAME

graphics_demos, bouncedemo, framedemo, jumpdemo, spheredemo, – graphics demonstration programs

SYNOPSIS

`/usr/demo/bouncedemo [-d dev] [-nx] [-r] [-q] /usr/demo/framedemo [-d dev] [-nx] [-r] [-q]`

`/usr/demo/jumpdemo [-c] [-d dev] [-nx] [-r] [-q]`

`/usr/demo/spheredemo [-d dev] [-nx] [-r] [-q]`

AVAILABILITY

These demos are available with the *Demos* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION**bouncedemo**

bouncedemo displays a bouncing square.

framedemo**framedemo**

displays a series of frames, each of which contains a 256 by 256 image one-bit-deep pixels (that is, the image is a square monochrome bitmap, with 256 bits on a side). **framedemo** looks for the frames in the files **frame.1** through **frame.n** in the current working directory, and displays them in numerical order. A set of sample frames is available in the directory `/usr/demo/globeframes/*`.

Interactive Commands

If you move the cursor onto the image surface, you can type certain commands to affect the rate at which the frames are displayed. The initial rate is one frame per second:

- f** Remove 1/20th of a second from the interval.
- F** Remove one second from the interval. **Ff** makes the interval as small as possible.
- s** Add 1/20th of a second.
- S** Add one second.

jumpdemo

jumpdemo simulates the famous *Star Wars* jump to light-speed-sequence using vector drawing. Colored stars are drawn on color surfaces.

spheredemo

spheredemo computes a random collection of shaded spheres. Colored spheres are drawn on color surfaces.

OPTIONS

- c** Rotate the color map to produce a sparkling effect.
- d surface** Run the demo on a surface other than the window or system console, for instance:

bouncedemo -d /dev/cgone0
- nx** Draw *x* items, or repeat a sequence *x* times.
- r** Retain the window. This allows the image to reappear when uncovered instead of restarting the demo.
- q** Quick exit. Useful for running several demos from within a shell script.

NAME

hack – replacement for *rogue*

SYNOPSIS

hack [**-d** *hackdir*] [**-s** *all* | *player* ...]

DESCRIPTION

hack is a display-oriented dungeons & dragons type game. Both display and command structure resemble *rogue*, although **hack** has twice as many monster types and requires three times as much memory.

Normally **hack** looks in `/usr/games/lib/hackdir` for the files listed below; this directory can be changed with the **-d** option. The **-s** option permits you to search the player record. Given the keyword **all**, **hack** lists all players; given the login name of a player, it lists all scores of that player.

FILES

record	top 100 list (start with an empty file)
news	changes or bugs (start with no news file)
data	information about objects and monsters
help	introductory information (no doubt outdated)
hh	compacted version of help
perm	empty file used for locking
rumors	texts for fortune cookies

NAME

hangman – computer version of the game hangman

SYNOPSIS

/usr/games/hangman

DESCRIPTION

In **hangman**, the computer picks a word from the on-line word list and you must try to guess it. The computer keeps track of which letters have been guessed and how many wrong guesses you have made on the screen in a graphic fashion.

FILES

/usr/dict/words on-line word list

NAME

hunt – a multiplayer multiterminal game

SYNOPSIS

`/usr/games/hunt[-m] [hostname] [-l name]`

DESCRIPTION

The object of the game **hunt** is to kill off the other players. There are no rooms, no treasures, and no monsters. Instead, you wander around a maze, find grenades, trip mines, and shoot down walls and players.

Your score is the ratio of number of kills to number of times you entered the game and is only kept for the duration of a single session of **hunt**. The more players you kill before you die, the better your score is.

hunt normally looks for an active game on the local network; if none is found, it starts one up on the local host. One may specify the location of the game by giving the *hostname* argument.

hunt only works on crt (vdt) terminals with at least 24 lines, 80 columns, and cursor addressing. The screen is divided in to 3 areas. On the right hand side is the status area. It shows you how much damage you've sustained, how many charges you have left, who's in the game, who's scanning (the asterisk in front of the name), who's cloaked (the plus sign in front of the name), and other players' scores. Most of the rest of the screen is taken up by your map of the maze, except for the 24th line, which is used for longer messages that do not fit in the status area.

hunt uses the same keys to move as **vi** does, for instance, **h,j,k,** and **l** for left, down, up, right respectively. To change which direction you're facing in the maze, use the upper case version of the movement key (for instance, **HJKL**).

Other commands are:

f	Fire (in the direction you're facing) (Takes 1 charge)
g	Throw grenade (in the direction you're facing) (Takes 9 charges)
F	Throw satchel charge (Takes 25 charges)
G	Throw bomb (Takes 49 charges)
o	Throw small slime bomb (Takes 15 charges)
O	Throw big slime bomb (Takes 30 charges)
s	Scan (where other players are) (Takes 1 charge)
c	Cloak (where you are) (Takes 1 charge)
^L	Redraw screen
q	Quit

Knowing what the symbols on the screen often helps:

- +	Walls
/\	Diagonal (deflecting) walls
#	Doors (dispersion walls)
;	Small mine
g	Large mine
:	Shot
o	Grenade
O	Satchel charge
@	Bomb
s	Small slime bomb
\$	Big slime bomb
> < ^ v	You facing right, left, up, or down

} {i! Other players facing right, left, up, or down
 * Explosion
 √
 -*E- Grenade and large mine explosion
 ∧

Satchel and bomb explosions are larger than grenades (5x5, 7x7, and 3x3 respectively).

Other helpful hints:

You can only fire in the direction you are facing.

You can only fire three shots in a row, then the gun must cool.

A shot only affects the square it hits.

Shots and grenades move 5 times faster than you do.

To stab someone,

you must face that player and move at them.

Stabbing does 3 points worth of damage and shooting does 5 points.

You start with 15 charges and get 5 more for every new player.

A grenade affects the nine squares centered about the square it hits.

A satchel affects the twenty-five squares centered about the square it hits.

A bomb affects the forty-nine squares centered about the square it hits.

One small mine and one large mine is placed in the maze for every new player.

A mine has a 5% probability of tripping when you walk directly at it;

50% when going sideways on to it; 95% when backing up on to it.

Tripping a mine costs you 5 points or 10 points respectively.

Defusing a mine is worth 1 charge or 9 charges respectively.

You cannot see behind you.

Scanning lasts for (20 times the number of players) turns.

Scanning takes 1 ammo charge, so do not waste all your charges scanning.

You get 2 more damage capacity points and 2 damage points taken away

whenever you kill someone.

Maximum typeahead is 5 characters.

A shot destroys normal (for instance, non-diagonal, non-door) walls.

Diagonal walls deflect shots and change orientation.

Doors disperse shots in random directions (up, down, left, right).

Diagonal walls and doors cannot be destroyed by direct shots but may

be destroyed by an adjacent grenade explosion.

Walls regenerate, reappearing in the order they were destroyed.

One percent of the regenerated walls will be diagonal walls or doors. When a wall is generated directly beneath a player, he is thrown in a random direction for a random period of time. When he lands, he sustains damage (up to 20 percent of the amount of damage he had before impact); that is, the less damage he had, the more nimble he is and therefore less likely to hurt himself on landing.

ENVIRONMENT

The environment variable HUNT is checked to get the player name. If you do not have this variable set, **hunt** will ask you what name you want to play under. You may also set up a single character keyboard map, but then you have to enumerate the options. For example:

```
setenv HUNT "name=Sneaky,mapkey=z0FfGg1f2g3F4G"
```

sets the player name to Sneaky, and the maps z to o, F to f, G to g, 1 to f, 2 to g, 3 to F, and 4 to G.

The *mapkey* option must be last.

It is a boring game if you are the only one playing.

OPTIONS

- m You enter the game as a monitor (you can see the action but you cannot play).
- l *name* Enter the game as player *name*.

FILES

/usr/games/lib/hunt.driver game coordinator

LIMITATIONS

hunt normally drives up the load average to be about (number_of_players + 0.5) greater than it would be without a **hunt** game executing. A limit of three players per host and nine players total is enforced by **hunt**.

BUGS

To keep up the pace, not everything is as realistic as possible.

NAME

life – John Conway's game of life

SYNOPSIS

`/usr/games/life`

AVAILABILITY

This game is available with the *Games* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

life is a program that plays John Conway's game of life. It only runs under `sunview(1)`.

When invoked, life will display a window with a small control panel at the top, and a large drawing area at the bottom. You can create pieces in the drawing area with the left button, and erase them with the middle button. When you select **Run** in the control panel, the pieces will begin to evolve, and the drawing region will update itself at a speed controlled by the slider labeled with **Fast** and **Slow**. life keeps track of all the pieces even if they are not visible. The scroll bars surrounding the drawing region can be used to see pieces that have moved out of view. There are some standard patterns that can be drawn by popping up a menu in the drawing subwindow.

The meaning of the items in the first row of the control panel (from left to right) are as follows. If you click on the picture which looks like a tic-tac-toe board, a grid will appear in the drawing region. If you click on **Step**, the mode will change from run mode (where the pieces update continuously) to step mode (where an update is only done when you click on **Step**). Following **Gen** is a number indicating the number of generations that have occurred. The button marked **Find** will scroll so that at least one piece is in view. This is useful when all the pieces disappear from view. The button marked **Clear** will clear the drawing region, but leave the other controls unchanged. **Reset** will reset all the panel controls, but will not erase any of the pieces, and **Quit** Exits the tool. The second row contains two sliders. The first controls the update speed when in run mode, the second controls the size of the pieces.

SEE ALSO

`sunview(1)`

NAME

mille – play Mille Bornes

SYNOPSIS

/usr/games/mille [file]

DESCRIPTION

mille plays a two-handed game reminiscent of the Parker Brother's game of Mille Bornes with you. The rules are described below. If a file name is given on the command line, the game saved in that file is started.

When a game is started up, the bottom of the score window will contain a list of commands. They are:

- P Pick a card from the deck. This card is placed in the 'P' slot in your hand.
- D Discard a card from your hand. To indicate which card, type the number of the card in the hand (or "P" for the just-picked card) followed by a carriage-return or space. The carriage-return or space is required to allow recovery from typos which can be very expensive, like discarding safeties.
- U Use a card. The card is again indicated by its number, followed by a carriage-return or space.
- O Toggle ordering the hand. By default off, if turned on it will sort the cards in your hand appropriately. This is not recommended for the impatient on slow terminals.
- Q Quit the game. This will ask for confirmation, just to be sure. Hitting DELETE (or RUBOUT) is equivalent.
- S Save the game in a file. If the game was started from a file, you will be given an opportunity to save it on the same file. If you don't wish to, or you did not start from a file, you will be asked for the file name. If you type a RETURN character without a name, the save will be terminated and the game resumed.
- R Redraw the screen from scratch. The command ^L (CTRL-L) will also work.
- W Toggle window type. This switches the score window between the startup window (with all the command names) and the end-of-game window. Using the end-of-game window saves time by eliminating the switch at the end of the game to show the final score. Recommended for hackers and other miscreants.

If you make a mistake, an error message will be printed on the last line of the score window, and a bell will beep.

At the end of each hand or game, you will be asked if you wish to play another. If not, it will ask you if you want to save the game. If you do, and the save is unsuccessful, play will be resumed as if you had said you wanted to play another hand/game. This allows you to use the "S" command to reattempt the save. (The game itself is a product of Parker Brothers, Inc.)

SEE ALSO

curses(3X)

CARDS

Here is some useful information. The number in brackets after the card name is the number of that card in the deck:

Hazard	Repair	Safety
Out of Gas [2]	Gasoline [6]	Extra Tank [1]
Flat Tire [2]	Spare Tire [6]	Puncture Proof [1]
Accident [2]	Repairs [6]	Driving Ace [1]
Stop [4]	Go [14]	Right of Way [1]
Speed Limit [3]	End of Limit [6]	

25 - [10], 50 - [10], 75 - [10], 100 - [12], 200 - [4]

RULES

Object: The point of game is to get a total of 5000 points in several hands. Each hand is a race to put down exactly 700 miles before your opponent does. Beyond the points gained by putting down milestones, there are several other ways of making points.

Overview: The game is played with a deck of 101 cards. *Distance* cards represent a number of miles traveled. They come in denominations of 25, 50, 75, 100, and 200. When one is played, it adds that many miles to the player's trip so far this hand. *Hazard* cards are used to prevent your opponent from putting down Distance cards. With the exception of the *speed limit* card, they can only be played if your opponent has a *Go* card on top of the Battle pile. The cards are *Out of Gas*, *Accident*, *Flat Tire*, *Speed Limit*, and *Stop*. *Remedy* cards fix problems caused by Hazard cards played on you by your opponent. The cards are *Gasoline*, *Repairs*, *Spare Tire*, *End of Limit*, and *Go*. *Safety* cards prevent your opponent from putting specific Hazard cards on you in the first place. They are *Extra Tank*, *Driving Ace*, *Puncture Proof*, and *Right of Way*, and there are only one of each in the deck.

Board Layout: The board is split into several areas. From top to bottom, they are: **SAFETY AREA (unlabeled):** This is where the safeties are played. **HAND:** These are the cards in your hand. **BATTLE:** This is the Battle pile. All the Hazard and Remedy Cards are played here, except the *Speed Limit* and *End of Limit* cards. Only the top card is displayed, as it is the only effective one. **SPEED:** The Speed pile. The *Speed Limit* and *End of Limit* cards are played here to control the speed at which the player is allowed to put down miles. **MILEAGE:** Miles are placed here. The total of the numbers shown here is the distance traveled so far.

Play: The first pick alternates between the two players. Each turn usually starts with a pick from the deck. The player then plays a card, or if this is not possible or desirable, discards one. Normally, a play or discard of a single card constitutes a turn. If the card played is a safety, however, the same player takes another turn immediately.

This repeats until one of the players reaches 700 points or the deck runs out. If someone reaches 700, they have the option of going for an *Extension*, which means that the play continues until someone reaches 1000 miles.

Hazard and Remedy Cards: Hazard Cards are played on your opponent's Battle and Speed piles. Remedy Cards are used for undoing the effects of your opponent's nastiness.

Go (Green Light) must be the top card on your Battle pile for you to play any mileage, unless you have played the *Right of Way* card (see below).

Stop is played on your opponent's *Go* card to prevent them from playing mileage until they play a *Go* card.

Speed Limit is played on your opponent's Speed pile. Until they play an *End of Limit* they can only play 25 or 50 mile cards, presuming their *Go* card allows them to do even that.

End of Limit is played on your Speed pile to nullify a *Speed Limit* played by your opponent.

Out of Gas is played on your opponent's *Go* card. They must then play a *Gasoline* card, and then a *Go* card before they can play any more mileage.

Flat Tire is played on your opponent's *Go* card. They must then play a *Spare Tire* card, and then a *Go* card before they can play any more mileage.

Accident is played on your opponent's *Go* card. They must then play a *Repairs* card, and then a *Go*

card before they can play any more mileage.

Safety Cards: Safety cards prevent your opponent from playing the corresponding Hazard cards on you for the rest of the hand. It cancels an attack in progress, and *always entitles the player to an extra turn.*

Right of Way prevents your opponent from playing both *Stop* and *Speed Limit* cards on you. It also acts as a permanent *Go* card for the rest of the hand, so you can play mileage as long as there is not a Hazard card on top of your Battle pile. In this case only, your opponent can play Hazard cards directly on a Remedy card besides a *Go* card.

Extra Tank When played, your opponent cannot play an *Out of Gas* on your Battle Pile.

Puncture Proof When played, your opponent cannot play a *Flat Tire* on your Battle Pile.

Driving Ace When played, your opponent cannot play an *Accident* on your Battle Pile.

Distance Cards: Distance cards are played when you have a *Go* card on your Battle pile, or a Right of Way in your Safety area and are not stopped by a Hazard Card. They can be played in any combination that totals exactly 700 miles, except that *you cannot play more than two 200 mile cards in one hand.* A hand ends whenever one player gets exactly 700 miles or the deck runs out. In that case, play continues until neither someone reaches 700, or neither player can use any cards in their hand. If the trip is completed after the deck runs out, this is called *Delayed Action.*

Coup Fouré: This is a French fencing term for a counter-thrust move as part of a parry to an opponents attack. In Mille Bornes, it is used as follows: If an opponent plays a Hazard card, and you have the corresponding Safety in your hand, you play it immediately, even *before* you draw. This immediately removes the Hazard card from your Battle pile, and protects you from that card for the rest of the game. This gives you more points (see "Scoring" below).

Scoring: Scores are totaled at the end of each hand, whether or not anyone completed the trip. The terms used in the Score window have the following meanings:

Milestones Played: Each player scores as many miles as they played before the trip ended.

Each Safety: 100 points for each safety in the Safety area.

All 4 Safeties: 300 points if all four safeties are played.

Each Coup Fouré: 300 points for each Coup Fouré accomplished.

The following bonus scores can apply only to the winning player.

Trip Completed: 400 points bonus for completing the trip to 700 or 1000.

Safe Trip: 300 points bonus for completing the trip without using any 200 mile cards.

Delayed Action: 300 points bonus for finishing after the deck was exhausted.

Extension: 200 points bonus for completing a 1000 mile trip.

Shut-Out: 500 points bonus for completing the trip before your opponent played any mileage cards.

Running totals are also kept for the current score for each player for the hand (**Hand Total**), the game (**Overall Total**), and number of games won (**Games**).

NAME

monop – Monopoly game

SYNOPSIS

/usr/games/monop [*filename*]

DESCRIPTION

monop is reminiscent of the Parker Brother's game Monopoly, and monitors a game between 1 to 9 users. It is assumed that the rules of Monopoly are known. The game follows the standard rules, with the exception that, if a property would go up for auction and there are only two solvent players, no auction is held and the property remains unowned.

The game, in effect, lends the player money, so it is possible to buy something which you cannot afford. However, as soon as a person goes into debt, he must "fix the problem", that is, make himself solvent, before play can continue. If this is not possible, the player's property reverts to his debtee, either a player or the bank. A player can resign at any time to any person or the bank, which puts the property back on the board, unowned.

Any time that the response to a question is a *string*, for instance a name, place or person, you can type ? to get a list of valid answers. It is not possible to input a negative number, nor is it ever necessary.

USAGE**Commands**

quit: Quit game. This allows you to quit the game. It asks you if you are sure.

print Print board. This prints out the current board. The columns have the following meanings (column headings are the same for the **where**, **own holdings**, and **holdings** commands):

Name	The first ten characters of the name of the square
Own	The <i>number</i> of the owner of the property.
Price	The cost of the property (if any)
Mg	This field has a '*' in it if the property is mortgaged
#	If the property is a Utility or Railroad, this is the number of such owned by the owner. If the property is land, this is the number of houses on it.
Rent	Current rent on the property. If it is not owned, there is no rent.

where: where players are: Tells you where all the players are. A '*' indicates the current player.

own holdings :

List your own holdings, that is, money, get-out-of-jail-free cards, and property.

holdings:

Holdings list. Look at anyone's holdings. It will ask you whose holdings you wish to look at. When you are finished, type **done**.

shell: Shell escape. Escape to a shell. When the shell dies, the program continues where you left off.

mortgage:

Mortgage property. Sets up a list of mortgageable property, and asks which you wish to mortgage.

unmortgage:

Unmortgage property. Unmortgage mortgaged property.

buy: Buy houses. Sets up a list of monopolies on which you can buy houses. If there is more than one, it asks you which you want to buy for. It then asks you how many for each piece of property, giving the current amount in parentheses after the property name. If you build in an unbalanced manner (a disparity of more than one house within the same monopoly), it asks you to re-input things.

sell: Sell houses. Sets up a list of monopolies from which you can sell houses. it operates in an

analogous manner to **buy**

- card:** Card for jail. Use a get-out-of-jail-free card to get out of jail. If you are not in jail, or you do not have one, it tells you so.
- pay:** Pay for jail. Pay \$50 to get out of jail, from whence you are put on Just Visiting. Difficult to do if you are not there.
- trade:** This allows you to trade with another player. It asks you whom you wish to trade with, and then asks you what each wishes to give up. You can get a summary at the end, and, in all cases, it asks for confirmation of the trade before doing it.
- resign:** Resign to another player or the bank. If you resign to the bank, all property reverts to its virgin state, and get-out-of-jail free cards revert to the deck.
- save:** Save game. Save the current game in a file for later play. You can continue play after saving, either by adding the file in which you saved the game after the **monop** command, or by using the **restore** command (see below). It will ask you which file you wish to save it in, and, if the file exists, confirm that you wish to overwrite it.
- restore:**
Restore game. Read in a previously saved game from a file. It leaves the file intact.
- roll:** Roll the dice and move forward to your new location. If you simply hit the RETURN key instead of a command, it is the same as typing *roll*.

FILES

`/usr/games/lib/cards.pck` chance and community chest cards

BUGS

No command can be given an argument instead of a response to a query.

NAME

moo - guessing game

SYNOPSIS

/usr/games/moo

DESCRIPTION

moo is a guessing game imported from England. The computer picks a number consisting of four distinct decimal digits. The player guesses four distinct digits being scored on each guess. A "cow" is a correct digit in an incorrect position. A "bull" is a correct digit in a correct position. The game continues until the player guesses the number (a score of four bulls).

NAME

number – convert Arabic numerals to English

SYNOPSIS

/usr/games/number

DESCRIPTION

number copies the standard input to the standard output, changing each decimal number to a fully spelled out version.

NAME

primes – print all primes larger than some given number

SYNOPSIS

`/usr/games/primes [number]`

DESCRIPTION

`primes` reads a number from the standard input and prints all primes larger than the given number. If *number* is given as an argument, it uses that number rather than reading one from the standard input.

BUGS

It obviously cannot print *all* primes larger than some given number. It will not behave very sensibly when it overflows an int.

NAME

quiz – test your knowledge

SYNOPSIS

`/usr/games/quiz [-ifilename] [-t] [category1 category2]`

DESCRIPTION

quiz gives associative knowledge tests on various subjects. It asks items chosen from *category1* and expects answers from *category2*. If no categories are specified, **quiz** gives instructions and lists the available categories.

quiz tells a correct answer whenever you type a bare newline. At the end of input, upon interrupt, or when questions run out, **quiz** reports a score and terminates.

The `-t` flag specifies ‘tutorial’ mode, where missed questions are repeated later, and material is gradually introduced as you learn.

The `-i` flag causes the named file to be substituted for the default index file. The lines of these files have the syntax:

```
line      = category newline | category ':' line
category = alternate | category 'l' alternate
alternate = empty | alternate primary
primary   = character | '[' category ']' | option
option    = '{' category '}'
```

The first category on each line of an index file names an information file. The remaining categories specify the order and contents of the data in each line of the information file. Information files have the same syntax. Backslash ‘\’ is used as with `sh(1)` to quote syntactically significant characters or to insert transparent newlines into a line. When either a question or its answer is empty, **quiz** will refrain from asking it.

FILES

`/usr/games/quiz.k/*`

BUGS

The construct ‘a|ab’ doesn’t work in an information file. Use ‘a{b}’.

NAME

rain – animated raindrops display

SYNOPSIS

/usr/games/rain

DESCRIPTION

rain's display is modeled after the VAX/VMS program of the same name. The terminal has to be set for 9600 baud to obtain the proper effect.

As with all programs that use **termcap**, the **TERM** environment variable must be set (and exported) to the type of the terminal being used.

FILES

/etc/termcap

NAME

random – select lines randomly from a file

SYNOPSIS

`/usr/games/random [-er] [divisor]`

DESCRIPTION

random acts as a text filter, randomly selecting lines from its standard input to write to the standard output. The probability that a given line is selected is normally 1/2; if a *divisor* is specified, it is treated as a floating-point number, and the probability is 1/*divisor* instead.

OPTIONS

- e** Don't read the standard input or write to the standard output. Instead, exit with a random exit status between 0 and 1, or between 0 and *divisor*-1 if *divisor* is specified.
- r** Don't buffer the output. If **-r** is not used, output is buffered in blocks, or line-buffered if the standard output is a terminal.

NAME

robots – fight off villainous robots

SYNOPSIS

`/usr/games/robots [-sjta] [scorefile]`

DESCRIPTION

robots pits you against evil robots, who are trying to kill you (which is why they are evil). Fortunately for you, even though they are evil, they are not very bright and have a habit of bumping into each other, thus destroying themselves. In order to survive, you must get them to kill each other off, since you have no offensive weaponry.

Since you are stuck without offensive weaponry, you are endowed with one piece of defensive weaponry: a teleportation device. When two robots run into each other or a junk pile, they die. If a robot runs into you, you die. When a robot dies, you get 10 points, and when all the robots die, you start on the next field. This keeps up until they finally get you.

Robots are represented on the screen by a '+', the junk heaps from their collisions by a '*', and you (the good guy) by a '@'.

The commands are:

h	move one square left
l	move one square right
k	move one square up
j	move one square down
y	move one square up and left
u	move one square up and right
b	move one square down and left
n	move one square down and right
.	(also space) do nothing for one turn
HJKLBNYU	
	run as far as possible in the given direction
>	do nothing for as long as possible
t	teleport to a random location
w	wait until you die or they all do
q	quit
^L	redraw the screen

All commands can be preceded by a count.

If you use the 'w' command and survive to the next level, you will get a bonus of 10% for each robot which died after you decided to wait. If you die, however, you get nothing. For all other commands, the program will save you from typos by stopping short of being eaten. However, with 'w' you take the risk of dying by miscalculation.

Only five scores are allowed per user on the score file. If you make it into the score file, you will be shown the list at the end of the game. If an alternate score file is specified, that will be used instead of the standard file for scores.

OPTIONS

-s	Do not play, just show the score file.
-j	Jump, when you run, don't show any intermediate positions; only show things at the end. This is useful on slow terminals.

- t Teleport automatically when you have no other option. This is a little disconcerting until you get used to it, and then it is very nice.
- a Advance into the higher levels directly, skipping the lower, easier levels.

FILES

/usr/games/lib/robots_roll the score file

BUGS

Bugs? You *crazy*, man!?

NAME

snake, snscore – display chase game

SYNOPSIS

`/usr/games/snake [-wn] [-ln]`

`/usr/games/snscore`

DESCRIPTION

snake is a display-based game which must be played on a CRT terminal from among those supported by `vi(1)`. The object of the game is to make as much money as possible without getting eaten by the snake. The `-l` and `-w` options allow you to specify the length and width of the field. By default the entire screen (except for the last column) is used.

You are represented on the screen by an `I`. The snake is 6 squares long and is represented by `S`'s. The money is `$`, and an exit is `#`. Your score is posted in the upper left hand corner.

You can move around using the same conventions as `vi(1)`, the `h`, `j`, `k`, and `l` keys work, as do the arrow keys. Other possibilities include:

- sefc** These keys are like `hjkl` but form a directed pad around the `d` key.
- HJKL** These keys move you all the way in the indicated direction to the same row or column as the money. This does *not* let you jump away from the snake, but rather saves you from having to type a key repeatedly. The snake still gets all his turns.
- SEFC** Likewise for the upper case versions on the left.
- ATPB** These keys move you to the four edges of the screen. Their position on the keyboard is the mnemonic, for example, `P` is at the far right of the keyboard.
- x** This lets you quit the game at any time.
- p** Points in a direction you might want to go.
- w** Space warp to get out of tight squeezes, at a price.
- !** Shell escape
- ^Z** Suspend the snake game, on systems which support it. Otherwise an interactive shell is started up.

To earn money, move to the same square the money is on. A new `$` will appear when you earn the current one. As you get richer, the snake gets hungrier. To leave the game, move to the exit (`#`).

A record is kept of the personal best score of each player. Scores are only counted if you leave at the exit, getting eaten by the snake is worth nothing.

As in pinball, matching the last digit of your score to the number which appears after the game is worth a bonus.

To see who wastes time playing snake, run `/usr/games/snscore .`

FILES

`/usr/games/lib/snakerawscores` database of personal bests
`/usr/games/lib/snake.log` log of games played

BUGS

When playing on a small screen, it's hard to tell when you hit the edge of the screen.

The scoring function takes into account the size of the screen. A perfect function to do this equitably has not been devised.

NAME

sunview_demos, canvas_demo, cursor_demo – Window-System demonstration programs

SYNOPSIS

/usr/demo/canvas_demo

/usr/demo/cursor_demo

AVAILABILITY

These demos are available with the *SunView 1 Demos* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION**canvas_Demo**

canvas_demo demonstrates the capabilities of the canvas subwindow package. It consists of two subwindows: a control panel and a canvas. By adjusting the items on the control panel, you can manipulate the attributes of the canvas, and see the results.

cursor_Demo

cursor_demo demonstrates what you can do with cursors. A single control panel is provide for adjusting the various cursor attributes. As you adjust the items on the control panel, the panel's cursor changes in appearance.

NAME

trek – trekkie game

SYNOPSIS

`/usr/games/trek [[-a] filename]`

DESCRIPTION

trek is a game of space glory and war. Below is a summary of commands. For complete documentation, see *Trek* by Eric Allman.

If a filename is given, a log of the game is written onto that file. If the `-a` flag is given before the filename, that file is appended to, not truncated.

The game will ask you what length game you would like. Valid responses are “short”, “medium”, and “long”. You may also type “restart”, which restarts a previously saved game. You will then be prompted for the skill, to which you must respond “novice”, “fair”, “good”, “expert”, “commodore”, or “impossible”. You should normally start out with a novice and work up.

In general, throughout the game, if you forget what is appropriate the game will tell you what it expects if you just type in a question mark.

COMMAND SUMMARY

abandon	capture
cloak up/down	
computer request; ...	damages
destruct	dock
help	impulse course distance
lrscan	move course distance
phasers automatic amount	
phasers manual amt1 course1 spread1 ...	
torpedo course [yes] angle/no	
ram course distance	rest time
shell	shields up/down
srscan [yes/no]	
status	terminate yes/no
undock	visual course
warp warp_factor	

NAME

worm -- play the growing worm game

SYNOPSIS

`/usr/games/worm [size]`

DESCRIPTION

In *worm*, you are a little worm, your body is the `o`'s on the screen and your head is the `@`. You move with the `hjkl` keys (as in the game *snake*). If you don't press any keys, you continue in the direction you last moved. The upper case `HJKL` keys move you as if you had pressed several (9 for `HL` and 5 for `JK`) of the corresponding lower case key (unless you run into a digit, then it stops).

On the screen you will see a digit; if your worm eats the digit it will grow longer, the actual amount longer depends on which digit it was that you ate. The object of the game is to see how long you can make the worm grow.

The game ends when the worm runs into either the sides of the screen, or itself. The current score (how much the worm has grown) is kept in the upper left corner of the screen.

The optional argument, if present, is the initial length of the worm.

BUGS

If the initial length of the worm is set to less than one or more than 75, various strange things happen.

NAME

worms - animate worms on a display terminal

SYNOPSIS

`/usr/games/worms [-field] [-length #] [-number #] [-trail]`

DESCRIPTION

`-field` makes a "field" for the worm(s) to eat; `-trail` causes each worm to leave a trail behind it. You can figure out the rest by yourself.

FILES

`/etc/termcap`

SEE ALSO

Snails by Karl Heuer

BUGS

The lower-right-hand character position will not be updated properly on a terminal that wraps at the right margin.

Terminal initialization is not performed.

NAME

wump – the game of hunt-the-wumpus

SYNOPSIS

/usr/games/wump

DESCRIPTION

wump plays the game of 'Hunt the Wumpus.' A Wumpus is a creature that lives in a cave with several rooms connected by tunnels. You wander among the rooms, trying to shoot the Wumpus with an arrow, meanwhile avoiding being eaten by the Wumpus and falling into Bottomless Pits. There are also Super Bats which are likely to pick you up and drop you in some random room.

The program asks various questions which you answer one per line; it will give a more detailed description if you want.

This program is based on one described in *People's Computer Company*, 2, 2 (November 1973).

NAME

miscellaneous – miscellaneous useful information pages

DESCRIPTION

This section contains miscellaneous documentation, mostly in the area of text processing macro packages for *troff*(1).

LIST OF MISC. TABLES

Name	Appears on Page	Description
ascii	ascii(7)	map of ASCII character set
eqnchar	eqnchar(7)	special character definitions for eqn
filesystem	filesystem(7)	filesystem organization
hier	hier(7)	file system hierarchy
man	man(7)	macros to format Reference Manual pages
me	me(7)	macros for formatting papers
ms	ms(7)	text formatting macros

NAME

ascii — map of ASCII character set

SYNOPSIS

cat /usr/pub/ascii

DESCRIPTION

/usr/pub/ascii is a map of the ASCII character set, to be printed as needed. It contains octal and hexadecimal values for each character. While not included in that file, a chart of decimal values is also shown here.

Octal — Character

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	

Hexadecimal — Character

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	

Decimal—Character

0	NUL	1	SOH	2	STX	3	ETX	4	EOT	5	ENQ	6	ACK	7	BEL
8	BS	9	HT	10	NL	11	VT	12	NP	13	CR	14	SO	15	SI
16	DLE	17	DC1	18	DC2	19	DC3	20	DC4	21	NAK	22	SYN	23	ETB
24	CAN	25	EM	26	SUB	27	ESC	28	FS	29	GS	30	RS	31	US
32	SP	33	!	34	"	35	#	36	\$	37	%	38	&	39	'
40	(41)	42	*	43	+	44	,	45	-	46	.	47	/
48	0	49	1	50	2	51	3	52	4	53	5	54	6	55	7
56	8	57	9	58	:	59	;	60	<	61	=	62	>	63	?
64	@	65	A	66	B	67	C	68	D	69	E	70	F	71	G
72	H	73	I	74	J	75	K	76	L	77	M	78	N	79	O
80	P	81	Q	82	R	83	S	84	T	85	U	86	V	87	W
88	X	89	Y	90	Z	91	[92	\	93]	94	^	95	_
96	`	97	a	98	b	99	c	100	d	101	e	102	f	103	g
104	h	105	i	106	j	107	k	108	l	109	m	110	n	111	o
112	p	113	q	114	r	115	s	116	t	117	u	118	v	119	w
120	x	121	y	122	z	123	{	124		125	}	126	~	127	DEL

FILES

`/usr/pub/ascii`

Online chart of octal and hexadecimal values for the ASCII character set.

NAME

eqnchar – special character definitions for eqn

SYNOPSIS

eqn /usr/pub/eqnchar [*filename*] | troff [*options*]

neqn /usr/pub/eqnchar [*filename*] | nroff [*options*]

DESCRIPTION

eqnchar contains troff(1) and nroff(1) character definitions for constructing characters that are not available on the Graphic Systems typesetter. These definitions are primarily intended for use with eqn(1) and neqn(1). It contains definitions for the following characters

<i>ciplus</i>	\oplus	//	//	<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>	\approx	<i>hbar</i>	\hbar	<i>bullet</i>	\bullet
<i>>wig</i>	\succ	<i>ppd</i>	\dagger	<i>prop</i>	\propto
<i><wig</i>	\prec	<i><-></i>	\leftrightarrow	<i>empty</i>	\emptyset
<i>=wig</i>	\equiv	<i><=></i>	\Leftrightarrow	<i>member</i>	\in
<i>star</i>	$*$	<i> <</i>	\nless	<i>nomem</i>	\notin
<i>bigstar</i>	\bigstar	<i> ></i>	\ngtr	<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>	\subseteq
<i>andsign</i>	\wedge	<i>3dot</i>	\vdots	<i>subset</i>	\subset
<i>=del</i>	\equiv	<i>thf</i>	\therefore	<i>supset</i>	\supset
<i>oppA</i>	\nless	<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

eqn(1), neqn(1), nroff(1), troff(1)

NAME

filesystem – file system organization

SYNOPSIS

/
/usr

DESCRIPTION

The SunOS file system tree is organized for easy administration. Distinct areas within the file system tree are provided for files that are private to one machine, files that can be shared by multiple machines of a common architecture, files that can be shared by all machines, and home directories. This organization allows the sharable files to be stored on one machine, while being accessed by many machines using a remote file access mechanism such as Sun's Network File System (NFS). Grouping together similar files makes the file system tree easier to upgrade and manage.

The file system tree consists of a root file system and a collection of mountable file systems. The **mount(8)** program attaches mountable file systems to the file system tree at mount points (directory entries) in the root file system, or other previously mounted file systems. Two file systems, / (the root) and /usr, must be mounted in order to have a fully functional system. The root file system is mounted automatically by the kernel at boot time; the /usr file system is mounted by the */etc/rc.boot* script, which is run as part of the booting process.

The root file system contains files that are unique to each machine; it can not be shared among machines. The root file system contains the following directories:

- /dev** Character and block special files. Device files provide hooks into hardware devices or operating system facilities. The **MAKEDEV(8)** command builds device files in the /dev directory. Typically, device files are built to match the kernel and hardware configuration of the machine.
- /etc** Various configuration files and system administration databases that are machine specific. You can think of /etc as the "home directory" of a machine, defining its "identity." Executable programs are no longer kept in /etc.
- /home** Mount points for home directories. This directory may be arranged so that shared user files are placed under the directory */home/machine-name* on machines serving as file servers. Machines may then be locally configured with mount points under /home for all of the file servers of interest, with the name of the mount point being the name of the file server.
- /mnt** A generic mount point. This is an empty directory available for temporarily mounting file systems on.
- /sbin** Executable programs that are needed in the boot process before /usr is mounted. /sbin contains *only* those programs that are needed in order to mount the /usr file system: **hostname(1)**, **ifconfig(8C)**, **init(8)**, **mount(8)**, and **sh(1)**. After /usr is mounted, the full complement of utilities are available.
- /tmp** Temporary files that are deleted at reboot time.
- /var** Files, such as log files, that are unique to a machine but that can grow to an arbitrary ("variable") size.
- /var/adm** System logging and accounting files.
- /var/preserve**
Backup files for **vi(1)** and **ex(1)**.
- /var/spool** Subdirectories for files used in printer spooling, mail delivery, **cron(8)**, **at(1)**, etc.
- /var/tmp** Transitory files that are not deleted at reboot time.

Because it is desirable to keep the root file system small, larger file systems are often mounted on /var and /tmp.

The file system mounted on `/usr` contains architecture-dependent and architecture-independent shareable files. The subtree rooted at `/usr/share` contains architecture-independent shareable files; the rest of the `/usr` tree contains architecture-dependent files. By mounting a common remote file system, a group of machines with a common architecture may share a single `/usr` file system. A single `/usr/share` file system can be shared by machines of any architecture. A machine acting as a file server may export many different `/usr` file systems to support several different architectures and operating system releases. Clients usually mount `/usr` read-only to prevent their accidentally modifying any shared files. The `/usr` file system contains the following subdirectories:

<code>/usr/5bin</code>	System V executables.
<code>/usr/5include</code>	System V include files.
<code>/usr/5lib</code>	System V library files.
<code>/usr/bin</code>	Executable programs. The bulk of the system utilities are located here.
<code>/usr/dict</code>	Dictionary databases.
<code>/usr/etc</code>	Executable system administration programs.
<code>/usr/games</code>	Executable game programs and data.
<code>/usr/include</code>	Include files.
<code>/usr/lib</code>	Program libraries and various architecture-dependent databases.
<code>/usr/pub</code>	Various data files.
<code>/usr/ucb</code>	Executable programs descended from the Berkeley Software Distribution.
<code>/usr/share</code>	Subtree for architecture-independent shareable files.
<code>/usr/share/man</code>	Subdirectories for the on-line reference manual pages.
<code>/usr/share/lib</code>	Architecture-independent databases.

A machine with disks may export root file systems, swap files and `/usr` file systems to diskless or partially-disked machines, which mount these into the standard file system hierarchy. The standard directory tree for exporting these file systems is:

<code>/export</code>	The root of the exported file system tree.
<code>/export/exec/architecture-name</code>	The exported <code>/usr</code> file system supporting <i>architecture-name</i> for the current release.
<code>/export/exec/architecture-name.release-name</code>	The exported <code>/usr</code> file system supporting <i>architecture-name</i> for SunOS <i>release-name</i> .
<code>/export/share</code>	The exported common <code>/usr/share</code> directory tree.
<code>/export/root/hostname</code>	The exported root file system for <i>hostname</i> .
<code>/export/swap/hostname</code>	The exported swap file for <i>hostname</i> .
<code>/export/var/hostname</code>	The exported <code>/var</code> directory tree for <i>hostname</i> .
<code>/export/dump/hostname</code>	The exported dump file for <i>hostname</i> .
<code>/export/crash/hostname</code>	The exported crash dump directory for <i>hostname</i> .

Changes from Previous Releases

The file system layout described here is quite a bit different from the layout employed previous to release 4.0 of SunOS. For compatibility with earlier releases of SunOS, and other versions of the UNIX system, symbolic links are provided for various files and directories linking their previous names to their current locations. The symbolic links provided include:

/bin → **/usr/bin** All programs previously located in **/bin** are now in **/usr/bin**.
/lib → **/usr/lib** All files previously located in **/lib** are now in **/usr/lib**.
/usr/adm → **/var/adm** The entire **/usr/adm** directory has been moved to **/var/adm**.
/usr/spool → **/var/spool** The entire **/usr/spool** directory has been moved to **/var/spool**.
/usr/tmp → **/var/tmp** The **/usr/tmp** directory has been moved to **/var/tmp**.
/etc/termcap → **/usr/share/lib/termcap**
/usr/5lib/terminfo → **/usr/share/lib/terminfo**
/usr/lib/me → **/usr/share/lib/me**
/usr/lib/ms → **/usr/share/lib/ms**
/usr/lib/tmac → **/usr/share/lib/tmac**
/usr/man → **/usr/share/man**

The following program binaries have been moved from **/etc** to **/usr/etc** with symbolic links to them left in **/etc**: **arp**, **clri**, **cron**, **chown**, **chroot**, **config**, **dkinfo**, **dmesg**, **dump**, **fastboot**, **fasthalt**, **fsck**, **halt**, **ifconfig**, **link**, **mkfs**, **mknod**, **mount**, **ncheck**, **news**, **pstat**, **rdump**, **reboot**, **renice**, **restore**, **rmt**, **rrestore**, **shutdown**, **umount**, **update**, **unlink**, and **vipw**.

In addition, some files and directories have been moved with no symbolic link left behind in the old location:

<i>Old Name</i>	<i>New Name</i>
/etc/biod	/usr/etc/biod
/etc/fsirand	/usr/etc/fsirand
/etc/getty	/usr/etc/getty
/etc/in.rlogind	/usr/etc/in.rlogind
/etc/in.routed	/usr/etc/in.routed
/etc/in.rshd	/usr/etc/in.rshd
/etc/inetd	/usr/etc/inetd
/etc/init	/usr/etc/init
/etc/nfsd	/usr/etc/nfsd
/etc/portmap	/usr/etc/portmap
/etc/rpc.lockd	/usr/etc/rpc.lockd
/etc/rpc.statd	/usr/etc/rpc.statd
/etc/ypbind	/usr/etc/ypbind
/usr/lib/sendmail.cf	/etc/sendmail.cf
/usr/preserve	/var/preserve
/usr/lib/aliases	/etc/aliases
/stand	/usr/stand
/etc/yp	/var/yp

Note: with this new file system organization, the approach to repairing a broken file system changes. One must mount **/usr** before doing an **fsck(8)**, for example. If the mount point for **/usr** has been destroyed, **/usr** can be mounted temporarily on **/mnt** or **/tmp**. If the root file system on a standalone system is so badly damaged that none of these mount points exist, or if **/sbin/mount** has been corrupted, the only way to repair it may be to re-install the root file system.

SEE ALSO

at(1), ex(1), hostname(1), sh(1), vi(1), intro(4), nfs(4P), hier(7), ifconfig(8C), init(8), MAKEDEV(8), mount(8), fsck(8), rc(8)

NAME

hier – file system hierarchy

DESCRIPTION

The following outline gives a quick tour through a typical SunOS file system hierarchy:

```

/      root directory of the file system
/dev/  devices (Section 4)
MAKEDEV
        shell script to create special files
MAKEDEV.local
        site specific part of MAKEDEV
console main system console, console(4S)
drum    paging device, drum(4)
*mem   memory special files, mem(4S)
null   null file or data sink, null(4)
pty[p-z]*
        pseudo terminal controllers, pty(4)
tty[ab] CPU serial ports, zs(4S)
tty[0123][0-f]
        MTI serial ports mti(4S)
tty[hijk][0-f]
        ALM-2 serial ports mcp(4S)
tty[p-z]*
        pseudo terminals, pty(4)
vme*   VME bus special files, mem(4S)
win   window system special files, win(4S)
xy*   disks, xy(4S)
rxy*  raw disk interfaces, xy(4S)
...
/etc/  system-specific maintenance and data files
dumpdates
        dump history, dump(8)
exports table of file systems exportable with NFS, exports(5)
fstab  file system configuration table, fstab(5)
group  group file, group(5)
hosts  host name to network address mapping file, hosts(5)
hosts.equiv
        list of trusted systems, hosts.equiv(5)
motd  message of the day, login(1)
mtab  mounted file table, mtab(5)
networks
        network name to network number mapping file, networks(5)
passwd password file, passwd(5)
phones private phone numbers for remote hosts, as described in phones(5)
printcap
        table of printers and capabilities, printcap(5)
protocols
        protocol name to protocol number mapping file, protocols(5)
rc    shell program to bring the system up multiuser
rc.boot startup file run at boot time
rc.local site dependent portion of rc
remote names and description of remote hosts for tip(1C), remote(5)
services
        network services definition file, services(5)

```

ttytab database of terminal information used by **getty(8)**
 ...

/export/ directory of exported files and file systems for clients, including swap files, root, and **/usr** file systems

/home/ directory of mount points for remote-mounted home directories and shared file systems

user home (initial working) directory for *user*

.profile set environment for **sh(1)**, **environ(5V)**

.project
 what you are doing (used by **(finger(1))**)

.cshrc startup file for **cs(1)**

.exrc startup file for **ex(1)**

.plan what your short-term plans are (used by **finger(1)**)

.rhosts host equivalence file for **rlogin(1C)**

.mailrc startup file for **mail(1)**

calendar
 user's datebook for **calendar(1)**

...

/lost+found directory for connecting detached files for **fsck(8)**

/mnt/ mount point for file systems mounted temporarily

/sbin/ executable programs needed to mount **/usr/**

hostname

ifconfig

init

mount

sh

/tmp/ temporary files, usually on a fast device, see also **/var/tmp/**

ctm* used by **cc(1V)**

e* used by **ed(1)**

...

/var/ directory of files that tend to grow or vary in size

adm/ administrative log files

lastlog record of recent logins, **utmp(5)**

lpacct line printer accounting **lpr(1)**

messages
 system messages

tracct phototypesetter accounting, **troff(1)**

utmp table of currently logged in users, **utmp(5)**

vaacct, vpaacct
 varian and versatec accounting **vtroff(1)**, **pac(8)**

wtmp login history, **utmp(5)**

...

preserve/
 editor temporaries preserved here after crashes/hangups

spool/ delayed execution files

cron/ used by **cron(8)**

lpd/ used by **lpr(1)**

lock present when line printer is active

cf* copy of file to be printed, if necessary

df* control file for print job

tf* transient control file, while **lpr** is working

mail/ mailboxes for **mail(1)**

name mail file for user *name*
name.lock lock file while *name* is receiving mail
mqueue/ mail queue for **sendmail(8)**
secretmail/ like **mail/**, but used by **xsend(1)**
uucp/ work files and staging area for **uucp(1C)**
LOGFILE summary log
LOG.* log file for one transaction
 ...
tmp/ temporary files, to keep **/tmp/** small
raster used by **plot(1G)**
stm* used by **sort(1V)**
 ...
yp/ Yellow Pages database files, **ypfiles(5)**
/usr/ general-purpose directory, usually a mounted file system
bin/ utility programs
as assembler, **as(1)**
cc C compiler executive, c.f. **/usr/lib/ccom**, **/usr/lib/cpp**, **/usr/lib/c2**
cs the C-shell, **cs(1)**
sh the Bourne shell, **sh(1)**
 ...
demo/ demonstration programs
diag/ system tests and diagnostics
dict/ word lists, etc.
spellhist history file for **spell(1)**
words principal word list, used by **look(1)**
 ...
etc/ system administration programs; c.f. section 8
catman update preformatted man pages, **catman(8)**
cron the clock daemon, **cron(8)**
dump file system backup program **dump(8)**
getty part of **login(1)**, **getty(8)**
in.comsat biff server (incoming mail daemon), **comsat(8C)**
init the parent of all processes, **init(8)**
mount **mount(8)**
yp/ Yellow Pages programs
ypinit build and install Yellow Pages database, **ypinit(8)**
yppush force propagation of a changed Yellow Pages map, **yppush(8)**
ypset point **ypbind** at a particular server, **ypset(8)**
 ...
 ...
games/
backgammon
lib/ library directory for game scores, etc.
quiz.k/ what **quiz(6)** knows
africa countries and capitals
index category index
 ...

```

...
hosts/ symbolic links to rsh(1C) for commonly accessed remote hosts
include/
    standard #include files
    a.out.h object file layout, a.out(5)
    images/ icon images
    machine/
        header files from /usr/share/sys/sys/machine; may be a symbolic link
    math.h intro(3M)
    net/ header files from /usr/share/sys/sys/net; may be a symbolic link
    nfs/ header files used in the Network File System (NFS)
    stdio.h standard I/O, intro(3S)
    sys/ kernel header files, c.f. /usr/share/sys/sys
    ...
lib/ object libraries, compiler program binaries, and other data
    ccom C compiler proper
    cpp C preprocessor
    c2 C code improver
    eign list of English words to be ignored by ptx(1)
    font/ fonts for troff(1)
        ftR Times Roman
        ftB Times Bold
        ...
    libc.a system calls, standard I/O, etc. (2,3,3S)
    libm.a math library, intro(3M)
    lint/ utility files for lint
        lint[12] subprocesses for lint(1V)
        llib-lc dummy declarations for /usr/lib/libc.a, used by lint(1V)
        llib-lm dummy declarations for /usr/lib/libm.a
        ...
    units conversion tables for units(1)
    uucp/ programs and data for uucp(1C)
        L.sys remote system names and numbers
        uucico the real copy program
        ...
    ...
local/ locally maintained software
old/ obsolete and unsupported programs
pub/ publicly readable data files
sccs/ binaries of programs that compose the source code control system (SCCS)
src/ system source code tree
stand/ standalone programs (not run under the Sun Operating System)
share/ architecture independent files
    lib/ architecture independent data files
        termcap
            description of terminal capabilities, termcap(5)
        tmac/ macros for troff(1)
            tmac.an
                macros for man(7)
            tmac.s macros for ms(7)
            ...
    ...

```

man/ on-line reference manual pages, **man(1)**
man?/ source files (**nroff(1)**) for sections 1 through 8 of the manual
as.1
 ...
cat?/ preformatted pages for sections 1 through 8 of the manual
 ...
sys/ SunOS kernel source and object modules
ucb/ binaries of programs developed at the University of California, Berkeley
ex line-oriented editor for experienced users, **ex(1)**
vi screen-oriented editor, **vi(1)**
 ...

/vmunix

the SunOS kernel binary

SEE ALSO

filesystem(7), **find(1)**, **finger(1)**, **grep(1V)**, **ls(1V)**, **rlogin(1C)**, **whatis(1)**, **whereis(1)**, **which(1)**, **ncheck(8)**

BUGS

The locations of files are subject to change without notice; the organization of your file system may vary.
 This list is incomplete.

NAME

`man` – macros to format Reference Manual pages

SYNOPSIS

`nroff -man filename...`

`troff -man filename...`

DESCRIPTION

These macros are used to lay out the reference pages in this manual.

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 `.SB` may be used 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.

Requests

<i>Request</i>	<i>Cause Break</i>	<i>If no Argument</i>	<i>Explanation</i>
<code>.B t</code>	no	<i>t</i> =n.t.l.*	Text is in bold font.
<code>.BI t</code>	no	<i>t</i> =n.t.l.	Join words, alternating bold and italic.
<code>.BR t</code>	no	<i>t</i> =n.t.l.	Join words, alternating bold and roman.
<code>.DT</code>	no	<code>.5i li...</code>	Restore default tabs.
<code>.HP i</code>	yes	<i>i</i> =p.i.*	Begin paragraph with hanging indent. Set prevailing indent to <i>i</i> .
<code>.I t</code>	no	<i>t</i> =n.t.l.	Text is italic.
<code>.IB t</code>	no	<i>t</i> =n.t.l.	Join words, alternating italic and bold.
<code>.IP x i</code>	yes	<i>x</i> =""	Same as <code>.TP</code> with tag <i>x</i> .
<code>.IR t</code>	no	<i>t</i> =n.t.l.	Join words, alternating italic and roman.
<code>.LP</code>	yes	-	Begin left-aligned paragraph. Set prevailing indent to <code>.5i</code> .
<code>.PD d</code>	no	<i>d</i> =4v	Set vertical distance between paragraphs.
<code>.PP</code>	yes	-	Same as <code>.LP</code> .
<code>.RE</code>	yes	-	End of relative indent. Restores prevailing indent.
<code>.RB t</code>	no	<i>t</i> =n.t.l.	Join words, alternating roman and bold.
<code>.RI t</code>	no	<i>t</i> =n.t.l.	Join words, alternating roman and italic.
<code>.RS i</code>	yes	<i>i</i> =p.i.	Start relative indent, increase indent by <i>i</i> . Sets prevailing indent to <code>.5i</code> for nested indents.
<code>.SB t</code>	no	-	Reduce size of text by 1 point, make text boldface.
<code>.SH t</code>	yes	<i>t</i> =n.t.l.	Section Heading.
<code>.SM t</code>	no	<i>t</i> =n.t.l.	Reduce size of text by 1 point.
<code>.SS t</code>	yes	<i>t</i> =n.t.l.	Section Subheading.
<code>.TH n s d f m</code>	yes	-	Begin reference page <i>n</i> , of of section <i>s</i> ; <i>d</i> is the date of the most recent change. If present, <i>f</i> is the left page footer; <i>m</i> is the main page (center) header. Sets prevailing indent and tabs to <code>.5i</code> .
<code>.TP i</code>	yes	<i>i</i> =p.i.	Begin indented paragraph, with the tag given on the next text line. Set prevailing indent to <i>i</i> .
			* n.t.l. = next text line; p.i. = prevailing indent
<code>.TX t p</code>	no	-	Resolve the title abbreviation <i>t</i> ; join to punctuation mark (or text) <i>p</i> .

Conventions

A typical manual page for a SunOS command or function is laid out as follows:

.TH TITLE [1-8]

The name of the command or function in upper-case, which serves as the title of the manual page. This is followed by the number of the section in which it appears.

.SH NAME name (or comma-separated list of names) – one-line summary

The name, or list of names, by which the command is called, followed by a dash and then a one-line summary of the action performed. All in roman font, this section contains no **troff(1)** commands or escapes, and no macro requests. It is used to generate the **whatis(1)** database.

.SH SYNOPSIS**Commands:**

The syntax of the command and its arguments, as typed on the command line. When in boldface, a word must be typed exactly as printed. When in italics, a word can be replaced with an argument that you supply. References to bold or italicized items are not capitalized in other sections, even when they begin a sentence.

Syntactic symbols appear in roman face:

- [] An argument, when surrounded by brackets is optional.
- | Arguments separated by a vertical bar are exclusive. You can supply only item from such a list.
- ... Arguments followed by an elipsis can be repeated. When an elipsis follows a bracketed set, the expression within the brackets can be repeated.

Functions:

If required, the data declaration, or **#include** directive, is shown first, followed by the function declaration. Otherwise, the function declaration is shown.

.SH DESCRIPTION

A narrative overview of the command or function's external behavior. This includes how it interacts with files or data, and how it handles the standard input, standard output and standard error. Internals and implementation details are normally omitted. This section attempts to provide a succinct overview in answer to the question, "what does it do?"

Literal text from the synopsis appears in boldface, as do literal filenames and references to items that appear elsewhere in the *SunOS Reference Manual*. Arguments are italicized.

If a command interprets either subcommands or an input grammar, its command interface or input grammar is normally described in a **USAGE** section, which follows the **OPTIONS** section. (The **DESCRIPTION** section only describes the behavior of the command itself, not that of subcommands.)

.SH OPTIONS

The list of options along with a description of how each affects the command's operation.

.SH FILES

A list of files associated with the command or function.

.SH SEE ALSO

A comma-separated list of related manual pages, followed by references to other published materials.

.SH DIAGNOSTICS

A list of diagnostic messages and an explanation of each.

.SH BUGS

A description of limitations, known defects, and possible problems associated with the command or function.

FILES

/usr/share/lib/tmac/tmac.an

SEE ALSO

man(1), nroff(1), troff(1), whatis(1)

Formatting Documents.

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**, **meqn**, **mefer**, and **tbl(1)** preprocessors for equations and tables is acceptable as input.

REQUESTS

In the following list, "initialization" refers to the first .pp, .lp, .ip, .np, .sh, or .uh macro. This list is incomplete.

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
.(xx	-	no	Begin indexed item in index <i>x</i>
.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
++ <i>m H</i>	-	no	Define paper section. <i>m</i> defines the part of the paper, and can be C (chapter), A (appendix), P (preliminary, for instance, 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 eqn or meqn .
.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.
.GE	-	yes	End <i>gremlin</i> picture.
.GS	-	yes	Begin <i>gremlin</i> picture.

.PE	-	yes	End <i>pic</i> picture.
.PS	-	yes	Begin <i>pic</i> picture.
.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)
.bu	-	yes	Begin bulleted paragraph
.bx <i>x</i>	no	no	Print <i>x</i> in a box (nofill only).
.ef 'x'y'z	****	no	Set even footer to <i>x y z</i>
.eh 'x'y'z	****	no	Set even header to <i>x y z</i>
.fo 'x'y'z	****	no	Set footer to <i>x y z</i>
.hx	-	no	Suppress headers and footers on next page.
.he 'x'y'z	****	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 'x'y'z	****	no	Set odd footer to <i>x y z</i>
.oh 'x'y'z	****	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.
.sm <i>x -</i>	no	no	Set <i>x</i> in a smaller pointsize.
.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 troff). (Nofill only).
.uh	-	yes	Like .sh but unnumbered.
.xp <i>x</i>	-	no	Print index <i>x</i> .

FILES

/usr/share/lib/tmac/tmac.e

/usr/share/lib/me/*

SEE ALSO

eqn(1), nroff(1), troff(1), refer(1), tbl(1)

Formatting Documents

NAME

ms – text formatting macros

SYNOPSIS

nroff –ms [options] filename ...

troff –ms [options] filename ...

DESCRIPTION

This package of **nroff**(1) and **troff**(1) 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**(1V). All external –ms macros are defined below.

Note: this –ms macro package is an extended version written at Berkeley and is a superset of the standard –ms macro packages as supplied by Bell Labs. Some of the Bell Labs macros have been removed; for instance, it is assumed that the user has little interest in producing headers stating that the memo was generated at Whippany Labs.

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: <i>n</i> =1 single, <i>n</i> =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**(1), **eqn**(1) and **refer**(1) preprocessors for equations, tables, and references is acceptable as input.

REQUESTS

Macro Name	Initial Value	Break? Reset?	Explanation
.AB x	–	y	begin abstract; if <i>x</i> =no do not label abstract
.AE	–	y	end abstract
.AI	–	y	author's institution
.AM	–	n	better accent mark definitions
.AU	–	y	author's name
.B x	–	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 x	–	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 x	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 x y	I	y	begin display with keep; <i>x</i> =I, L, C, B; <i>y</i> =indent
.ID y	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 x	–	n	even page footer <i>x</i> (3 part as for .tl)
.EH x	–	n	even page header <i>x</i> (3 part as for .tl)
.EN	–	y	end displayed equation produced by eqn

.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 first 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 first 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 TAB characters 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 tbl
.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 troff
.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 refer 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 .JP-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
<code>*Q</code>	quote (" in <code>nroff</code> , " in <code>troff</code>)
<code>*U</code>	unquote (" in <code>nroff</code> , " in <code>troff</code>)
<code>*-</code>	dash (-- in <code>nroff</code> , — in <code>troff</code>)
<code>*(MO</code>	month (month of the year)
<code>*(DY</code>	day (current date)
<code>**</code>	automatically numbered footnote
<code>*' </code>	acute accent (before letter)
<code>*`</code>	grave accent (before letter)
<code>*^</code>	circumflex (before letter)
<code>*,</code>	cedilla (before letter)
<code>*:</code>	umlaut (before letter)
<code>*_</code>	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.

FILES

`/usr/share/lib/tmac/tmac.s`
`/usr/share/lib/ms/ms.???`

SEE ALSO

`col(1V)`, `eqn(1)`, `nroff(1)`, `refer(1)`, `tbl(1)`, `troff(1)`

Formatting Documents

BUGS

Floating keeps and regular keeps are diverted to the same space, so they cannot be mixed together with predictable results.

NAME

intro – introduction to system maintenance and operation commands

DESCRIPTION

This section contains information related to system bootstrapping, operation and maintenance. It describes all the server processes and daemons that run on the system, as well as standalone (PROM monitor) programs.

Disk formatting and labeling is done by `format(8S)`. Bootstrapping of the system is described in `boot(8S)` and `init(8)`. The standard set of commands run by the system when it boots is described in `rc(8)`. Related commands include those that check the consistency of file systems, `fsck(8)`; those that mount and unmount file systems, `mount(8)`; add swap devices, `swapon(8)`; force completion of outstanding file system I/O, `sync(2)`; shutdown or reboot a running system `shutdown(8)`, `halt(8)`, and `reboot(8)`; and, set the time on a machine from the time on another machine `rdate(8)`.

Creation of file systems is discussed in `mkfs(8)` and `newfs(8)`. File system performance parameters can be adjusted with `tunefs(8)`. File system backups and restores are described in `dump(8)` and `restore(8)`.

Procedures for adding new users to a system are described in `adduser(8)`, using `vipw(8)` to lock the password file during editing. `crash(8S)` which describes what happens when the system crashes, `savecore(8)` and `analyze(8)`, which can be used to analyze system crash dumps. Occasionally useful as adjuncts to the `fsck(8)` file system repair program are `clri(8)`, `dcheck(8)`, `icheck(8)`, and `ncheck(8)`.

Configuring a new version of the kernel requires using the program `config(8)`; major system bootstraps often require the use of `mkproto(8)`. New devices are added to the `/dev` directory (once device drivers are configured into the system) using `makedev(8)` and `mknod(8)`. The `installboot(8S)` command can be used to install freshly compiled programs. The `catman(8)` command preformats the on-line manual pages.

Resource accounting is enabled by the `accton` command, and summarized by `sa(8)`. Login time accounting is performed by `ac(8)`. Disk quotas are managed using `quot(8)`, `quotacheck(8)`, `quotaon(8)`, and `repquota(8)`.

A number of servers and daemon processes are described in this section. The `update(8)` daemon forces delayed file system I/O to occur and `cron(8)` runs periodic events (such as removing temporary files from the disk periodically). The `syslogd(8)` daemon maintains the system error log. The `init(8)` process is the initial process created when the system boots. It manages the reboot process and creates the initial login prompts on the various system terminals, using `getty(8)`. The Internet super-server `inetd(8C)` invokes all other internet servers as needed. These servers include the remote shell servers `rshd(8C)` and `rexecd(8C)`, the remote login server `rlogind(8C)`, the FTP and TELNET daemons `ftpd(8C)`, and `telnetd(8C)`, the TFTP daemon `tftpd(8C)`, and the mail arrival notification daemon `comsat(8C)`. Other network daemons include the 'load average/who is logged in' daemon `rwhod(8C)`, the routing daemon `routed(8C)`, and the mail daemon `sendmail(8)`.

If network protocols are being debugged, then the protocol debugging trace program `trpt(8C)` is often useful. Remote magnetic tape access is provided by `rsh` and `rmt(8C)`. Remote line printer access is provided by `lpd(8)`, and control over the various print queues is provided by `lpc(8)`. Printer cost-accounting is done through `pac(8)`.

Network host tables may be gotten from the ARPA NIC using `gettable(8C)` and converted to UNIX-system-usable format using `htable(8)`.

RPC and NFS daemons

RPC and NFS daemons include:

portmap	used by RPC based services.
yplibd	used by the Yellow Pages to locate the Yellow Pages server.
biod	used by NFS clients to read ahead to, and write behind from, network file systems.
nfsd	the NFS server process that responds to NFS requests on NFS server machines.
ypserv	the Yellow Pages server, typically run on each NFS server.
rstatd	the server counterpart of the remote speedometer tools.

mountd the **mount** server that runs on NFS server machines and responds to requests by other machines to mount file systems.

rwalld used for broadcasting messages over the network.

LIST OF MAINTENANCE COMMANDS

Name	Appears on Page	Description
ac	ac(8)	login accounting
accton	sa(8)	system accounting
adbgen	adbgen(8)	generate adb script
adduser	adduser(8)	procedure for adding new users
arp	arp(8C)	address resolution display and control
audit	audit(8)	audit trail maintenance
audit_warn	audit_warn(8)	audit space low warning script
auditd	auditd(8)	audit daemon
automount	automount(8)	automatically mount NFS file systems
biod	nfsd(8)	NFS daemons
boot	boot(8S)	start the system kernel or a standalone program
bootparamd	bootparamd(8)	boot parameter server
captainfo	captainfo(8V)	convert a termcap description into a terminfo description
catman	catman(8)	create the cat files for the manual
chown	chown(8)	change owner
chroot	chroot(8)	change root directory for a command
client	client(8)	add/remove diskless systems
clri	clri(8)	clear i-node
comsat	comsat(8C)	biff server
config	config(8)	build system configuration files
crash	crash(8S)	what happens when the system crashes
cron	cron(8)	clock daemon
dcheck	dcheck(8)	file system directory consistency check
dinfo	dinfo(8)	report information about a disk's geometry and partitioning
dmesg	dmesg(8)	collect system diagnostic messages to form error log
dump	dump(8)	incremental file system dump
dumpfs	dumpfs(8)	dump file system information
edquota	edquota(8)	edit user quotas
eeprom	eeprom(8S)	Sun-3 EEPROM display and load utility
etherd	etherd(8C)	Ethernet statistics server
etherfind	etherfind(8C)	find packets on Ethernet
exportfs	exportfs(8)	export and unexport directories to NFS clients
fastboot	fastboot(8)	reboot/halt the system without checking the disks
fasthalt	fastboot(8)	reboot/halt the system without checking the disks
fingerd	fingerd(8C)	remote user information server
format	format(8S)	disk partitioning and maintenance utility
fparel	fparel(8)	Sun FPA online reliability tests
fpaversion	fpaversion(8)	print FPA version
fsck	fsck(8)	file system consistency check and interactive repair
fsirand	fsirand(8)	install random inode generation numbers
ftpd	ftpd(8C)	DARPA Internet File Transfer Protocol server
gettable	gettable(8C)	get DoD Internet format host table from a host
getty	getty(8)	set terminal mode
gpconfig	gpconfig(8)	initialize the Graphics Processor
grpck	grpck(8)	check group database entries
halt	halt(8)	stop the processor
htable	htable(8)	convert DoD Internet format host table

icheck	icheck(8)	file system storage consistency check
ifconfig	ifconfig(8C)	configure network interface parameters
inetd	inetd(8C)	Internet services daemon
infocmp	infocmp(8V)	compare or print out terminfo descriptions
init	init(8)	process control initialization
installboot	boot(8S)	start the system kernel or a standalone program
iostat	iostat(8)	report I/O statistics
ipallocald	ipallocald(8C)	Ethernet-to-IP address allocator
kadb	kadb(8S)	adb-like kernel and standalone-program debugger
keyenvoy	keyenvoy(8C)	talk to keyserver
keyserv	keyserv(8C)	server for storing public and private keys
kgmon	kgmon(8)	generate a dump of the operating system's profile buffers
ldconfig	ldconfig(8)	configure cache for ld.so
link	link(8)	exercise link and unlink system calls
lockd	lockd(8C)	network lock daemon
logintool	logintool(8)	graphic login interface
lpc	lpc(8)	line printer control program
lpd	lpd(8)	printer daemon
mailstats	mailstats(8)	print statistics collected by sendmail
MAKEDBM	makedbm(8)	make a Yellow Pages dbm file
MAKEDEV	makedev(8)	make system special files
makekey	makekey(8)	generate encryption key
mc68881version	mc68881version(8)	print the MC68881 mask number and approximate clock rate
mconnect	mconnect(8)	connect to SMTP mail server socket
mkfs	mkfs(8)	construct a file system
mknod	mknod(8)	build special file
mkproto	mkproto(8)	construct a prototype file system
modload	modload(8)	load a loadable module
modstat	modstat(8)	display status of loadable modules
modunload	modunload(8)	unload a loadable module
monitor	monitor(8S)	system ROM monitor
mount	mount(8)	mount and dismount filesystems
mountd	mountd(8C)	NFS mount request server
named	named(8C)	Internet domain name server
ncheck	ncheck(8)	generate names from i-numbers
ndbootd	ndbootd(8C)	ND boot block server
netconfig	netconfig(8C)	PNP boot service
netstat	netstat(8C)	show network status
newaliases	newaliases(8)	rebuild the data base for the mail aliases file
newfs	newfs(8)	construct a new file system
newkey	newkey(8)	create a new key in the publickey database
nfsd	nfsd(8)	NFS daemons
nfsstat	nfsstat(8C)	Network File System statistics
pac	pac(8)	printer/plotter accounting information
ping	ping(8C)	send ICMP ECHO_REQUEST packets to network hosts
pnps386	pnps386(8C)	PNP diskless boot service
pnpsboot	pnpsboot(8C)	PNP diskless boot service
pnpd	pnpd(8C)	PNP daemon
portmap	portmap(8C)	DARPA port to RPC program number mapper
praudit	praudit(8)	print contents of an audit trail file
pstat	pstat(8)	print system facts
pwck	pwck(8)	check password database entries
pwdauthd	pwdauthd(8C)	server for authenticating passwords

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
rarpd	rarpd(8C)	DARPA Reverse Address Resolution Protocol service
rc.boot	rc(8)	command scripts for auto-reboot and daemons
rc.local	rc(8)	command scripts for auto-reboot and daemons
rc	rc(8)	command scripts for auto-reboot and daemons
rdate	rdate(8C)	set system date from a remote host
rdump	dump(8)	incremental file system dump
reboot	reboot(8)	restart the operating system
renice	renice(8)	alter priority of running processes
repquota	repquota(8)	summarize quotas for a file system
restore	restore(8)	incremental file system restore
rex	rex(8C)	RPC-based remote execution server
rexecd	rexecd(8C)	remote execution server
rlogind	rlogind(8C)	remote login server
rmail	rmail(8C)	handle remote mail received via uucp
rmt	rmt(8C)	remote magtape protocol module
route	route(8C)	manually manipulate the routing tables
routed	routed(8C)	network routing daemon
rpcinfo	rpcinfo(8C)	report RPC information
rquotad	rquotad(8C)	remote quota server
rrestore	restore(8)	incremental file system restore
rshd	rshd(8C)	remote shell server
rstatd	rstatd(8C)	kernel statistics server
rusersd	rusersd(8C)	network username server
rwalld	rwalld(8C)	network rwall server
rwhod	rwhod(8C)	system status server
sa	sa(8)	system accounting
savecore	savecore(8)	save a core dump of the operating system
sendmail	sendmail(8)	send mail over the internet
setup_client	setup_client(8)	create or remove a nfs client on a 4.0 server.
setup_exec	setup_exec(8)	install architecture-dependent executable files
showmount	showmount(8)	show all remote mounts
shutdown	shutdown(8)	close down the system at a given time
spray	spray(8C)	spray packets
sprayd	sprayd(8C)	spray server
statd	statd(8C)	network status monitor
sticky	sticky(8)	persistent text and append-only directories
suninstall	suninstall(8)	SunOS software installation program
swapon	swapon(8)	specify additional device for paging and swapping
sysdiag	sysdiag(8)	system diagnostics
syslogd	syslogd(8)	log system messages
talkd	talkd(8C)	server for talk program
telnetd	telnetd(8C)	DARPA TELNET protocol server
tftpd	tftpd(8C)	DARPA Trivial File Transfer Protocol server
tic	tic(8V)	terminfo compiler
timed	timed(8C)	DARPA Time server
tnamed	tnamed(8C)	DARPA Trivial name server
trpt	trpt(8C)	transliterate protocol trace
tunefs	tunefs(8)	tune up an existing file system
umount	mount(8)	mount and dismount filesystems

unconfigure	unconfigure(8)	reset the network configuration for a system
unlink	link(8)	exercise link and unlink system calls
update	update(8)	periodically update the super block
uuclean	uuclean(8C)	uucp spool directory clean-up
vipw	vipw(8)	edit the password file
vmstat	vmstat(8)	report virtual memory statistics
ypbind	ypserv(8)	Yellow Pages server and binder processes
ypinit	ypinit(8)	build and install Yellow Pages database
ypmake	ypmake(8)	rebuild Yellow Pages database
yppasswdd	yppasswdd(8C)	server for modifying Yellow Pages password file
yppoll	yppoll(8)	what version of a YP map is at a YP server host
yppush	yppush(8)	force propagation of a changed YP map
ypserv	ypserv(8)	Yellow Pages server and binder processes
ypset	ypset(8)	point ypbind at a particular server
ypupdated	ypupdated(8C)	server for changing yp information
ypwhich	ypwhich(8)	what machine is the YP server?
ypxfr	ypxfr(8)	transfer YP map from a YP server to here
zdump	zdump(8)	time zone dumper
zic	zic(8)	time zone compiler

NAME

ac – login accounting

SYNOPSIS

/usr/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.

The accounting file **/var/adm/wtmp** is maintained by **init(8)** and **login(1)**. 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.

OPTIONS

- w** Specify an alternate *wtmp* file.
- p** Print individual totals; without this option, only totals are printed.
- d** 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, **/var/adm/wtmp** is used.

FILES

/var/adm/wtmp

SEE ALSO

login(1), **utmp(5)**, **init(8)**, **sa(8)**

NAME

adbgen – generate adb script

SYNOPSIS

`/usr/lib/adb/adbgen filename.adb ...`

DESCRIPTION

adbgen makes it possible to write adb(1) scripts that do not contain hard-coded dependencies on structure member offsets. The input to **adbgen** is a file named *filename.adb* which contains **adbgen** header information, then a null line, then the name of a structure, and finally an adb script. **adbgen** only deals with one structure per file; all member names are assumed to be in this structure. The output of **adbgen** is an adb script in *filename*. **adbgen** operates by generating a C program which determines structure member offsets and sizes, which in turn generates the adb script.

The header lines, up to the null line, are copied verbatim into the generated C program. Typically these include C `#include` statements to include the header files containing the relevant structure declarations.

The adb script part may contain any valid adb commands (see adb(1)), and may also contain **adbgen** requests, each enclosed in `{}`s. Request types are:

- 1 Print a structure member. The request form is `{member format}`. *member* is a member name of the *structure* given earlier, and *format* is any valid adb format request. For example, to print the `p_pid` field of the *proc* structure as a decimal number, you would write `{p_pid,d}`.
- 2 Reference a structure member. The request form is `{*member,base}`. *member* is the member name whose value is desired, and *base* is an adb register name which contains the base address of the structure. For example, to get the `p_pid` field of the *proc* structure, you would get the *proc* structure address in an adb register, say `<f`, and write `{*p_pid,<f}`.
- 3 Tell **adbgen** that the offset is ok. The request form is `{OFFSETOK}`. This is useful after invoking another adb script which moves the *adb dot*.
- 4 Get the size of the *structure*. The request form is `{SIZEOF}`. **adbgen** replaces this request with the size of the structure. This is useful in incrementing a pointer to step through an array of structures.
- 5 Get the offset to the end of the structure. The request form is `{END}`. This is useful at the end of the structure to get adb to align the *dot* for printing the next structure member.

adbgen keeps track of the movement of the *adb dot* and emits adb code to move forward or backward as necessary before printing any structure member in a script. **adbgen**'s model of the behavior of *adb*'s *dot* is simple: it is assumed that the first line of the script is of the form *struct_address/adb text* and that subsequent lines are of the form *+/adb text*. This causes the *adb dot* to move in a sane fashion. **adbgen** does not check the script to ensure that these limitations are met. **adbgen** also checks the size of the structure member against the size of the adb format code and warns you if they are not equal.

EXAMPLE

If there were an include file *x.h* which contained:

```
struct x {
    char    *x_cp;
    char    x_c;
    int     x_i;
};
```

Then an **adbgen** file (call it *script.adb*) to print it would be:

```
#include "x.h"
x
./"x_cp"16t"x_c"8t"x_i"n{x_cp,X}{x_c,C}{x_i,D}
```

After running **adbgen** the output file **script** would contain:

```
16t"x_c"8t"x_i"nXC+D"" /"x_cp"16t"x_c"8t"x_i"nXC+D
```

To invoke the script you would type:

```
x$<script
```

FILES

/usr/lib/adb/* **adb** scripts for debugging the kernel

SEE ALSO

adb(1), **kadb(8S)**

Debugging Tools

BUGS

adb syntax is ugly; there should be a higher level interface for generating scripts.

Structure members which are bit fields cannot be handled because C will not give the address of a bit field. The address is needed to determine the offset.

DIAGNOSTICS

Warnings about structure member sizes not equal to **adb** format items and complaints about badly formatted requests. The C compiler complains if you reference a structure member that does not exist. It also complains about **&** before array names; these complaints may be ignored.

NAME

adduser – procedure for adding new users

DESCRIPTION

To add an account for a new user, the system administrator (or super-user):

- Create an entry for the new user in the system password files.
- Create a home directory for the user, and change ownership so the new user owns that directory.
- Optionally set up skeletal dot files for the new user (.cshrc, .login, .profile...).
- If the account is on a system running the YP name service, take additional measures.

USAGE**Making an Entry in the Password File**

To add an entry for the new login name on a local host, first edit the `/etc/passwd` file — inserting a line for the new user. This must be done with the password file locked, for instance, by using `vipw(8)`, and the insertion must be made above the line containing the string:

```
+::0:0:::
```

This line is used to indicate that additional accounts can be found in the YP.

To add an entry for the new login name on to the YP, add an identical line to the file `/etc/passwd` on the YP master server, and run `make(1)` in the directory `/var/yp` (see `yppmake(8)` for details) to propagate the change.

The new user is assigned a group and user ID number. User ID numbers (or `userids`, or `UIDs`) should be unique for each user and consistent across the NFS domain, since they control access to files. Group ID numbers (or `groupids`, or `GIDs`) need not be unique. Typically, users working on similar projects will be assigned to the same group. The system staff is group 10 for historical reasons, and the super-user is in this group.

An entry for a new user “francine” would look like this:

```
francine::235:20:& Featherstonehaugh:/usr/francine:/bin/csh
```

Fields in each password-file entry are delimited by colons, and have the following meanings:

- Login name (“francine”). The login name is limited to eight characters in length.
- Encrypted password or the string `##name` if encrypted passwords are stored in the password adjunct file. Typically, if passwords are to be stored in the main password file, this field is left empty, so no password is needed when the user first logs in. If security demands a password, it should be assigned by running `passwd(1)` immediately after exiting the editor. The number of significant characters in a password is eight. (See `passwd(1)`.)
- User ID. The UID is a number which identifies that user uniquely in the system. Files owned by the user have this number stored in their data blocks, and commands such as `ls(1V)` use it to look up the owner’s login name. For this reason, you cannot randomly change this number. See `passwd(5)` for more information.
- Group ID. The UID number identifies the group to which the user belongs by default (although the user may belong to additional groups as well). All files that the user creates have this number stored in their data blocks, and commands such as `ls(1V)` use it to look up the group name. Group names and assignments are listed in the file `/etc/group` (which is described in `group(5)`) or in the YP group map.
- This field is called the GCOS field (from earlier implementation of the operating system) and is traditionally used to hold the user’s full name. Some installations have other information encoded in this field. From this information we can tell that Francine’s real name is ‘Francine Featherstonehaugh’. The `&` in the entry is shorthand for the user’s login name.

- User's home directory. This is the directory in which that user is "positioned" when they log in.
- Initial shell which this user will see on login. If this field is empty, sh(1) is used as the initial shell.

An entry for a new user "francine" would look like this:

```
francine:::::lo:ad,+dw
```

Fields in each password adjunct file entry are delimited by colons, and have the following meanings:

- Login name ("francine"). This name must match the login name in the password file.
- Encrypted password. Typically, this field is left empty when adding the line using the editor. passwd(1) should be run immediately after exiting the editor.
- The next three fields are the minimum label, the maximum label, and the default label. These fields should be left empty, since they are reserved for future use.
- The next two fields are for the always-audit flags and the never-audit flags. Always-audit flags specify which events guaranteed to be audited for that user. Never-audit flags specify which events are guaranteed not to be audited for that user. For a description of audit flags, see audit_data(5).

Making a Home Directory

As shown in the password file entry above, the name of Francine's home directory is to be /usr/francine. This directory must be created using mkdir(1), and Francine must be given ownership of it using chown(8), in order for her profile files to be read and executed, and to have control over access to it by other users:

```
example# mkdir /usr/francine
example# /usr/etc/chown francine /usr/francine
```

If running under NFS, the mkdir(1) and chown(8) commands must be performed on the NFS server.

Setting Up Skeletal Profile Files

New users often need assistance in setting up their profile files to initialize the terminal properly, configure their search path, and perform other desired functions at startup. Providing them with skeletal profile files saves time and interruptions for both the new user and the system administrator.

Such files as .profile (if they use /usr/bin/sh as the shell), or .cshrc and .login (if they use /usr/bin/csh as the shell), can include commands that are performed automatically at each login, or whenever a shell is invoked, such as tset(1). The ownership of these files must be changed to belong to the new user, either by running su(1) before making copies, or by using chown(8).

FILES

```
/etc/passwd          password file
/etc/group            group file
/etc/yp/src/passwd
~/.cshrc
~/.login
~/.profile
```

SEE ALSO

csh(1), ls(1v), make(1), mkdir(1), passwd(1), sh(1), su(1), tset(1), audit(2), audit_control(5), audit_data(5), group(5), passwd(5), passwd.adjunct(5), audit(8), auditd(8), chown(8), vipw(8), ypmake(8)

Network Programming

NAME

arp – address resolution display and control

SYNOPSIS

arp *hostname*

arp -a [*vmunix* [*kmem*]]

arp -d *hostname*

arp -s *hostname ether_address* [**temp**] [**pub**] [**trail**]

arp -f *filename*

DESCRIPTION

The **arp** program displays and modifies the Internet-to-Ethernet address translation tables used by the address resolution protocol (**arp(4P)**).

With no flags, the program displays the current ARP entry for *hostname*. The host may be specified by name or by number, using Internet dot notation.

OPTIONS

- a** Display all of the current ARP entries by reading the table from the file *kmem* (default */dev/kmem*) based on the kernel file *vmunix* (default */vmunix*).
- d** Delete an entry for the host called *hostname*. This option may only be used by the super-user.
- s** Create an ARP entry for the host called *hostname* with the Ethernet address *ether_address*. The Ethernet address is given as six hex bytes separated by colons. The entry will be permanent unless the word **temp** is given in the command. If the word **pub** is given, the entry will be published, for instance, this system will respond to ARP requests for *hostname* even though the host-name is not its own. The word **trail** indicates that trailer encapsulations may be sent to this host.
- f** Read the file named *filename* and set multiple entries in the ARP tables. Entries in the file should be of the form

hostname ether_address [**temp**] [**pub**] [**trail**]

with argument meanings as given above.

SEE ALSO

arp(4P), **ifconfig(8C)**

NAME

audit – audit trail maintenance

SYNOPSIS

```
audit [ -n|-s|-t ]  
audit -d username  
audit -u username audit_event_state
```

AVAILABILITY

This program is available with the *Security* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

The **audit** command is the general administrator's interface to kernel auditing. The process audit state for a user can be temporarily or permanently altered. The audit daemon may be notified to read the contents of the **audit_control** file and re-initialize the current audit directory to the first directory listed in the **audit_control** file, or to open a new audit file in the current audit directory specified in the **audit_control** file as last read by the audit daemon. Auditing may also be terminated/disabled.

OPTIONS

- n Signal audit daemon to close the current audit file and open a new audit file in the current audit directory.
- s Signal audit daemon to read audit control file. The audit daemon stores the information internally.
- t Signal audit daemon to disable auditing and die.
- d *username*
Change the process audit state of all processes owned by *username*. This new process audit state is constructed from the system and user audit values as specified in the **audit_control** and **passwd.adjunct** files respectively.
- u *username audit_event_state*
Set the process audit state from *audit_event_state* for all current processes owned by *username*. See **audit_control(5)** for the format of the system audit value. The process audit state is one argument. Enclose the audit event state in quotes, or do not use SPACE characters in the process audit state specification. A new login session reconstructs the process audit state from the audit flags in the **audit_control** and **passwd.adjunct** files.

SEE ALSO

audit(2), **setuseraudit(2)**, **getauditflags(3)**, **getfauditflags(3)**, **audit_control(5)**, **passwd.adjunct(5)**

NAME

`audit_warn` – audit space low warning script

SYNOPSIS

`/etc/security/audit/audit_warn` logfile

DESCRIPTION

The `audit_warn` shell script is used to take appropriate action when audit filesystem space is running low. This script is run when the audit filesystem free disk space drops below the value *minfree* as specified in the file `audit_control`. The script is passed the name of the current audit log file.

The `audit_warn` script included with the installation tape issues the command "audit -n", which tells the audit daemon to switch to the next audit directory.

SEE ALSO

`auditd(8)`, `audit(8)`, `audit.log(5)`, `audit_control(5)`

NAME

auditd – audit daemon

SYNOPSIS

/etc/auditd [*username*]

DESCRIPTION

The **auditd** program is the daemon process that drives audit log generation. **auditd** runs as root unless the *username* parameter specifies another valid user ID. Use of this parameter is recommended to insure that **auditd** works over NFS and that the audit data is secure.

After reading the **audit_control** file, **auditd** opens an audit log file in the first directory specified. If there is an error opening this file, the daemon tries successive directories until successful. Then the daemon invokes the **auditsvc(2)** system call to initiate audit record logging to the audit log file. The system call does not return until:

- disk space is low on the audit filesystem
- there is an error writing the audit log file or
- the daemon receives a signal

auditd simply ignores most signals and re-issues the **auditsvc()** system call. However, for **SIGHUP**, the daemon re-reads the **audit_control** file, closes the current audit log file, and opens a new audit log file based on the new directory list.

If the **auditsvc()** system call returns because of low disk space, **auditd** invokes the shell script **audit_warn(8)** with the name of the current audit log file, then returns to the **auditsvc()** system call to continue auditing using the same audit file.

When the **auditsvc()** system call returns because of an error, **auditd** recovers from the problem by closing the current audit log file and opening a new audit log file in the next directory in the list. This recovers from most errors, such as lack of disk space, or file server crashes.

Should the audit daemon run out of audit directories, it attempts to recover. It suspends itself for a few seconds, and then re-reads the **audit_control** file. It then tries again to open audit logs in the specified audit directory list.

SEE ALSO

auditsvc(2), **audit_control(5)**, **audit.log(5)**, **audit(8)**, **audit_warn(8)**

NAME

automount – automatically mount NFS file systems

SYNOPSIS

```
automount [ -mnT ] [ -tl duration ] [ -tm interval ] [ -tw interval ]
    [ directory mapname [ -mount-options ] ] ...
```

DESCRIPTION

automount is a daemon that will automatically and transparently mount an NFS file system whenever a file or directory within in that system is opened. **automount** forks a daemon, which appears to be an NFS server to the kernel; lookups on the specified *directory* are intercepted by this daemon, which uses the map contained in *mapname* to determine a server, exported file system, and appropriate mount options for a given file system. The named map can either be a file on the local system, or a Yellow Pages map. *directory* is a full pathname starting with a '/

When supplied, **-mount-options** consists of the leading **-** and a comma-separate list of **mount(8)** options; if mount options are specified in the map, however, those in the map take precedence.

Once mounted, members of the *directory* are made available using a symbolic link to the real mount point within a temporary directory.

If *directory* does not exist, the daemon creates it, and then removes it automatically when the daemon exits.

Since the name-to-location binding is dynamic, updates to a Yellow Pages map are transparent to the user. This obviates the need to “pre-mount” shared file systems for applications that have “hard coded” references to files. It also obviates the need to maintain records of which hosts must be mounted for what applications.

Maps

automount looks first for the indicated *mapname* in a file by that name. If there is no such file, it looks for a YP map by that name.

An automount map is composed of a list of mappings, with one mapping per line. Each mapping is composed of the following fields:

```
basename [ -mount-options ] location [...]
```

where *basename* is the name of a subdirectory within the *directory* specified in the **automount** command line (not a relative pathname). The *location* field consists of an entry of the form:

```
host:directory[:subdir]
```

where *host* is the name of the host from which to mount the file system, *directory* is the pathname of the directory to mount, and *subdir*, when supplied, is the name of a subdirectory to which the symbolic link is made. This can be used to prevent duplicate mounts in cases where multiple directories in the same remote file system are accessed.

The contents of a YP map can be included within a map by adding an entry of the form:

```
+mapname
```

A mapping can be continued across line breaks using a \ as the last character before the NEWLINE. Comments begin with a # and end at the subsequent NEWLINE.

If more than one *location* is supplied, there is no guarantee as to which location will be used; the first location to respond to the mount request gets mounted. The *mount-options* field can be used to supply options to the **mount(8)** command for the mounted file system.

Special Maps

There are two special maps currently available. The `-hosts` map uses the Yellow Pages `hosts.byname` map to locate a remote host when the hostname is specified as a subdirectory of *directory*. This map specifies mounts of all exported file systems from any host. For instance, if the following `automount` command is already in effect:

```
automount /net -hosts
```

then a reference to `/net/hermes/usr` would initiate an automatic mount of all file systems from `hermes` that `automount` can mount; references to a directory under `/net/hermes` will refer to the corresponding directory on `hermes`. The `-passwd` map uses the `passwd(5)` database to attempt to locate the home directory of a user. For instance, if the following `automount` command is already in effect:

```
automount /homes -passwd
```

then if the home directory shown in the `passwd` entry for the user *username* has the form `/dir/server/username`, and `server` matches the host system on which that directory resides, references to files in `/homes/username` result in the file system containing that directory being mounted if necessary, and all such references will refer to that user's home directory.

Configuration

`automount` normally consults the `auto.master` Yellow Pages configuration database for a list of initial *directory* to *mapname* pairs, and sets up automatic mounts for them in addition to those given on the command line; if there are duplications, the command-line arguments take precedence. (Note that this database contains arguments to the `automount` command, rather than mappings, and that `automount` does *not* look for an `auto.master` file on the local host.)

OPTIONS

- `-m` Suppress initialization of *directory-mapname* pairs listed in the `auto.master` Yellow Pages database.
- `-n` Disable dynamic mounts. With this option, references through the `automount` daemon only succeed when the target filesystem has been previously mounted. This can be used to prevent NFS servers from cross-mounting each other.
- `-T` Trace. Expand each NFS call and display it on the standard output.
- `-tl duration`
Specify a *duration*, in seconds, that a looked up name remains cached when not in use. The default is 5 minutes.
- `-tm interval`
Specify an *interval*, in seconds, between attempts to mount a filesystem. The default is 30 seconds.
- `-tw interval`
Specify an *interval*, in seconds, between attempts to dismount filesystems that have exceeded their cached times. The default is 1 minute.

EXAMPLE

```
tutorial# automount -m /net -hosts
```

Provide `automount` access to the exported file systems of any host in the Yellow Pages `hosts.byname` database, by prefixing the pathname with `/net/hostname/`:

```
tutorial% ls /net/hermes/usr/src ...
```

FILES

`/tmp_mnt` directory under which filesystems are dynamically mounted

SEE ALSO

`mount(8)`

BUGS

Shell filename expansion does not apply to objects not currently mounted or cached. For instance, in the above example, the command `ls /net/*` might not list `hermes` as a subdirectory of `/net`.

NAME

boot – start the system kernel, or a standalone program

SYNOPSIS

```
>b [ device [ (c, u, p) ] ] [ filename ] [ -a ] boot-flags
>b?
>b!
```

DESCRIPTION

The boot program is started by the PROM monitor and loads the kernel, or another executable program, into memory.

The form **b?** displays all boot devices and their device arguments.

The form **b!** boots, but does not perform a RESET.

USAGE**Booting Standalone**

When booting standalone, the boot program (/boot) is brought in by the PROM from the file system. This program contains drivers for all devices.

Booting a Sun-3 System Over the Network

When booting over the network, the Sun-3 system PROM obtains a version of the boot program from a server using the Trivial File Transfer Protocol (TFTP). The client broadcasts a RARP request containing its Ethernet address. A server responds with the client's Internet address. The client then sends a TFTP request for its boot program to that server (or if that fails, it broadcasts the request). The filename requested (unqualified — not a pathname) is the hexadecimal, uppercase representation of the client's Internet address, for example:

```
Using IP Address    192.9.1.17 = C0090111
```

When the Sun server receives the request, it looks in the directory /tftpboot for *filename*. That file is typically a symbolic link to the client's boot program, normally **boot.sun3** in the same directory. The server invokes the TFTP server, **tftpd(8C)**, to transfer the file to the client.

When the file is successfully read in by the client, the boot program jumps to the load-point and loads **vmunix** (or a standalone program). In order to do this, the boot program makes a broadcast RARP request to find the client's IP address, and then makes a second broadcast request to a **bootparamd(8)** bootparams daemon, for information necessary to boot the client. The bootparams daemon obtains this information either from a local /etc/bootparams database file, or from a Yellow Pages (YP) map. The boot program sends two requests to the bootparams daemon, the first, **whoami**, to obtain its hostname, and the second, **getfile**, to obtain the name of the client's server and the pathname of the client's root partition.

The boot program then performs a **mount(8)** operation to mount the client's root partition, after which it can read in and execute any program within that partition by pathname (including a symbolic link to another file within that same partition). Typically, it reads in the file /vmunix. If the program is not read in successfully, boot responds with a short diagnostic message.

Booting a Sun-2, Sun-4, or Sun386i System Over the Network

Sun-2, Sun-4 and Sun386i systems boot over the network in a similar fashion. However, the filename requested from a server must have a suffix that reflects the system architecture of the machine being booted. For these systems, the requested filename has the form:

```
ip-address.arch
```

where *ip-address* is the machine's Internet Protocol (IP) address in hex, and *arch* is a suffix representing its architecture. (Only Sun-3 systems may omit the *arch* suffix.) These filenames are restricted to 14 characters for compatibility with System V and other operating systems. Therefore, the architecture suffix is limited to 5 characters; it must be in upper case. At present, the following suffixes are recognized: **SUN2** for Sun-2 system, **SUN3** for Sun-3 system, **SUN4** for Sun-4 system, **S386** for Sun386i system, and **PCNFS** for PC-NFS.

Note: a Sun-2 system boots from its server using one extra step. It broadcasts an ND request which is intercepted by the user-level `ndbootd(8)` server. This server sends back a standalone program that carries out the same TFTP request sequence as is done for all the other systems.

System Startup

Once the system is loaded and running, the kernel performs some internal housekeeping, configures its device drivers, and allocates its internal tables and buffers. The kernel then starts process number 1 to run `init(8)`, which performs file system housekeeping, starts system daemons, initializes the system console, and begins multiuser operation. Some of these activities are omitted when `init` is invoked with certain *boot-flags*. These are typically entered as arguments to the boot command, and passed along by the kernel to `init`.

OPTIONS

<i>device</i>	One of:
	ie Intel Ethernet
	ec 3Com Ethernet
	le Lance Ethernet (Sun 3-50 system)
	sd SCSI disk
	st SCSI 1/4" tape
	mt Tape Master 9-track 1/2" tape
	xt Xylogics 1/2" tape
	xy Xylogics 440/450 disk
c	Controller number, 0 if there is only one controller for the indicated type of device.
u	Unit number, 0 if only there is only one driver.
<i>filename</i>	Name of a standalone program in the selected partition, such as <code>stand/diag</code> or <code>vmunix</code> . Note: <i>filename</i> is relative to the root of the selected device and partition. It never begins with '/' (backslash). If <i>filename</i> is not given, the boot program uses a default value (normally <code>vmunix</code>). This is stored in the <code>vmunix</code> variable in the <code>boot</code> executable file supplied by Sun, but can be patched to indicate another standalone program loaded using <code>adb(1)</code> .
-a	Prompt interactively for the device and name of the file to boot. For more information on how to boot from a specific device, refer to <i>Installing the SunOS</i> .
<i>boot-flags</i>	The boot program passes all <i>boot-flags</i> to the kernel or standalone program. They are typically arguments to that program or, as with those listed below, arguments to programs that it invokes.
	-b Pass the -b flag through the kernel to <code>init(8)</code> to skip execution of the <code>/etc/rc.local</code> script.
	-h Halt after loading the system.
	-s Pass the -s flag through the kernel to <code>init(8)</code> for single-user operation.
	-i <i>initname</i> Pass the -i <i>initname</i> to the kernel to tell it to run <i>initname</i> as the first program rather than the default <code>/single/init</code> .

FILES

<code>/boot</code>	standalone boot program
<code>/tftpboot/????????</code>	symbolic link to the boot program for a client
<code>/tftpboot/boot.sun3</code>	programs to boot from the client's root partition
<code>/usr/etc/in.tftpd</code>	TFTP server
<code>/usr/mdcc/installboot</code>	program to install boot blocks from a remote host
<code>/vmunix</code>	
<code>/usr/boot</code>	
<code>/etc/bootparams</code>	

SEE ALSO

adb(1), tftp(1), bootparamd(8), init(8), mount(8), ndbootd(8C), rc(8), reboot(8), tftpd(8C), kadb(8S), monitor(8S)

*Installing the SunOS
System and Network Administration*

BUGS

On the Sun-2 system, the PROM passes in the default name `vmunix`, overriding the the boot program's patchable default.

NAME

bootparamd – boot parameter server

SYNOPSIS

`/usr/etc/rpc.bootparamd [-d]`

DESCRIPTION

bootparamd is a server process that provides information to diskless clients necessary for booting. It consults either the **bootparams** database or the `/etc/bootparams` file if the Yellow Pages service is not running.

bootparamd can be invoked either by **inetd(8C)** or by the user.

OPTIONS

-d Display the debugging information.

FILES

`/etc/bootparams`

SEE ALSO

inetd(8C)

NAME

C2conv, C2unconv – convert system to or from C2 security

SYNOPSIS

C2conv

C2unconv

AVAILABILITY

This program is available with the *Security* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

C2conv converts a standard SunOS system to operate with C2-level security.

The program prompts for information regarding installation base, client systems (if the system is a SunDisk server), audit devices (if it is an audit file server), and names of file systems (if it is a remote audit server). The program also requests certain information for the `audit_control(5)` file; default values may be used for audit flags and for the "minfree" value. Finally, it requests the user ID of person (or list of persons) to notify (by `mail(1)`) when C2 administrative tasks are required. The default ID is `root` for the host being converted.

Once it has this information, **C2conv** uses it to set up the necessary files for a C2 secure system, reporting on its progress as it proceeds.

C2unconv backs out the changes made to `/etc/passwd` and `/etc/group`. It does not back out changes to other files.

FILES

`/etc/passwd`

`/etc/group`

`/etc/fstab`

SEE ALSO

`audit_control(5)`

NAME

captoinfo – convert a termcap description into a terminfo description

SYNOPSIS

captoinfo [*-v ...*] [*-V*] [*-1*] [*-w width*] *filename ...*

DESCRIPTION

captoinfo converts the **termcap(5)** the terminal description entries given in *filename* into **terminfo(5V)** source entries, and writes them to the standard output along with any comments found in that file. A description that is expressed as relative to another description (as specified in the **termcap** *tc=* capability) is reduced to the minimum superset before being written.

If no *filename* is given, then the environment variable **TERMCAP** is used for the filename or entry. If **TERMCAP** is a full pathname to a file, only the terminal-name is specified in the environment variable **TERM** is extracted from that file. If that environment variable is not set, then the file */etc/termcap* is read.

OPTIONS

- v** Verbose. Print tracing information on the standard error as the program runs. Additional **-v** options increase the level of detail.
- V** Version. Display the version of the program on the standard error and exit.
- 1** Print fields one-per-line. Otherwise, fields are printed several to a line, to a maximum width of 60 characters.
- w width**
Change the output to *width* characters.

CAVEATS

Certain **termcap** defaults are assumed to be true. The bell character (**terminfo** *bel*) is assumed to be **^G**. The linefeed capability (**termcap** *nl*) is assumed to be the same for both **cursor_down** and **scroll_forward** (**terminfo** *cu***d1** and *ind*, respectively.) Padding information is assumed to belong at the end of the string.

The algorithm used to expand parameterized information for **termcap** fields such as **cursor_position** (**termcap** *cm*, **terminfo** *cup*) can sometimes produce a string that may not be optimal. In particular, the rarely used **termcap** operation *%n* produces strings that are especially long. Most occurrences of these non-optimal strings will be flagged with a warning message and may need to be recoded by hand.

The short two-letter name at the beginning of the list of names in a **termcap** entry, a hold-over from an earlier version of the system, has been removed.

FILES

*/usr/share/lib/terminfo/?/**
compiled terminal description database

/etc/termcap

SEE ALSO

curses(3V), **termcap(5)**, **terminfo(5V)**, **infocmp(8V)**, **tic(8V)**

DIAGNOSTICS

tgetent failed with return code n

The **termcap** entry is not valid. In particular, check for an invalid **'tc='** entry.

unknown type given for the termcap code cc.

The **termcap** description had an entry for *cc* whose type was not boolean, numeric or string.

wrong type given for the boolean (numeric, string) termcap code cc.

The boolean **termcap** entry *cc* was entered as a numeric or string capability.

the boolean (numeric, string) termcap code cc is not a valid name.

An unknown **termcap** code was specified.

tgetent failed on TERM=term.

The terminal type specified could not be found in the **termcap** file.

TERM=term: cap *cc* (info *ii*) is

The termcap code was specified as a null string. The correct way to cancel an entry is with an '@', as in ':bs@:'. Giving a null string could cause incorrect assumptions to be made by the software which uses termcap or terminfo.

a function key for *cc* was specified, but it already has the value

vv. When parsing the *ko* capability, the key *cc* was specified as having the same value as the capability *cc*, but the key *cc* already had a value assigned to it.

the unknown termcap name *cc* was specified in the *ko* termcap capability.

A key was specified in the *ko* capability which could not be handled.

the *vi* character *v* (info *ii*) has the value *xx*, but *ma* gives *n*.

The *ma* capability specified a function key with a value different from that specified in another setting of the same key.

the unknown *vi* key *v* was specified in the *ma* termcap capability.

A *vi*(1) key unknown to captainfo was specified in the *ma* capability.

Warning: termcap *sg* (*nn*) and termcap *ug* (*nn*) had different values.

terminfo assumes that the *sg* (now *xmc*) and *ug* values were the same.

Warning: the string produced for *ii* may be inefficient.

The parameterized string being created should be rewritten by hand.

Null termname given.

The terminal type was null. This is given if the environment variable **TERM** is not set or is null.

cannot open *filename* for reading.

The specified file could not be opened.

NAME

catman – create the cat files for the manual

SYNOPSIS

/usr/etc/catman [**-nptw**] [**-M directory**] [**-T tmac.an**] [*sections*]

DESCRIPTION

catman creates the preformatted versions of the on-line manual from the **nroff(1)** input files. Each manual page is examined and those whose preformatted versions are missing or out of date are recreated. If any changes are made, **catman** recreates the **whatis** database.

If there is one parameter not starting with a '-', it is taken to be a list of manual sections to look in. For example

```
catman 123
```

only updates manual sections 1, 2, and 3.

If an unformatted source file contains only a line of the form '**.so manx/yyy.x**', a symbolic link is made in the **catx** or **fmtx** directory to the appropriate preformatted manual page. This feature allows easy distribution of the preformatted manual pages among a group of associated machines with **rdist(1)**, since it makes the directories of preformatted manual pages self-contained and independent of the unformatted entries.

OPTIONS

- n** Do not (re)create the **whatis** database.
- p** Print what would be done instead of doing it.
- t** Create **troffed** entries in the appropriate **fmt** subdirectories instead of **nroffing** into the **cat** subdirectories.
- w** Only create the **whatis** database. No manual reformatting is done.
- M** Update manual pages located in the specified directory (**/usr/share/man** by default).
- T** Use **tmac.an** in place of the standard manual page macros.

ENVIRONMENT

TROFF The name of the formatter to use when the **-t** flag is given. If not set, '**troff**' is used.

FILES

/usr/share/man	default manual directory location
/usr/share/man/man?/*.*	raw (nroff input) manual sections
/usr/share/man/cat?/*.*	preformatted nroffed manual pages
/usr/share/man/fmt?/*.*	preformatted troffed manual pages
/usr/share/man/whatis	whatis database location
/usr/lib/makewhatis	command script to make whatis database

SEE ALSO

man(1), **nroff(1)**, **rdist(1)**, **troff(1)**, **whatis(1)**

DIAGNOSTICS

man?/xxx? (.so'ed from **man?/yyy?**): No such file or directory

The file outside the parentheses is missing, and is referred to by the file inside them.

target of .so in man?/xxx? must be relative to **/usr/man**

catman only allows references to filenames that are relative to the directory **/usr/share/man**.

opendir:man?: No such file or directory

A harmless warning message indicating that one of the directories **catman** normally looks for is missing.

***.*: No such file or directory**

A harmless warning message indicating **catman** came across an empty directory.

NAME

chown – change owner

SYNOPSIS

`/usr/etc/chown [-fR] owner[.group] filename ...`

DESCRIPTION

chown changes the owner of the *filenames* to *owner*. The owner may be either a decimal user ID or a login name found in the password file. An optional *group* may also be specified. The group may be either a decimal group ID or a group name found in the GID file.

Only the super-user can change owner, in order to simplify accounting procedures.

OPTIONS

- f** Do not report errors.
- R** Recursively descend into directories setting the ownership of all files in each directory encountered. When symbolic links are encountered, their ownership is changed, but they are not traversed.

FILES

`/etc/passwd` password file

SEE ALSO

`chgrp(1)`, `chown(2)`, `group(5)`, `passwd(5)`

NAME

chroot – change root directory for a command

SYNOPSIS

/usr/etc/chroot newroot command

DESCRIPTION

The given command is executed *relative* to the new root. The meaning of any initial '/' (slashes) in path names is changed for a command and any of its children to *newroot*. Furthermore, the initial working directory is *newroot*.

Input and output redirections on the command line are made with respect to the *original* root:

chroot newroot command >x

creates the file *x* relative to the original root, not the new one.

This command is restricted to the super-user.

The new root path name is always relative to the current root: even if a **chroot** is already in effect; the *newroot* argument is relative to the current root of the running process.

SEE ALSO

chdir(2)

BUGS

One should exercise extreme caution when referring to special files in the new root file system.

NAME

client – add or remove diskless Sun386i systems

SYNOPSIS

client [**-a arch**] [**-h hostid**] [**-o os**] [**-q**] [**-t minutes**] **add bootserver client etheraddress ipaddress**
client remove client
client modify client [diskful | diskless | slave]

AVAILABILITY

Sun386i systems only.

DESCRIPTION

client can be used to manually add and remove diskless clients of a PNP boot server. After successful completion of the command, the diskless client can boot. Only users in the *networks* group (group 12) on the boot server are allowed to change configurations using this utility. **client** can be invoked from any system on the network.

The boot server of a system is the only machine truly required for that system to boot to the point of allowing user logins; it must accordingly provide name, booting, and time services. Diskless clients can provide none of these services themselves. Diskful clients, however, can provide most of their own boot services. Network clients only need name and time services from the network, and can use any boot server.

To add a diskless client, use the **add** operation. To remove a diskless, diskful, or network client, use the **remove** operation. To change a system's network role, use the **modify** operation.

A server can reject a configuration request if it is disallowed by the contents of the *bootserver*s map (e.g., too many clients would be configured, or too little free space would be left on the server), or if no system software for the client is available.

OPTIONS

- a arch** Specifies the architecture code of the client; it defaults to *s386*. (Note: architecture codes are different from architecture names. Architecture codes are used in diskless booting, and are at most five characters in length, while architecture names can be longer.)
- h hostid** Specifies the host ID of the client; if supplied, it is used as the root password for the system. It defaults to the null string.
- o os** Specifies the operating system; defaults to 'unix'. This is currently used only to construct the system's *publickey* data, where applicable; this is never done if the system has no *hostid* specified.
- q** Quiet. Displays only error messages.
- t minutes** Sets the RPC timeout to the number of minutes indicated; this defaults to 15 minutes. If the bootserver takes more time than this to complete, **client** will exit. Unless the server has already completed setup, but not yet sent status to **client**, this will cause the bootserver to back out of the setup, deallocating all assigned resources.

SEE ALSO

pnpd(8C), **netconfig(8C)**, **publickey(5)**

BUGS

Unless the *hostid* is assigned, the root filesystem for the diskless client is not set up beyond copying the *proto* and *boot* files into it. This means that **netconfig** will often handle other parts of the setup.

NAME

clri – clear inode

SYNOPSIS

/usr/etc/clri filesystem i-number...

DESCRIPTION

Note: **clri** has been superceded for normal file system repair work by **fsck(8)**.

clri writes zeros on the inodes 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 inode 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 inode which does appear in a directory, care should be taken to track down the entry and remove it. Otherwise, when the inode 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 inode, so the whole cycle is likely to be repeated again and again.

SEE ALSO

icheck(8) **fsck(8)**

BUGS

If the file is open, **clri** is likely to be ineffective.

NAME

comsat – biff server

SYNOPSIS

/usr/etc/in.comsat

DESCRIPTION

comsat is the server process which listens for reports of incoming mail and notifies users who have requested to be told when mail arrives. It is invoked as needed by **inetd(8C)**, and times out if inactive for a few minutes.

comsat listens on a datagram port associated with the **biff(1)** service specification (see **services(5)**) for one line messages of the form

user@mailbox-offset

If the *user* specified is logged in to the system and the associated terminal has the owner execute bit turned on (by a 'biff y'), the *offset* is used as a seek offset into the appropriate mailbox file and the first 7 lines or 560 characters of the message are printed on the user's terminal. Lines which appear to be part of the message header other than the From, To, Date, or Subject lines are not printed when displaying the message.

FILES

/etc/utmp to find out who's logged on and on what terminals

SEE ALSO

biff(1), **services(5)**, **inetd(8C)**

BUGS

The message header filtering is prone to error.

The notification should appear in a separate window so it does not mess up the screen.

NAME

`config` – build system configuration files

SYNOPSIS

`/etc/config [-fgnp] [-o obj_dir] config_file`

DESCRIPTION

`config` does the preparation necessary for building a new system kernel with `make(1)`. The *config_file* named on the command line describes the kernel to be made in terms of options you want in your system, size of tables, and device drivers to be included. When you run `config`, it uses several input files located in the current directory (typically the `conf` subdirectory of the system source including your *config_file*). The format of this file is described below.

If the directory named `./config_file` does not exist, `config` will create one. One of `config`'s output files is a makefile which you use with `make(1)` to build your system.

You use `config` as follows. Run `config` from the `conf` subdirectory of the system source (in a typical Sun environment, from `/usr/include/sys/conf`):

```
example# /etc/config config_file
Doing a "make depend"
example# cd ../config_file
example# make
... lots of output...
```

While `config` is running watch for any errors. Never use a kernel which `config` has complained about; the results are unpredictable. If `config` completes successfully, you can change directory to the `./config_file` directory, where it has placed the new makefile, and use `make` to build a kernel. The output files placed in this directory include `ioconf.c`, which contains a description of I/O devices attached to the system; `mbglue.s`, which contains short assembly language routines used for vectored interrupts, a makefile, which is used by `make` to build the system; a set of header files (*device_name.h*) which contain the number of various devices that may be compiled into the system; and a set of swap configuration files which contain definitions for the disk areas to be used for the root file system, swapping, and system dumps.

Now you can install your new kernel and try it out.

OPTIONS

- `-f` Set up the makefile for fast builds. This is done by building a `vmunix.o` file which includes all the `.o` files which have no source. This reduces the number of files which have to be `stated` during a system build. This is done by prelinking all the files for which no source exists into another file which is then linked in place of all these files when the kernel is made. This makefile is faster because it does not `stat` the object files during the build.
- `-g` Get the current version of a missing source file from its SCCS history, if possible.
- `-n` Do not do the 'make depend'. Normally `config` will do the 'make depend' automatically. If this option is used `config` will print 'Don't forget to do a "make depend"' before completing as a reminder.
- `-p` Configure the system for profiling (see `kgmon(8)` and `gprof(1)`).
- `-o obj_dir`
Use `./obj_dir` instead of `./OBJ` as the directory to find the object files when the corresponding source file is not present in order to generate the files necessary to compile and link your kernel.

USAGE

Input Grammar

In the following descriptions, a number can be a decimal integer, a whole octal number or a whole hexadecimal number. Hex and octal numbers are specified to `config` in the same way they are specified to the C compiler, a number starting with `0x` is a hex number and a number starting with just a `0` is an octal number.

Comments are begin with a # character, and end at the next NEWLINE. Lines beginning with TAB characters are considered continuations of the previous line. Lines of the configuration file can be one of two basic types. First, there are lines which describe general things about your system:

machine "type"

This system is to run on the machine type specified. Only one machine type can appear in the config file. The legal *types* for a Sun system are **sun2**, **sun3**, **sun4**, and **sun386**. Note: the double quotes around *type* are part of the syntax, and must be included.

cpu "type"

This system is to run on the cpu type specified. More than one cpu type can appear in the config file. Legal *types* for a sun2 machine are noted in the annotated config file in *Installing the SunOS*.

ident name

Give the system identifier — a name for the machine or machines that run this kernel. Note that *name* must be enclosed in double quotes if it contains both letters and digits. Also, note that if *name* is **GENERIC**, you need not include the 'options **GENERIC**' clause in order to specify 'swap generic'.

maxusers number

The maximum expected number of simultaneously active user on this system is *number*. This number is used to size several system data structures.

options optlist

Compile the listed options into the system. Options in this list are separated by commas. A line of the form:

options FUNNY, HAHA

yields

-DFUNNY -DHAHA

to the C compiler. An option may be given a value, by following its name with = (equal sign) then the value enclosed in (double) quotes. None of the standard options use such a value.

In addition, options can be used to bring in additional files if the option is listed in the **files** files. All options should be listed in upper case. In this case, no corresponding *option.h* will be created as it would be using the corresponding *pseudo-device* method.

config sysname config_clauses...

Generate a system with name *sysname* and configuration as specified in *config-clauses*. The *sysname* is used to name the resultant binary image and per-system swap configuration files. The *config_clauses* indicate the location for the root file system, one or more disk partitions for swapping and paging, and a disk partition to which system dumps should be made. All but the root device specification may be omitted; **config** will assign default values as described below.

root A root device specification is of the form 'root on *xy0d*'. If a specific partition is omitted — for example, if only root on *xy0* is specified — the 'a' partition is assumed. When a generic system is being built, no root specification should be given; the root device will be defined at boot time by prompting the console.

swap To specify a swap partition, use a clause of the form: 'swap on *partition*'. Swapping areas may be almost any size. Partitions used for swapping are sized at boot time by the system; to override dynamic sizing of a swap area the number of sectors in the swap area can be specified in the config file. For example, 'swap on *xy0b* size 99999' would configure a swap partition with 99999 sectors. If **swap generic** or no *partition* is specified with **on**, partition *b* on the root device is used. For dataless clients, use 'swap on type *nfs*'.

dumps The location to which system dumps are sent may be specified with a clause of the form 'dumps on *xyz*'. If no dump device is specified, the first swap partition specified is used. If a device is specified without a particular partition, the 'b' partition is assumed. If a generic configuration is to be built, no dump device should be specified; the dump device will be assigned to the swap device dynamically configured at boot time. Dumps are placed at the end of the partition specified. Their size and location is recorded in global kernel variables *dumpsiz*e and *dumpl*o, respectively, for use by *savecore*(8).

Device names specified in configuration clauses are mapped to block device major numbers with the *devices.machine*, where *machine* is the machine type previously specified in the configuration file. If a device name to block device major number mapping must be overridden, a device specification may be given in the form 'major *x* minor *y*'.

The second group of lines in the configuration file describe which devices your system has and what they are connected to (for example, a Xylogics 450 Disk Controller at address 0xee40 in the Multibus I/O space). These lines have the following format:

```
dev_type dev_name at con_dev more_info
```

dev_type is either *controller*, *disk*, *tape*, *device*, or *pseudo-device*. These types have the following meanings:

controller	A disk or tape controller.
disk or tape	Devices connected to a controller.
device	Something "attached" to the main system bus, like a cartridge tape interface.
pseudo-device	A software subsystem or driver treated like a device driver, but without any associated hardware. Current examples are the pseudo-tty driver and various network subsystems. For pseudo-devices, <i>more_info</i> may be specified as an integer, that gives the value of the symbol defined in the header file created for that device, and is generally used to indicate the number of instances of the pseudo-device to create.

dev_name is the standard device name and unit number (if the device is not a *pseudo-device*) of the device you are specifying. For example, *xyz0* is the *dev_name* for the first Xylogics controller in a system; *ar0* names the first quarter-inch tape controller.

con_dev is what the device you are specifying is connected to. It is either *nexus?*, a bus type, or a controller. There are several bus types which are used by *config* and the kernel.

The different possible bus types are:

obmem	On board memory
obio	On board io
mbmem	Multibus memory (<i>sun2</i> system only)
mbio	Multibus io (<i>sun2</i> system only)
vme16d16 (vme16)	16 bit VMEbus/ 16 bit data
vme24d16 (vme24)	24 bit VMEbus/ 16 bit data
vme32d16	32 bit VMEbus/ 16 bit data (<i>sun3</i> system only)
vme16d32	16 bit VMEbus/ 32 bit data (<i>sun3</i> system only)
vme24d32	24 bit VMEbus/ 32 bit data (<i>sun3</i> system only)
vme32d32 (vme32)	32 bit VMEbus/ 32 bit data (<i>sun3</i> system only)

All of these bus types are declared to be connected to *nexus*. The devices are hung off these buses. If the bus is wildcarded, then the autoconfiguration code will determine if it is appropriate to probe for the device on the machine that it is running on. If the bus is numbered, then the autoconfiguration code will only look for that device on machine type *N*. In general, the Multibus and VMEbus bus types are always wildcarded.

more_info is a sequence of the following:

csr address	Specify the address of the csr (command and status registers) for a device. The csr addresses specified for the device are the addresses within the bus type specified. The csr address must be specified for all controllers, and for all devices connected to a main system bus.
drive number	For a disk or tape, specify which drive this is.
flags number	These flags are made available to the device driver, and are usually read at system initialization time.
priority level	For devices which interrupt, specify the interrupt level at which the device operates.
vector intr number [intr number . . .]	For devices which use vectored interrupts on VMEbus systems, <i>intr</i> specify the vectored interrupt routine and <i>number</i> the corresponding vector to be used (0x40-0xFF).

A ? may be substituted for a number in two places and the system will figure out what to fill in for the ? when it boots. You can put question marks on a *con_dev* (for example, at virtual '?'), or on a drive number (for example, drive '?'). This allows redundancy, as a single system can be built which will boot on different hardware configurations.

The easiest way to understand config files is to look at a working one and modify it to suit your system. Good examples are provided in *Installing the SunOS*.

FILES

Files in */usr/include/sys/conf* which may be useful for developing the *config_file* used by config are:

GENERIC	These are generic configuration files for either a Sun-2 or Sun-3 system. They contain all possible device descriptions lines for the particular architecture.
README	File describing how to make a new kernel.

As shipped from Sun, the files used by */etc/config* as input are in the */usr/include/sys/conf* directory:

<i>config_file</i>	System-specific configuration file
makefile.sun[23]	Generic prototype makefile for Sun-[23] systems
files	List of common files required to build a basic kernel
files.sun[23]	List of files for a Sun-[23] specific kernel
devices.sun[23]	Name to major device mapping file for Sun-[23] systems

/etc/config places its output files in the *./config_file* directory:

mbglue.s	Short assembly language routines used for vectored interrupts
ioconf.c	Describes I/O devices attached to the system
makefile	Used with make(1) to build the system
device_name.h	a set of header files (various <i>device_name</i> 's) containing devices which can be compiled into the system

SEE ALSO

gprof(1), **make(1)**, **kgmon(8)**, **savecore(8)**

The SYNOPSIS portion of each device entry in Section 4 of this manual.

Installing the SunOS
System and Network Administration

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 displays 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)`. Unless some unexpected inconsistency is encountered in the state of the file systems due to hardware or software failure, the system will then resume multiuser operations.

The system has a large number of internal consistency checks; if one of these fails, it will panic with a very short message indicating which one failed.

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, you can run the program `savecore(8)` to preserve a copy of this core image and kernel namelist 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, as described in *Debugging Tools*

The most common cause of system failures is hardware failure, which can reflect itself in different ways.

Here are some messages that you may encounter, with some hints as to causes. In each case there is a possibility that a hardware or software error produced the message in some unexpected way.

FILES

<code>/vmunix</code>	the system kernel
<code>/etc/rc.local</code>	script run when the local system starts up

SEE ALSO

`adb(1)`, `analyze(8)`, `reboot(8)` `sa(8)`, `savecore(8)`

Debugging Tools

DIAGNOSTICS**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 should not be a panic, but until the data structure is fixed, involved, running out of entries causes a crash. If this happens, you should make the timeout table bigger by changing the value of `nccallout` in the `param.c` file, and then rebuild your system.

trap type type, pid process-id, pc = program-counter, sr = status-register, context context-number

A unexpected trap has occurred within the system; typical trap types are:

- Bus error
- Address error
- Illegal instruction
- Divide by zero
- Chk instruction
- Trapv instruction
- Privilege violation
- Trace
- 1010 emulator trap
- 1111 emulator trap
- Stack format error

- Uninitialized interrupt
- Spurious interrupt

The favorite trap types in system crashes are “Bus error” or “Address error”, indicating a wild reference. The *process-id* is the ID of the process running at the time of the fault, *program-counter* is the hexadecimal value of the program counter, *status-register* is the hexadecimal value of the status register, and *context-number* is the context that the process was running in. 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.

NAME

cron – clock daemon

SYNOPSIS

/usr/etc/cron

DESCRIPTION

cron executes commands at specified dates and times. Regularly scheduled commands can be specified according to instructions found in **crontab** files in the directory **/var/spool/cron/crontabs**. Users can submit their own **crontab** files using the **crontab(1)** command. Commands that are to be executed only once may be submitted using the **at(1)** command.

cron only examines **crontab** files and **at** command files during process initialization and when a file changes using **crontab** or **at**. This reduces the overhead of checking for new or changed files at regularly scheduled intervals.

Since **cron** never exits, it should only be executed once. This is normally done by running **cron** from the initialization process through the file **/etc/rc**; see **init(8)**. **/var/spool/cron/FIFO** is a FIFO file that **crontab** and **at** use to communicate with **cron**; it is also used as a lock file to prevent the execution of more than one **cron**.

FILES

/var/spool/cron main cron directory
/var/spool/cron/FIFO FIFO for sending messages to cron
/var/spool/cron/crontabs
 directory containing crontab files

SEE ALSO

at(1), **crontab(1)**, **sh(1)**, **init(8)**, **syslogd(8)**

DIAGNOSTICS

cron logs various errors to the system log daemon, **syslogd(8)**, with a facility code of **cron**. The messages are listed here, grouped by severity level.

Err Severity

Can't create /var/spool/cron/FIFO: reason

cron was unable to start up because it could not create **/var/spool/cron/FIFO**.

Can't access /var/spool/cron/FIFO: reason

cron was unable to start up because it could not access **/var/spool/cron/FIFO**.

Can't open /var/spool/cron/FIFO: reason

cron was unable to start up because it could not open **/var/spool/cron/FIFO**.

Can't start cron - another cron may be running (/var/spool/cron/FIFO exists)

cron found that **/var/spool/cron/FIFO** already existed when it was started; this normally means that **cron** had already been started, but it may mean that an earlier **cron** terminated abnormally without removing **/var/spool/cron/FIFO**.

Can't stat /var/spool/cron/FIFO: reason

cron could not get the status of **/var/spool/cron/FIFO**.

Can't change directory to directory:reason

cron could not change to the directory **directory**.

Can't read directory:reason

cron could not read the directory **directory**.

error reading message: reason

An error occurred when **cron** tried to read a control message from **/var/spool/cron/FIFO**.

message received — bad format

A message was successfully read by **cron** from `/var/spool/cron/FIFO`, but the message was not of a form recognized by **cron**.

SIGTERM

received **cron** was told to terminate by having a SIGTERM signal sent to it.

cron could not unlink /var/spool/cron/FIFO: reason

cron was told to terminate, but it was unable to unlink `/var/spool/cron/FIFO` before it terminated.

******* CRON ABORTED *******

cron terminated, either due to an error or because it was told to.

Can't open queuedefs file file:reason

cron could not open a *queuedefs* file.

I/O error reading queuedefs file file:reason

An I/O error occurred while **cron** was reading a *queuedefs* file.

Using default queue definitions

An error occurred while trying to read a *queuedefs* file; the default queue definitions will be used.

Can't allocate numberbytes of space

An internal error occurred in **cron** while trying to allocate memory.

Info Severity**queue queue max run limit reached**

There were more jobs running or to be run in the queue *queue* than the maximum number specified. **cron** will wait until one of the currently-running jobs completes before starting to run a new one.

MAXRUN (25) procs reached

There were more than 25 jobs running or to be run by **cron**. **cron** will wait until one of the currently-running jobs completes before starting to run a new one.

***** cron started *****

cron started running.

> CMD: command

A **cron** job was started. *Command* is the command to be run.

> user pid queue time

A **cron** job was started for user *user*, in queue *queue*, with process ID *pid*, at the date and time *time*.

< user pid queue time status

A **cron** job completed for user *user*, in queue *queue*, with process ID *pid*, at the date and time *time*. If the command terminated with a non-zero exit status or a signal, *status* indicates the exit status or signal.

Notice Severity**Can't fork**

An attempt to fork (2) to run a new job failed; **cron** will attempt again after a 30-second delay.

Warning Severity**Can't stat queuedefs file file:reason**

cron could not get the status of a *queuedefs* file in order to determine whether it has changed. **cron** will assume it has changed and will reread it.

NAME

dcheck – file system directory consistency check

SYNOPSIS

/usr/etc/dcheck [*-i numbers*] [*filesystem*]

DESCRIPTION

Note: **dcheck** has been superceded for normal consistency checking by **fsck(8)**.

dcheck reads the directories in a file system and compares the link-count in each inode with the number of directory entries by which it is referenced. If the file system is not specified, **dcheck** checks a set of default file systems.

dcheck is fastest if the raw version of the special file is used, since the i-list is read in large chunks.

OPTIONS

-i numbers

numbers is 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.

FILES

Default file systems vary with installation.

SEE ALSO

fs(5), **fsck(8)**, **clri(8)**, **icheck(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.

Inode numbers less than 2 are invalid.

NAME

devnm – device name

SYNOPSIS

/etc/devnm [*name*]...

DESCRIPTION

devnm identifies the special file associated with the mounted file system where each *name* argument resides. This command can be used to construct a mount table entry for the root file system.

EXAMPLE

If **/usr** is mounted on **/dev/dsk/c1d0s2**, then the command:

/etc/devnm /usr

produces:

/dev/dsk/c1d0s2 usr

FILES

/dev/dsk/*

/etc/mtab

SEE ALSO

mount(8), fstab(5)

NAME

dkinfo – report information about a disk's geometry and partitioning

SYNOPSIS

/usr/etc/dkinfo *disk* [*partition*]

DESCRIPTION

dkinfo gives the total number of cylinders, heads, and sectors or tracks on the specified *disk*, and gives this information along with the starting cylinder for the specified *partition*. If no *partition* is specified on the command line, **dkinfo** reports on all partitions.

The *disk* specification here is a disk name of the form *xxn*, where *xx* is the controller device abbreviation (ip, xy, etc.) and *n* is the disk number. The *partition* specification is simply the letter used to identify that partition in the standard UNIX system nomenclature. For example, **'/usr/etc/dkinfo xy0'** reports on the first disk in a system controlled by a Xylogics controller; **'/usr/etc/dkinfo xy0g'** reports on the seventh partition of such a disk.

EXAMPLE

A request for information on my local disk, an 84 MByte disk controlled by a Xylogics 450 controller, might look like this:

```
#/usr/etc/dkinfo xy0
xy0: Xylogics 450 controller at addr ee40, unit # 0
586 cylinders 7 heads 32 sectors/track
a: 15884 sectors (70 cyls, 6 tracks, 12 sectors)
starting cylinder 0
b: 33440 sectors (149 cyls, 2 tracks)
starting cylinder 71
c: 131264 sectors (586 cyls)
starting cylinder 0
d: No such device or address
e: No such device or address
f: No such device or address
g: 81760 sectors (365 cyls)
starting cylinder 221
h: No such device or address
#
```

FILES

/dev/rxxnp

SEE ALSO

dkio(4S), **format(8)**

NAME

dmesg – collect system diagnostic messages to form error log

SYNOPSIS

/usr/etc/dmesg [-]

DESCRIPTION

Note: **dmesg** is obsoleted by **syslogd(8)** for maintenance of the system error log.

dmesg looks in a system buffer for recently printed diagnostic messages and prints them on the standard output. The messages are those printed or logged by the system when errors occur. 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.

FILES

/var/adm/msgbuf scratch file for memory of '-' option

SEE ALSO

syslogd(8)

NAME

dump, rdump – incremental file system dump

SYNOPSIS

/usr/etc/dump [*options* [*arguments*]] *filesystem*

DESCRIPTION

dump backs up all files in *filesystem*, or files changed after a certain date, to magnetic tape; on Sun386i systems, **dump** works on both magnetic tape and diskettes. *options* is a string that specifies **dump** options, as shown below. Any *arguments* supplied for specific options are given as subsequent words on the command line, in the same order as that of the *options* listed.

If no *options* are given, the default is **9u**.

OPTIONS

- 0–9** The “dump level.” All files in the *filesystem* that have been modified since the last **dump** at a lower dump level are copied to the volume. For instance, if you did a “level 2” dump on Monday, followed by a “level 4” dump on Tuesday, a subsequent “level 3” dump on Wednesday would contain all files modified or added since the “level 2” (Monday) backup. A “level 0” dump copies the entire filesystem to the dump volume.
- b factor** Blocking factor. Specify the blocking factor for tape writes. The default is 20 blocks per write. Note: the blocking factor is specified in terms of 512 bytes blocks, for compatibility with **tar(1)**. The default blocking factor for tapes of density 6250BPI and greater is 64. The default blocking factor for cartridge tapes (c option specified) is 126. The highest blocking factor available with most tape drives is 126.
- c** Cartridge. Use a cartridge instead of the standard half-inch reel. This sets the density to 1000BPI and the blocking factor to 126. The length is set to 425 feet. (This option is incompatible with the **d** option, unless you specify a density of 1000BPI with that option).
- d bpi** Tape density. The density of the tape, expressed in BPI, is taken from *bpi*. This is used to keep a running tab on the amount of tape used per reel. The default density is 1600 except for cartridge tape. Unless a higher density is specified explicitly, **dump** uses its default density — even if the tape drive is capable of higher-density operation (for instance, 6250BPI). Note: the density specified should correspond to the density of the tape device being used, or **dump** will not be able to handle end-of-tape properly. The **d** option is not compatible with the **D** option.
- D** Diskette. Specify diskette as the dump media.
- f dump-file** Dump file. Use *dump-file* as the file to dump to, instead of */dev/rmt8*. If *dump-file* is specified as ‘-’, dump to the standard output. If the filename argument is of the form *machine:device*, dump to a remote machine. Since **dump** is normally run by *root*, the name of the local machine must appear in the *.rhosts* file of the remote machine. If the filename argument is of the form *user@machine:device*, **dump** will attempt to execute as the specified user on the remote machine. The specified user must have a *.rhosts* file on the remote machine that allows root from the local machine. If **dump** is called as **rdump**, the dump device defaults to *dumphost:/dev/rmt8*. To direct the output to a desired remote machine, set up an alias for *dumphost* in the file */etc/hosts*.
- n** Notify. When this option is specified, if **dump** requires attention, it sends a terminal message (similar to **wall(1)**) to all operators in the “operator” group.
- s size** Specify the *size* of the volume being dumped to. When the specified size is reached, **dump** waits for you to change the volume. **dump** interprets the specified size as the length in feet for tapes, and cartridges and as the number of 1024 byte blocks for diskettes. The following are defaults:
- | | |
|-----------|--|
| tape | 2300 feet |
| cartridge | 425 feet |
| diskette | 1422 blocks (Corresponds to a 1.44 Mb diskette, with one cylinder reserved for bad block information.) |

- t tracks** Specify the number of tracks for a cartridge tape. On all Sun-2 systems the default is 4 tracks, although some Sun-2 systems have 9 track drives. On all other machines the default is 9 tracks. The **t** option is not compatible with the **D** option.
- u** Update the dump record. Add an entry to the file `/etc/dumpdates`, for each filesystem successfully dumped that includes the filesystem name, date, and dump level. This file can be edited by the super-user.
- w** List the filesystems that need backing up. This information is gleaned from the files `/etc/dumpdates` and `/etc/fstab`. When the **w** option is used, all other options are ignored. After reporting, **dump** exits immediately.
- W** Like **w**, but includes all filesystems that appear in `/etc/dumpdates`, along with information about their most recent dump dates and levels. Filesystems that need backing up are highlighted.

FILES

<code>/dev/rmt8</code>	default unit to dump to
<code>dumphost:/dev/rmt8</code>	default remote unit to dump to if called as <code>rdump</code>
<code>/etc/dumpdates</code>	dump date record
<code>/etc/fstab</code>	dump table: file systems and frequency
<code>/etc/group</code>	to find group <i>operator</i>
<code>/etc/hosts</code>	

SEE ALSO

`tar(1)`, `wall(1)`, `dump(5)`, `fstab(5)`, `restore(8)`, `shutdown(8)`

DIAGNOSTICS

While running, **dump** emits many verbose messages.

Exit Codes

- 0 Normal exit.
- 1 Startup errors encountered.
- 3 Abort – no checkpoint attempted.

BUGS

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 is recommended that incremental dumps also be performed with the system running in single-user mode.

dump does not support multi-file multi-volume tapes.

NOTES

Operator Intervention

dump requires operator intervention on these conditions: end of volume, end of dump, volume write error, volume 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** option, **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**, as appropriate.

Since backing up a disk can involve a lot of time and effort, **dump** checkpoints at the start of each volume. If writing that volume fails for some reason, **dump** will, with operator permission, restart itself from the checkpoint after a defective volume has been replaced.

dump reports periodically, and in verbose fashion. Each report includes estimates of the percentage of the dump completed and how long it will take to complete the dump.

Suggested Dump Schedule

It is vital to perform full, "level 0", dumps at regular intervals. When performing a full dump, bring the machine down to single-user mode using `shutdown(8)`. While preparing for a full dump, it is a good idea to clean the tape drive and heads.

Incremental dumps allow for convenient backup and recovery on a more frequent basis of active files, with a minimum of media and time. However there are some tradeoffs. First, the interval between backups should be kept to a minimum (once a day at least). To guard against data loss as a result of a media failure (a rare, but possible occurrence), it is a good idea to capture active files on (at least) two sets of dump volumes. Another consideration is the desire to keep unnecessary duplication of files to a minimum to save both operator time and media storage. A third consideration is the ease with which a particular backed-up version of a file can be located and restored. The following four-week schedule offers a reasonable trade-off between these goals.

	<i>Sun</i>	<i>Mon</i>	<i>Tue</i>	<i>Wed</i>	<i>Thu</i>	<i>Fri</i>
<i>Week 1:</i>	Full	5	5	5	5	3
<i>Week 2:</i>		5	5	5	5	3
<i>Week 3:</i>		5	5	5	5	3
<i>Week 4:</i>		5	5	5	5	3

Although the Tuesday — Friday incrementals contain “extra copies” of files from Monday, this scheme assures that any file modified during the week can be recovered from the previous day’s incremental dump.

Process Priority of dump

dump uses multiple processes to allow it to read from the disk and write to the media concurrently. Due to the way it synchronizes between these processes, any attempt to run **dump** with a **nice** (process priority) of ‘-5’ or better will likely make **dump** run *slower* instead of faster.

NAME

dumpfs – dump file system information

SYNOPSIS

/usr/etc/dumpfs device

DESCRIPTION

dumpfs prints out the super block and cylinder group information for the file system or special device specified. The listing is very long and detailed. This command is useful mostly for finding out certain file system information such as the file system block size and minimum free space percentage.

SEE ALSO

fs(5), fsck(8), newfs(8), tune2fs(8)

NAME

edquota – edit user quotas

SYNOPSIS

/usr/etc/edquota [-p proto-user] usernames...

/usr/etc/edquota -t

DESCRIPTION

edquota is a quota editor. One or more users may be specified on the command line. For each user a temporary file is created with an ASCII representation of the current disk quotas for that user and an editor is then invoked on the file. The quotas may then be modified, new quotas added, etc. Upon leaving the editor, **edquota** reads the temporary file and modifies the binary quota files to reflect the changes made.

The editor invoked is **vi(1)** unless the **EDITOR** environment variable specifies otherwise.

Only the super-user may edit quotas. (In order for quotas to be established on a file system, the root directory of the file system must contain a file, owned by root, called **quotas**. See **quotaon(8)** for details.)

OPTIONS

- p** Duplicate the quotas of the prototypical user specified for each user specified. This is the normal mechanism used to initialize quotas for groups of users.
- t** Edit the soft time limits for each file system. If the time limits are zero, the default time limits in **<ufs/quota.h>** are used. Time units of **sec(onds)**, **min(utes)**, **hour(s)**, **day(s)**, **week(s)**, and **month(s)** are understood. Time limits are printed in the greatest possible time unit such that the value is greater than or equal to one.

FILES

quotas	quota file at the file system root
/etc/mtab	mounted file systems

SEE ALSO

quota(1), **vi(1)**, **quotactl(2)**, **quotacheck(8)**, **quotaon(8)**, **repquota(8)**

BUGS

The format of the temporary file is inscrutable.

NAME

`eeeprom` – EEPROM display and load utility

SYNOPSIS

`eeeprom` [`-i`] [`-`] [`-f filename`] [`field [=value]`] ...

`eeeprom` [`-i`] [`-c`] [`-f filename`]

AVAILABILITY

Not available for Sun-2 systems.

DESCRIPTION

`eeeprom` displays or changes the values of fields in the EEPROM. It processes fields in the order given. When processing a *field* accompanied by a *value*, `eeeprom` makes the indicated alteration to the EEPROM; otherwise it displays the *field*'s value. When given no field specifiers, `eeeprom` displays the values of all EEPROM fields. A `-` flag specifies that fields and values are to be read from stdin (one *field* or *field=value* per line).

`eeeprom` verifies the EEPROM checksums and complains if they are incorrect; if the `-i` flag is specified, erroneous checksums are ignored. If the `-c` flag is specified, all incorrect checksums are recomputed and corrected in the EEPROM.

OPTIONS

- `-i` Ignore bad checksums.
- `-f filename`
Use *filename* as the EEPROM device.
- `-c` Correct bad checksums.
- `-` Read field names and values from stdin.

The field names and their possible values are:

<code>hwupdate</code>	a valid date (including "today" and "now")
<code>memsize</code>	8 bit integer (megabytes of memory on machine)
<code>memtest</code>	8 bit integer (megabytes of memory to test)
<code>scrsz</code>	"1024x1024", "1152x900", "1600x1280", or "1440x1440"
<code>watchdog_reboot</code>	"true" or "false"
<code>default_boot</code>	"true" or "false"
<code>bootdev</code>	<code>%c%c (%x,%x,%x)</code>
<code>kbdtype</code>	8 bit integer (0 for all Sun keyboards)
<code>keyclick</code>	"true" or "false"
<code>console</code>	"b&w" or "ttya" or "ttyb" or "color"
<code>custom_logo</code>	"true" or "false"
<code>banner</code>	banner string
<code>diagdev</code>	<code>%c%c (%x,%x,%x)</code> - diagnostic boot device
<code>diagpath</code>	diagnostic boot path
<code>ttya_no_rtsdtr</code>	"true" or "false"
<code>ttyb_no_rtsdtr</code>	"true" or "false"
<code>columns</code>	number of columns on screen (8-bit integer)
<code>rows</code>	number of rows on screen (8-bit integer)

FILES

`/dev/eeeprom`

SEE ALSO

`<mon/eeeprom.h>`

NAME

etherd – Ethernet statistics server

SYNOPSIS

/usr/etc/rpc.etherd interface

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

etherd is a server which puts *interface* into promiscuous mode, and keeps summary statistics of all the packets received on that interface. It responds to RPC requests for the summary. You must be root to run **etherd**.

interface is a networking interface such as *ie0*, *ie1*, *ec0*, *ec1* and *le0*.

traffic(1C) displays the information obtained from **etherd** in graphical form.

SEE ALSO

traffic(1C)

NAME

etherfind – find packets on Ethernet

SYNOPSIS

etherfind [*-nprtuvx*] [*-c count*] [*-i interface*] *expression*

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

etherfind prints out the headers of packets on the ethernet that match the boolean *expression*. When an internet packet is fragmented into more than one ethernet packet, all fragments except the first are marked with an asterisk. You must be root to invoke etherfind.

OPTIONS

- n** Do not convert host addresses and port numbers to names.
- p** Normally, the selected interface is put into promiscuous mode, so that etherfind has access to all packets on the ethernet. However, when the **-p** flag is used, the interface will not go promiscuous.
- r** RPC mode: treat each packet as an RPC message, printing the program and procedure numbers.
- t** Timestamps: precede each packet listing with a time value in seconds and hundredths of seconds since the first packet.
- u** Make the output line buffered.
- v** Verbose mode: print out some of the fields of TCP and UDP packets.
- x** Dump the header in hex, in addition to the line printed for each packet by default.
- c count** Exit after receiving *count* packets. This is sometimes useful for dumping a sample of ethernet traffic to a file for later analysis.
- i interface** etherfind listens on *interface*. The program netstat(8C) when invoked with the **-i** flag lists all the interfaces that a machine has.

expression

The syntax of *expression* is similar to that used by find(1). Here are the allowable primaries.

- dst destination**
True if the destination field of the packet is *destination*, which may be either an address or a name.
- src source**
True if the source field of the packet is *source*, which may be either an address or a name.
- between host1 host2**
True if either the source of the packet is *host1* and the destination *host2*, or the source is *host2* and the destination *host1*.
- dstnet destination**
True if the destination field of the packet has a network part of *destination*, which may be either an address or a name.
- srcnet source**
True if the source field of the packet has a network part of *source*, which may be either an address or a name.

- srcport *port***
True if the packet has a source port value of *port*. It must be either *udp* or *tcp* (see *tcp(4P)*), *udp(4P)*). The *port* can be a number or a name used in */etc/services*.
- dstport *port***
True if the packet has a destination port value of *port*. The *port* can be a number or a name.
- less *length***
True if the packet has a length less than or equal to *length*.
- greater *length***
True if the packet has a length greater than or equal to *length*.
- proto *protocol***
True if the packet is an ip packet (see *ip(4P)*) of protocol type *protocol*. *Protocol* can be a number or one of the names *icmp*, *udp*, *nd*, or *tcp*.
- byte *byte op value***
True if byte number *byte* of the packet is in relation *op* to *value*. Legal values for *op* are *+*, *<*, *>*, *&*, and *|*. Thus *4=6* is true if the fourth byte of the packet has the value 6, and *20&0xf* is true if byte twenty has one of its four low order bits nonzero.
- broadcast**
True if the packet is a broadcast packet.
- arp** True if the packet is a arp packet (see *arp(4P)*).
- rarp** True if the packet is a rarp packet.
- ip** True if the packet is an ip packet.

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

A parenthesized group of primaries and operators (parentheses are special to the Shell and must be escaped).

The negation of a primary ('!' is the unary *not* operator).

Concatenation of primaries (the *and* operation is implied by the juxtaposition of two primaries).

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

EXAMPLE

To find all packets arriving at or departing from sundown

```
example% etherfind -src sundown -o -dst sundown
example%
```

SEE ALSO

find(1), *traffic(1C)*, *arp(4P)*, *ip(4P)*, *nit(4P)* *tcp(4P)*, *udp(4P)*, *netstat(8C)*

BUGS

The syntax is painful.

NAME

exportfs – export and unexport directories to NFS clients

SYNOPSIS

`/usr/etc/exportfs [-avu] [-o options] [directory]`

DESCRIPTION

`exportfs` makes a local directory (or file) available for mounting over the network by NFS clients. It is normally invoked at boot time by the `/etc/rc.local` script, and uses information contained in the `/etc/exports` file to export a *directory* (which must be specified as a full pathname). The super-user can run `exportfs` at any time to alter the list or characteristics of exported directories. Directories that are currently exported are listed in the file `/etc/xtab`.

With no options or arguments, `exportfs` prints out the list of directories currently exported.

OPTIONS

- a** All. Export all directories listed in `/etc/exports`, or if `-u` is specified, unexport all of the currently exported directories.
- v** Verbose. Print each directory as it is exported or unexported.
- u** Unexport the indicated directories.
- i** Ignore the options in `/etc/exports`. Normally, `exportfs` will consult `/etc/exports` for the options associated with the exported directory.

-o options

Specify a comma-separated list of optional characteristics for the directory being exported. *options* can be selected from among:

ro Export the directory read-only. If not specified, the directory is exported read-write.

rw=hostname[:hostname]...

Export the directory read-mostly. Read-mostly means exported read-only to most machines, but read-write to those specified. If not specified, the directory is exported read-write to all.

anon=uid

If a request comes from an unknown user, use *uid* as the effective user ID. Note: root users (uid 0) are always considered "unknown" by the NFS server, unless they are included in the "root" option below. The default value for this option is -2. Setting the value of "anon" to -1 disables anonymous access. Note that by default secure NFS accepts insecure requests as anonymous, and those wishing for extra security can disable this feature by setting "anon" to -1.

root=hostname[:hostname]...

Give root access only to the root users from a specified *hostname*. The default is for no hosts to be granted root access.

access=client[:client]...

Give mount access to each *client* listed. A *client* can either be a hostname, or a netgroup (see `netgroup(5)`). Each *client* in the list is first checked for in the `/etc/netgroup` database, and then the `/etc/hosts` database. The default value allows any machine to mount the given directory.

secure Require clients to use a more secure protocol when accessing the directory.

FILES

<code>/etc/exports</code>	static export information
<code>/etc/xtab</code>	current state of exported directories
<code>/etc/netgroup</code>	

SEE ALSO

exports(5), netgroup(5)

WARNINGS

You cannot export a directory that is either a parent- or a sub-directory of one that is currently exported and *within the same filesystem*. It would be illegal, for example, to export both `/usr` and `/usr/local` if both directories resided in the same disk partition.

NAME

extract_unbundled – extract and execute unbundled-product installation scripts

SYNOPSIS

extract_unbundled [*-ddevice* [*-rremote-host*]] [*-DEFAULT*]

DESCRIPTION

extract_unbundled extracts and executes the installation scripts from release tapes for Sun unbundled software products. If no options are specified, it prompts for input as to the tape device, or remote host-name from which to the software is to be installed. For information about installing a specific product, refer to the installation manual that accompanies that product.

OPTIONS

-ddevice

Install from the indicated tape drive, such as *st0*, *mt0* or *ar0*.

-rremote_host

Install from the device given in the *-d* option on the indicated remote host.

-DEFAULT

Execute the installation script using all default values. Otherwise the installation script prompts for any optional values.

NAME

fastboot, fasthalt – reboot/halt the system without checking the disks

SYNOPSIS

/usr/etc/fastboot [*boot-options*]

/usr/etc/fasthalt [*halt-options*]

DESCRIPTION

fastboot and **fasthalt** are shell scripts that reboot and halt the system without checking the file systems. This is done by creating a file **/fastboot**, then invoking the **reboot(8)** program. The system startup script, **/etc/rc**, looks for this file and, if present, skips the normal invocation of **fsck(8)**.

FILES

/usr/etc/fastboot
/etc/rc

SEE ALSO

fsck(8), **halt(8)**, **init(8)**, **rc(8)**, **reboot(8)**

NAME

fingerd – remote user information server

SYNOPSIS

/usr/etc/in.fingerd

DESCRIPTION

fingerd implements the server side of the Name/Finger protocol, specified in RFC 742. The Name/Finger protocol provides a remote interface to programs which display information on system status and individual users. The protocol imposes little structure on the format of the exchange between client and server. The client provides a single “command line” to the finger server which returns a printable reply.

fingerd waits for connections on TCP port 79. Once connected it reads a single command line terminated by a <RETURN-LINE-FEED> which is passed to **finger(1)**. **fingerd** closes its connections as soon as the output is finished.

If the line is null (only a RETURN-LINEFEED is sent) then **finger** returns a “default” report that lists all people logged into the system at that moment.

If a user name is specified (for instance, eric<RETURN-LINE-FEED>) then the response lists more extended information for only that particular user, whether logged in or not. Allowable “names” in the command line include both “login names” and “user names”. If a name is ambiguous, all possible derivations are returned.

SEE ALSO

finger(1)

Harrenstien, Ken, *NAME/FINGER*, RFC 742, Network Information Center, SRI International, Menlo Park, Calif., December 1977.

BUGS

Connecting directly to the server from a TIP or an equally narrow-minded TELNET-protocol user program can result in meaningless attempts at option negotiation being sent to the server, which will foul up the command line interpretation. **fingerd** should be taught to filter out IAC's and perhaps even respond negatively (IAC *will not*) to all option commands received.

NAME

format – disk partitioning and maintenance utility

SYNOPSIS

format [*-f command-file*] [*-l log-file*] [*-x data-file*] [*-d disk-name*] [*-t disk_type*]
 [*-p partition-name*] [*-s*] *diskname* ...

DESCRIPTION

format enables you to format, label, repair and analyze disks on your Sun computer. Unlike previous disk maintenance programs, **format** runs under SunOS. Because there are limitations to what can be done to the system disk while the system is running, **format** is also supported within the memory-resident system environment. For most applications, however, running **format** under SunOS is the more convenient approach.

If no *disk-list* is present, **format** uses the disk list defined in the data file specified with the *-x* option. If that option is omitted, the data file defaults to *format.dat* in the current directory, or else */etc/format.dat*.

OPTIONS*-f command-file*

Take command input from *command-file* rather than the standard input. The file must contain commands that appear just as they would if they had been entered from the keyboard. With this option, **format** does not issue *continue?* prompts.

-l log-file

Log a transcript of the **format** session to the indicated *log-file*, including the standard input, the standard output and the standard error.

-x data-file

Use the disk list contained in *data-file*.

-d disk_name

Specify which disk should be made current upon entry into the program. The disk is specified by its logical name (for instance, *-xy0*). This can also be accomplished by specifying a single disk in the disk list.

-t disk-type

Specify the type of disk which is current upon entry into the program. A disk's type is specified by name in the data file. This option can only be used if a disk is being made current as described above.

-p partition-name

Specify the partition table for the disk which is current upon entry into the program. The table is specified by its name as defined in the data file. This option can only be used if a disk is being made current, and its type is either specified or available from the disk label.

-s

Silent. Suppress all of the standard output. Error messages are still displayed. This is generally used in conjunction with the *-f* option.

FILES

/etc/format.dat default data file

NAME

fparel - Sun FPA online reliability tests

SYNOPSIS

fparel [**-pn**] [**-v**]

DESCRIPTION

fparel is a command to execute the Sun FPA online confidence and reliability test program. **fparel** tests about 90% of the functions of the FPA board, and tests all FPA contexts not in use by other processes. **fparel** runs without disturbing other processes that may be using the FPA. **fparel** can only be run by the super-user.

After a successful pass, **fparel** writes

time, date: Sun FPA Passed. The contexts tested are: 0, 1, ... 31

to the file `/var/adm/diaglog`.

If a pass fails, **fparel** writes

time, date: Sun FPA failed

along with the test name and context number that failed, to the file `/var/adm/diaglog`. **fparel** then broadcasts the message

time, date: Sun FPA failed, disabled, service required

to all users of the system. Next, **fparel** causes the kernel to disable the FPA. Once the kernel disables the FPA, the system must be rebooted to make it accessible.

The file `/etc/rc.local` should contain an entry to cause **fparel** to be invoked upon reboot to be sure that the FPA remains unaccessible in cases where rebooting doesn't correct the problem. See `rc(8)`.

The `crontab(5)` file for root should contain an entry indicating that `cron(8)` is to run **fparel** daily, such as:

```
7 2 * * * /usr/etc/fpa/fparel
```

which causes **fparel** to run at seven minutes past two, every day. See `cron(8)` and `crontab(5)` for details.

OPTIONS

- pn** Perform *n* passes. Default is *n*=1. **-p0** means perform 2147483647 passes.
- v** Run in verbose mode with detailed test results to the standard output.

FILES

`/var/adm/diaglog` Log of **fparel** diagnostics.
`/etc/rc.local`
`/var/spool/cron/crontabs/root`
`/usr/etc/fpa/*` directory containing FPA microcode, data files, and loader

SEE ALSO

`fpaversion(8)`, `crontab(5)`, `cron(8)`, `rc(8)`

NAME

fpaersion, fpa_download – print FPA version, load microcode

SYNOPSIS

fpaersion [**-hlqv**] [**-t** [**cdhimprstvxCIMS**]]

DESCRIPTION

fpaersion performs various tests on the FPA (floating point accelerator). With no arguments, it prints the version number of the microcode and constants that are currently installed on **/dev/fpa**, and performs a quick test to ensure proper operation.

OPTIONS

- h** Help. Print command-line summary.
- l** Loop through tests infinitely.
- q** Quiet output. Print out only error messages.
- v** Verbose output.
- t** Specify certain tests:
 - c** Command register format instructions.
 - d** Double precision format instructions.
 - h** Help. Print summary of test specifiers.
 - i** Imask register.
 - m** Mode register.
 - p** Simple pipe sequencing.
 - r** User registers for all contexts.
 - s** Single precision format instructions.
 - t** Status generation.
 - v** Print version number and date of microcode and constants.
 - x** Extended format instructions.
 - C** Check checksum for microcode, mapping RAM, and constant RAM.
 - M** Command register format matrix instructions.
 - S** Shadow registers.

FILES

/dev/fpa physical FPA device
/usr/etc/fpa/fpamicro* microcode binaries for specific versions
/usr/etc/fpa/fpa_constants microcode data file
/usr/etc/fpa/fpa_download microcode loader

SEE ALSO

fparel(8), **sysdiag(8)**

NAME

fsck – file system consistency check and interactive repair

SYNOPSIS

/usr/etc/fsck -p [filesystem ...]

/usr/etc/fsck [-b block#] [-w] [-y] [-n] [filesystem] ...

DESCRIPTION

The first form of *fsck* preens a standard set of file systems or the specified file systems. It is normally used in the */etc/rc* script during automatic reboot. In this case, *fsck* reads the table */etc/fstab* to determine the file systems to check. It inspects disks in parallel, taking maximum advantage of I/O overlap to check the file systems as quickly as possible.

Normally, the root file system is checked in pass 1; other root-partition file systems are checked in pass 2. Small file systems on separate partitions are checked in pass 3, while larger ones are checked in passes 4 and 5.

Only partitions marked in */etc/fstab* with a file system type of “4.2” and a non-zero pass number are checked.

fsck corrects innocuous inconsistencies such as: unreferenced inodes, too-large link counts in inodes, missing blocks in the free list, blocks appearing in the free list and also in files, or incorrect counts in the super block, automatically. It displays a message for each inconsistency corrected that identifies the nature of, and file system on which, the correction is to take place. After successfully correcting a file system, *fsck* prints the number of files on that file system, the number of used and free blocks, and the percentage of fragmentation.

If *fsck* encounters other inconsistencies that it cannot fix automatically, it exits with an abnormal return status (and the reboot fails).

If sent a QUIT signal, *fsck* will finish the file system checks, then exit with an abnormal return status that causes the automatic reboot to fail. This is useful when you wish to finish the file system checks, but do not want the machine to come up multiuser.

Without the *-p* option, *fsck* audits and interactively repairs inconsistent conditions on file systems. In this case, it asks for confirmation before attempting any corrections. Inconsistencies other than those mentioned above can often result in some loss of data. The amount and severity of data lost can be determined from the diagnostic output.

The default action for each correction is to wait for the operator to respond either **yes** or **no**. If the operator does not have write permission on the file system, *fsck* will default to a *-n* (no corrections) action.

If no file systems 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. Incorrect directory sizes.
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. If the *lost+found* directory does not exist, it is created. If there is insufficient space its size is increased.

A file system may be specified by giving the name of the cooked or raw device on which it resides, or by giving the name of its mount point. If the latter is given, *fsck* finds the name of the device on which the file system resides by looking in */etc/fstab*.

Checking the raw device is almost always faster.

OPTIONS

- 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.
- w** Check writable file systems only.
- y** Assume a yes response to all questions asked by *fsck*; this should be used with extreme caution, as it is a free license to continue, even after severe problems are encountered.
- n** Assume a no response to all questions asked by *fsck*; do not open the file system for writing.

FILES

/etc/fstab contains default list of file systems to check

DIAGNOSTICS

The diagnostics produced by *fsck* are fully enumerated and explained in *System and Network Administration*.

EXIT STATUS

- 0** Either no errors detected or all errors were corrected.
- 4** Root file system errors were corrected. The system must be rebooted.
- 8** Some uncorrected errors exist on one or more of the file systems checked, there was a syntax error, or some other operational error occurred.
- 12** A signal was caught during processing.

SEE ALSO

fstab(5), *fs(5)*, *newfs(8)*, *mkfs(8)*, *crash(8S)*, *reboot(8)*

System and Network Administration

BUGS

There should be some way to start a *fsck -p* at pass *n*.

NAME

fsirand – install random inode generation numbers

SYNOPSIS

fsirand [**-p**] *special*

DESCRIPTION

fsirand installs random inode generation numbers on all the inodes on device *special*, and also installs a filesystem ID in the superblock. This helps increase the security of filesystems exported by NFS.

fsirand must be used only on an unmounted filesystem that has been checked with **fsck(8)**. The only exception is that it can be used on the root filesystem in single-user mode, if the system is immediately re-booted afterwards.

OPTIONS

-p Print out the generation numbers for all the inodes, but do not change the generation numbers.

SEE ALSO

fsck(8)

NAME

ftpd – DARPA Internet File Transfer Protocol server

SYNOPSIS

`/usr/etc/in.ftpd [-dl] [-timeout] host.socket`

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

ftpd is the DARPA Internet File Transfer Protocol (FTP) server process. The server is invoked by the Internet daemon **inetd**(8C) each time a connection to the FTP service (see **services**(5)) is made, with the connection available as descriptor 0 and the host and socket the connection originated from (in hex and decimal respectively) as argument.

Inactive connections are timed out after 60 seconds.

If the **-d** option is specified, debugging information is logged to the system log daemon, **syslogd**(8).

If the **-l** option is specified, each FTP session is logged to **syslogd**.

The FTP server will timeout an inactive session after 15 minutes. If the **-t** option is specified, the inactivity timeout period will be set to *timeout*.

The FTP server currently supports the following FTP requests; case is not distinguished.

Request	Description
ABOR	abort previous command
ACCT	specify account (ignored)
ALLO	allocate storage (vacuously)
APPE	append to a file
CDUP	change to parent of current working directory
CWD	change working directory
DELE	delete a file
HELP	give help information
LIST	give list files in a directory (“ls -lg”)
MKD	make a directory
MODE	specify data transfer <i>mode</i>
NLST	give name list of files in directory (“ls”)
NOOP	do nothing
PASS	specify password
PASV	prepare for server-to-server transfer
PORT	specify data connection port
PWD	print the current working directory
QUIT	terminate session
RETR	retrieve a file
RMD	remove a directory
RNFR	specify rename-from file name
RNTO	specify rename-to file name

STOR	store a file
STOU	store a file with a unique name
STRU	specify data transfer <i>structure</i>
TYPE	specify data transfer <i>type</i>
USER	specify user name
XCUP	change to parent of current working directory
XCWD	change working directory
XMKD	make a directory
XPWD	print the current working directory
XRMD	remove a directory

The remaining FTP requests specified in RFC 959 are recognized, but not implemented.

The FTP server will abort an active file transfer only when the ABOR command is preceded by a Telnet "Interrupt Process" (IP) signal and a Telnet "Synch" signal in the command Telnet stream, as described in RFC 959.

ftpd interprets file names according to the "globbing" conventions used by **csh**(1). This allows users to utilize the metacharacters '* ? [] {}'.

ftpd authenticates users according to three rules.

- 1) The user name must be in the password data base, **/etc/passwd**, and not have a null password. In this case a password must be provided by the client before any file operations may be performed.
- 2) The user must have a standard shell returned by **getusershell**(3).
- 3) If the user name is "anonymous" or "ftp", an anonymous FTP account must be present in the password file (user "ftp"). In this case the user is allowed to log in by specifying any password (by convention this is given as the client host's name).

In the last case, **ftpd** takes special measures to restrict the client's access privileges. The server performs a **chroot**(2) command to the home directory of the "ftp" user. In order that system security is not breached, it is recommended that the "ftp" subtree be constructed with care; the following rules are recommended.

- ~B ftp Make the home directory owned by "ftp" and unwritable by anyone.
- ~ftp/bin Make this directory owned by the super-user and unwritable by anyone. The program **ls**(1V) must be present to support the list commands. This program should have mode 111.
- ~ftp/etc Make this directory owned by the super-user and unwritable by anyone. The files **passwd**(5) and **group**(5) must be present for the **ls** command to work properly. These files should be mode 444.
- ~ftp/pub Make this directory mode 777 and owned by "ftp". Users should then place files which are to be accessible via the anonymous account in this directory.

DIAGNOSTICS

ftpd logs various errors to the system log daemon, **syslogd**, with a facility code of **daemon**. The messages are listed here, grouped by severity level.

Err Severity

getpeername failed: *reason*
A **getpeername**(2) call failed.

getsockname failed: *reason*
A **getsockname**(2) call failed.

signal failed: *reason*
A **signal**(3) call failed.

setsockopt failed: *reason*

A **setsockopt** call (see **getsockopt(2)**) failed.

ioctl failed: *reason*

A **ioctl(2)** call failed.

directory: *reason*

ftpd did not have write permission on the directory *directory* in which a file was to be created by the **STOU** command.

Info Severity

These messages are logged only if the **-I** flag is specified.

FTPD: connection from *host* at *time*

A connection was made to **ftpd** from the host *host* at the date and time *time*.

FTPD: User *user* timed out after *timeout* seconds at *time*

The user *user* was logged out because they hadn't entered any commands after *timeout* seconds; the logout occurred at the date and time *time*.

Debug Severity

These messages are logged only if the **-d** flag is specified.

FTPD: command: *command*

A command line containing *command* was read from the FTP client.

lost connection

The FTP client dropped the connection.

<— *replycode*

<— *replycode*—

A reply was sent to the FTP client with the reply code *replycode*. The next message logged will include the message associated with the reply. If a **-** follows the reply code, the reply is continued on later lines.

SEE ALSO

ftp(1C), **getsockopt(2)**, **getusershell(3)**, **syslogd(8)**

Postel, Jon, and Joyce Reynolds, *File Transfer Protocol (FTP)*, RFC 959, Network Information Center, SRI International, Menlo Park, Calif., October 1985.

BUGS

The anonymous account is inherently dangerous and should be avoided when possible.

The server must run as the super-user to create sockets with privileged port numbers. It maintains an effective user id of the logged in user, reverting to the super-user only when binding addresses to sockets. The possible security holes have been extensively scrutinized, but are possibly incomplete.

NAME

gettable – get DoD Internet format host table from a host

SYNOPSIS

/usr/etc/gettable *host*

DESCRIPTION

gettable is a simple program used to obtain the DoD Internet host table from a “hostname” server. The indicated *host* is queried for the table. The table, if retrieved, is placed in the file **hosts.txt**.

gettable operates by opening a TCP connection to the port indicated in the service specification for “hostname”. A request is then made for “ALL” names and the resultant information is placed in the output file.

gettable is best used in conjunction with the **htable(8)** program which converts the DoD Internet host table format to that used by the network library lookup routines.

SEE ALSO

intro(3N), **htable(8)**

Harrenstien, Ken, Mary Stahl, and Elizabeth Feinler, *HOSTNAME Server*, RFC 953, Network Information Center, SRI International, Menlo Park, Calif., October 1985.

BUGS

Should allow requests for only part of the database.

NAME

getty – set terminal mode

SYNOPSIS

`/usr/etc/getty [type [tty]]`

Sun386i SYSTEM SYNOPSIS

`/usr/etc/getty [-n] [type [tty]]`

DESCRIPTION

getty, which is invoked by **init(8)**, opens and initializes a *tty* line, reads a login name, and invokes **login(1)**.

The *tty* argument is the name of the character-special file in `/dev` that corresponds to the terminal. If there is no *tty* argument, or the argument is `'-'`, the *tty* line is assumed to be opened as file descriptor 0.

The *type* argument, if supplied, is used as an index into the **gettytab(5)** database—to determine the characteristics of the line. If this argument is absent, or if there is no such entry, the default entry is used. If there is no `/etc/gettytab` file, a set of system-supplied defaults is used.

When the indicated entry is located, **getty** clears the terminal screen, prints a banner heading, and prompts for a login name. Usually, either the banner or the login prompt includes the system's hostname.

Next, **getty** prompts for a login and reads the login name, one character at a time. When it receives a NULL character (which is assumed to be the result pressing the BREAK, or “interrupt” key), **getty** switches to the entry **gettytab** entry named in the *nx* field. It reinitializes the line to the new characteristics, and then prompts for a login once again. This mechanism typically is used to cycle through a set of line speeds (baud rates) for each terminal line. For instance, a rotary dialup might have entries for the speeds: 300, 1200, 150, and 110 baud, with each *nx* field pointing to the next one in succession.

The user terminates login input line with a NEWLINE or RETURN character. The latter is preferable; it sets up the proper treatment of RETURN characters (see **tty(4)**). **getty** checks to see if the terminal has only upper-case alphabetical characters. If all alphabetical characters in the login name are in upper case, the system maps them along with all subsequent upper-case input characters to lower-case internally; they are displayed in upper case for the benefit of the terminal. To force recognition of an upper-case character, the shell allows them to be quoted (typically by preceding each with a backslash, `'\'`).

Finally, **getty** calls **login(1)** with the login name as an argument.

getty can be set to time out after a certain interval; this hangs up dial-up lines if the login name is not entered in time.

Sun386i SYSTEM DESCRIPTION

For Sun386i system, the value of *type* is the constant **Sun**, for the console frame buffer.

Sun386i SYSTEM OPTIONS

-n invoke the full screen login program **logintool(8)**, and optionally the “New User Accounts” feature. May only be used on a frame buffer. Unless removed from the console entry in `/etc/ttytab`, this option is in effect by default.

FILES

`/etc/gettytab`

SEE ALSO

login(1), **ioctl(2)**, **tty(4)**, **gettytab(5)**, **ttytab(5)**, **init(8)**, **logintool(8)**

DIAGNOSTICS

ttyxx: No such device or address.

ttyxx: No such file or directory.

A terminal which is turned on in the **ttys** file cannot be opened, likely because the requisite lines are either not configured into the system, the associated device was not attached during boot-time system configuration, or the special file in `/dev` does not exist.

NAME

gpconfig – initialize the Graphics Processor

SYNOPSIS

/usr/etc/gpconfig gpunit [[-b] [-f] fbunit...]

DESCRIPTION

gpconfig binds **cgtwo** frame buffers to the GP, (Graphics Processor) and loads and starts the appropriate microcode in the GP. For example, the command line:

```
/usr/etc/gpconfig gpone0 cgtwo0 cgtwo1
```

will bind the frame buffer boards **cgtwo0** and **cgtwo1** to the Graphics Processor **gpone0**. The devices **/dev/gpone0a** and **/dev/gpone0b** will then refer to the combination of **gpone** and **cgtwo0** or **cgtwo1** respectively.

The same **cgtwo** frame buffer cannot be bound to more than one GP.

All **cgtwo** frame buffer boards bound to a GP must be configured to the same width and height.

The standard version of the file **/etc/rc.local** contains the following **gpconfig** command line:

```
/usr/etc/gpconfig gpone0 -f -b cgtwo0
```

This binds **gpone0** and **cgtwo0** as **gpone0a**, causes **gpone0a** to use the Graphics Buffer Board if it is present, and redirects **/dev/fb** to be **/dev/gpone0a**. If another configuration is desired, edit the command line in **/etc/rc.local** to do the appropriate thing.

It is inadvisable to run the **gpconfig** command while the GP is being used. Unpredictable results may occur. If it is necessary to change the frame buffer bindings to the GP (or to stop using the GP altogether), bring the system down gently, boot single user, edit the **gpconfig** line in the **/etc/rc.local** file, and bring the system back up multiuser.

OPTIONS

- b** Configure the GP to use the Graphics Buffer as well. Currently only one GP-to-frame-buffer binding is allowed to use the graphics buffer at a time. Only the last **-b** option in the command line takes effect.
- f** Redirect **/dev/fb** to the device formed by binding **gpunit** with **fbunit**. Only the last **-f** option in the command line takes effect.

FILES

```
/dev/cgtwo[0-9]
/dev/fb
/dev/gpone[0-3][abcd]
/usr/lib/gp1cg2.1024.unicode
/usr/lib/gp1cg2.1152.unicode
/etc/rc.local
```

SEE ALSO

cgtwo(4S), **gpone(4S)**

NAME

grpck – check group database entries

SYNOPSIS

/usr/etc/grpck [*filename*]

DESCRIPTION

Note: Optional Software (System V Option). Refer to *Installing the SunOS* for information on how to install this command.

grpck checks that a file in **group(5)** does not contain any errors; it checks the **/etc/group** file by default.

FILES

/etc/group

DIAGNOSTICS**Too many/few fields**

An entry in the group file does not have the proper number of fields.

No group name

The group name field of an entry is empty.

Bad character(s) in group name

The group name in an entry contains characters other than lower-case letters and digits.

Invalid GID

The group ID field in an entry is not numeric or is greater than 65535.

Null login name

A login name in the list of login names in an entry is null.

Login name not found in password file

A login name in the list of login names in an entry is not in the password file.

SEE ALSO

groups(1), group(5), passwd(5)

NAME

halt – stop the processor

SYNOPSIS

/usr/etc/halt [**-nqy**]

DESCRIPTION

halt writes out any information pending to the disks and then stops the processor.

halt normally logs the system shutdown to the system log daemon, **syslogd(8)**, and places a shutdown record in the login accounting file **/var/adm/wtmp**. These actions are inhibited if the **-n** or **-q** options are present.

OPTIONS

- n** Prevent the *sync* before stopping.
- q** Do a quick halt. No graceful shutdown is attempted.
- y** Halt the system, even from a dialup terminal.

FILES

/var/adm/wtmp login accounting file

SEE ALSO

reboot(8), **shutdown(8)**, **syslogd(8)**

NAME

htable – convert DoD Internet format host table

SYNOPSIS

/usr/etc/htable filename

DESCRIPTION

htable converts a host table in the format specified by RFC 952 to the format used by the network library routines. Three files are created as a result of running **htable**: **hosts**, **networks**, and **gateways**. The **hosts** file is used by the **gethostent(3N)** routines in mapping host names to addresses. The **networks** file is used by the **getnetent(3N)** routines in mapping network names to numbers. The **gateways** file is used by the routing daemon in identifying “passive” Internet gateways; see **routed(8C)** for an explanation.

If any of the files **localhosts**, **localnetworks**, or **localgateways** are present in the current directory, the file’s contents is prepended to the output file without interpretation. This allows sites to maintain local aliases and entries which are not normally present in the master database.

htable is best used in conjunction with the **gettable(8C)** program which retrieves the DoD Internet host table from a host.

FILES

localhosts
localnetworks
localgateways

SEE ALSO

intro(3N), **gethostent(3N)**, **getnetent(3N)**, **gettable(8C)**, **routed(8C)**

Harrenstien, Ken, Mary Stahl, and Elizabeth Feinler, *DoD Internet Host Table Specification*, RFC 952, Network Information Center, SRI International, Menlo Park, Calif., October 1985.

BUGS

Does not properly calculate the **gateways** file.

NAME

icheck – file system storage consistency check

SYNOPSIS

/usr/etc/icheck [**-s**] [**-b numbers**] [*filesystem*]

DESCRIPTION

Note: **icheck** has been superceded 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; that is, not in any file nor in the free list.

With the **-s** option **icheck** ignores the actual free list and reconstructs a new one by rewriting the superblock 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 superblock will not continue to be used. Notice also that the words in the superblock which indicate the size of the free list and of the i-list are believed. If the superblock has been curdled these words will have to be patched. The **-s** option suppresses the normal output reports.

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

fs(5), **clri(8)**, **dcheck(8)**, **fsck(8)**, **ncheck(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 superblocks 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

ifconfig – configure network interface parameters

SYOPSIS

```
/etc/ifconfig interface [ address_family ] [ address [ dest_address ] ] [ parameters ] [ netmask mask ]
    [ broadcast address ] [ metric n ]

/etc/ifconfig interface [ protocol_family ]
```

DESCRIPTION

ifconfig is used to assign an address to a network interface and/or to configure network interface parameters. **ifconfig** must be used at boot time to define the network address of each interface present on a machine; it may also be used at a later time to redefine an interface's address or other operating parameters. Used without options, **ifconfig** displays the current configuration for a network interface. If a protocol family is specified, **ifconfig** will report only the details specific to that protocol family. Only the super-user may modify the configuration of a network interface.

The *interface* parameter is a string of the form "name unit", for example **ie0**.

Since an interface may receive transmissions in differing protocols, each of which may require separate naming schemes, the parameters and addresses are interpreted according to the rules of some address family, specified by the *address_family* parameter. The address family currently supported is **inet**. If no address family is specified, **inet** is assumed.

For the DARPA Internet family (**inet**), the address is either a host name present in the host name data base (see **hosts(5)**) or in the Yellow Pages map **hosts**, or a DARPA Internet address expressed in the Internet standard "dot notation". Typically, an Internet address specified in dot notation will consist of your system's network number and the machine's unique host number. A typical Internet address is **192.9.200.44**, where **192.9.200** is the network number and **44** is the machine's host number.

If the *dest_address* parameter is supplied in addition to the *address* parameter, it specifies the address of the correspondent on the other end of a point to point link.

OPTIONS

The following *parameters* may be set with **ifconfig**:

- up** Mark an interface "up". This may be used to enable an interface after an "ifconfig down." It happens automatically when setting the first address on an interface. If the interface was reset when previously marked down, the hardware will be re-initialized.
- down** Mark an interface "down". When an interface is marked "down", the system will not attempt to transmit messages through that interface. If possible, the interface will be reset to disable reception as well. This action does not automatically disable routes using the interface.
- trailers** (**inet** only) Enable the use of a "trailer" link level encapsulation when sending (default — really?). If a network interface supports trailer encapsulation, the system will, when possible, encapsulate outgoing messages in a manner which minimizes the number of memory to memory copy operations performed by the receiver. This feature is machine-dependent, and therefore not recommended. On networks that support the Address Resolution Protocol (see **arp(4P)**; currently, only 10 Mb/s Ethernet), this flag indicates that the system should request that other systems use trailer encapsulation when sending to this host. Similarly, trailer encapsulations will be used when sending to other hosts that have made such requests.
- trailers** Disable the use of a "trailer" link level encapsulation.
- arp** Enable the use of the Address Resolution Protocol in mapping between network level addresses and link level addresses (default). This is currently implemented for mapping between DARPA Internet addresses and 10Mb/s Ethernet addresses.
- arp** Disable the use of the Address Resolution Protocol.

- metric *n*** Set the routing metric of the interface to *n*, default 0. The routing metric is used by the routing protocol (routed(8c)). Higher metrics have the effect of making a route less favorable; metrics are counted as addition hops to the destination network or host.
- netmask *mask*** (inet only) Specify how much of the address to reserve for subdividing networks into sub-networks. The mask includes the network part of the local address and the subnet part, which is taken from the host field of the address. The mask can be specified as a single hexadecimal number with a leading 0x, with a dot-notation Internet address, or with a pseudo-network name listed in the network table networks(5). The mask contains 1's for the bit positions in the 32-bit address which are to be used for the network and subnet parts, and 0's for the host part. The mask should contain at least the standard network portion, and the subnet field should be contiguous with the network portion. If a + (plus sign) is given for the netmask value, then the network number is looked up in the netmasks.byaddr map of the Yellow Pages (or in the /etc/netmasks) file if not running Yellow Pages.
- broadcast *address*** (inet only) Specify the address to use to represent broadcasts to the network. The default broadcast address is the address with a host part of all 0's.

EXAMPLE

If your workstation is not attached to an Ethernet, the ie0 interface should be marked "down" as follows:

```
ifconfig ie0 down
```

FILES

/etc/netmasks

SEE ALSO

intro(3N), netmasks(5), netstat(8C), rc(8)

DIAGNOSTICS

Messages indicating the specified interface does not exist, the requested address is unknown, or the user is not privileged and tried to alter an interface's configuration.

NAME

inetd – Internet services daemon

SYNOPSIS

/usr/etc/inetd [**-d**] [*configuration-file*]

DESCRIPTION

inetd, the Internet services daemon, is normally run at boot time by the */etc/rc.local* script. When started **inetd** reads its configuration information from *configuration-file*, the default being */etc/inetd.conf*. See *inetd.conf(5)* for more information on the format of this file. It listens for connections on the Internet addresses of the services that its configuration file specifies. When a connection is found, it invokes the server daemon specified by that configuration file for the service requested. Once a server is finished, **inetd** continues to listen on the socket (except in some cases which will be described below).

Rather than having several daemon processes with sparsely distributed requests each running concurrently, **inetd**, reduces the load on the system by invoking Internet servers only as they are needed.

inetd itself provides a number of simple TCP-based services. These include **echo**, **discard**, **chargen** (character generator), **daytime** (human readable time), and **time** (machine readable time, in the form of the number of seconds since midnight, January 1, 1900). For details of these services, consult the appropriate RFC, as listed below, from the Network Information Center.

inetd rereads its configuration file whenever it receives a hangup signal, **SIGHUP**. New services can be activated, and existing services deleted or modified in between whenever the file is reread.

SEE ALSO

inetd.conf(5), **comsat(8C)**, **ftpd(8C)**, **rexecd(8C)**, **rlogind(8C)**, **rshd(8C)**, **telnetd(8C)**, **tftpd(8C)**

Postel, Jon, "Echo Protocol," RFC 862, Network Information Center, SRI International, Menlo Park, Calif., May 1983.

Postel, Jon, "Discard Protocol," RFC 863, Network Information Center, SRI International, Menlo Park, Calif., May 1983.

Postel, Jon, "Character Generator Protocol," RFC 864, Network Information Center, SRI International, Menlo Park, Calif., May 1983.

Postel, Jon, "Daytime Protocol," RFC 867, Network Information Center, SRI International, Menlo Park, Calif., May 1983.

Postel, Jon, and Ken Harrenstien, "Time Protocol," RFC 868, Network Information Center, SRI International, Menlo Park, Calif., May 1983.

NAME

infocmp – compare or print out terminfo descriptions

SYNOPSIS

```
infocmp [ -cdnILCruvV1 ] [ -sd ] [ -si ] [ -sl ] [ -sc ] [ -w width ] [ -A directory ] [ -B directory ]
      [ termname ... ]
```

DESCRIPTION

infocmp compares a binary **terminfo(5V)** entry with other terminfo entries, rewrites a terminfo description to take advantage of the *use=field*, or prints out a terminfo description from the corresponding binary file in a variety of formats. It displays boolean fields first, then numeric fields, then string fields.

It can also convert a terminfo entry to a **termcap(5)** entry; the **-C** flag causes **infocmp** to perform this conversion. Some termcap variables are not supported by terminfo, but those that can be derived from terminfo variables are displayed. Not all terminfo capabilities are translated either; only those that are allowed in a termcap entry are normally displayed. Specifying the **-r** option eliminates this restriction, allowing all capabilities to be displayed in termcap form.

Because padding is collected at the beginning of a capability, not all capabilities are displayed. Since mandatory padding is not supported by terminfo and termcap strings are not as flexible, it is not always possible to convert a terminfo string capability into an equivalent working termcap capability. Also, a subsequent conversion of the termcap file back into terminfo format will not necessarily reproduce the original source; **infocmp** attempts to convert parameterized strings, and comments out those that it can not.

Some common terminfo parameter sequences, their termcap equivalents, and some terminal types which commonly have such sequences, are:

Terminfo	Termcap	Representative Terminals
%p1%c	%.	adm
%p1%d	%d	hp, ANSI standard, vt100
%p1%'x'%'%+%c	%+x	concept
%i	%i	ANSI standard, vt100
%p1%?'%'x'%'%>%t%p1%'y'%'%+%;	%>xy	concept
%p2 is printed before %p1	%r	hp

If no *termname* arguments are given, the environment variable **TERM** is used for all expected *termname* arguments.

OPTIONS

Default Options

If no options are specified and either zero or one *termname* is specified, the **-I** option is assumed to be in effect. If more than one *termname* is specified, the **-d** option is assumed.

Comparison Options

infocmp compares the description of the first terminal *termname* with each of the descriptions for terminals listed in subsequent *termname* arguments. If a capability is defined for only one of the terminals, the value returned will depend on the type of the capability: **F** for boolean variables, **-1** for integer variables, and **NULL** for string variables.

- c** Produce a list of capabilities common to both entries. Capabilities that are not set are ignored. This option can be used as a quick check to see if the **-u** option is worth using.
- d** Produce a list of capabilities that differ between descriptions.
- n** Produce a list of capabilities in neither entry.

Source Listing Options

The **-I**, **-L**, and **-C** options produce a source listing for each terminal named.

- I** Use the terminfo names.
- L** Use the long C variable name listed in *<term.h>*.

- C Display only those capabilities that have **termcap** equivalents, using the **termcap** names and displaying them in **termcap** form whenever possible.

The source produced by the **-C** option may be used directly as a **termcap** entry, but not all of the parameterized strings may be changed to the **termcap** format. All padding information for strings is collected together and placed at the beginning of the string where **termcap** expects it. Mandatory padding (padding information with a trailing '/') will become optional.

- r When using **-C**, display all capabilities, not just those capabilities that have **termcap** equivalents.
- u Produce a **terminfo** source description for the first named terminal which is relative to the descriptions given by the entries for all terminals named subsequently on the command line, by analyzing the differences between them, and producing a description with **use=** fields for the other terminals. In this manner, it is possible to retrofit generic **terminfo** entries into a terminal's description. Or, if two similar terminals exist, but were coded at different times or by different people so that each description is a full description, using **infocmp** will show what can be done to change one description to be relative to the other.

A capability is displayed with an at-sign (@) if it no longer exists in the first terminal, but one of the other terminal entries contains a value for it. A capability's value gets printed if the value in the first *termname* is not found in any of the other *termname* entries, or if the first of the other *termname* entries has a different value for that capability.

The order of the other *termname* entries is significant. Since the **terminfo** compiler **tic(8V)** does a left-to-right scan of the capabilities, specifying two **use=** entries that contain differing entries for the same capabilities will produce different results, depending on the order in which they are given. **infocmp** flags any such inconsistencies between the other *termname* entries as they are found.

Alternatively, specifying a capability after a **use=** entry that contains it, will cause the second specification to be ignored. Using **infocmp** to recreate a description can be a useful check to make sure that everything was specified correctly in the original.

Specifying superfluous **use=** slows down the comparison, but is not fatal; **infocmp** flags superfluous **use=** fields.

Sorting Options

- sd Sort fields in the order that they are stored in the **terminfo** database.
- si Sort fields by **terminfo** name.
- sl Sort fields by the long C variable name.
- sc Sort fields by the **termcap** name.

If no sorting option is given, fields are sorted alphabetically by the **terminfo** name within each type, except in the case of the **-C** or the **-L** options, which cause the sorting to be done by the **termcap** name or the long C variable name, respectively.

Changing Databases

The location of the compiled **terminfo** database is taken from the environment variable **TERMINFO**. If the variable is not defined, or if the terminal is not found in that location, the system **terminfo** database, usually in **/usr/share/lib/terminfo**, is used. The options **-A** and **-B** may be used to override this location. With these options, it is possible to compare descriptions for a terminal with the same name located in two different databases. This is useful for comparing descriptions for the same terminal created by different people.

- A Set **TERMINFO** for the first *termname* argument.
- B Set **TERMINFO** for the remaining *termname* arguments.

Other Options

- v Print out tracing information on the standard error.
- V Print out the version of the program in use on the standard error and exit.
- 1 Print fields out one to a line. Otherwise, fields are printed several to a line to a maximum width of 60 characters.
- w *width*
Change the output to *width* characters.

FILES

`/usr/share/lib/terminfo/?/*`
compiled terminal description database
`/usr/5include/term.h`

SEE ALSO

`curses(3V)`, `termcap(5)`, `terminfo(5V)`, `tic(8V)`

DIAGNOSTICS**malloc is out of space!**

There was not enough memory available to process all the terminal descriptions requested. Run `infocmp` in several smaller stages (with fewer *termname* arguments).

use= order dependency found:

A value specified in one relative terminal specification was different from that in another relative terminal specification.

'use=*term*' did not add anything to the description.

A relative terminal name did not contribute anything to the final description.

must have at least two terminal names for a comparison to be done.

The `-u`, `-d` and `-c` options require at least two terminal names.

NAME

init – process control initialization

SYNOPSIS

/usr/etc/init [**-bs**]

DESCRIPTION

init is invoked inside the operating system as the last step in the boot procedure. It normally runs the sequence of commands in the script **/etc/rc.boot** (see **rc(8)**) to check the file system. If passed the **-b** flag from the boot program, **init** skips this step. If the file system check succeeds or is skipped, **init** runs the commands in **/etc/rc** and **/etc/rc.local** to begin multiuser operation; otherwise it commences single-user operation by giving the super-user a shell on the console. It is possible to pass the **-s** parameter from the boot program to **init** so that single-user operation is commenced immediately.

Whenever a single-user shell is created, and the system is running as a secure system, the **init** program demands the super-user password. This is to prevent an ordinary user from invoking a single-user shell and thereby circumventing the system's security. Logging out (for instance, by entering an EOT) causes **init** to proceed with a multi-user boot. The super-user password is demanded whenever the system is running secure as determined by **issecure(3)**, or the terminal is labeled "secure" in **/etc/ttytab**.

Whenever single-user operation is terminated (for instance by killing the single-user shell) **init** runs the scripts mentioned above.

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/ttytab** and executes a command for each terminal specified in the file. This command will usually be **/usr/etc/getty**. **getty(8)** opens and initializes the terminal line, reads the user's name and invokes **login(1)** to log in the user and execute the shell.

Ultimately the shell will terminate because it received EOF, either explicitly, as a result of hanging up, or from the user logging out. The main path of **init**, which has been waiting for such an event, wakes up and removes the appropriate entry from the file **/etc/utmp**, which records current users. **init** then makes an entry in **/var/adm/wtmp**, which maintains a history of logins and logouts. The **/var/adm/wtmp** entry is made only if a user logged in successfully on the line. Then the appropriate terminal is reopened and the command for that terminal is reinvoked.

init catches the *hangup* signal (SIGHUP) and interprets it to mean that the file **/etc/ttytab** should be read again. The shell process on each line which used to be active in **/etc/ttytab** 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 terminal lines without rebooting the system by changing **/etc/ttytab** and sending a *hangup* signal to the **init** process: use **'kill -HUP 1'**.

init terminates multi-user operations and resumes single-user mode if sent a terminate (SIGTERM) signal: use **'kill -TERM 1'**. If there are processes outstanding which are deadlocked (due to hardware or software failure), **init** does not wait for them all to die (which might take forever), but times out after 30 seconds and prints a warning message.

init ceases to create new processes, and allows the system to slowly die away, when sent a terminal stop (SIGTSTP) signal: use **'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)**.

Whenever it reads **/etc/ttytab**, **init** will normally write out an old-style **/etc/ttys** file reflecting the contents of **/etc/ttytab**. This is required in order that programs built on earlier versions of SunOS that read the **/etc/ttys** file (for example, programs using the **ttyslot(3)** routine, such as **shelltool(1)**) may continue to run. If it is not required that such programs run, **/etc/ttys** may be made a link (hard or symbolic) to **/etc/ttytab** and **init** will not write to **/etc/ttys**.

init's role is so critical that if it dies, the system will reboot itself automatically. If, at bootstrap time, the **init** program cannot be located, the system will print an error message and panic.

DIAGNOSTICS

command failing, sleeping.

A process being started to service a line is exiting quickly each time it is started. This is often caused by a ringing or noisy terminal line. `init` will sleep for 30 seconds, then continue trying to start the process.

WARNING: Something is hung (won't 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`
`/var/adm/wtmp`
`/etc/ttytab`
`/etc/rc`
`/etc/rc.local`
`/etc/rc.boot`
`/usr/etc/getty`

SEE ALSO

`kill(1)`, `login(1)`, `sh(1)`, `shelltool(1)`, `issecure(3)`, `ttyslot(3)`, `ttytab(5)`, `getty(8)`, `halt(8)`, `rc(8)`, `reboot(8)`, `shutdown(8)`

NAME

installboot – install bootblocks in a disk partition

SYNOPSIS

/usr/mdec/installboot [-vlt] bootfile protobootblk bootdevice

DESCRIPTION

The **boot(8S)** program is loaded from disk by bootblock code which resides in the bootblock area of a disk partition. In order for the bootblock code to read the boot program (usually **/boot**) it is necessary for it to know the block numbers occupied by the boot program. Previous versions of the bootblock code could find **/boot** by interpreting the file system on the partition from which it was being booted, but this is no longer so.

installboot plugs the block numbers of the boot program into a table in the bootblock code, and writes the modified bootblock code onto the disk. Note carefully that **installboot** must be run every time the boot program is reinstalled, since in general, the block list of the boot program will change each time it is written.

bootfile is the name of the boot program, usually **/boot**. *protobootblk* is the name of the bootblock code into which the block numbers of the boot program are to be inserted. The file read in must have an **a.out(5)** header, but it will be written out to the device with the header removed. *bootdevice* is the name of the disk device onto which the bootblock code is to be installed.

You can see how **installboot** works by making the destination a regular file instead of a device, and examining the result with **od(1V)**.

OPTIONS

- v** Verbose. Display detailed information about the size of the boot program, etc.
- l** Print out the list of block numbers of the boot program.
- t** Test. Display various internal test messages.

EXAMPLE

To install the bootblocks onto the root partition on a Xylogics disk:

```
example% cd /usr/mdec
```

```
example% installboot -vlt /boot bootxy /dev/xy0a
```

For an SD disk, you would use **bootsd** and **/dev/sd0a**, respectively, in place of **bootxy** and **/dev/xy0a**.

SEE ALSO

od(1V), **init(8)**, **boot(8S)**, **bootparamd(8)**, **kadb(8S)**, **ndbootd(8C)**, **monitor(8S)**, **rc(8)**, **reboot(8)**

System and Network Administration

Installing the SunOS

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 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 fiftieth 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 approximate average seek times are calculated 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(8)

NAME

ipallocald – Ethernet-to-IP address allocator

SYNOPSIS

/usr/etc/rpc.ipallocald

AVAILABILITY

Sun386i systems only.

DESCRIPTION

ipallocald is a daemon that determines or temporarily allocates IP addresses within a network segment. It has complete knowledge of the hosts listed in the yellow pages, and, if the system is running the name server, of any hosts listed in internet domain tables automatically accessed on that host through the standard library *gethostbyaddr()* call.

This protocol uses DES authentication (the Sun Secure RPC protocol) to restrict access to this function. The only clients privileged to allocate addresses are those whose net IDs are in the *networks* group. For machine IDs, the machine must be a YP server.

The daemon uses permanent entries in the */etc/ethers* and */etc/hosts* files when they exist and are usable. In other cases, such as when a system is new to the network, *ipallocald* will enter a temporary mapping in a local cache. Entries in the cache are removed when there have been no references to a given entry in a 24-hour period. This cache survives system crashes so that IP addresses will remain consistent.

The daemon also provides corresponding IP address to name mapping.

ipallocald refuses to allocate addresses on networks not listed in the *netrange* file, or for which no free address is available.

FILES

/etc/ez/ipallocald.cache
/etc/ez/ipallocald.netrange

SEE ALSO

pnp(3R), ipallocal(3R), ipallocald(8C), pnpboot(8C), netconfig(8C)

NAME

kadb – adb-like kernel and standalone-program debugger

SYNOPSIS

> **b kadb** [-d] [*boot-flags*]

DESCRIPTION

kadb is an interactive debugger that is similar in operation to **adb**(1), and runs as a standalone program under the PROM monitor. You can use **kadb** to debug the kernel, or to debug any standalone program.

Unlike **adb**, **kadb** runs in the same supervisor virtual address space as the program being debugged — although it maintains a separate context. The debugger runs as a *coprocess* that cannot be killed (no ‘:k’) or rerun (no ‘:r’). There is no signal control (no ‘:i’, ‘:t’, or ‘:s’), although the keyboard facilities (CTRL-C, CTRL-S, and CTRL-Q) are simulated.

While the kernel is running under **kadb**, the abort sequence (L1-A or BREAK) drops the system into **kadb** for debugging — as will a system panic. When running other standalone programs under **kadb**, the abort sequence will pass control to the PROM monitor. **kadb** is then invoked from the monitor by jumping to the starting address for **kadb** found in `/usr/include/debug/debug.h` (currently this can be done for both Sun-2 and Sun-3 system machines with the monitor command ‘g fd0000’, and with the monitor command ‘g fe005000’ for Sun386i systems). **kadb**’s user interface is similar to **adb**. Note: **kadb** prompts with

kadb>

Most **adb** commands function in **kadb** as expected. Typing an abort sequence in response to the prompt returns you to the PROM monitor, from which you can examine control spaces that are not accessible within **adb** or **kadb**. The PROM monitor command **c** will return control to **kadb**. As with “adb -k”, **\$p** works when debugging kernels (by actually mapping in new user pages). The verbs **?** and **/** are equivalent in **kadb**, since there is only one address space in use.

OPTIONS

kadb is booted from the PROM monitor as a standalone program. If you omit the **-d** flag, **kadb** automatically loads and runs **vmunix** from the filesystem **kadb** was loaded from. The **kadb vmunix** variable can be patched to change the default program to be loaded.

-d Interactive startup. Prompts with
kadb:

for a file to be loaded. From here, you can enter a boot sequence line to load a standalone program. Boot flags entered in response to this prompt are included with those already set and passed to the program. If you type a RETURN only, **kadb** loads **vmunix** from the filesystem that **kadb** was loaded from.

boot-flags

You can specify boot flags as arguments when invoking **kadb**. Note: **kadb** always sets the **-d** (debug) boot flag, and passes it to the program being debugged.

USAGE

Refer to **adb** in *Debugging Tools*.

Kernel Macros

As with **adb**, kernel macros are supported. With **kadb**, however, the macros are compiled into the debugger itself, rather than being read in from the filesystem. The **kadb** command **\$M** lists macros known to **kadb**.

Setting Breakpoints

Self-relocating programs such as the SunOS kernel need to be relocated before breakpoints can be used. To set the first breakpoint for such a program, start it with ‘:s’; **kadb** is then entered after the program is relocated (when the system initializes its interrupt vectors). Thereafter, ‘:s’ single-steps as with **adb**. Otherwise, use ‘:c’ to start up the program.

Sun386i System Commands

The Sun386i system version of **kadb** has the following additional commands. Note, for the general syntax of **adb** commands, see **adb(1)**.

- :i** Read a byte (with the INB instruction) in from the port at *address*.
- :o** Send a byte (with the OUTB instruction) containing *count* out through the port at *address*.
- :p** Like **:b** in **adb(1)**, but sets a breakpoint using the hardware debug register instead of the breakpoint instruction. The advantage of using **:p** is that when setting breakpoints with the debug register it is not necessary to have write access to the breakpoint location. Four (4) breakpoints can be set with the hardware debug registers.
- \$\$** Switch I/O from the console to the serial port or vice versa.
- [** Like **:e** in **adb(1)**, but requires only one keystroke and no RETURN character.
-]** Like **:s** in **adb(1)**, but requires only one keystroke and no RETURN character.

Automatic Rebooting with kadb

You can set up your workstation to automatically reboot **kadb** by patching the *vmunix* variable in **/boot** with the string **kadb**. (Refer to **adb** in *Debugging Tools* for details on how to patch executables.)

FILES

/vmunix
/boot
/kadb
/usr/include/debug/debug.h

SEE ALSO

adb(1), **boot(8S)**
Debugging Tools
Writing Device Drivers

BUGS

There is no floating-point support, except on Sun386i systems.

kadb cannot reliably single-step over instructions that change the status register.

When sharing the keyboard with the operating system the monitor's input routines can leave the keyboard in a confused state. If this should happen, disconnect the keyboard momentarily and then reconnect it. This forces the keyboard to reset as well as initiating an abort sequence.

Most of the bugs listed in **adb(1)** also apply to **kadb**.

NAME

keyenvoy – talk to keyserver

SYNOPSIS

keyenvoy

DESCRIPTION

keyenvoy is used by some RPC programs to talk to the key server, **keyserv(8C)**. The key server will not talk to anything but a root process, and **keyenvoy** is a set-uid root process that acts as an intermediary between a user process that wishes to talk to the key server and the key server itself.

This program cannot be run interactively.

SEE ALSO

keyserv(8C)

NAME

keyserv – server for storing public and private keys

SYNOPSIS

keyserv [**-n**]

DESCRIPTION

keyserv is a daemon that is used for storing the private encryption keys of each user logged into the system. These encryption keys are used for accessing secure network services such as secure NFS. When a user logs in to the system, the **login(1)** program uses the login password to decrypt the user's encryption key stored in the Yellow Pages, and then gives the decrypted key to the **keyserv** daemon to store away.

Normally, root's key is read from the file **/etc/.rootkey** when the daemon starts up. This is useful during power-fail reboots when no one is around to type a password, yet you still want the secure network services to operate normally.

OPTIONS

-n Do not read root's key from **/etc/.rootkey**. Instead, prompt the user for the password to decrypt root's key stored in the Yellow Pages and then store the decrypted key in **/etc/.rootkey** for future use. This option is useful if the **/etc/.rootkey** file ever gets out of date or corrupted.

FILES

/etc/.rootkey

SEE ALSO

login(1), **publickey(5)**

NAME

kgmon – generate a dump of the operating system's profile buffers

SYNOPSIS

/usr/etc/kgmon [**-bhpr**] [*filesystem*] [*memory*]

DESCRIPTION

kgmon is a tool used when profiling the operating system. When no arguments are supplied, **kgmon** indicates the state of operating system profiling as running, off, or not configured (see **config(8)**). If the **-p** flag is specified, **kgmon** extracts profile data from the operating system and produces a **gmon.out** file suitable for later analysis by **gprof(1)**.

OPTIONS

- b** Resume the collection of profile data.
- h** Stop the collection of profile data.
- p** Dump the contents of the profile buffers into a **gmon.out** file.
- r** Reset all the profile buffers. If the **-p** flag is also specified, the **gmon.out** file is generated before the buffers are reset.

If neither **-b** nor **-h** is specified, the state of profiling collection remains unchanged. For example, if the **-p** flag is specified and profile data is being collected, profiling is momentarily suspended, the operating system profile buffers are dumped, and profiling is immediately resumed.

FILES

/vmunix	the default system
/dev/kmem	the default memory
gmon.out	

SEE ALSO

gprof(1), **config(8)**

DIAGNOSTICS

Users with only read permission on **/dev/kmem** cannot change the state of profiling collection. They can get a **gmon.out** file with the warning that the data may be inconsistent if profiling is in progress.

NAME

ldconfig – link-editor configuration

SYNOPSIS

/usr/etc/ldconfig [*directory ...*]

DESCRIPTION

ldconfig is used to configure a performance-enhancing cache for the run-time link-editor, **ld.so**. It is run from **/etc/rc.local** and periodically via **cron** to avoid linking with stale libraries. It should be also be run manually when a new shared object (e.g., a shared library) is installed on the system.

When invoked with no arguments, a default set of directories are built into the cache – these are the directories searched by default by the link editors. Additional directories may be specified on the command line.

FILES

/etc/ld.so.cache holds the cached data.

SEE ALSO

ld(1)

NAME

link, unlink – exercise link and unlink system calls

SYNOPSIS

/usr/etc/link filename1 filename2

/usr/etc/unlink filename

DESCRIPTION

link and **unlink** perform their respective system calls on their arguments, abandoning all error checking.

SEE ALSO

rm(1), link(2), unlink(2)

WARNINGS

Only the super-user can unlink a directory, in which case the files it contains are lost. The files can, however, be recovered from the file system's **lost+found** directory after performing an **fsck**.

If you have write permission on the directory in which *filename* resides, **unlink** removes that file without warning, regardless of its ownership.

NAME

lockd – network lock daemon

SYNOPSIS

/etc/rpc.lockd [*-t timeout*] [*-g graceperiod*]

DESCRIPTION

lockd processes lock requests that are either sent locally by the kernel or remotely by another lock daemon. **lockd** forwards lock requests for remote data to the server site's lock daemon through the RPC/XDR(3N) package. **lockd** then requests the status monitor daemon, **statd(8C)**, for monitor service. The reply to the lock request will not be sent to the kernel until the status daemon and the server site's lock daemon have replied.

If either the status monitor or server site's lock daemon is unavailable, the reply to a lock request for remote data is delayed until all daemons become available.

When a server recovers, it waits for a grace period for all client site lockds to submit reclaim requests. Client site lockds, on the other hand, are notified by the **statd** of the server recovery and promptly resubmit previously granted lock requests. If a lockd fails to secure a previously granted lock at the server site, the lockd sends SIGLOST to a process.

OPTIONS

-t timeout

Use *timeout* (seconds) as the interval instead of the default value (15 seconds) to retransmit lock request to the remote server.

-g graceperiod

Use *graceperiod* (seconds) as the grace period duration instead of the default value (45 seconds).

SEE ALSO

fcntl(2V), **lockf(3)**, **signal(3)**, **statd(8C)**

NAME

logintool – graphic login interface

AVAILABILITY

Sun386i systems only.

DESCRIPTION

logintool is invoked by **getty(8)** to display a full screen window for logging in. It cannot be run from the shell. It is more attractive than the traditional 'login:' prompt, and also provides help for the person without a username and information about the workstation.

logintool is normally invoked on the console by **getty(8)**, and works only on a frame buffer.

If the "newlogin" policy in the "policies" YP map is set to "unrestricted," then **logintool** may create new user accounts in the Yellow Pages. The account resides on the local system if it is diskful, or on the system's boot server if the local system is diskless.

FILES

/usr/share/lib/lez/login

SEE ALSO

getty(8)

NAME

lpc – line printer control program

SYNOPSIS

/usr/etc/lpc [*command* [*parameter...*]]

DESCRIPTION

lpc controls the operation of the printer, or of multiple printers, as described in the */etc/printcap* database. **lpc** commands can be used to start or stop a printer, disable or enable a printer's spooling queue, rearrange the order of jobs in a queue, or display the status of each printer—along with its spooling queue and printer daemon.

With no arguments, **lpc** runs interactively, prompting with **lpc>**. If arguments are supplied, **lpc** interprets the first as a *command* to execute; each subsequent argument is taken as a *parameter* for that command. The standard input can be redirected so that **lpc** reads commands from a file.

USAGE**Commands**

Commands may be abbreviated to an unambiguous substring. Note: the *printer* parameter is specified just by the name of the printer (as *lw*), not as you would specify it to **lpr(1)** or **lpq(1)** (not as *-Plw*).

? [*command*]...

help [*command*]...

Display a short description of each command specified in the argument list, or, if no arguments are given, a list of the recognized commands.

abort [all| [*printer* ...]]

Terminate an active spooling daemon on the local host immediately and then disable printing (preventing new daemons from being started by **lpr(1)**) for the specified printers. The **abort** command can only be used by the super-user.

clean [all| [*printer* ...]]

Remove all files with names beginning with *cf*, *tf*, or *df* from the specified printer queue(s) on the local machine. The **clean** command can only be used by the super-user.

disable [all| [*printer* ...]]

Turn the specified printer queues off. This prevents new printer jobs from being entered into the queue by **lpr(1)**. The **disable** command can only be used by the super-user.

down [all| [*printer* ...] [*message*]

Turn the specified printer queue off, disable printing and put *message* in the printer status file. The message doesn't need to be quoted, the remaining arguments are treated like **echo(1V)**. This is normally used to take a printer down and let others know why (**lpq(1)** indicates that the printer is down, as does the **status** command).

enable [all| [*printer* ...]]

Enable spooling on the local queue for the listed printers, so that **lpr(1)** can put new jobs in the spool queue. The **enable** command can only be used by the super-user.

exit

quit Exit from **lpc**.

restart [all| [*printer* ...]]

Attempt to start a new printer daemon. This is useful when some abnormal condition causes the daemon to die unexpectedly leaving jobs in the queue. **lpq(1)** reports that there is no daemon present when this condition occurs. This command can be run by any user.

start [all| [*printer* ...]]

Enable printing and start a spooling daemon for the listed printers. The **start** command can only be used by the super-user.

status [*all* | [*printer ...*]]

Display the status of daemons and queues on the local machine. This command can be run by any user.

stop [*all* | [*printer ...*]]

Stop a spooling daemon after the current job completes and disable printing. The **stop** command can only be used by the super-user.

topq *printer* [*job# ...*] [*user ...*]

Move the print job(s) specified by *job#* or those job(s) belonging to *user* to the top (head) of the printer queue. The **topq** command can only be used by the super-user.

up [*all* | [*printer ...*]] Enable everything and start a new printer daemon. Undoes the effects of **down**.

FILES

/etc/printcap	printer description file
/var/spool/*	spool directories
/var/spool/*/lock	lock file for queue control

SEE ALSO

lpq(1), **lpr(1)**, **lprm(1)**, **printcap(5)**, **lpd(8)**

DIAGNOSTICS

?Ambiguous command

The abbreviation you typed matches more than one command.

?Invalid command

You typed a command or abbreviation that was not recognized.

?Privileged command

You used a command can be executed only by the super-user.

NAME

lpd – printer daemon

SYNOPSIS

`/usr/lib/lpd [-l] [-L logfile] [port#]`

DESCRIPTION

lpd is the line printer daemon (spool area handler). It is normally invoked at boot time from the `rc(8)` script, making a single pass through the `printcap(5)` file to find out about the existing printers and printing any files left after a crash. It then accepts requests to print files in a queue, transfer files to a spooling area, display a queue's status, or remove jobs from a queue. In each case, it forks a child process for each request, and continues to listen for subsequent requests.

The Internet port number used to communicate with other processes is normally obtained with `getservent(3N)`, but can be specified with the `port#` argument.

OPTIONS

-l Log valid requests received from the network. This can be useful for debugging purposes.

-L *logfile*

Change the file used for writing error conditions to *logfile*. The default is to report a message using the `syslog(3)` facility.

OPERATION**Access Control**

Access control is provided by two means. First, all requests must come from one of the machines listed in either the file `/etc/hosts.equiv` or `/etc/hosts.lpd`. Second, if the `rs` capability is specified in the `printcap` entry, `lpr(1)` requests are only be honored for users with accounts on the printer host.

Lock File

The lock file in each spool directory is used to prevent multiple daemons from becoming active, and to store information about the daemon process for `lpr(1)`, `lpq(1)`, and `lprm(1)`.

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)`.

Control Files

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 that indicates what to do with the remainder of the line.

J	Job name to print on the burst page.
C	Classification line on the burst page.
L	Literal. This line contains identification information from the password file, and causes a burst page to be printed.
T	Title string for page headings printed by <code>pr(1V)</code> .
H	Hostname of the machine where <code>lpr(1)</code> was invoked.
P	Person. Login name of the person who invoked <code>lpr(1)</code> . This is used to verify ownership by <code>lprm(1)</code> .
M	Send mail to the specified user when the current print job completes.
f	Formatted File, the name of a file to print that is already formatted.
l	Like f , but passes control characters along, and does not make page breaks.
p	Name of a file to print using <code>pr(1V)</code> as a filter.
t	Troff File. The file contains <code>troff(1)</code> output (cat phototypesetter commands).
n	Ditroff File. The file contains device independent troff output.
d	DVI File. The file contains T _E X output (DVI format from Stanford).
g	Graph File. The file contains data produced by <code>plot(3X)</code> .

c	Cifplot File. The file contains data produced by <i>cifplot</i> .
v	The file contains a raster image.
r	The file contains text data with FORTRAN carriage control characters.
1	Troff Font R. The name of a font file to use instead of the default.
2	Troff Font I. The name of the font file to use instead of the default.
3	Troff Font B. The name of the font file to use instead of the default.
4	Troff Font S. The name of the font file to use instead of the default.
W	Width. Changes the page width (in characters) used by <i>pr</i> (1V) and the text filters.
I	Indent. Specify the number of characters by which to indent the output.
U	Unlink. The name of file to remove upon completion of printing.
N	Filename. The name of the file being printed, or a blank for the standard input (when <i>lpr</i> (1) is invoked in a pipeline).

Data Files

If a file can not be opened, an error message is logged using the LOG_LPR facility of *syslog*(3). *lpd* will try up to 20 times to reopen a file it expects to be there, after which it proceeds to the next file or job.

Minfree File

The file *minfree* in each spool directory contains the number of disk blocks to leave free so that the line printer queue won't completely fill the disk. The *minfree* file can be edited with your favorite text editor.

FILES

/etc/printcap	printer description file
/var/spool/*	spool directories
/var/spool/*/minfree	minimum free space to leave
/dev/lp*	line printer devices
/dev/printer	socket for local requests
/etc/hosts.equiv	hosts allowed equivalent host access
/etc/hosts.lpr	hosts allowed printer access only

SEE ALSO

lpr(1), *lpq*(1), *lprm*(1), *printcap*(5), *lpc*(8), *pac*(8)

NAME

mailstats – print statistics collected by sendmail

SYNOPSIS

/usr/etc/mailstats [filename]

DESCRIPTION

mailstats prints out the statistics collected by the **sendmail** program on mailer usage. These statistics are collected if the file indicated by the **S** configuration option of **sendmail** exists. The **mailstats** program first prints the time that the statistics file was created and the last time it was modified. It will then print a table with one row for each mailer specified in the configuration file. The first column is the mailer number, followed by the symbolic name of the mailer. The next two columns refer to the number of messages received by *sendmail*, and the last two columns refer to messages sent by *sendmail*. The number of messages and their total size (in 1024 byte units) is given. No numbers are printed if no messages were sent (or received) for any mailer.

You might want to add an entry to */var/spool/cron/crontab/root* to reinitialize the statistics file once a night. Copy */dev/null* into the statistics file or otherwise truncate it to reset the counters.

FILES

/etc/sendmail.st default statistics file
/etc/sendmail.cf sendmail configuration file
/var/spool/cron/crontab/root
/dev/null

SEE ALSO

sendmail(8)

BUGS

Mailstats should read the configuration file instead of having a hard-wired table mapping mailer numbers to names.

NAME

makedbm – make a Yellow Pages dbm file

SYNOPSIS

```
makedbm [ -b ] [ -i yp_input_file ] [ -o yp_output_name ] [ -d yp_domain_name ]
  [ -m yp_master_name ] infile outfile
makedbm [ -u dbmfilename ]
```

DESCRIPTION

makedbm takes *infile* and converts it to a pair of files in **ndbm(3)** format, namely *outfile.pag* and *outfile.dir*. Each line of the input file is converted to a single **dbm** record. All characters up to the first TAB or SPACE form the key, and the rest of the line is the data. If a line ends with '\', then the data for that record is continued on to the next line. It is left for the clients of the Yellow Pages to interpret #; **makedbm** does not itself treat it as a comment character. *infile* can be '-', in which case the standard input is read.

makedbm is meant to be used in generating **dbm** files for the Yellow Pages, and it generates a special entry with the key *yp_last_modified*, which is the date of *infile* (or the current time, if *infile* is '-').

OPTIONS

- b** Interdomain. Propagate a map to all servers using the interdomain name server **named(8C)**.
- i** Create a special entry with the key *yp_input_file*.
- o** Create a special entry with the key *yp_output_name*.
- d** Create a special entry with the key *yp_domain_name*.
- m** Create a special entry with the key *yp_master_name*. If no master host name is specified, *yp_master_name* will be set to the local host name.
- u** Undo a **dbm** file. That is, print out a **dbm** file one entry per line, with a single space separating keys from values.

EXAMPLE

It is easy to write shell scripts to convert standard files such as */etc/passwd* to the key value form used by **makedbm**. For example,

```
#!/bin/awk -f
BEGIN { FS = ":"; OFS = "\t"; }
{ print $1, $0 }
```

takes the */etc/passwd* file and converts it to a form that can be read by **makedbm** to make the Yellow Pages file *passwd.byname*. That is, the key is a username, and the value is the remaining line in the */etc/passwd* file.

SEE ALSO

yppasswd(1), **ndbm(3)**, **named(8C)**

NAME

makedev, MAKEDEV – make system special files

SYNOPSIS

/dev/MAKEDEV device-name ...

DESCRIPTION

MAKEDEV is a shell script normally used to install special files. It resides in the */dev* directory, as this is the normal location of special files. Arguments to MAKEDEV are usually of the form *device-name?* where *device-name* is one of the supported devices listed in section 4 of the manual and '?' is a logical unit number (0-9). A few special arguments create assorted collections of devices and are listed below.

std Create the *standard* devices for the system; for example, */dev/console*, */dev/tty*.

local Create those devices specific to the local site. This request runs the shell file */dev/MAKEDEV.local*. Site specific commands, such as those used to setup dialup lines as "ttyd?" should be included in this file.

Since all devices are created using *mknod(8)*, this shell script is useful only to the super-user.

DIAGNOSTICS

Either self-explanatory, or generated by one of the programs called from the script. Use *sh -x MAKEDEV* in case of trouble.

SEE ALSO

intro(4), *config(8)*, *mknod(8)*

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(1)** and **crypt(1)**). Usually **makekey**'s input and output will be pipes.

SEE ALSO

crypt(1), **ed(1)**

NAME

`mc68881version` – print the MC68881 mask number and approximate clock rate

SYNOPSIS

`/usr/etc/mc68881version`

AVAILABILITY

Sun-2, Sun-3, and Sun-4 systems only.

DESCRIPTION

`mc68881version` determines whether an MC68881 or MC68882 floating-point coprocessor is available, and if so, determines its apparent mask number and approximate clock rate and prints them on the standard output. The clock rate is derived by timing floating-point operations with `getrusage(2)` and is thus somewhat variable.

SEE ALSO

`getrusage(2)`

NAME

mconnect – connect to SMTP mail server socket

SYNOPSIS

/usr/etc/mconnect [**-p** *port*] [**-r**] [*hostname*]

DESCRIPTION

mconnect opens a connection to the mail server on a given host, so that it can be tested independently of all other mail software. If no host is given, the connection is made to the local host. Servers expect to speak the Simple Mail Transfer Protocol (SMTP) on this connection. Exit by typing the **quit** command. Typing EOF will send an end of file to the server. An interrupt closes the connection immediately and exits.

OPTIONS

-p *port* Specify the port number instead of the default SMTP port (number 25) as the next argument.
-r “Raw” mode: disable the default line buffering and input handling. This gives you a similar effect as **telnet** to port number 25, not very useful.

FILES

/usr/lib/sendmail.hf help file for SMTP commands

SEE ALSO

sendmail(8)

Postel, Jonathan B *Simple Mail Transfer Protocol*, RFC821 August 1982, SRI Network Information Center

NAME

mkfile – create a file

SYNOPSIS

mkfile [**-nv**] *size*[**k|b|m**] *filename* ...

DESCRIPTION

mkfile creates one or more files that are suitable for use as NFS-mounted swap areas. The sticky bit is set, and the file is padded with zeroes by default. The default *size* is in bytes, but it can be flagged as kilobytes, blocks, or megabytes, with the **k**, **b**, or **m** suffixes, respectively.

OPTIONS

- n** Create an empty *filename*. The size is noted, but disk blocks aren't allocated until data is written to them.
- v** Verbose. Report the names and sizes of created files.

NAME

mkfs – construct a file system

SYNOPSIS

```
/usr/etc/mkfs [ -N ] special size [ nsect ] [ ntrack ] [ blksize ] [ fragsize ] [ ncpg ] [ minfree ]
[ rps ] [ nbpi ] [ opt ] [ apc ] [ rot ]
```

DESCRIPTION

Note: file systems are normally created with the **newfs(8)** command.

mkfs constructs a file system by writing on the special file *special* unless the **-N** flag has been specified. 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 inodes is calculated as a function of the file system size. No boot program is initialized by **mkfs** (see **newfs(8)**).

OPTIONS

The optional arguments allow fine tune control over the parameters of the file system.

nsect The number of sectors per track on the disk. The default is **32**.

ntrack The number of tracks per cylinder on the disk. The default is **16**.

blksize The primary block size for files on the file system. It must be a power of two, currently selected from **4096** or **8192** (the default).

fragsize 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**. The default is **1024**.

ncpg The number of disk cylinders per cylinder group. This number must be in the range **1** to **32**. The default is **16**.

minfree 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%**.

rps The rotational speed of the disk, in revolutions per second. The default is **60**.

nbpi The number of bytes for which one inode block is allocated. This parameter is currently set at one inode block for every **2048** bytes.

opt Space or time optimization preference; **s** specifies optimization for space, **t** specifies optimization for time. The default is **t**.

apc The number of alternates per cylinder (SCSI devices only). The default is **0**.

rot 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.

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

dir(5), **fs(5)**, **fsck(8)**, **newfs(8)**, **tunefs(8)**

System and Network Administration

McKusick, Joy, Leffler; *A Fast File System for UNIX*,

NOTES

newfs(8) is much to be preferred for most routine uses.

NAME

mknod – build special file

SYNOPSIS

/usr/etc/mknod filename [c] [b] major minor

/usr/etc/mknod filename p

DESCRIPTION

mknod makes a special file. The first argument is the *filename* of the entry. In the first form, the second argument 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 (for example, unit, drive, or line number). Only the super-user is permitted to invoke this form of the **mknod** command.

In the second form, **mknod** makes a named pipe (FIFO).

The first form of **mknod** is only for use by system configuration people. Normally you should use **/dev/MAKEDEV** instead when making special files.

SEE ALSO

mknod(2), **makedev(8)**

NAME

mkproto – construct a prototype file system

SYNOPSIS

/usr/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 SPACE or NEW-LINE characters. 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(1V)**.

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 /usr/bin/sh
      ken  d—755 6 1
      $
      b0   b—644 3 1 0 0
      c0   c—644 3 1 0 0
      $
$

```

SEE ALSO

chmod(1V), **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

modload – load a Sun386i module

SYNOPSIS

modload *filename* [**-conf** *config_file*] [**-entry** *entry_point*] [**-exec** *exec_file*] [**-o** *output_file*]
[**-nolink**] [**-A** *vmunix_file*]

AVAILABILITY

Sun386i systems only.

DESCRIPTION

modload loads a loadable module into a running system. The input file *filename* is an object file (.o file).

OPTIONS**-conf** *config_file*

Use this configuration file to configure the loadable driver being loaded. The commands in this file are the same as those that the **config(8)** program recognizes. There are two additional commands, **blockmajor** and **charmajor**, shown in the configuration file example below.

-entry *entry_point*

This is the module entry point. This is passed by **modload** to **ld(1)** when the module is linked. The default module entry point name is 'xxx init'.

-exec *exec_file*

This is the name of a shell script or executable image file that will be executed if the module is successfully loaded. It is always passed the module id and module type as the first two arguments. For loadable drivers, the third and fourth arguments are the block major and character major numbers respectively. For a loadable system call, the third argument is the system call number.

-o *output_file*

This is the name of the output file that is produced by the linker. If this option is omitted, then the output file name is *filename*> without the '.o'.

-nolink This option can be used if **modload** has already been issued once and the output file already exists. One must take care that neither the kernel nor the module have changed.

-A *vmunix_file*

This is the file that is passed to the linker to resolve module references to kernel symbols. The default is */vmunix*. The symbol file must be for the currently running kernel or the module is likely to crash the system.

EXAMPLE

```

controller      fdc0 at atmem csr 0x001000 irq 6 priority 3
controller      fdc2 at atmem csr 0x002000 irq 5 priority 2
disk            fd0 at fdc0 drive 0
disk            fd0 at fdc0 drive 1
disk            fd0 at fdc0 drive 2
device          fd0 at fdc2 drive 0 csr 0x003000 irq 4 priority 2
disk            fd0 at fdc2 drive 1
blockmajor 51
charmajor 52

```

SEE ALSO

ld(1), **modunload(8)**, **modstat(8)**

NAME

modstat – display status of Sun386i modules

SYNOPSIS

modstat [**-id** *module_id*]

AVAILABILITY

Sun386i systems only.

DESCRIPTION

modstat displays the status of the loaded modules.

OPTIONS

-id *module_id*
Display status of only this module.

SEE ALSO

modload(8), **modunload(8)**

NAME

modunload – unload a Sun386i module

SYNOPSIS

modunload **-id** *module_id* [**-exec** *exec_file*]

AVAILABILITY

Sun386i systems only.

DESCRIPTION

modunload unloads a loadable module from a running system. The *module_id* is the ID of the module as shown by **modstat(8)**.

OPTIONS

-exec *exec_file*

This is the name of a shell script or executable image file that will be executed before the module is unloaded. It is always passed the module id and module type as the first two arguments. For loadable drivers, the third and fourth arguments are the block major and character major numbers respectively. For a loadable system call, the third argument is the system call number.

SEE ALSO

modload(8), **modstat(8)**

NAME

monitor – system ROM monitor

SYNOPSIS

L1-A

BREAK

DESCRIPTION

The CPU board of the Sun workstation contains an EPROM (or set of EPROMs), called the *monitor*, that controls the system during startup. The monitor tests the system before attempting to boot the operating system. If you interrupt the boot procedure by holding down L1 while typing a or A on the workstation keyboard (or BREAK if the console is a dumb terminal) the monitor issues the prompt:

>

and accepts commands interactively.

USAGE

Commands

- +|- Increment or decrement the current address and display the contents of the new location.
- ^C *source destination n*
 (CTRL-C) Copy, byte-by-byte a block of length *n* from the *source* address to the *destination* address.
- ^I *program* (CTRL-I) Display the compilation date and location of *program*.
- ^T *virtual_address*
 (CTRL-T) Display the physical address to which *virtual_address* is mapped.
- a [*n*] [*action*]. . . (Sun-2 and Sun-3 systems only)
 Open A-register (cpu address register) *n*, and perform indicated actions. The number *n* can be any value from 0 to 7, inclusive. The default value is 0. A hexadecimal *action* argument assigns the value you supply to the register *n*. A non-hex *action* terminates command input.
- b [!] [*device* [(*c,u,p*)]] [*pathname*] [*arguments_list*]
b[?] Reset appropriate parts of the system and bootstrap a program. A '!' (preceding the *device* argument) prevents the system reset from occurring. Programs can be loaded from various devices (such as a disk, tape or Ethernet). 'b' with no arguments will cause a default boot, either from a disk, or from an Ethernet controller. 'b?' displays all boot devices and their *device* arguments, where *device* is one of:
 - ie Intel Ethernet
 - le Lance Ethernet (Sun-2, Sun-3, Sun-4 systems only)
 - sd SCSI disk
 - st SCSI 1/4" tape
 - mt Tape Master 9-track 1/2" tape (Sun-2, Sun-3, Sun-4 systems only)
 - xd Xylogics 7053 disk (Sun-2, Sun-3, Sun-4 systems only)
 - xt Xylogics 1/2" tape (Sun-2, Sun-3, Sun-4 systems only)
 - xy Xylogics 440/450 disk (Sun-2, Sun-3, Sun-4 systems only)
 - fd Diskette (Sun386i system only)
- c* A controller number (0 if only one controller),
- u* A unit number (0 if only one driver), and
- p* A partition.
- pathname* A pathname for a program such as /stand/diag. /vmunix is the default.
- arguments_list*
 A list of up to seven arguments to pass to the program being booted.

c [*virtual_address*]

Resume execution of a program. When given, *virtual_address* is the address at which execution will resume. The default is the current PC (EIP on Sun386i systems). Registers are restored to the values shown by the **a**, **d**, and **r** commands (for Sun-2 and Sun-3 systems), or by the **d** and **r** commands (for Sun-4 systems), or by the **d** command (for Sun386i systems).

d [*window_number*] (Sun-4 systems only)

Display (dump) the state of the processor. The processor state is observable only after:

- An unexpected trap was encountered.
- A user program dropped into the monitor (by calling *abortent*).
- The user manually entered the monitor by typing L1-A or BREAK.

The display consists of the following:

- The special registers: PSR, PC, nPC, TBR, WIM and Y
- Eight global registers, and
- 24 window registers (8 *in*, 8 *local*, and 8 *out*), corresponding to one of the 7 available windows. If a Floating-Point Unit is on board, its status register along with its 32 floating-point registers are also shown.

window_number

Display the indicated *window_number*, which can be any value between 0 and 6, inclusive. If no window is specified and the PSR's current window pointer contains a valid window number, registers from the window that was active just prior to entry into the monitor are displayed. Otherwise, registers from window 0 are displayed.

d (Sun386i systems only)

Display (dump) the state of the processor. This display consists of the registers, listed below:

Processor Registers:	EAX, ECX, EDX, ESI, EDI, ESP, EBP, EFLAGS, EIP
Segment Registers:	ES, CS, SS, DS, FS, GS
Memory Management Registers:	GDTR, LDTR, IDTR, TR
Control Registers:	CR0, CR2, CR3
Debug Registers:	DR0, DR1, DR2, DR3, DR6, DR7
Test Registers:	TR6, TR7

The processor's state is observable only after an unexpected trap, a user program has "dropped" into the monitor (by calling monitor function *abortent*) or the user has manually "broken" into the monitor (by typing L1-A on the Workstation console, or BREAK on the dumb terminal's keyboard).

d [*n*] [*action*]... (Sun-2 and Sun-3 systems only)

Open D-register (cpu data register) *n*, and perform indicated actions. The number *n* can be any value from 0 to 7, inclusive. The default is 0. See the **a** command for a description of *action*.

e [*virtual_address*] [*action*]...

Open the 16 bit word at *virtual_address* (default zero). On Sun-2, Sun-3, and Sun-4 systems, the address is interpreted in the address space defined by the **s** command. See the **a** command for a description of *action*.

f *virtual_address1 virtual_address2 pattern [size]* (Sun-3 and Sun-4 systems only)

Fill the bytes, words or long words from *virtual_address1* (lower) to *virtual_address2* (higher) with the constant, *pattern*. The *size* argument can take one of the following values

- b** byte format (the default)
- w** word format
- l** long word format

For example, the following command fills the address block from 0x1000 to 0x2000 with the word pattern, 0xABCD:

```
f 1000 2000 ABCD W
```

g [*vector*] [*argument*]

g [*virtual_address*] [*argument*]

Goto (jump to) a predetermined or default routine (first form), or to a user-specified routine (second form). The value of *argument* is passed to the routine. If the *vector* or *virtual_address* argument is omitted, the value in the PC is used as the address to jump to.

To set up a predetermined routine to jump to, a user program must, prior to executing the monitor's **g** command, set the variable `*romp->v_vector_cmd` to be equal to the virtual address of the desired routine. Predetermined routines need not necessarily return control to the monitor.

The default routine, defined by the monitor, prints the user-supplied *vector* according to the format supplied in *argument*. This format can be one of:

- `%x` hexadecimal
- `%d` decimal

g0 (Sun-2, Sun-3, and Sun-4 only)

When the monitor is running as a result of the system being interrupted, force a panic and produce a crash dump.

g4

When the monitor is running as a result of the system being interrupted, force a kernel stack trace.

h (Sun-3 and Sun-4 and Sun386i systems)

Display the help menu for monitor commands and their descriptions. To return to the monitor's basic command level, press ESCAPE or **q** before pressing RETURN.

i [*cache_data_offset*] [*action*]... (Sun-3/200 series and Sun-4 systems only)

Modify cache data RAM command. Display and/or modify one or more of the cache data addresses. See the **a** command for a description of *action*.

j [*cache_tag_offset*] [*action*]... (Sun-3/200 series and Sun-4 systems only)

Modify cache tag RAM command. Display and/or modify the contents of one or more of the cache tag addresses. See the **a** command for a description of *action*.

k [*reset_level*]

Reset the system. If *reset_level* is:

- 0** CPU reset only (Sun-2 and Sun-3 systems). Reset VMEbus, interrupt registers, video monitor (Sun-4 systems). This is the default. Reset video (Sun386i systems).
- 1** Software reset.
- 2** Power-on reset. Resets and clears the memory. Runs the EPROM-based diagnostic self test, which can take several minutes, depending upon how much memory is being tested.

kb Display the system banner.

- l** [*virtual_address*] [*action*] ...
 Open the long word (32 bit) at memory address *virtual_address* (default zero). On Sun-2, Sun-3 and Sun-4 systems, the address is interpreted in the address space defined by the *s* command (below). See the *a* command for a description of *action*.
- m** [*virtual_address*] [*action*] ...
 Open the segment map entry that maps *virtual_address* (default zero). On Sun-2, Sun-3 and Sun-4 systems, the address is interpreted in the address space defined by the *s* command. Not supported on Sun386i. See the *a* command for a description of *action*.
- nd** (Sun386i systems only)
ne
ni Disable, enable, or invalidate the cache, respectively
- o** [*virtual_address*] [*action*] ...
 Open the byte location specified by *virtual_address* (default zero). On Sun-2, Sun-3 and Sun-4 systems, the address is interpreted in the address space defined by the *s* command. See the *a* command for a description of *action*.
- p** [*virtual_address*] [*action*] ...
 Open the page map entry that maps *virtual_address* (default zero) in the address space defined by the *s* command. See the *a* command for a description of *action*.
- p** [*port_address*] [[*nonhex_char* [*hex_value*] | *hex_value*] ...] (Sun386i systems only)
 Display or modify the contents of one or more port I/O addresses in byte mode. Each port address is treated as a 8-bit unit. The optional *port_address*, argument, which is a 16-bit quantity, specifies the initial port I/O address. See the *e* command for argument descriptions.
- q** [*eeeprom_offset*] [*action*] ... (Sun-3 and Sun-4 systems only)
 Open the EEPROM *eeeprom_offset* (default zero) in the EEPROM address space. All addresses are referenced from the beginning or base of the EEPROM in physical address space, and a limit check is performed to insure that no address beyond the EEPROM physical space is accessed. On Sun386i systems, open the NVRAM *nvrasm_offset* (default zero). This command is used to display or modify configuration parameters, such as: the amount of memory to test during self test, whether to display a standard or custom banner, if a serial port (A or B) is to be the system console, etc. See the *a* command for a description of *action*.
- r** [*reg_name*] [[*nonhex_char* [*hex_value*] | *hex_value*] ...] (Sun386i systems only)
 Display or modify one or more of the processor registers. If *reg_name* is specified (2 or 3 characters from the above list), that register is displayed first. The default is EAX. See note on register availability under the command *d* (for Sun386i systems). See the *e* command for argument descriptions.
- s** [*step_count*] (Sun386i systems only)
 Single step the execution of the interrupted program. The *step_count* argument specifies the number of single steps to execute before displaying the monitor prompt. The default is 1.
- r** [*register_number*] [*action*] ... (Sun-2 and Sun-3 systems only)
 Display and/or modify the register indicated. *register_number* can be one of:
- CA 68020 Cache Address Register
 - CC 68020 Cache Control Register
 - CX 68020 System and User Context
 - DF Destination Function code
 - IS 68020 Interrupt Stack Pointer
 - MS 68020 Master Stack Pointer
 - PC Program Counter
 - SC 68010 System Context

SF Source Function code
 SR Status Register
 SS 68010 Supervisor Stack Pointer
 UC 68010 User Context
 US User Stack Pointer
 VB Vector Base

Alterations to these registers (except SC and UC) do not take effect until the next **c** command is executed. See the **a** command for a description of *action*.

r [*register_number*] (Sun-4 systems only)
r [*register_type*]
r [*w window_number*]

Display and/or modify one or more of the IU or FPU registers.

A hexadecimal *register_number* can be one of:

0x00—0x0f window(0,i0)—window(0,i7), window(0,i0)—window(0,i7)
 0x16—0x1f window(1,i0)—window(1,i7), window(1,i0)—window(1,i7)
 0x20—0x2f window(2,i0)—window(2,i7), window(2,i0)—window(2,i7)
 0x30—0x3f window(3,i0)—window(3,i7), window(3,i0)—window(3,i7)
 0x40—0x4f window(4,i0)—window(4,i7), window(4,i0)—window(4,i7)
 0x50—0x5f window(5,i0)—window(5,i7), window(5,i0)—window(5,i7)
 0x60—0x6f window(6,i0)—window(6,i7), window(6,i0)—window(6,i7)
 0x70—0x77 g0, g1, g2, g3, g4, g5, g6, g7
 0x78—0x7d PSR, PC, nPC, WIM, TBR, Y
 0x7e—0x9e FSR, f0—f31

Register numbers can only be displayed after an unexpected trap, a user program has entered the monitor using the *abortent* function, or the user has entered the monitor by manually typing **L1-A** or **BREAK**.

If a *register_type* is given, the first register of the indicated type is displayed. *register_type* can be one of:

f floating-point
g global
s special

If **w** and a *window_number* (0—6) are given, the first *in*-register within the indicated window is displayed. If *window_number* is omitted, the window that was active just prior to entering the monitor is used. If the PSR's current window pointer is invalid, window 0 is used.

s [*code*] (Sun-2 and Sun-3 systems only)

Set or query the address space to be used by subsequent memory access commands. *code* is one of:

0 undefined
1 user data space
2 user program space
3 user control space
4 undefined
5 supervisor data space
6 supervisor program space
7 supervisor control space

If *code* is omitted, **s** displays the current address space.

s [*asi*] (Sun-4 systems only)

Set or display the Address Space Identifier. With no argument, **s** displays the current Address Space Identifier. The *asi* value can be one of:

0x2	control space
0x3	segment table
0x4	Page table
0x8	user instruction
0x9	supervisor instruction
0xa	user data
0xb	supervisor data
0xc	flush segment
0xd	flush page
0xe	flush context
0xf	cache data

t [*program*] (Sun-3 systems only)

Trace the indicated standalone *program*. Works only with programs that do not affect interrupt vectors.

u [*echo*]

u [*port*] [*options*] [*u*]

u [*u*] [*virtual_address*]

With no arguments, display the current I/O device characteristics including: current input device, current output device, BAUD rates for serial ports A and B, an input-to-output echo indicator, and virtual addresses of mapped UART devices. With arguments, set or configure the current I/O device. With the **u** argument (*uu...*), set the I/O device to be the *virtual_address* of a UART device currently mapped.

echo Can be either **e** to enable input to be echoed to the output device, or **ne**, to indicate that input is not echoed.

port Assign the indicated *port* to be the current I/O device. *port* can be one of:

- a** serial port A
- b** serial port B (except on Sun386i systems)
- k** the workstation keyboard
- s** the workstation screen

baud_rate Any legal BAUD rate.

options can be any combination of:

- i** input
- o** output
- u** UART
- e** echo input to output
- ne** do not echo input
- r** reset indicated serial port (a and b ports only)

If either **a** or **b** is supplied, and no *options* are given, the serial port is assigned for both input and output. If **k** is supplied with no *options*, it is assigned for input only. If **s** is supplied with no *options*, it is assigned for output only.

v *virtual_address1 virtual_address2 [size]* (Sun-3 and Sun-4 systems only)

Display the contents of *virtual_address1* (lower) *virtual_address2* (higher) in the format specified by *size*:

- b** byte format (the default)
- w** word format
- l** long word format

Enter return to pause for viewing; enter another return character to resume the display. To terminate the display at any time, press the space bar.

For example, the following command displays the contents of virtual address space from address 0x1000 to 0x2000 in word format:

```
v 1000 2000 W
```

w [*virtual_address*] [*argument*] (Sun-3 and Sun-4 systems only)

Set the execution vector to a predetermined or default routine. Pass *virtual_address* and *argument* to that routine.

To set up a predetermined routine to jump to, a user program must, prior to executing the monitor's **w** command, set the variable `*romp->v_vector_cmd` to be equal to the virtual address of the desired routine. Predetermined routines need not necessarily return control to the monitor.

The default routine, defined by the monitor, prints the user-supplied *vector* according to the format supplied in *argument*. This format can be one of:

- %x** hexadecimal
- %d** decimal

x (Sun-3 and Sun-4 systems only)

Display a menu of extended tests. These diagnostics permit additional testing of such things as the I/O port connectors, video memory, workstation memory and keyboard, and boot device paths.

y c *context_number* (Sun-4 systems only)

y p|s *context_number virtual_address*

Flush the indicated context, context page, or context segment.

- c** flush context *context_number*
- p** flush the page beginning at *virtual_address* within context *context_number*
- s** flush the segment beginning at *virtual_address* within context *context_number*

z [*number*] [*breakpoint_virtual_address*] [*type*] [*len*] (Sun386i systems only)

Set or reset breakpoints for debugging. With no arguments, this command displays the existing breakpoints. The *number* argument is a values from 0 to 3, corresponding to the processor debug registers, DR0 to DR3, respectively. Up to 4 distinct breakpoints can be specified. If *number* is not specified then the monitor chooses a breakpoint number. The *breakpoint_virtual_address* argument specifies the breakpoint address. The *type* argument can be one of:

- x** Instruction Execution breakpoint (the default)
- m** for Data Write only breakpoint
- r** Data Reads and Writes only breakpoint.

The *len* argument can be one of: 'b', 'w', or 'l', corresponding to the breakpoint field length of byte, word, or long-word, respectively. The default is 'b'. Since the breakpoints are set in the on-chip registers, an instruction breakpoint can be placed in ROM code or in code shared by several tasks. If the *number* argument is specified but not *breakpoint_virtual_address*, the corresponding breakpoint is reset.

z [*virtual_address*] (Sun-3 systems only)

Set a breakpoint at *virtual_address* in the address space selected by the **s** command.

FILES

/vmunix

NAME

mount, umount – mount and dismount filesystems

SYNOPSIS

```

/usr/etc/mount [ -p ]
/usr/etc/mount -a[fnv] [ -t type ]
/usr/etc/mount [ -fnrv ] [ -t type ] [ -o options ] filesystem directory
/usr/etc/mount [ -vfn ] [ -o options ] filesystem | directory

/usr/etc/umount [ -t type ] [ -h host ]
/usr/etc/umount -a[v]
/usr/etc/umount [ -v ] filesystem | directory ...

```

DESCRIPTION

mount attaches a named *filesystem* to the filesystem hierarchy at the pathname location *directory*, which must already exist. If *directory* has any contents prior to the **mount** operation, these remain hidden until the *filesystem* is once again unmounted. If *filesystem* is of the form *host:pathname*, it is assumed to be an NFS filesystem (type *nfs*).

umount unmounts a currently mounted filesystem, which can be specified either as a *directory* or a *filesystem*.

mount and **umount** maintain a table of mounted filesystems in */etc/mtab*, described in *fstab(5)*. If invoked without an argument, **mount** displays the contents of this table. If invoked with either a *filesystem* or *directory* only, **mount** searches the file */etc/fstab* for a matching entry, and mounts the filesystem indicated in that entry on the indicated directory.

MOUNT OPTIONS

- p** Print the list of mounted filesystems in a format suitable for use in */etc/fstab*.
- a** All. Attempt to mount all the filesystems described in */etc/fstab*. If a *type* argument is specified with **-t**, mount all filesystems of that type. Filesystems are not necessarily mounted in the order shown in */etc/fstab*.
- f** Fake an */etc/mtab* entry, but do not actually mount any filesystems.
- n** Mount the filesystem without making an entry in */etc/mtab*.
- v** Verbose. Display a message indicating each filesystem being mounted.
- t type** Specify a filesystem type. The accepted types are *4.2*, and *nfs*; see *fstab(5)* for a description of these types.
- r** Mount the specified filesystem read-only. This is a shorthand for:

```
mount -o ro filesystem directory
```

Physically write-protected and magnetic-tape filesystems must be mounted read-only. Otherwise errors occur when the system attempts to update access times, even if no write operation is attempted.

-o options

Specify filesystem *options* —list of comma-separated words from the list below. Some options are valid for all filesystem types, while others apply to a specific type only.

options valid on *all* filesystems:

rw ro	Read/write or read-only.
suid nosuid	Setuid execution allowed or disallowed.
grpuid	Create files with BSD semantics for the propagation of the group ID. Under this option, files inherit the GID of the directory in which they are created, regardless of the directory's set-GID bit.

noauto Do not mount this filesystem that is currently mounted read-only. If the filesystem is not currently mounted, an error results.

The default is 'rw,suid'.

options specific to 4.2 filesystems:

quota|noquota Usage limits are enforced, or are not enforced. The default is **noquota**.

options specific to nfs (NFS) filesystems:

bg|fg If the first attempt fails, retry in the background, or, in the foreground.
retry=*n* The number of times to retry the mount operation.
rsize=*n* Set the read buffer size to *n* bytes.
wsize=*n* Set the write buffer size to *n* bytes.
timeo=*n* Set the NFS timeout to *n* tenths of a second.
retrans=*n* The number of NFS retransmissions.
port=*n* The server IP port number.
soft|hard Return an error if the server does not respond, or continue the retry request until the server responds.
intr Allow keyboard interrupts on hard mounts.
secure Use a more secure protocol for NFS transactions.
acregmin=*n* Hold cached attributes for at least *n* seconds after file modification.
acregmax=*n* Hold cached attributes for no more than *n* seconds after file modification.
acdirmin=*n* Hold cached attributes for at least *n* seconds after directory update.
acdirmax=*n* Hold cached attributes for no more than *n* seconds after directory update.
actimeo=*n* Set *min* and *max* times for regular files and directories to *n* seconds.

Regular defaults are:

```
fg, retry=10000, timeo=7, retrans=3, port=NFS_PORT, hard, \
acregmin=3, acregmax=60, acdirmin=30, acdirmax=60
```

Defaults for *rsize* and *wsize* are set internally by the system kernel.

UMOUNT OPTIONS

- h *host*** Unmount all filesystems listed in */etc/mstab* that are remote-mounted from *host*.
- t *type*** Unmount all filesystems listed in */etc/mstab* that are of a given *type*.
- a** Unmount all filesystems currently mounted (as listed in */etc/mstab*).
- v** Verbose. Display a message indicating each filesystem being unmounted.

NFS FILESYSTEMS

Background vs. Foreground

Filesystems mounted with the **bg** option indicate that **mount** is to retry in the background if the server's mount daemon (**mountd(8c)**) does not respond. **mount** retries the request up to the count specified in the **retry=*n*** option. Once the filesystem is mounted, each NFS request made in the kernel waits **timeo=*n*** tenths of a second for a response. If no response arrives, the time-out is multiplied by 2 and the request is retransmitted. When the number of retransmissions has reached the number specified in the **retrans=*n*** option, a filesystem mounted with the **soft** option returns an error on the request; one mounted with the **hard** option prints a warning message and continues to retry the request.

Read-Write vs. Read-Only

Filesystems that are mounted **rw** (read-write) should use the **hard** option.

Interrupting Processes With Pending NFS Requests

The **intr** option allows keyboard interrupts to kill a process that is hung while waiting for a response on a hard-mounted filesystem.

Secure Filesystems

The `secure` option must be given if the server requires secure mounting for the filesystem.

File Attributes

The attribute cache retains file attributes on the client. Attributes for a file are assigned a time to be flushed. If the file is modified before the flush time, then the flush time is extended by the time since the last modification (under the assumption that files that changed recently are likely to change soon). There is a minimum and maximum flush time extension for regular files and for directories. Setting `actimeo=n` extends flush time by *n* seconds for both regular files and directories.

SYSTEM V COMPATIBILITY**System V File-Creation Semantics**

Ordinarily, when a file is created its GID is set to the effective GID of the calling process. This behavior may be overridden on a per-directory basis, by setting the set-GID bit of the parent directory; in this case, the GID is set to the GID of the parent directory (see `open(2V)` and `mkdir(2)`). Files created on filesystems that are mounted with the `grpuid` option will obey BSD semantics; that is, the GID is unconditionally inherited from that of the parent directory.

EXAMPLES

To mount a local disk:	<code>mount /dev/xy0g /usr</code>
To fake an entry for <code>nd</code> root:	<code>mount -ft 4.2 /dev/nd0 /</code>
To mount all 4.2 filesystems:	<code>mount -at 4.2</code>
To mount a remote filesystem:	<code>mount -t nfs serv:/usr/src /usr/src</code>
To mount a remote filesystem:	<code>mount serv:/usr/src /usr/src</code>
To hard mount a remote filesystem:	<code>mount -o hard serv:/usr/src /usr/src</code>
To save current mount state:	<code>mount -p > /etc/fstab</code>

FILES

<code>/etc/mtab</code>	table of mounted filesystems
<code>/etc/fstab</code>	table of filesystems mounted at boot

SEE ALSO

`mkdir(2)`, `mount(2)`, `unmount(2)`, `open(2V)`, `fstab(5)`, `mtab(5)`, `mountd(8C)`, `nfsd(8)`

BUGS

Mounting filesystems full of garbage crashes the system.

If the directory on which a filesystem is to be mounted is a symbolic link, the filesystem is mounted on *the directory to which the symbolic link refers*, rather than being mounted on top of the symbolic link itself.

NAME

mountd – NFS mount request server

SYNOPSIS

/usr/etc/rpc.mountd [**-n**]

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

mountd is an RPC server that answers file system mount requests. It reads the file **/etc/xtab**, described in **exports(5)**, to determine which file systems are available for mounting by which machines. It also provides information as to what file systems are mounted by which clients. This information can be printed using the **showmount(8)** command.

The **mountd** daemon is normally invoked by **rc(8)**.

OPTIONS

-n Do not check that the clients are root users. Though this option makes things slightly less secure, it does allow older versions (pre-3.0) of client NFS to work.

FILES

/etc/xtab

SEE ALSO

exports(5), **rc(8)**, **showmount(8)**

NAME

named – Internet domain name server

SYNOPSIS

```
/usr/etc/in.named [ -d level ] [ -p port ] [[ -b ] bootfile ]
```

DESCRIPTION

filenamed is the Internet domain name server. It is used by hosts on the DARPA Internet to provide access to the Internet distributed naming database. See RFC 1034 and RFC 1035 for more details. With no arguments *filenamed* reads */etc/named.boot* for any initial data, and listens for queries on a privileged port.

OPTIONS

-d level Print debugging information. *level* is a number indicating the level of messages printed.

-p port Use a different *port* number.

-b bootfile
Use *bootfile* rather than */etc/named.boot*.

EXAMPLE

```

;
;      boot file for name server
;
; type          domain          source file or host
;
domain         berkeley.edu
primary        berkeley.edu  named.db
secondary      cc.berkeley.edu 10.2.0.78 128.32.0.10
cache          named.ca

```

The **domain** line specifies that **berkeley.edu** is the domain of the given server.

The **primary** line states that the file **named.db** contains authoritative data for **berkeley.edu**. The file **named.db** contains data in the master file format, described in RFC 1035, except that all domain names are relative to the origin; in this case, **berkeley.edu** (see below for a more detailed description).

The **secondary** line specifies that all authoritative data under **cc.berkeley.edu** is to be transferred from the name server at **10.2.0.78**. If the transfer fails it will try **128.32.0.10**, and continue for up to 10 tries at that address. The secondary copy is also authoritative for the domain.

The **cache** line specifies that data in **named.ca** is to be placed in the cache (typically such data as the locations of root domain servers). The file **named.ca** is in the same format as **named.db**.

The master file consists of entries of the form:

```

$INCLUDE <filename>
$ORIGIN <domain>
<domain> <opt_ttl> <opt_class> <type> <resource_record_data>

```

where *domain* is '.' for the root, '@' for the current origin, or a standard domain name. If *domain* is a standard domain name that does not end with '.', the current origin is appended to the domain. Domain names ending with '.' are unmodified.

The *opt_ttl* field is an optional integer number for the time-to-live field. It defaults to zero.

The *opt_class* field is currently one token, 'IN' for the Internet.

The *type* field is one of the following tokens; the data expected in the *resource_record_data* field is in parentheses.

A A host address (dotted quad).
NS An authoritative name server (domain).
MX A mail exchanger (domain).

CNAME

The canonical name for an alias (domain).

SOA Marks the start of a zone of authority (5 numbers). (see RFC 1035).

MB A mailbox domain name (domain).

MG A mail group member (domain).

MR A mail rename domain name (domain).

NULL A null resource record (no format or data).

WKS A well know service description (not implemented yet).

PTR A domain name pointer (domain).

HINFO Host information (cpu_type OS_type).

MINFO Mailbox or mail list information (request_domain error_domain).

FILES

<code>/etc/named.boot</code>	name server configuration boot file
<code>/etc/named.pid</code>	the process ID
<code>/var/tmp/named.run</code>	debug output
<code>/var/tmp/named_dump.db</code>	dump of the name servers database

SEE ALSO

kill(1), **signal(3)**, **resolver(3)**, **resolve.conf(5)**

Mockapetris, Paul, *Domain Names - Concepts and Facilities*, RFC 1034, Network Information Center, SRI International, Menlo Park, Calif., November 1987.

Mockapetris, Paul, *Domain Names - Implementation and Specification*, RFC 1035, Network Information Center, SRI International, Menlo Park, Calif., November 1987.

Mockapetris, Paul, *Domain System Changes and Observations*, RFC 973, Network Information Center, SRI International, Menlo Park, Calif., January 1986.

Partridge, Craig, *Mail Routing and the Domain System*, RFC 974, Network Information Center, SRI International, Menlo Park, Calif., January 1986.

NOTES

The following signals have the specified effect when sent to the server process using the **kill(1)** command.

SIGHUP Causes server to read `named.boot` and reload database.

SIGQUIT

Dumps current data base and cache to `/var/tmp/named_dump.db`.

SIGEMT Turns on debugging; each subsequent **SIGEMT** increments debug level.

SIGFPE Turns off debugging completely.

NAME

ncheck – generate names from i-numbers

SYNOPSIS

`/usr/etc/ncheck [-i numbers] [-as] [filesystem]`

DESCRIPTION

Note: For most normal file system maintenance, the function of ncheck is subsumed by fsck(8).

ncheck with no argument generates a pathname versus i-number list of all files on a set of default file systems. Names of directory files are followed by ‘.’

A file system may be specified by the optional *filesystem* argument.

The report is in no useful order, and probably should be sorted.

OPTIONS

-i *numbers*

Report only those files whose *i-numbers* follow.

-a Print the names ‘.’ and ‘..’, which are ordinarily suppressed.

-s Report only special files and files with set-user-ID mode. This is intended to discover concealed violations of security policy.

SEE ALSO

sort(1V), dcheck(8), fsck(8), ickcheck(8)

DIAGNOSTICS

When the filesystem structure is improper, ‘??’ denotes the “parent” of a parentless file and a pathname beginning with ‘...’ denotes a loop.

NAME

ndbootd – ND boot block server

SYNOPSIS

ndbootd [**-dv**]

DESCRIPTION

ndbootd sends boot blocks to diskless Sun-2 system clients that request them using the (now obsolete) ND protocol. This server uses the boot block contained in the file **/tftpboot/sun2.bb**. A client must appear in the **ethers(5)** and **hosts(5)** databases, in order for the request to be served. In determining whether to serve the client, **ndbootd** checks the **/tftpboot** directory for a file whose name is the client's IP address in hexadecimal notation. For example, if the file **/tftpboot/C00901AD** exists, the machine at IP address 192.9.1.173 can be served. This file normally contains the boot program that is sent to the client by **tftpd(8C)**.

Only root can invoke **ndbootd**.

OPTIONS

- d** Debug. Display information about ignored packets, retransmissions, and address translation.
- v** Verbose. Show a detailed listing of packets sent and received, etc.

If either option is used, all output is sent to the invoking terminal. Otherwise, error output (if any) appears on the console.

FILES

/tftpboot	bootfiles directory
/tftpboot/sun2.bb	boot blocks
/tftpboot/????????	boot programs for clients

SEE ALSO

ethers(5), **hosts(5)**, **boot(8S)**, **tftpd(8C)**

NAME

netconfig – PNP boot service

SYNOPSIS

/single/netconfig [*-e*] [*-n*]

AVAILABILITY

Sun386i systems only.

DESCRIPTION

netconfig is used both for automatic installation of new diskful systems, and during routine booting of all systems. The sequence of actions taken by **netconfig** depends on which of these situations is in effect, but it always sets the hostname, domainname, time, timezone, and interface IP address. If the system is newly installed on the network, it does more, perhaps interrogating the user about system configuration.

netconfig is invoked with the *-e* option from the */etc/rc.boot* script.

Invoked without options, *netconfig* may perform PNP set up, including set up of files, passwords, and secure RPCs. Unless *-n* is specified, it writes */etc/net.conf*, which is read later by *rc.boot*. This includes the VERBOSE flag, derived from NVRAM data, which controls the verbosity of the commands in *rc.boot*.

Routine Booting

Boot servers use information stored locally in YP maps rather than acquiring it over the network, except that they get the time from the *timehost* system if it is up. The following describes the steps taken by boot clients: diskful clients, diskless clients, and network clients.

Boot clients first invoke *rarp* to acquire an IP address. This is followed by a ICMP Netmask request to obtain the IP subnetwork mask, and then a PNP_WHOAMI RPC to determine the system's name, YP domain, and time zone. Then the systems clock is set using the RFC 868 time service. If PNP_WHOAMI fails, a PNP_SETUP sequence is followed by set up of */etc/passwd* and other files.

OPTIONS

- e* Check shell environment variables. This option is specified during routine boot. HOSTNAME and DOMAINNAME are used to determine if the system is a YP server using local YP maps. Otherwise, if NETWORKED is YES, **netconfig** probes the network for network configuration. MUST_SETUP requires writing */etc/passwd* and other files for setup in restricted network environments.
- n* Used in conjunction with *'-e'*, this does not probe the network for anything but just sets the hostname and domainname of the system from the environment variables HOSTNAME and DOMAINNAME respectively. Does not write the */etc/net.conf* file.

FILES

/var/yp/domainname/netmasks
/var/yp/domainname/hosts

SEE ALSO

pnpboot(8C), **pnpd(8C)**, **rarpd(8C)** **pnp(3R)**

NAME

netstat – show network status

SYNOPSIS

netstat [**-aAn**] [**-f** *address_family*] [**system**] [**core**]

netstat [**-n**] [**-s**] [**-m** | **-i** | **-h** | **-r**] [**-f** *address_family*] [**system**] [**core**]

netstat [**-n**] [**-I** *interface*] *interval* [**system**] [**core**]

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

netstat displays the contents of various network-related data structures in various formats, depending on the options you select.

The first form of the command displays a list of active sockets for each protocol. The second form selects one from among various other network data structures. The third form displays running statistics of packet traffic on configured network interfaces; the *interval* argument indicates the number of seconds in which to gather statistics between displays.

The default value for the **system** argument is */vmunix*; for *core*, the default is */dev/kmem*.

OPTIONS

- a** Show the state of all sockets; normally sockets used by server processes are not shown.
- A** Show the address of any protocol control blocks associated with sockets; used for debugging.
- f** *address_family*
Limit statistics or address control block reports to those of the specified *address_family*, which can be one of:
 - inet** For the AF_INET address family, or
 - unix** For the AF_UNIX family.
- h** Show the state of the IMP host table. (This does not work in an environment where the IMP host tables do not exist.)
- i** Show the state of interfaces that have been auto-configured. Interfaces that are statically configured into a system, but not located at boot time, are not shown.
- I** *interface*
Highlight information about the indicated *interface* in a separate column; the default (for the third form of the command) is the interface with the most traffic since the system was last rebooted. *interface* can be any valid interface listed in the system configuration file, such as *ie0* or *le0*.
- m** Show the statistics recorded by management routines for the network's private buffer pool.
- n** Show network addresses as numbers. netstat normally displays addresses as symbols. This option may be used with any of the display formats.
- r** Show the routing tables. (When **-s** is also present, show routing statistics instead.)
- s** Show per-protocol statistics. When used with the **-r** option, show routing statistics.
- t** Replace queue length information with timer information.

DISPLAYS**Active Sockets (First Form)**

The display for each active socket shows the local and remote address, the send and receive queue sizes (in bytes), the protocol, and the internal state of the protocol.

The symbolic format normally used to display socket addresses is either:

hostname.port

when the name of the host is specified, or:

network.port

if a socket address specifies a network but no specific host. Each *hostname* and *network* is shown according to its entry in the */etc/hosts* or the */etc/networks* file, as appropriate.

If the network or hostname for an address is not known (or if the *-n* option is specified), the numerical network address is shown. Unspecified, or "wildcard", addresses and ports appear as "*". (For more information regarding the Internet naming conventions, refer to *inet(3N)*).

TCP Sockets

The possible state values for TCP sockets are as follows:

CLOSED	Closed: the socket is not being used.
LISTEN	Listening for incoming connections.
SYN_SENT	Actively trying to establish connection.
SYN_RECEIVED	Initial synchronization of the connection under way.
ESTABLISHED	Connection has been established.
CLOSE_WAIT	Remote shut down: waiting for the socket to close.
FIN_WAIT_1	Socket closed, shutting down connection.
CLOSING	Closed, then remote shutdown: awaiting acknowledgement.
LAST_ACK	Remote shut down, then closed: awaiting acknowledgement.
FIN_WAIT_2	Socket closed, waiting for shutdown from remote.
TIME_WAIT	Wait after close for remote shutdown retransmission.

Network Data Structures (Second Form)

The form of the display depends upon which of the *-m*, *-i*, *-h* or *-r*, options you select. (If you specify more than one of these options, *netstat* selects one in the order listed here.)

Routing Table Display

The routing table display lists the available routes and the status of each. Each route consists of a destination host or network, and a gateway to use in forwarding packets. The *flags* column shows the status of the route (U if "up"), whether the route is to a gateway (G), and whether the route was created dynamically by a redirect (D).

Direct routes are created for each interface attached to the local host; the gateway field for such entries shows the address of the outgoing interface.

The *refcnt* column gives the current number of active uses per route. (Connection-oriented protocols normally hold on to a single route for the duration of a connection, whereas connectionless protocols obtain a route while sending to the same destination.)

The *use* column displays the number of packets sent per route.

The *interface* entry indicates the network interface utilized for the route.

Cumulative Traffic Statistics (Third Form)

When the *interval* argument is given, *netstat* displays a table of cumulative statistics regarding packets transferred, errors and collisions, the network addresses for the interface, and the maximum transmission unit ("mtu"). The first line of data displayed, and every 24th line thereafter, contains cumulative statistics from the time the system was last rebooted. Each subsequent line shows incremental statistics for the *interval* (specified on the command line) since the previous display.

SEE ALSO

hosts(5), *networks(5)*, *protocols(5)*, *services(5)*, *iostat(8)*, *trpt(8C)*, *vmstat(8)*

BUGS

The notion of errors is ill-defined. Collisions mean something else for the IMP.

The kernel's tables can change while netstat is examining them, creating incorrect or partial displays.

NAME

newaliases – rebuild the data base for the mail aliases file

SYNOPSIS

newaliases

DESCRIPTION

newaliases rebuilds the random access data base for the mail aliases file **/etc/aliases**. It is run automatically by **sendmail(8)** (in the default configuration) whenever a message is sent.

FILES

/etc/aliases

SEE ALSO

aliases(5), **sendmail(8)**

NAME

newfs – construct a new file system

SYNOPSIS

`/usr/etc/newfs [-nNv] [mkfs-options] block-special-file`

DESCRIPTION

newfs is a “friendly” front-end to the **mkfs(8)** program. On Sun systems, the disk type is determined by reading the disk label for the specified *block-special-file*.

block-special-file is the name of a block special device residing in `/dev`. If you want to make a file system on `sd0`, you can specify `sd0` `rsd0` or `/dev/rsd0`; if you only specify `sd0`, **newfs** will find the proper device.

newfs then calculates the appropriate parameters to use in calling **mkfs**, builds the file system by forking **mkfs** and, if the file system is a root partition, installs the necessary bootstrap programs in its initial 16 sectors.

OPTIONS

- n** Do not install the bootstrap programs.
- N** Print out the file system parameters without actually creating the file system.
- v** Verbose. **newfs** prints out its actions, including the parameters passed to **mkfs**.

mkfs-options

Options that override the default parameters passed to **mkfs(8)** are:

-b *block-size*

The block size of the file system in bytes.

-c *#cylinders/group*

The number of cylinders per cylinder group in a file system. The default value used is 16.

-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.

-f *frag-size*

The fragment size of the file system in bytes.

-i *bytes/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.

-m *free-space%*

The percentage of space reserved from normal users; the minimum free space threshold. The default value used is 10%.

-o *optimization*

(*space* or *time*). The file system can either be instructed to try to minimize the time spent allocating blocks, or to try to minimize the space fragmentation on the disk. If the minimum free space threshold (as specified by the **-m** option) is less than 10%, the default is to optimize for space; if the minimum free space threshold is greater than or equal to 10%, the default is to optimize for time.

-r *revolutions/minute*

The speed of the disk in revolutions per minute (normally 3600).

-s *size* The size of the file system in sectors.

-t #tracks/cylinder

The number of tracks per cylinders on the disk.

FILES

/usr/etc/mkfs to actually build the file system
/usr/mdec for boot strapping programs */dev*

SEE ALSO

fs(5), fsck(8), mkfs(8), tuneefs(8)

System and Network Administration

NAME

newkey – create a new key in the publickey database

SYNOPSIS

newkey [**-h** *hostname*] [**-u** *username*]

DESCRIPTION

newkey is normally run by the network administrator on the YP master machine in order to establish public keys for users and super-users on the network. These keys are needed for using secure RPC or secure NFS.

newkey will prompt for the login password of the given username and then create a new public/secret key pair in */etc/publickey* encrypted with the login password of the given user.

Use of this program is not required: users may create their own keys using **chkey** (1).

OPTIONS

-u *username* Create a new public key for the given username. Prompts for the Yellow Pages password of the given username.

-h *hostname* Create a new public key for the super-user at the given hostname. Prompts for the root password of the given hostname.

SEE ALSO

keylogin(1), **chkey**(1), **publickey**(5), **keyserv**(8)

NAME

nfsd, *biod* – NFS daemons

SYNOPSIS

/usr/etc/nfsd [*nservers*]

/usr/etc/biod [*nservers*]

DESCRIPTION

nfsd starts the daemons that handle client filesystem requests. *nservers* is the number of file system request daemons to start. This number should be based on the load expected on this server. Four seems to be a good number.

biod starts *nservers* asynchronous block I/O daemons. This command is used on a NFS client to buffer cache handle read-ahead and write-behind. The magic number for *nservers* in here is also four.

When a file that is opened by a client is unlinked (by the server), a file with a name of the form *.nfsXXX* (where *XXX* is a number) is created by the client. When the open file is closed, the *.nfsXXX* file is removed. If the client crashes before the file can be closed, the *.nfsXXX* file is not removed.

FILES

.nfsXXX client machine pointer to an open-but-unlinked file

SEE ALSO

exports(5), *mountd*(8C)

NAME

nfsstat – Network File System statistics

SYNOPSIS

nfsstat [**-csnrz**]

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

nfsstat displays statistical information about the NFS (Network File System) and RPC (Remote Procedure Call), interfaces to the kernel. It can also be used to reinitialize this information. If no options are given the default is

nfsstat -csnr

That is, display everything, but reinitialize nothing.

OPTIONS

- c** Display client information. Only the client side NFS and RPC information will be printed. Can be combined with the **-n** and **-r** options to print client NFS or client RPC information only.
- s** Display server information.
- n** Display NFS information. NFS information for both the client and server side will be printed. Can be combined with the **-c** and **-s** options to print client or server NFS information only.
- r** Display RPC information.
- z** Zero (reinitialize) statistics. This option is for use by the super-user only, and can be combined with any of the above options to zero particular sets of statistics after printing them.

DISPLAYS

The server RPC display includes the fields:

calls	total number of RPC calls received
badcalls	total number of calls rejected
nullrecv	number of times no RPC packet was available when trying to receive
badlen	number of packets that were too short
xdr call	number of packets that had a malformed header

The server NFS display shows the number of NFS calls received (**calls**) and rejected (**badcalls**), and the counts and percentages for the various calls that were made.

The client RPC display includes the following fields:

calls	total number of RPC calls sent
badcalls	total of calls rejected by a server
retrans	number of times a call had to be retransmitted
badxid	number of times a reply did not match the call
timeout	number of times a call timed out
wait	number of times a call had to wait on a busy CLIENT handle
newcred	number of times authentication information had to be refreshed

The client NFS display shows the number of calls sent and rejected, as well as the number of times a CLIENT handle was received (**nclget**), the number of times a call had to sleep while awaiting a handle (**nclsleep**), as well as a count of the various calls and their respective percentages.

FILES

/vmunix	system namelist
/dev/kmem	kernel memory

NAME

nslookup – query name servers interactively

SYNOPSIS

nslookup [*-l*] [*address*]

DESCRIPTION

nslookup is an interactive program to query ARPA Internet domain name servers. The user can contact servers to request information about a specific host or print a list of hosts in the domain.

OPTIONS

-l Use the local host's name server instead of the servers in */etc/resolve.conf*. (If */etc/resolve.conf* does not exist or does not contain server information, the **-l** option does not have any effect).

address Use the name server on the host machine with the given Internet address.

USAGE**Overview**

The Internet domain name-space is tree-structured, with four top-level domains at present:

COM	commercial establishments
EDU	educational institutions
GOV	government agencies
MIL	MILNET hosts

If you are looking for a specific host, you need to know something about the host's organization in order to determine the top-level domain it belongs to. For instance, if you want to find the Internet address of a machine at UCLA, do the following:

1. Connect with the root server using the **root** command. The root server of the name space has knowledge of the top-level domains.
2. Since UCLA is a university, its domain name is **ucla.edu**. Connect with a server for the **ucla.edu** domain with the command **serverucla.edu**. The response will print the names of hosts that act as servers for that domain. Note: the root server does not have information about **ucla.edu**, but knows the names and addresses of hosts that do. Once located by the root server, all future queries will be sent to the UCLA name server.
3. To request information about a particular host in the domain (for instance, **locus**), just type the host name. To request a listing of hosts in the UCLA domain, use the **ls** command. The **ls** command requires a domain name (in this case, **ucla.edu**) as an argument.

Note: if you are connected with a name server that handles more than one domain, all lookups for host names must be fully specified with its domain. For instance, the domain **harvard.edu** is served by **seismo.css.gov**, which also services the **css.gov** and **cornell.edu** domains. A lookup request for the host **aiken** in the **harvard.edu** domain must be specified as **aiken.harvard.edu**. However, the

set domain= name

and

set defname

commands can be used to automatically append a domain name to each request.

After a successful lookup of a host, use the **finger** command to see who is on the system, or to finger a specific person. To get other information about the host, use the

set querytype= value

command to change the type of information desired and request another lookup. (**finger** requires the type to be A.)

Commands

Commands may be interrupted at any time by typing `^C`. To exit, type `^D` (EOF). The command line length must be less than 80 characters. Note: an unrecognized command will be interpreted as a host name.

host [*server*]

Look up information for *host* using the current default server or using *server* if it is specified.

server domain

lserver domain

Change the default server to *domain*. *lserver* uses the initial server to look up information about *domain* while *server* uses the current default server. If an authoritative answer can't be found, the names of servers that might have the answer are returned.

root Changes the default server to the server for the root of the domain name space. Currently, the host `sri-nic.arpa` is used; this command is a synonym for '`lserver sri-nic.arpa`'.) The name of the root server can be changed with the `set root` command.

finger [*name*]

Connect with the finger server on the current host, which is defined by a previous successful lookup for a host's address information (see the `set querytype=A` command). As with the shell, output can be redirected to a named file using `>` and `>>`.

ls [`-ah`]

List the information available for *domain*. The default output contains host names and their Internet addresses. The `-a` option lists aliases of hosts in the domain. The `-h` option lists CPU and operating system information for the domain. As with the shell, output can be redirected to a named file using `>` and `>>`. When output is directed to a file, hash marks are printed for every 50 records received from the server.

view*filename*

Sort and list the output of the `ls` command with `more(1)`.

help

? Print a brief summary of commands.

set*keyword*[=*value*]

This command is used to change state information that affects the lookups. Valid keywords are:

all Prints the current values of the various options to `set`. Information about the current default server and host is also printed.

[no]deb[*ug*]

Turn debugging mode on. A lot more information is printed about the packet sent to the server and the resulting answer. The default is `nodebug`.

[no]def[*name*]

Append the default domain name to every lookup. The default is `nodename`.

do[*main*]=*filename*

Change the default domain name to *name*. The default domain name is appended to all lookup requests if `defname` option has been set. The default is the value in `/etc/resolve.conf`.

q[*querytype*]=*value*

Change the type of information returned from a query to one of:

A The host's Internet address (the default).

CNAME

The canonical name for an alias.

HINFO The host CPU and operating system type.

MD The mail destination.

MX The mail exchanger.
MB The mailbox domain name.
MG The mail group member.
MINFO The mailbox or mail list information.

(Other types specified in the RFC883 document are valid, but are not very useful.)

[no]recurse

Tell the name server to query other servers if it does not have the information. The default is *recurse*.

ret[ry]=*count*

Set the number of times to retry a request before giving up to *count*. When a reply to a request is not received within a certain amount of time (changed with **set timeout**), the request is resent. The default is *count* is 2.

ro[ot]=*host*

Change the name of the root server to *host*. This affects the **root** command. The default root server is *sri-nic.arpa*.

t[timeout]=*interval*

Change the time-out for a reply to *interval* seconds. The default *interval* is 10 seconds.

[no]v[c]

Always use a virtual circuit when sending requests to the server. The default is *novc*.

DIAGNOSTICS

If the lookup request was not successful, an error message is printed. Possible errors are:

Time-out

The server did not respond to a request after a certain amount of time (changed with **set timeout= value**) and a certain number of retries (changed with **set retry= value**).

No information

Depending on the query type set with the **set querytype** command, no information about the host was available, though the host name is valid.

Non-existent domain

The host or domain name does not exist.

Connection refused

Network is unreachable

The connection to the name or finger server could not be made at the current time. This error commonly occurs with **finger** requests.

Server failure

The name server found an internal inconsistency in its database and could not return a valid answer.

Refused

The name server refused to service the request.

The following error should not occur and it indicates a bug in the program.

Format error

The name server found that the request packet was not in the proper format.

FILES

/etc/resolve.conf initial domain name and name server addresses.

SEE ALSO

resolver(3), **resolve.conf(5)**, **named(8C)**, RFC 882, RFC 883

NAME

pac – printer/plotter accounting information

SYNOPSIS

`/usr/etc/pac [-cps] [-Pprinter] [-pprice] [filename...]`

DESCRIPTION

pac reads the printer/plotter accounting files, accumulating the number of pages (the usual case) or feet (for raster devices) of paper consumed by each user, and printing out how much each user consumed in pages or feet and dollars. If any *filenames* are specified, then statistics are only printed for those users; usually, statistics are printed for every user who has used any paper.

OPTIONS

- c Sorted the output by cost; usually the output is sorted alphabetically by name.
- pprice Use the value *price* for the cost in dollars instead of the default value of 0.02.
- r Reverse the sorting order.
- s Summarize the accounting information on the summary accounting file; this summary is necessary since on a busy system, the accounting file can grow by several lines per day.
- Pprinter Do accounting for the named *printer*. Normally, accounting is done for the default printer (site dependent) or the value of the PRINTER environment variable is used.

FILES

<code>/var/adm/?acct</code>	raw accounting files
<code>/var/adm/?_sum</code>	summary accounting files

BUGS

The relationship between the computed price and reality is as yet unknown.

NAME

ping – send ICMP ECHO_REQUEST packets to network hosts

SYNOPSIS

`/usr/etc/ping host [timeout]`

`/usr/etc/ping [-s] [-rv] host [packet-size] [count]`

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

ping utilizes the ICMP protocol's mandatory ECHO_REQUEST datagram to elicit an ICMP ECHO_RESPONSE from the specified *host*, or network gateway. ECHO_REQUEST datagrams, or "pings," have an IP and ICMP header, followed by a *struct* timeval, and then an arbitrary number of bytes to pad out the packet. If *host* responds, **ping** will print *host* is alive on the standard output and exit. Otherwise after *timeout* seconds, it will write **no answer from host**. The default value of *timeout* is 20 seconds.

When the `-s` flag is specified, **ping** sends one datagram per second, and prints one line of output for every ECHO_RESPONSE that it receives. No output is produced if there is no response. In this second form, **ping** computes round trip times and packet loss statistics; it displays a summary of this information upon termination or timeout. The default datagram packet size is 64 bytes, or you can specify a size with the *packet-size* command-line argument. If an optional *count* is given, **ping** sends only that number of requests.

When using **ping** for fault isolation, first 'ping' the local host to verify that the local network interface is running.

OPTIONS

- `-r` Bypass the normal routing tables and send directly to a host on an attached network. If the host is not on a directly-attached network, an error is returned. This option can be used to **ping** a local host through an interface that has been dropped by the router daemon, see **routed(8C)**.
- `-v` Verbose output. List any ICMP packets, other than ECHO_RESPONSE, that are received.

SEE ALSO

icmp(4P), **ifconfig(8C)**, **netstat(8C)**, **rpcinfo(8C)**, **spray(8C)**

NAME

pnpboot, pnp.s386 – pnp diskless boot service

SYNOPSIS — Sun386i

/tftpboot/pnp.s386

AVAILABILITY

Sun386i systems only.

DESCRIPTION

pnp.s386 is a level 2 boot program that requests actions necessary to set up a diskless workstation on the network.

The PNP diskless boot service is used by diskless workstations at installation time to locate a server that will configure the diskless client.

The last steps of the level 1 boot (from the PROM) are to load the level 2 program through **rarpd(8C)** and **tftpd(8C)**. The first step in the boot sequence is RARP to acquire an IP address. This is followed by TFTP service calls to acquire the **pnp.sun*** program file needed for the client's architecture. A **PNP_ACQUIRE** RPC is then broadcast to locate a server willing to configure the diskless client.

A **PNP_SETUP** is issued to the server which returns one of three statuses: success, failure, or **in_progress**. As long as the server responds with a status of **in_progress** the client will periodically issue a **PNP_POLL** until the status changes to either success or failure.

The last step is to reboot the client. This goes through a RARP, TFTP, BOOT sequence, with the boot using the normal **boot.sun*** file and **bootparamd(8)** service.

The system will have been set up using the IP address returned in the first step and a system name will have been assigned.

FILES

/tftpboot/pnp.sun*

SEE ALSO

bootparam(3R), **bootparams(5)**, **boot(8S)**, **bootparamd(8)**, **netconfig(8C)**, **pnpd(8C)**, **ipallocald(8C)**, **rarpd(8C)**, **tftpd(8C)**

NAME

pnpd – PNP daemon

SYNOPSIS

/usr/etc/rpc.pnpd

AVAILABILITY

Sun386i systems only.

DESCRIPTION

pnpd is used during routine booting of systems to determine their network configuration, and by new systems to configure themselves on a network. **pnpd** adds and removes diskless clients of the boot server on which it is running. The **pnpd** daemon is normally invoked in **rc.local**. The RPCs are used by **netconfig(8C)**, **pnp.s386** (see **pnpboot(8C)**), and **client(8)**.

The **bootserver** YP map specifies limits on server capacity and default swap size.

FILES

/export/exec/arch
symbolic link to **/export/exec/arch.release**
/export/exec/arch.release
symbolic link to **/usr** for the architecture
/export/exec/arch.release/boot
root binaries

SEE ALSO

pnp(3R), **client(8)**, **ipallocald(8C)**, **netconfig(8C)**, **pnpboot(8C)**

NAME

portmap – DARPA port to RPC program number mapper

SYNOPSIS

/usr/etc/rpc.portmap

DESCRIPTION

portmap is a server that converts RPC program numbers into DARPA protocol port numbers. It must be running in order to make RPC calls.

When an RPC server is started, it will tell **portmap** what port number it is listening to, and what RPC program numbers it is prepared to serve. When a client wishes to make an RPC call to a given program number, it will first contact **portmap** on the server machine to determine the port number where RPC packets should be sent.

Normally, standard RPC servers are started by **inetd(8C)**, so **portmap** must be started before **inetd** is invoked.

SEE ALSO

inetd.conf(5), **rpcinfo(8)**, **inetd(8)**

BUGS

If **portmap** crashes, all servers must be restarted.

NAME

praudit – print contents of an audit trail file

SYNOPSIS

praudit [**-lrs**] [**-ddel**] [*filename ...*]

AVAILABILITY

This program is available with the *Security* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

praudit reads the listed *filenames* (or standard input, if no *files* are specified) and interprets the data as audit trail records as defined in **audit_control(5)**. By default, times, security labels, user and group IDs are converted to their ASCII representation. Record type and event fields are converted to long ASCII representation.

OPTIONS

- l** Print records one line per record. The record type and event fields are always converted to their short ASCII representation.
- r** Print records in their raw form. Times, security labels, user IDs, group IDs, record types, and events are displayed as integers. Currently, in SunOS 4.0, labels are not used and are displayed as zero in this mode. This option and the **-s** option are exclusive. If both are used, a format usage error message is output.
- s** Print records in their short form. All numeric fields are converted to ASCII and displayed. The short ASCII representations for the record type and event fields are used. Security labels are displayed in their short representation. Again, labels are not currently used. This option and the **-r** option are exclusive. If both are used, a format usage error message is output.
- ddel** Use *del* as the field delimiter instead of the default delimiter, which is the comma. If *del* has special meaning for the shell, it must be quoted. The maximum size of a delimiter is four characters.

FILES

/etc/passwd

SEE ALSO

audit(2), **setuseraudit(2)**, **getauditflags(3)**, **audit_control(5)**

NAME

pstat – print system facts

SYNOPSIS

`/usr/etc/pstat [-afipSsT] [-u pid] [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 `/vmunix` unless *system* is specified.

OPTIONS

-a Under **-p**, describe all process slots rather than just active ones.

-f Print the open file table with these headings:

LOC	The memory address of this table entry.
TYPE	The type of object the file table entry points to.
FLG	Miscellaneous state variables encoded thus: <ul style="list-style-type: none"> R open for reading W open for writing A open for appending S shared lock present X exclusive lock present I signal prgp when data ready
CNT	Number of processes that know this open file.
MSG	Number of references from message queue.
DATA	The location of the vnode table entry or socket for this file.
OFFSET	The file offset (see <code>lseek(2)</code>).

-i Print the inode table including the associated vnode entries with these headings:

ILOC	The memory address of this table entry.
IFLAG	Miscellaneous inode state variables encoded thus: <ul style="list-style-type: none"> A inode access time must be corrected C inode change time must be corrected L inode is locked R inode is being referenced U update time (<code>fs(5)</code>) must be corrected W wanted by another process (<code>L</code> flag is on)
IDEVICE	Major and minor device number of file system in which this inode resides.
INO	I-number within the device.
MODE	Mode bits in octal, see <code>chmod(2)</code> .
NLK	Number of links to this inode.
UID	User ID of owner.
SIZE/DEV	Number of bytes in an ordinary file, or major and minor device of special file.
VFLAG	Miscellaneous vnode state variables encoded thus: <ul style="list-style-type: none"> R root of its file system S shared lock applied E exclusive lock applied Z process is waiting for a shared or exclusive lock
CNT	Number of open file table entries for this vnode.
SHC	Reference count of shared locks on the vnode.
EXC	Reference count of exclusive locks on the vnode (this may be '> 1' if, for example, a file descriptor is inherited across a fork).
TYPE	Vnode file type, either VNON (no type), VREG (regular), VDIR (directory), VBLK (block device), VCHR (character device), VLNK (symbolic link), VSOCK (socket), VFIFO (named pipe), or VBAD (bad).

-p Print process table for active processes with these headings:

LOC	The memory address of this table entry.
S	Run state encoded thus: <ul style="list-style-type: none"> 0 no process 1 awaiting an event 2 (abandoned state) 3 runnable 4 being created 5 being terminated 6 stopped (by signal or under trace)
F	Miscellaneous state variables, ORed together (hexadecimal): <ul style="list-style-type: none"> 0000001 loaded 0000002 a system process (scheduler or page-out daemon) 0000004 locked for swap out 0000008 swapped out during process creation 0000010 process is being traced 0000020 tracing parent has been told that process is stopped 0000040 user settable lock in memory 0000080 in page-wait 0000100 prevented from swapping during fork(2) 0000200 will restore old mask after taking signal 0000400 exiting 0000800 doing physical I/O 0001000 process resulted from a vfork(2) which is not yet complete 0002000 another flag for vfork(2) 0004000 process has no virtual memory, as it is a parent in the context of vfork(2) 0008000 process is demand paging pages from its executable image vnode 0010000 process has advised of sequential VM behavior with vadvise(2) 0020000 process has advised of random VM behavior with vadvise(2) 0080000 process is a session process group leader 0100000 process is tracing another process 0200000 process needs a profiling tick 0400000 process is scanning descriptors during select 4000000 process has done record locks 8000000 process is having its system calls traced
PRI	Scheduling priority, see <code>getpriority(2)</code> .
SIG	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 <code>getpriority(2)</code> .
PGRP	Process number of root of process group.
PID	The process ID number.
PPID	The process ID of parent process.
RSS	Resident set size — the number of physical page frames allocated to this process.
SRSS	RSS at last swap (0 if never swapped).

SIZE The size of the process image. That is, the sum of the data and stack segment sizes, not including the sizes of any shared libraries.

WCHAN Wait channel number of a waiting process.

LINK Link pointer in list of runnable processes.

-S Print the streams table with these headings:

LOC The memory address of this table entry.

WRQ The address of this stream's write queue.

VNODE The address of this stream's vnode.

DEVICE Major and minor device number of device to which this stream refers.

PGRP This stream's process group number.

FLG Miscellaneous stream state variables encoded thus:

- I waiting for `ioctl()` to finish
- R read/recvmmsg is blocked
- W write/putmsg is blocked
- P priority message is at stream head
- H device has been "hung up" (`M_HANGUP`)
- O waiting for open to finish
- M stream is linked under multiplexor
- D stream is in message-discard mode
- N stream is in message-nondiscard mode
- E fatal error has occurred (`M_ERROR`)
- T waiting for queue to drain when closing
- 2 waiting for previous `ioctl()` to finish before starting new one
- 3 waiting for acknowledgment for `ioctl()`
- B stream is in non-blocking mode
- A stream is in asynchronous mode
- o stream uses old-style no-delay mode
- S stream has had `TOSTOP` set
- C `VTIME` clock running
- V `VTIME` timer expired
- r collision on `select()` for reading
- w collision on `select()` for writing
- e collision on `select()` for exceptional condition

The queues on the write and read sides of the stream are listed for each stream. Each queue is printed with these headings:

NAME The name of the module or driver for this queue.

COUNT The approximate number of bytes on this queue.

FLG Miscellaneous state variables encoded thus:

- E queue is enabled to run
- R someone wants to get from this queue when it becomes non-empty
- W someone wants to put on this queue when it drains
- F queue is full
- N queue should not be enabled automatically by a `putq`

MINPS The minimum packet size for this queue.

MAXPS The maximum packet size for this queue, or `INF` if there is no maximum.

HIWAT The high-water mark for this queue.

LOWAT The low-water mark for this queue.

- s** Print information about swap space usage:
- allocated:** The amount of swap space (in bytes) allocated to private pages.
 - reserved:** The number of swap space bytes not currently allocated, but claimed by memory mappings that have not yet created private pages.
 - used:** The total amount of swap space, in bytes, that is either allocated or reserved.
 - available:** The total swap space, in bytes, that is currently available for future reservation and allocation.
- T** Print the number of used and free slots in the several system tables. This is useful for checking to see how full system tables have become if the system is under heavy load. Shows both used and cached inodes.
- u *pid*** Print information about the process with ID *pid*.

FILES

/vmunix **namelist**
/dev/kmem **default source of tables**

SEE ALSO

ps(1), chmod(2), fork(2), lseek(2), getpriority(2), stat(2), vadvice(2), vfork(2), fs(5), iostat(8), vmstat(8)

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

pwck – check password database entries

SYNOPSIS

`/usr/etc/pwck [file]`

DESCRIPTION

Note: Optional Software (System V Option). Refer to *Installing the SunOS* for information on how to install this command.

pwck checks that a file in `passwd(5)` does not contain any errors; it checks the `/etc/passwd` file by default.

FILES

`/etc/passwd`

DIAGNOSTICS**Too many/few fields**

An entry in the password file does not have the proper number of fields.

No login name

The login name field of an entry is empty.

Bad character(s) in login name

The login name in an entry contains characters other than lower-case letters and digits.

First char in login name not lower case alpha

The login name in an entry does not begin with a lower-case letter.

Login name too long

The login name in an entry has more than 8 characters.

Invalid UID

The user ID field in an entry is not numeric or is greater than 65535.

Invalid GID

The group ID field in an entry is not numeric or is greater than 65535.

No login directory

The login directory field in an entry is empty.

Login directory not found

The login directory field in an entry refers to a directory that does not exist.

Optional shell file not found.

The login shell field in an entry refers to a program or shell script that does not exist.

No netgroup name

The entry is a Yellow Pages entry referring to a netgroup, but no netgroup is present.

Bad character(s) in netgroup name

The netgroup name in a Yellow Pages entry contains characters other than lower-case letters and digits.

First char in netgroup name not lower case alpha

The netgroup name in a Yellow pages entry does not begin with a lower-case letter.

SEE ALSO

`group(5)`, `passwd(5)`

NAME

pwdauthd – server for authenticating passwords

SYNOPSIS

/usr/etc/rpc.pwdauthd

AVAILABILITY

This program is available with the *Security* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

pwdauthd is a server that determines authentication for users and groups. It handles authentication requests from **pwdauth(3)** and **grpauth**. Communication to and from **pwdauthd** is by means of RPC calls. The server is passed a *filename* and a *password*. It returns an integer value that specifies whether the *password* is valid. The possible return values are **PWA_VALID** if the name is valid, **PWA_INVALID** if the name is invalid, and **PWA_UNKNOWN** if validity cannot be determined because no adjunct files are present.

If **pwdauthd** is serving **pwdauth**, it determines whether the **passwd.adjunct** file exists. If not, it returns **PWA_UNKNOWN**. In this case, **pwdauth** knows to check the **/etc/passwd** file. Otherwise, the server calls **getpwanam** (see **getpwaent(3)**) to get the entry for *filename* in either the local or the Yellow Pages file for **passwd.adjunct**. If the encrypted password guess matches the encrypted password from the file, **pwdauthd** returns **PWA_VALID**. If the passwords do not match, it returns **PWA_INVALID**.

If **pwdauthd** is serving **grpauth**, it determines whether the **group.adjunct** file exists. If not, it returns **PWA_UNKNOWN**. In this case, **grpauth** knows to check the **/etc/group** file. Otherwise, the server calls **getgranam** (see **getgraent(3)**) to get the entry for *filename* in either the local or the Yellow Pages file for **group.adjunct**. If the encrypted password guess matches the encrypted password from the file, **pwdauthd** returns **PWA_VALID**. If the passwords do not match, it returns **PWA_INVALID**.

SEE ALSO

getpwaent(3), **getgraent(3)**, **pwdauth(3)**

NAME

quot – summarize file system ownership

SYNOPSIS

`/usr/etc/quot [-acfhnv] [filesystem]`

DESCRIPTION

quot displays the number of blocks (1024 bytes) in the named *filesystem* currently owned by each user.

OPTIONS

- a** Generate a report for all mounted file systems.
- c** Display three columns giving file size in blocks, number of files of that size, and cumulative total of blocks in that size or smaller file.
- f** Display count of number of files as well as space owned by each user.
- h** Estimate the number of blocks in the file — this doesn't account for files with holes in them.
- n** Run the pipeline `ncheck filesystem | sort +0n | quot -n filesystem` to produce a list of all files and their owners.
- v** Display three columns containing the number of blocks not accessed in the last 30, 60, and 90 days.

FILES

<code>/etc/mtab</code>	mounted file systems
<code>/etc/passwd</code>	to get user names

SEE ALSO

`du(1V)`, `ls(1V)`

NAME

quotacheck – file system quota consistency checker

SYNOPSIS

`/usr/etc/quotacheck [-v] [-p] filesystem...`

`/usr/etc/quotacheck [-apv]`

DESCRIPTION

quotacheck examines each file system, builds a table of current disk usage, and compares this table against that stored in the disk quota file for the file system. If any inconsistencies are detected, both the quota file and the current system copy of the incorrect quotas are updated (the latter only occurs if an active file system is checked).

quotacheck expects each file system to be checked to have a quota file named *quotas* in the root directory. If none is present, **quotacheck** will ignore the file system.

quotacheck is normally run at boot time from the `/etc/rc.local` file, see `rc(8)`, before enabling disk quotas with `quotaon(8)`.

quotacheck accesses the raw device in calculating the actual disk usage for each user. Thus, the file systems checked should be quiescent while **quotacheck** is running.

OPTIONS

- `-v` Indicate the calculated disk quotas for each user on a particular file system. **quotacheck** Normally reports only those quotas modified.
- `-a` Check all the file systems indicated in `/etc/fstab` to be read-write with disk quotas.
- `-p` Run parallel passes on the required file systems, using the pass numbers in `/etc/fstab` in an identical fashion to `fsck(8)`.

FILES

<code>quotas</code>	quota file at the file system root
<code>/etc/mstab</code>	mounted file systems
<code>/etc/fstab</code>	default file systems

SEE ALSO

`quotactl(2)`, `quotaon(8)`, `rc(8)`

NAME

quotaon, quotaoff – turn file system quotas on and off

SYNOPSIS

/usr/etc/quotaon [-v] *filesystem...*

/usr/etc/quotaon [-av]

/usr/etc/quotaoff [-v] *filesystem...*

/usr/etc/quotaoff [-av]

DESCRIPTION OF QUOTAON

quotaon announces to the system that disk quotas should be enabled on one or more file systems. The file systems specified must be mounted at the time. The file system quota files must be present in the root directory of the specified file system and be named *quotas*.

DESCRIPTION OF QUOTAOFF

quotaoff announces to the system that file systems specified should have any disk quotas turned off.

OPTIONS TO QUOTAON

- a All file systems in */etc/fstab* marked read-write with quotas will have their quotas turned on. This is normally used at boot time to enable quotas.
- v Display a message for each file system where quotas are turned on.

OPTIONS TO QUOTAOFF

- a Force all file systems in */etc/fstab* to have their quotas disabled.
- v Display a message for each file system affected.

These commands update the status field of devices located in */etc/mtab* to indicate when quotas are on or off for each file system.

FILES

quotas	quota file at the file system root
/etc/mtab	mounted file systems
/etc/fstab	default file systems

SEE ALSO

quotactl(2), fstab(5), mtab(5)

NAME

rarpd – DARPA Reverse Address Resolution Protocol service

SYNOPSIS

/usr/etc/rarpd if hostname

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

rarpd starts a daemon that responds to Reverse Address Resolution Protocol (Reverse ARP) requests. The daemon forks a copy of itself, and requires root privileges.

The Reverse ARP protocol is used by machines at boot time to discover their (32 bit) IP address given their (48 bit) Ethernet address. In order for the request to be answered, a machine's name-to-IP-address entry must exist in the */etc/hosts* file and its name-to-Ethernet-address entry must exist in the */etc/ethers* file. Furthermore, the server that runs the **rarpd** daemon must have entries in both files. Note that if the server machine is using the Yellow Pages service, the server's files are ignored, and the appropriate Yellow Pages maps queried.

The first argument, *if*, is one of the interface parameter strings (listed in *boot(8S)*), in the form of "name unit", for example *ie0*. The second argument, *hostname*, is the interface's corresponding host name. The *if*, *hostname* pair should be the same as the arguments passed to the *ifconfig(8)* command. As with *ifconfig*, **rarpd** must be invoked for each interface that the server wishes to support. Therefore a gateway machine may invoke the **rarpd** multiple times, for example:

```
/usr/etc/rarpd ie0 host  
/usr/etc/rarpd ie1 host-backbone
```

FILES

/etc/ethers
/etc/hosts

SEE ALSO

boot(8S), *ifconfig(8C)*, *ipallocald(8C)*, *netconfig(8C)*, *ethers(5)*, *hosts(5)*, *policies(5)*, *netconfig(8C)*, *ipallocald(8C)*

Finlayson, Ross, Timothy Mann, Jeffrey Mogul, and Marvin Theimer, *A Reverse Address Resolution Protocol*, RFC 903, Network Information Center, SRI International, Menlo Park, Calif., June 1984.

NAME

rc, **rc.boot**, **rc.local** – command scripts for auto-reboot and daemons

SYNOPSIS

/etc/rc
/etc/rc.boot
/etc/rc.local

DESCRIPTION

rc and **rc.boot** are command scripts that are invoked by **init(8)** to perform file system housekeeping and to start system daemons. **rc.local** is a script for commands that are pertinent only to a specific site or client machine.

rc.boot sets the machine name, and then, if coming up multi-user, runs **fsck(8)** with the **-p** option. This “preens” the disks of minor inconsistencies resulting from the last system shutdown and checks for serious inconsistencies caused by hardware or software failure. If **fsck(8)** detects a serious disk problem, it returns an error and **init(8)** brings the system up in single-user mode. When coming up single-user, when **init(8)** is invoked by **fastboot(8)**, or when it is passed the **-b** flag from **boot(8S)**, functions performed in the **rc.local** file, including this disk check, are skipped.

Next, **rc** runs. If the system came up single-user, **rc** runs when the single-user shell terminates (see **init(8)**). It mounts 4.2 filesystems and spawns a shell for **/etc/rc.local**, which mounts NFS filesystems, and starts local daemons. After **rc.local** returns, **rc** starts standard daemons, preserves editor files, clears **/tmp**, starts system accounting (if applicable), starts the network (where applicable), and if enabled, runs **savecore(8)** to preserve the core image after a crash.

Sun386i SYSTEM DESCRIPTION

These files operate as described above with the following variations:

fsck(8) is invoked with the **-y** option to prevent users being put in single-user mode by happenstance.

rc.boot invokes **netconfig(8C)** to configure the system for the network before booting. **netconfig** is invoked before the **/usr** filesystem is mounted, because **/usr** might be mounted from a server. **netconfig** writes **/etc/net.conf** unless the **-n** option is specified, controlling system booting.

The file **/etc/net.conf** stores these environment variables: The **VERBOSE** environment variable controls the verbosity of the messages from the **rc** script; its value is taken from **NVRAM**. The **NETWORKED** environment variable controls whether services useful only on a networked system are started in **/etc/rc.local**. The **PNP** environment variable, set up during initial system installation, controls whether local network configuration information is used or whether that information comes from the network. (Using automatic system installation causes all systems except boot servers to get this information from the network, facilitating network reconfiguration.) The **HOSTNAME** and **DOMAINNAME** environment variables, used together, help determine if this system is a boot server or, with **PNP** set to **no**, control the host name and domain name.

rc.boot dynamically loads device drivers.

rc invokes any programs found in **/var/recover** to clean up any operations partially completed when the system crashed or was shut down.

rc.local starts the automounter.

FILES

/etc/rc
/etc/rc.boot
/etc/rc.local
/etc/net.conf
/var/recover/*
/var/yp/*
/tmp

SEE ALSO**automount(8), boot(8S), fastboot(8), init(8), reboot(8), savecore(8), netconfig(8C)**

NAME

rdate – set system date from a remote host

SYNOPSIS

/usr/ucb/rdate hostname

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

rdate sets the local date and time from the **hostname** given as argument. You must be super-user on the local system. Typically **rdate** can be inserted as part of your **/etc/rc.local** startup script.

FILES

/etc/rc.local

SEE ALSO

timed(8C)

BUGS

Could be modified to accept a list of hostnames and try each until a valid date returned. Better yet would be to write a real date server that accepted broadcast requests.

NAME

reboot – restart the operating system

SYNOPSIS

`/usr/etc/reboot [-dnq] [boot arguments]`

DESCRIPTION

reboot executes the **reboot(2)** system call to restart the kernel. The kernel is loaded into memory by the PROM monitor, which transfers control to it. See **boot(8S)** for details.

Although **reboot** can be run by the super-user at any time, **shutdown(8)** is normally used first to warn all users logged in of the impending loss of service. See **shutdown(8)** for details.

reboot performs a **sync(1)** operation on the disks, and then a multiuser reboot is initiated. See **init(8)** for details.

reboot normally logs the reboot to the system log daemon, **syslogd(8)**, and places a shutdown record in the login accounting file `/var/adm/wtmp`. These actions are inhibited if the `-n` or `-q` options are present.

Power Fail and Crash Recovery

Normally, the system will reboot itself at power-up or after crashes.

OPTIONS

- `-d` Dump system core before rebooting.
- `-n` Avoid the **sync(1)**. It can be used if a disk or the processor is on fire.
- `-q` Quick. Reboots quickly and ungracefully, without first shutting down running processes.

Boot Arguments

If a boot argument string is given, it is passed to the boot command in the PROM monitor. The string must be quoted if it contains spaces or other characters that could be interpreted by the shell. If the first character of the boot argument string is a minus sign '-' the string must be preceded by an option terminator string '--' For example: 'reboot -- -s' to reboot and come up single user, 'reboot vmunix.test' to reboot to a new kernel. See **boot(8S)** for details.

FILES

`/var/adm/wtmp` login accounting file

SEE ALSO

sync(1), **reboot(2)**, **fsck(8)**, **halt(8)**, **init(8)**, **shutdown(8)**, **syslogd(8)**, **boot(8S)**, **crash(8S)**

NAME

renice – alter priority of running processes

SYNOPSIS

/usr/etc/renice priority pid...

/usr/etc/renice priority [-p pid...] [-g pgrp...] [-u username...]

DESCRIPTION

renice alters the scheduling priority of one or more running processes.

OPTIONS

By default, the processes to be affected are specified by their process IDs. *priority* is the new priority value.

-p pid ...

Specify a list of process IDs.

-g pgrp ...

Specify a list of process group IDs. The processes in the specified process groups have their scheduling priority altered.

-u user ...

Specify a list of user IDs or usernames. All processes owned by each *user* have their scheduling altered.

Users other than the super-user may only alter the priority of processes they own, and can only monotonically increase their “nice value” within the range 0 to 20. (This prevents overriding administrative fiats.) The super-user may alter the priority of any process and set the priority to any value in the range -20 — 20. Useful priorities are: 19 (the affected processes will run only when nothing else in the system wants to), 0 (the “base” scheduling priority) and any negative value (to make things go very fast).

If no **who** parameter is specified, the current process (alternatively, process group or user) is used.

FILES

/etc/passwd to map user names to user ID's

SEE ALSO

pstat(8)

BUGS

If you make the priority very negative, then the process cannot be interrupted.

To regain control you must make the priority greater than zero.

Users other than the super-user cannot increase scheduling priorities of their own processes, even if they were the ones that decreased the priorities in the first place.

NAME

repquota – summarize quotas for a file system

SYNOPSIS

/usr/etc/repquota [-v] *filesystem...*

/usr/etc/repquota [-av]

DESCRIPTION

repquota prints a summary of the disc usage and quotas for the specified file systems. For each user the current number of files and amount of space (in kilobytes) is printed, along with any quotas created with **edquota(8)**.

OPTIONS

-a Report on all file systems indicated in **/etc/fstab** to be read-write with quotas.

-v Report all quotas, even if there is no usage.

Only the super-user may view quotas which are not their own.

FILES

quotas quota file at the file system root

/etc/fstab default file systems

SEE ALSO

edquota(8), **quota(1)**, **quotacheck(8)**, **quotactl(2)**, **quotaon(8)**

NAME

restore, rrestore – incremental file system restore

SYNOPSIS

/usr/etc/restore options [filename...]

DESCRIPTION

restore restores files from backup tapes created with the **dump(8)** command. *options* is a string of at least one of the options listed below, along with any modifiers and arguments you supply. Remaining arguments to **restore** are the names of files (or directories whose files) are to be restored to disk. Unless the **h** modifier is in effect, a directory name refers to the files it contains, and (recursively) its subdirectories and the files they contain.

OPTIONS

- i** Interactive. After reading in the directory information from the tape, **restore** invokes an interactive interface that allows you to browse through the dump tape's directory hierarchy, and select individual files to be extracted. See **Interactive Commands**, below, for a description of available commands.
- r** Restore the entire tape. Load the tape's full contents into the current directory. This option should only be used to restore a complete dump tape onto a clear filesystem, or to restore an incremental dump tape after a full "level 0" restore. For example:


```
example# /usr/etc/newfs /dev/rxy0g
example# /usr/etc/mount /dev/xy0g /mnt
example# cd /mnt
example# restore r
```

is a typical sequence to restore a "level 0" dump. Another restore can be done to get an incremental dump in on top of this.
- R** Resume restoring. **restore** requests a particular tape of a multivolume set from which to resume a full restore (see the **r** option above). This allows **restore** to start from a checkpoint when it is interrupted in the middle of a full restore.
- t** Table of contents. List each *filename* that appears on the tape. If no *filename* argument is given, the root directory is listed. This results in a list of all files on the tape, unless the **h** modifier is in effect. (The **t** option replaces the function of the old **dumpdir** program).
- x** Extract the named files from the tape. If a named file matches a directory whose contents were written onto the tape, and the **h** modifier is not in effect, the directory is recursively extracted. The owner, modification time, and mode are restored (if possible). If no *filename* argument is given, the root directory is extracted. This results in the entire tape being extracted unless the **h** modifier is in effect.

Modifiers

Some of the following modifiers take arguments that are given as separate words on the command line. When more than one such modifier appears within *options*, the arguments must appear in the same order as the modifiers that they apply to.

- c** Convert the contents of the dump tape to the new filesystem format.
- d** Debug. Turn on debugging output.
- h** Extract the actual directory, rather than the files that it references. This prevents hierarchical restoration of complete subtrees from the tape.
- m** Extract by inode numbers rather than by filename to avoid regenerating complete pathnames. This is useful if only a few files are being extracted.
- v** Verbose. **restore** displays the name of each file it restores, preceded by its file type.
- y** Do not ask whether to abort the restore in the event of tape errors. **restore** tries to skip over the bad tape block(s) and continue as best it can.

b factor

Blocking factor. Specify the blocking factor for tape reads. By default, **restore** will attempt to figure out the block size of the tape. Note: a tape block is 512 bytes.

f dump-file

Use *dump-file* instead of */dev/rmt?* as the file to restore from. If *dump-file* is specified as '-', **restore** reads from the standard input. This allows, **dump(8)** and **restore** to be used in a pipeline to dump and restore a file system:

```
example# dump 0f - /dev/rxy0g |
```

If the name of the file is of the form *machine:device* the restore is done from the specified machine over the network using **rmt(8C)**. Since **restore** is normally run by root, the name of the local machine must appear in the *.rhosts* file of the remote machine. If the file is specified as *user@machine:device*, **restore** will attempt to execute as the specified user on the remote machine. The specified user must have a *.rhosts* file on the remote machine that allows root from the local machine. If **restore** is called as **rrestore**, the tape defaults to **dumphost:/dev/rmt8**. To direct the input from a desired remote machine, set up an alias for **dumphost** in the file */etc/hosts*.

s n Skip to the *n*'th file when there are multiple dump files on the same tape. For example, the command:

```
example# restore xfs /dev/nrar0 5
```

would position you at the fifth file on the tape.

USAGE

Interactive Commands

restore enters interactive mode when invoked with the **i** option. Interactive commands are reminiscent of the shell. For those commands that accept an argument, the default is the current directory.

ls [*directory*] List files in *directory* or the current directory, represented by a '.' (period). Directories are appended with a '/' (backslash). Entries marked for extraction are prefixed with a '*' (asterisk). If the verbose option is in effect, inode numbers are also listed.

cd *directory*

Change to directory *directory* (within the dump-tape).

pwd Print the full pathname of the current working directory.

add [*filename*]

Add the current directory, or the named file or directory *directory* to the list of files to extract. If a directory is specified, add that directory and its files (recursively) to the extraction list (unless the **h** modifier is in effect).

delete [*filename*]

Delete the current directory, or the named file or directory from the list of files to extract. If a directory is specified, delete that directory and all its descendents from the extraction list (unless the **h** modifier is in effect). The most expedient way to extract a majority of files from a directory is to add that directory to the extraction list, and then delete specific files to omit.

extract Extract all files on the extraction list from the dump tape. **restore** asks which volume the user wishes to mount. The fastest way to extract a small number of files is to start with the last tape volume and work toward the first.

verbose Toggle the status of the **v** modifier. While **v** is in effect, the **ls** command lists the inode numbers of all entries, and **restore** displays information about each file as it is extracted.

help Display a summary of the available commands.

quit **restore** exits immediately, even if the extraction list is not empty.

FILES

/dev/rmt8	the default tape drive
dumphost:/dev/rmt8	the default tape drive if called as rrestore
/tmp/rstdir*	file containing directories on the tape
/tmp/rstmode*	owner, mode, and timestamps for directories
/restoresymtable	information passed between incremental restores

SEE ALSO

dump(8), mkfs(8), mount(8), newfs(8), rmt(8C)

BUGS

restore can get confused when doing incremental restores from dump tapes that were made on active file systems.

A “level 0” dump must be done after a full restore. Because **restore** runs in user mode, it has no control over inode allocation; this means that **restore** repositions the files, although it does not change their contents. Thus, a full dump must be done to get a new set of directories reflecting the new file positions, so that later incremental dumps will be correct.

DIAGNOSTICS

restore complains about bad option characters.

Read errors result in complaints. If **y** has been specified, or the user responds **y**, **restore** will attempt to continue.

If the dump extends over more than one tape, **restore** asks the user to change tapes. If the **x** or **i** option has been specified, **restore** also asks which volume the user wishes to mount.

There are numerous consistency checks that can be listed by **restore**. Most checks are self-explanatory or can “never happen”. Common errors are given below.

Converting to new file system format.

A dump tape created from the old file system has been loaded. It is automatically converted to the new file system format.

filename: not found on tape

The specified file name was listed in the tape directory, but was not found on the tape. This is caused by tape read errors while looking for the file, and from using a dump tape created on an active file system.

expected next file *inumber*, got *inumber*

A file that was not listed in the directory showed up. This can occur when using a dump tape created on an active file system.

Incremental tape too low

When doing an incremental restore, a tape that was written before the previous incremental tape, or that has too low an incremental level has been loaded.

Incremental tape too high

When doing incremental restore, a tape that does not begin its coverage where the previous incremental tape left off, or one that has too high an incremental level has been loaded.

Tape read error while restoring *filename*

Tape read error while skipping over inode *inumber*

Tape read error while trying to resynchronize

A tape read error has occurred.

If a file name is specified, then its contents are probably partially wrong. If an inode is being skipped or the tape is trying to resynchronize, then no extracted files have been corrupted, though files may not be found on the tape.

resync restore, skipped *num* blocks

After a tape read error, restore may have to resynchronize itself. This message lists the number of blocks that were skipped over.

NAME

rexd – RPC-based remote execution server

SYNOPSIS

/usr/etc/rpc.rexd

DESCRIPTION

rexd is the Sun RPC server for remote program execution. This daemon is started by **inetd**(8) whenever a remote execution request is made,

For noninteractive programs, the standard file descriptors are connected directly to TCP connections. Interactive programs involve pseudo-terminals, in a fashion that is similar to the login sessions provided by **rlogin**(1). This daemon may use the NFS to mount file systems specified in the remote execution request.

FILES

/dev/ttypn	pseudo-terminals used for interactive mode
/etc/passwd	authorized users
/tmp/dbrexd?????	temporary mount points for remote file systems.

SEE ALSO

on(1), **rex**(3), **exports**(5), **inetd**(8), **inetd.conf**(5)

DIAGNOSTICS

Diagnostic messages are normally printed on the console, and returned to the requestor.

BUGS

Should be better access control.

RESTRICTIONS

Root cannot execute commands using **rex**d client programs such as **on**(1).

NAME

rexecd – remote execution server

SYNOPSIS

/usr/etc/in.rexecd host.port

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

rexecd is the server for the rexec(3N) routine. The server provides remote execution facilities with authentication based on user names and encrypted passwords. It is invoked automatically as needed by inetd(8C), and then executes the following protocol:

- 1) The server reads characters from the socket up to a null (\0) byte. The resultant string is interpreted as an ASCII number, base 10.
- 2) 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.
- 3) A null terminated user name of at most 16 characters is retrieved on the initial socket.
- 4) A null terminated, encrypted, password of at most 16 characters is retrieved on the initial socket.
- 5) 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.
- 6) rexecd then validates the user as is done at login time and, if the authentication was successful, changes to the user's home directory, and establishes the user and group protections of the user. If any of these steps fail the connection is aborted with a diagnostic message returned.
- 7) 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 rexecd.

SEE ALSO

inetd(8C)

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 7 above upon successful completion of all the steps prior to the command execution).

username too long

The name is longer than 16 characters.

password too long

The password is longer than 16 characters.

command too long

The command line passed exceeds the size of the argument list (as configured into the system).

Login incorrect.

No password file entry for the user name existed.

Password incorrect.

The wrong password was supplied.

No remote directory.

The chdir command to the home directory failed.

Try again.

A fork by the server failed.

/usr/bin/sh: ...

The user's login shell could not be started.

BUGS

Indicating '**Login incorrect**' as opposed to '**Password incorrect**' is a security breach which allows people to probe a system for users with null passwords.

A facility to allow all data exchanges to be encrypted should be present.

NAME

rlogind – remote login server

SYNOPSIS

/usr/etc/in.rlogind host.port

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 is invoked by **inetd(8C)** when a remote login connection is established, and executes the following protocol:

- 1) The server checks the client's source port. If the port is not in the range 0-1023, the server aborts the connection. The client's address and port number are passed as arguments to **rlogind** by **inetd** in the form *host.port* with *host* in hex and *port* in decimal.
- 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(5V)**.

SEE ALSO

inetd(8C)

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.

/usr/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

rmail – handle remote mail received via uucp

SYNOPSIS

rmail *recipient...*

DESCRIPTION

rmail interprets incoming mail received through **uucp(1C)**, collapsing “From” lines in the form generated by **binmail(1)** into a single line of the form *return-path!sender*, and passing the processed mail on to **sendmail(8)**.

rmail is explicitly designed for use with **uucp(1C)** and **sendmail(8)**.

SEE ALSO

binmail(1), **uucp(1C)**, **sendmail(8)**

NAME

rmt – remote magtape protocol module

SYNOPSIS

/usr/etc/rmt

DESCRIPTION

rmt is a program used by the remote dump and restore programs in manipulating a magnetic tape drive through an interprocess communication connection. rmt is normally started up with an rexec(3N) or rcmd(3N) call.

The rmt program accepts requests specific to the manipulation of magnetic tapes, performs the commands, then responds with a status indication. All responses are in ASCII and in one of two forms. Successful commands have responses of

*A*number\n

where *number* is an ASCII representation of a decimal number. Unsuccessful commands are responded to with

*E*error-number\n*e*rror-message\n

where *error-number* is one of the possible error numbers described in intro(2) and *error-message* is the corresponding error string as printed from a call to perror(3). The protocol is comprised of the following commands (a space is present between each token):

- | | |
|--------------------------|--|
| S | Return the status of the open device, as obtained with a MTIOCGET ioctl call. If the operation was successful, an “ack” is sent with the size of the status buffer, then the status buffer is sent (in binary). |
| C device | Close the currently open device. The <i>device</i> specified is ignored. |
| I operation count | Perform a MTIOCOP ioctl(2) command using the specified parameters. The parameters are interpreted as the ASCII representations of the decimal values to place in the <i>mt_op</i> and <i>mt_count</i> fields of the structure used in the ioctl call. The return value is the <i>count</i> parameter when the operation is successful. |
| L whence offset | Perform an lseek(2) operation using the specified parameters. The response value is that returned from the lseek call. |
| O device mode | Open the specified <i>device</i> using the indicated <i>mode</i> . <i>device</i> is a full pathname and <i>mode</i> is an ASCII representation of a decimal number suitable for passing to open(2V). If a device had already been opened, it is closed before a new open is performed. |
| Rcount | Read <i>count</i> bytes of data from the open device. rmt performs the requested read(2V) and responds with <i>Acount-read</i> \n if the read was successful; otherwise an error in the standard format is returned. If the read was successful, the data read is then sent. |
| W count | Write data onto the open device. rmt reads <i>count</i> bytes from the connection, aborting if a premature EOF is encountered. The response value is that returned from the write(2V) call. |

Any other command causes rmt to exit.

DIAGNOSTICS

All responses are of the form described above.

SEE ALSO

intro(2), ioctl(2), lseek(2), open(2V), read(2V), write(2V), perror(3), rcmd(3N), rexec(3N), mtio(4), dump(8), restore(8)

BUGS

People tempted to use this for a remote file access protocol are discouraged.

NAME

route – manually manipulate the routing tables

SYNOPSIS

`/usr/etc/route [-fhn] add | delete destination [gateway [metric]]`

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

route manually manipulate the network routing tables normally maintained by the system routing daemon, **routed(8C)**. **route** allows the super-user to operate directly on the routing table for the specific host or network indicated by *destination*. The *gateway* argument, if present, indicates the network gateway to which packets should be addressed. The *metric* argument indicates the number of “hops” to the *destination*. The default value for *metric* is 0.

The **add** command instructs **route** to add a route to *destination*. **delete** deletes a route.

Routes to a particular host are distinguished from those to a network by interpreting the Internet address associated with the *destination* parameter. If the destination has a “local address part” of `INADDR_ANY`, then the route is assumed to be to a network; otherwise, it is presumed to be a route to a host. If the route is to a destination connected by a gateway, the *metric* parameter should be greater than 0. All symbolic names specified for a *destination* or *gateway* are looked up in the hosts database, `/etc/hosts`. If this lookup fails, then the name is looked up in the networks database, `/etc/networks`.

OPTIONS

- f** Flush the routing tables of all gateway entries. If this is used in conjunction with one of the subcommands described below, **route** flushes the gateways before performing the subcommand.
- h** Treat the destination as a host.
- n** Treat the destination as a network.

FILES

`/etc/hosts`
`/etc/networks`

SEE ALSO

ioctl(2), **routing(4N)**, **routed(8C)**

BUGS

The change operation is not implemented, one should add the new route, then delete the old one.

DIAGNOSTICS

add destination address:gateway addresss flags value

The specified route is being added to the tables. The values printed are from the routing table entry supplied in the `ioctl(2)` call.

delete destination address:gateway addresss flags value

The specified route is being deleted.

destination address done

When the **-f** flag is specified, each routing table entry deleted is indicated with a message of this form.

Network is unreachable

An attempt to add a route failed because the gateway listed was not on a directly-connected network. Give the next-hop gateway instead.

not in table

A delete operation was attempted for an entry that is not in the table.

routing table overflow

An add operation was attempted, but the system was unable to allocate memory to create the new entry.

NAME

routed – network routing daemon

SYNOPSIS

`/etc/in.routed` [`-qstv`] [`logfile`]

DESCRIPTION

routed is invoked at boot time to manage the network routing tables. The routing daemon uses a variant of the Xerox NS Routing Information Protocol in maintaining up to date kernel routing table entries.

In normal operation **routed** listens on udp(4P) socket 520 (decimal) for routing information packets. If the host is an internetwork router, it periodically supplies copies of its routing tables to any directly connected hosts and networks.

When **routed** is started, it uses the `SIOCGIFCONF` ioctl(2) to find those directly connected interfaces configured into the system and marked “up” (the software loopback interface is ignored). If multiple interfaces are present, it is assumed the host will forward packets between networks. **routed** then transmits a *request* packet on each interface (using a broadcast packet if the interface supports it) and enters a loop, listening for *request* and *response* packets from other hosts.

When a *request* packet is received, **routed** formulates a reply based on the information maintained in its internal tables. The *response* packet generated contains a list of known routes, each marked with a “hop count” metric (a count of 16, or greater, is considered “infinite”). The metric associated with each route returned provides a metric *relative to the sender*.

request packets received by **routed** are used to update the routing tables if one of the following conditions is satisfied:

- (1) No routing table entry exists for the destination network or host, and the metric indicates the destination is “reachable” (that is, the hop count is not infinite).
- (2) The source host of the packet is the same as the router in the existing routing table entry. That is, updated information is being received from the very internetwork router through which packets for the destination are being routed.
- (3) The existing entry in the routing table has not been updated for some time (defined to be 90 seconds) and the route is at least as cost effective as the current route.
- (4) The new route describes a shorter route to the destination than the one currently stored in the routing tables; the metric of the new route is compared against the one stored in the table to decide this.

When an update is applied, **routed** records the change in its internal tables and generates a *response* packet to all directly connected hosts and networks. **routed** waits a short period of time (no more than 30 seconds) before modifying the kernel’s routing tables to allow possible unstable situations to settle.

In addition to processing incoming packets, **routed** also periodically checks the routing table entries. If an entry has not been updated for 3 minutes, the entry’s metric is set to infinity and marked for deletion. Deletions are delayed an additional 60 seconds to insure the invalidation is propagated throughout the internet.

Hosts acting as internetwork routers gratuitously supply their routing tables every 30 seconds to all directly connected hosts and networks.

Supplying the `-s` option forces **routed** to supply routing information whether it is acting as an internetwork router or not. The `-q` option is the opposite of the `-s` option. If the `-t` option is specified, all packets sent or received are printed on the standard output. In addition, **routed** will not divorce itself from the controlling terminal so that interrupts from the keyboard will kill the process. Any other argument supplied is interpreted as the name of file in which **routed**’s actions should be logged. This log contains information about any changes to the routing tables and a history of recent messages sent and received which are related to the changed route. The `-v` option allows a logfile to be created showing the changes made to the routing tables with a timestamp.

In addition to the facilities described above, **routed** supports the notion of “distant” *passive* and *active* gateways. When **routed** is started up, it reads the file `/etc/gateways` to find gateways which may not be identified using the `SIOGIFCONF ioctl`. Gateways specified in this manner should be marked *passive* if they are not expected to exchange routing information, while gateways marked *active* should be willing to exchange routing information (that is, they should have a **routed** process running on the machine). *Passive* gateways are maintained in the routing tables forever and information regarding their existence is included in any routing information transmitted. *Active* gateways are treated equally to network interfaces. Routing information is distributed to the gateway and if no routing information is received for a period of the time, the associated route is deleted.

The `/etc/gateways` is comprised of a series of lines, each in the following format:

```
< net | host > filename1 gateway filename2 metric value < passive | active >
```

The `net` or `host` keyword indicates if the route is to a network or specific host.

filename1 is the name of the destination network or host. This may be a symbolic name located in `/etc/networks` or `/etc/hosts`, or an Internet address specified in “dot” notation; see `inet(3N)`.

filename2 is the name or address of the gateway to which messages should be forwarded.

Value is a metric indicating the hop count to the destination host or network.

The keyword `passive` or `active` indicates if the gateway should be treated as *passive* or *active* (as described above).

FILES

```
/etc/gateways      for distant gateways
/etc/networks
/etc/hosts
```

SEE ALSO

`ioctl(2)`, `inet(3N)`, `udp(4P)`

BUGS

The kernel’s routing tables may not correspond to those of **routed** for short periods of time while processes utilizing existing routes exit; the only remedy for this is to place the routing process in the kernel.

routed should listen to intelligent interfaces, such as an IMP, and to error protocols, such as ICMP, to gather more information.

NAME

rpcinfo – report RPC information

SYNOPSIS

```
rpcinfo -p [ host ]
rpcinfo [ -n portnum ] -u host program [ version ]
rpcinfo [ -n portnum ] -t host program [ version ]
rpcinfo -b program version
rpcinfo -d program version
```

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

rpcinfo makes an RPC call to an RPC server and reports what it finds.

OPTIONS

- p Probe the portmapper on *host*, and print a list of all registered RPC programs. If *host* is not specified, it defaults to the value returned by `hostname(1)`.
- u Make an RPC call to procedure 0 of *program* on the specified *host* using UDP, and report whether a response was received.
- t Make an RPC call to procedure 0 of *program* on the specified *host* using TCP, and report whether a response was received.
- n Use *portnum* as the port number for the *-t* and *-u* options instead of the port number given by the portmapper.
- b Make an RPC broadcast to procedure 0 of the specified *program* and *version* using UDP and report all hosts that respond.
- d Delete registration for the RPC service of the specified *program* and *version*. This option can be exercised only by the super-user.

The *program* argument can be either a name or a number.

If a *version* is specified, **rpcinfo** attempts to call that version of the specified *program*. Otherwise, **rpcinfo** attempts to find all the registered version numbers for the specified *program* by calling version 0 (which is presumed not to exist; if it does exist, **rpcinfo** attempts to obtain this information by calling an extremely high version number instead) and attempts to call each registered version. Note: the version number is required for *-b* and *-d* options.

EXAMPLES

To show all of the RPC services registered on the local machine use:

```
example% rpcinfo -p
```

To show all of the RPC services registered on the machine named **klaxon** use:

```
example% rpcinfo -p klaxon
```

To show all machines on the local net that are running the Yellow Pages service use:

```
example% rpcinfo -b ypserv 'version' | uniq
```

where 'version' is the current Yellow Pages version obtained from the results of the *-p* switch above.

To delete the registration for version 1 of the **walld** service use:

```
example% rpcinfo -d walld 1
```

SEE ALSO

rpc(5), portmap(8C)

RPC Programming Guide in Network Programming

BUGS

In releases prior to SunOS 3.0, the Network File System (NFS) did not register itself with the portmapper; **rpcinfo** cannot be used to make RPC calls to the NFS server on hosts running such releases.

NAME

rquotad – remote quota server

SYNOPSIS

/usr/etc/rpc.rquotad

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

rquotad is an **rpc(4)** server which returns quotas for a user of a local file system which is mounted by a remote machine over the NFS. The results are used by **quota(1)** to display user quotas for remote file systems. The **rquotad** daemon is normally invoked by **inetd(8C)**.

FILES

quotas quota file at the file system root

SEE ALSO

quota(1), nfs(4), rpc(4), services(5), inetd(8C)

NAME

rshd – remote shell server

SYNOPSIS

/usr/etc/in.rshd host.port

DESCRIPTION

rshd is the server for the rcmd(3N) routine and, consequently, for the rsh(1C) program. The server provides remote execution facilities with authentication based on privileged port numbers.

rshd is invoked by inetd(8C) each time a shell service is requested, and executes the following protocol:

- 1) The server checks the client's source port. If the port is not in the range 0-1023, the server aborts the connection. The client's host address (in hex) and port number (in decimal) are the argument passed to rshd.
- 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 the lookup fails, the connection is terminated. If the chdir fails, it does a chdir to / (root). 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.

FILES

/etc/hosts.equiv

SEE ALSO

rsh(1C), rcmd(3N), syslogd(8)

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.

DIAGNOSTICS

The following 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.

Permission denied.

The authentication procedure described above failed.

Can't make pipe.

The pipe needed for the `stderr`, was not created.

Try again.

A *fork* by the server failed.

/usr/bin/sh:...

The user's login shell could not be started.

In addition, daemon's status messages and internal diagnostics are logged to the appropriate system log using the `syslogd(8)` facility.

NAME

rstatd – kernel statistics server

SYNOPSIS

/usr/etc/rpc.rstatd

DESCRIPTION

rstatd is a server which returns performance statistics obtained from the kernel. These statistics are graphically displayed by **perfmeter(1)**. The **rstatd** daemon is normally invoked by **inetd(8C)**.

Systems with disk drivers to be monitored by this daemon must be configured so as to report disk (**_dk_xfer**) statistics.

SEE ALSO

perfmeter(1), **services(5)**, **inetd(8C)**

NAME

rusersd – network username server

SYNOPSIS

/usr/etc/rpc.rusersd

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

rusersd is a server that returns a list of users on the network. The **rusersd** daemon is normally invoked by **inetd(8C)**.

SEE ALSO

perfmeter(1), **rusers(1C)**, **services(5)**, **inetd(8C)**

NAME

rwalld – network rwall server

SYNOPSIS

/usr/etc/rpc.rwalld

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

rwalld is a server that handles **rwall(1C)** and **shutdown(2)** requests. It is implemented by calling **wall(1)** to all the appropriate network machines. The **rwalld** daemon is normally invoked by **inetd(8C)**.

SEE ALSO

rwall(1C), **wall(1)**, **shutdown(2)** **services(5)**, **inetd(8C)**,

NAME

rwod – system status server

SYNOPSIS

/usr/etc/rwod

AVAILABILITY

Due to its potential impact on network performance, this service is commented out of the */etc/rc.local* system initialization script. It is provided only for 4.3 BSD compatibility.

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

rwod is the server which maintains the database used by the **rwho(1C)** and **ruptime(1C)** programs. Its operation is predicated on the ability to *broadcast* messages on a network.

rwod 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 **rwod** servers' status messages, validating them, then recording them in a collection of files located in the directory */var/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 **rwod** are placed in files named *whod.hostname* in the directory */var/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.

FILES

/var/spool/rwho

DIAGNOSTICS

Status and diagnostic messages are logged to the appropriate system log using the **syslogd(8)** facility.

SEE ALSO

rwho(1C), **ruptime(1C)**, **w(1)**, **gethostname(2)**, **nlist(3)**, **utmp(5)**, **syslogd(8)**

BUGS

This service takes up progressively more network bandwidth as the number of hosts on the local net increases. For large networks, the cost becomes prohibitive. RPC-based services such as **rup(1C)** and **rusers(1C)** provide a similar function with greater efficiency.

rwhod should relay status information between networks. People often interpret the server dying as a machine going down.

NAME

sa, accton – system accounting

SYNOPSIS

```
/usr/etc/sa [ -abcdDfijkKlmnrstu ] [ -v[n] ] [ -S savacctfile ] [ -U usracctfile ] [ filename ]
/usr/etc/accton [ filename ]
```

DESCRIPTION

With an argument naming an existing *filename*, *accton* causes system accounting information for every process executed to be placed at the end of the file. If no argument is given, accounting is turned off.

sa reports on, cleans up, and generally maintains accounting files.

sa is able to condense the information in */var/adm/acct* into a summary file */var/adm/savacct* which contains a count of the number of times each command was called and the time resources consumed. This condensation is desirable because on a large system */var/adm/acct* can grow by 500K bytes per day. The summary file is normally read before the accounting file, so the reports include all available information.

If a file name is given as the last argument, that file will be treated as the accounting file; */var/adm/acct* is the default.

Output fields are labeled: **cpu** for the sum of user+system time (in minutes), **re** for real time (also in minutes), **k** for CPU-time averaged core usage (in 1k units), **avio** for average number of I/O operations per execution. With options fields labeled **tio** for total I/O operations, **k*sec** for CPU storage integral (kilo-core seconds), **u** and **s** for user and system CPU time alone (both in minutes) will sometimes appear.

sa also breaks out accounting statistics by user. This information is kept in the file */var/adm/usracct*.

OPTIONS

- a** Print all command names, even those containing unprintable characters and those used only once. By default, those are placed under the name '***other.'
- b** Sort output by sum of user and system time divided by number of calls. Default sort is by sum of user and system times.
- c** Besides total user, system, and real time for each command print percentage of total time over all commands.
- d** Sort by average number of disk I/O operations.
- D** Print and sort by total number of disk I/O operations.
- f** Force no interactive threshold compression with **-v** flag.
- i** Do not read in summary file.
- j** Instead of total minutes time for each category, give seconds per call.
- k** Sort by CPU-time average memory usage.
- K** Print and sort by CPU-storage integral.
- l** Separate system and user time; normally they are combined.
- m** Print number of processes and number of CPU minutes for each user.
- n** Sort by number of calls.
- r** Reverse order of sort.
- s** Merge accounting file into summary file */var/adm/savacct* when done.
- t** For each command report ratio of real time to the sum of user and system times.
- u** Superseding all other flags, print for each record in the accounting file the user ID and command name.

- v Followed by a number *n*, types the name of each command used *n* times or fewer. If *n* is not specified, it defaults to 1. Await a reply from the terminal; if it begins with *y*, add the command to the category '**junk**.' This is used to strip out garbage.
- S The following filename is used as the command summary file instead of `/var/adm/savacct`.
- U The following filename is used instead of `/var/adm/usracct` to accumulate the per-user statistics printed by the `-m` option.

FILES

<code>/var/adm/acct</code>	raw accounting
<code>/var/adm/savacct</code>	summary by command
<code>/var/adm/usracct</code>	summary by user ID

SEE ALSO

`acct(2)`, `acct(5)`, `ac(8)`

BUGS

`sa`'s execution time increases linearly with the magnitude of the largest positive user ID in `/etc/passwd`.

NAME

savecore — save a core dump of the operating system

SYNOPSIS

/usr/etc/savecore dirname [system-name]

DESCRIPTION

savecore saves a core dump of the kernel (assuming that one was made) and writes a reboot message in the shutdown log. It is meant to be called near the end of the */etc/rc.local* file after the system boots. However, it is not normally run by default. You must edit that file to enable it.

savecore checks the core dump to be certain it corresponds with the version of the operating system currently running. If it does, savecore saves the core image in the file *dirname/vmcore.n* and the kernel's namelist, in *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 is less free space on the filesystem containing *dirname* than the number obtained from the minfree file, the core dump is not saved. If the minfree file does not exist, savecore always writes out the core file (assuming that a core dump was taken).

savecore also logs a reboot message using facility LOG_AUTH (see syslog(3)) If the system crashed as a result of a panic, savecore logs the panic string too.

If the core dump was from a system other than */vmunix*, the name of that system must be supplied as *system-name*.

FILES

/vmunix the kernel
/etc/rc.local

SEE ALSO

syslog(3), sa(8), crash(8S)

BUGS

Can be fooled into thinking a core dump is the wrong size.

You must run savecore very soon after booting — before the swap space containing the crash dump is overwritten by programs currently running.

NAME

sendmail – send mail over the internet

SYNOPSIS

```
/etc/sendmail [ -ba ] [ -bd ] [ -bi ] [ -bm ] [ -bp ] [ -bs ] [ -bt ] [ -bv ] [ -bz ]
               [ -Cfile ] [ -dX ] [ -Ffullname ] [ -fname ] [ -hN ] [ -n ] [ -ox value ] [ -q[ time ] ]
               [ -rname ] [ -t ] [ -v ] [ address ... ]
```

DESCRIPTION

sendmail sends a message to one or more people, routing the message over whatever networks are necessary. **sendmail** does internetwork forwarding as necessary to deliver the message to the correct place.

sendmail is not intended as a user interface routine; other programs provide user-friendly front ends; **sendmail** is used only to deliver pre-formatted messages.

With no flags, **sendmail** reads its standard input up to an EOF, or a line with a single dot and sends a copy of the letter found there to all of the addresses listed. It determines the network to use based on the syntax and contents of the addresses.

Local addresses are looked up in the local aliases(5) file, or by using the Yellow Pages name service, and aliased appropriately. In addition, if there is a .forward file in a recipient's home directory, **sendmail** forwards a copy of each message to the list of recipients that file contains. Aliasing can be prevented by preceding the address with a backslash. Normally the sender is not included in alias expansions, for example, if 'john' sends to 'group', and 'group' includes 'john' in the expansion, then the letter will not be delivered to 'john'.

sendmail will also route mail directly to other known hosts in a local network. The list of hosts to which mail is directly sent is maintained in the file /usr/lib/mailhosts.

OPTIONS

- ba** Go into ARPANET mode. All input lines must end with a CR-LF, and all messages will be generated with a CR-LF at the end. Also, the "From:" and "Sender:" fields are examined for the name of the sender.
- bd** Run as a daemon, waiting for incoming SMTP connections.
- bi** Initialize the alias database.
- bm** Deliver mail in the usual way (default).
- bp** Print a summary of the mail queue.
- bs** Use the SMTP protocol as described in RFC 821. This flag implies all the operations of the **-ba** flag that are compatible with SMTP.
- bt** Run in address test mode. This mode reads addresses and shows the steps in parsing; it is used for debugging configuration tables.
- bv** Verify names only — do not try to collect or deliver a message. Verify mode is normally used for validating users or mailing lists.
- bz** Create the configuration freeze file.
- Cfile** Use alternate configuration file.
- dX** Set debugging value to *X*.
- Ffullname**
Set the full name of the sender.
- fname** Sets the name of the "from" person (that is, the sender of the mail). **-f** can only be used by "trusted" users (who are listed in the config file).
- hN** Set the hop count to *N*. The hop count is incremented every time the mail is processed. When it reaches a limit, the mail is returned with an error message, the victim of an aliasing loop.

- Mid** Attempt to deliver the queued message with message-id *id*.
- n** Don't do aliasing.
- ox value**
Set option *x* to the specified *value*. Options are described below.
- q[time]**
Processed saved messages in the queue at given intervals. If *time* is omitted, process the queue once. *Time* is given as a tagged number, with *s* being seconds, *m* being minutes, *h* being hours, *d* being days, and *w* being weeks. For example, **-q1h30m** or **-q90m** would both set the timeout to one hour thirty minutes.
- rname** An alternate and obsolete form of the **-f** flag.
- Rstring**
Go through the queue of pending mail and attempt to deliver any message with a recipient containing the specified string. This is useful for clearing out mail directed to a machine which has been down for awhile.
- t** Read message for recipients. "To:", "Cc:", and "Bcc:" lines will be scanned for people to send to. The "Bcc:" line will be deleted before transmission. Any addresses in the argument list will be suppressed.
- v** Go into verbose mode. Alias expansions will be announced, etc.

PROCESSING OPTIONS

There are also a number of processing options that may be set. Normally these will only be used by a system administrator. Options may be set either on the command line using the **-o** flag or in the configuration file. These are described in detail in the *Installation and Operation Guide*. The options are:

- Afile** Use alternate alias file.
- c** On mailers that are considered "expensive" to connect to, do not initiate immediate connection. This requires queueing.
- dx** Set the delivery mode to *x*. Delivery modes are **i** for interactive (synchronous) delivery, **b** for background (asynchronous) delivery, and **q** for queue only — that is, actual delivery is done the next time the queue is run.
- D** Run `newaliases(8)` to automatically rebuild the alias database, if necessary.
- ex** Set error processing to mode *x*. Valid modes are **m** to mail back the error message, **w** to "write" back the error message (or mail it back if the sender is not logged in), **p** to print the errors on the terminal (default), **q** to throw away error messages (only exit status is returned), and **e** to do special processing for the BerkNet. If the text of the message is not mailed back by modes **m** or **w** and if the sender is local to this machine, a copy of the message is appended to the file `dead.letter` in the sender's home directory.
- Fmode** The mode to use when creating temporary files.
- f** Save UNIX-system-style "From" lines at the front of messages.
- gN** The default group ID to use when calling mailers.
- Hfile** The SMTP help file.
- i** Do not take dots on a line by themselves as a message terminator.
- Ln** The log level.
- m** Send to "me" (the sender) also if I am in an alias expansion.
- o** If set, this message may have old style headers. If not set, this message is guaranteed to have new style headers (that is, commas instead of spaces between addresses). If set, an adaptive algorithm is used that will correctly determine the header format in most cases.

Q*queuedir*

Select the directory in which to queue messages.

r*timeout*

The timeout on reads; if none is set, **sendmail** will wait forever for a mailer.

S*file*

Save statistics in the named file.

s

Always instantiate the queue file, even under circumstances where it is not strictly necessary.

T*ime*

Set the timeout on messages in the queue to the specified time. After sitting in the queue for this amount of time, they will be returned to the sender. The default is three days.

t*stz,dtz*

Set the name of the time zone.

u*N*

Set the default user id for mailers.

If the first character of the user name is a vertical bar, the rest of the user name is used as the name of a program to pipe the mail to. It may be necessary to quote the name of the user to keep **sendmail** from suppressing the blanks from between arguments.

sendmail returns an exit status describing what it did. The codes are defined in *<sysexits.h>*

EX_OK	Successful completion on all addresses.
EX_NOUSER	User name not recognized.
EX_UNAVAILABLE	Catchall meaning necessary resources were not available.
EX_SYNTAX	Syntax error in address.
EX_SOFTWARE	Internal software error, including bad arguments.
EX_OSERR	Temporary operating system error, such as "cannot fork".
EX_NOHOST	Host name not recognized.
EX_TEMPFAIL	Message could not be sent immediately, but was queued.

If invoked as *newaliases*, **sendmail** rebuilds the alias database. If invoked as *mailq*, **sendmail** prints the contents of the mail queue.

FILES

Except for */etc/sendmail.cf*, these pathnames are all specified in */etc/sendmail.cf*. Thus, these values are only approximations.

<i>/etc/aliases</i>	raw data for alias names
<i>/etc/aliases.pag</i>	data base of alias names
<i>/etc/aliases.dir</i>	
<i>/usr/lib/mailhosts</i>	list of hosts to which mail can be sent directly
<i>/etc/sendmail.cf</i>	configuration file
<i>/etc/sendmail.fc</i>	frozen configuration
<i>/etc/sendmail.hf</i>	help file
<i>/etc/sendmail.st</i>	collected statistics
<i>/usr/bin/uux</i>	to deliver uucp mail
<i>/usr/bin/mail</i>	to deliver local mail
<i>/var/spool/mqueue/*</i>	temp files and queued mail
<i>~/forward</i>	list of recipients for forwarding messages

SEE ALSO

biff(1), **binmail(1)**, **mail(1)**, **aliases(5)**

System and Network Administration

Su, Zaw-Sing, and Jon Postel, *The Domain Naming Convention for Internet User Applications*, RFC 819, Network Information Center, SRI International, Menlo Park, Calif., August 1982.

Postel, Jon, *Simple Mail Transfer Protocol*, RFC 821, Network Information Center, SRI International, Menlo Park, Calif., August 1982.

Crocker, Dave, *Standard for the Format of ARPA-Internet Text Messages*, RFC 822, Network Information Center, SRI International, Menlo Park, Calif., August 1982.

NAME

`setup_client` – create or remove an NFS client

SYNOPSIS

```
/usr/etc/install/script/setup_client op clientname yp_type swapsize rootpath swappath
    dumppath homepath execpath arch
```

DESCRIPTION

`setup_client` adds an NFS client to a server, or to removes one. It can only be run by the super-user. It is also used by `suninstall(8)`.

The *op* argument indicates which operation to perform; it can be either **add** or **remove**, to indicate whether to add or remove a client. *clientname* is the hostname of the client. *yp_type* indicates the type of Yellow Pages server or service to provide to the client, if any; it can be one of **master**, **slave**, **client** or **none**. *swapsize* is the number of bytes reserved for client's swap file. *rootpath* is the pathname of parent directory in which various client root directories reside; *rootpath/clientname* is the pathname of the client's root directory. *swappath* is the pathname of parent directory in which various client swap files reside; *swappath/clientname* is the pathname of the client's swap file. *dumppath* is the parent pathname in which various client dump files reside; *dumppath/clientname* is the pathname of the client's dump file. *homepath* is the pathname of the (parent) directory in which the various home directories are to reside; it is the pathname of the directory that the client is to mount as **/home**. *execpath* is the full pathname of the directory in which the executables for the architecture specified by the *arch* argument. This is the directory that the client mounts as **/usr**. *arch* specifies the client's architecture (for instance, **sun4**, **sun3**...). `setup_client` with no arguments displays a usage message that includes the proper *arch* argument for each supported architecture.

USAGE

Before you add or remove a client, you must first make sure that the Internet and Ethernet addresses for *clientname* are listed in the YP hosts database (if the server is running the YP), or in the server's **/etc/hosts** and **/etc/ethers** databases, respectively (otherwise). Then, run `setup_client` with the **add** or **remove** operation. When adding a client, you must then bootstrap that client machine.

You cannot add a client to a server that does not support the specified architecture. The executable directory for that client's architecture must be present on the server. If this file is absent, an error results.

`setup_client` updates the **/etc/bootparams** file. If the server is a YP master, it updates local YP database. It *does not* propagate the local update to other YP servers. To propagate the updates, use the following commands:

```
example# cd /var/yp
example# make
```

If the server is running YP but is not a YP master, `setup_client` issues a warning to indicate that the database is out of date.

When *arch* is given as **sun2**, `suninstall` issues a reminder to run the **/usr/etc/ndbootd** daemon for booting Sun-2 systems.

`setup_client` creates *swappath/clientname* with the *size*, (number of bytes) you specify. You can append one of **K** or **k** to indicate kilobytes, **M** or **m** to indicate megabytes, or **B** or **b** to indicate 512-byte blocks, to *size*. Otherwise, *size* is taken to indicate an exact byte count.

`suninstall` updates the **/etc/exports** file to allow root access to each client's root file system. It exports the client's swap and dump partitions only to the client. Note: the system administrator should verify that the **/etc/exports** file contains correct information, and that file systems are exported to the correct users and groups. Refer to `exportfs(8)` for details on exporting file systems.

EXAMPLES

This example shows how to add a Sun-4 system NFS client to a server.

```
example# setup_client add frodo client 16M /exports/roots /exports/swaps /exports/dumps /home \
/exports/execs/sun4/4.0 sun4
```

To remove this client, you would merely substitute **remove** for **add** in the above example.

FILES

```
/etc/hosts
/etc/ethers
/usr/etc/ndbootd
/etc/bootparams
/etc/exports
```

SEE ALSO

exportfs(8), **setup_exec(8)** **suninstall(8)**

Installing the SunOS

DIAGNOSTICS

incorrect number of arguments

Check number and order of the arguments.

must be run as root (super-user).

You must be root to use **setup_client**.

invalid operation type “xx”.

Valid operations are **add** and **remove**.

ATTENTION: xxxxxxxx -> boot.sun? not created.

(Sun-3 systems only.) A symbolic link can not be created because the boot file does not exist.

ATTENTION: xxxxxxxx.SUN? -> boot.sun? not created.

(Other than Sun-3 systems.) A symbolic link can not be created because the boot file does not exist.

ATTENTION: /usr/etc/ndbootd needs to be running on server before bringing up “client”.

The Sun-2 system boot daemon must be running in order to bootstrap a Sun-2 system.

NAME

setup_exec – install architecture-dependent executables on a heterogeneous file server

SYNOPSIS

/usr/etc/install/setup_exec arch execpath

DESCRIPTION

setup_exec installs architecture-dependent executables from either from a local tape drive or a remote host. It is used to convert a standalone system or homogeneous file server to a heterogeneous file server. **setup_exec** is a forms-based utility that can be invoked directly, but it is also used by **suninstall(8)**. It can only be invoked by the super-user.

The *arch* argument specifies the machine architecture to install (for instance, *sun4*, *sun3*...). When run with no arguments, **setup_exec** displays a usage line that includes the proper format of the *arch* argument for each supported architecture. *execpath* is the full pathname of the directory in which to install the executables. When **setup_exec** is done, the *execpath* directory is ready to mount as */usr* by the heterogeneous server's NFS clients of the indicated *arch*.

setup_exec also updates the */etc/exports* file (see **exportfs(8)**) to export the executable directories it has installed. The system administrator should verify this file to make sure that the directory has been exported to the correct groups.

EXAMPLE

This example shows how to install a directory of executables for Sun-4 system clients running 4.0.

```
example# setup_exec sun4 /exports/execs/sun4/4.0
```

FILES

<i>/etc/hosts</i>	hosts database
<i>/etc/ethers</i>	database of hostnames and Ethernet addresses
<i>/etc/exports</i>	database of exported file systems
<i>/usr/etc/install/files/extractlist.arch</i>	record of extracted categories for the indicated architecture

SEE ALSO

exportfs(8), **setup_client(8)**, **suninstall(8)**

Installing the SunOS

DIAGNOSTICS**incorrect number of arguments**

Check the number and the order of arguments.

invalid architecture type “arch”.

You supplied a value for *arch* that is not supported.

invalid tape drive type “drive”.

Valid tape drive types are *local* and *remote*.

invalid tape type “tape”.

Valid tape types are *ar*, *st*, *mt*, and *xt*.

can't reach tapehost “tapehost”.

The IP address of *tapehost* is not in the hosts database, that is, the hosts YP database if the Yellow Pages are running, or the */etc/hosts* file otherwise.

Load release tape *n*

Mount the release tape specified on the screen and type RETURN to continue.

NAME

showmount – show all remote mounts

SYNOPSIS

`/usr/etc/showmount [-ade] [host]`

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

`showmount` lists all the clients that have remotely mounted a filesystem from *host*. This information is maintained by the `mountd(8C)` server on *host*, and is saved across crashes in the file `/etc/rmtab`. The default value for *host* is the value returned by `hostname(1)`.

OPTIONS

`-a` Print all remote mounts in the format
hostname:directory

where *hostname* is the name of the client, and *directory* is the root of the file system that has been mounted.

`-d` List directories that have been remotely mounted by clients.

`-e` Print the list of exported file systems.

FILES

`/etc/rmtab`

SEE ALSO

`hostname(1)`, `exports(5)`, `exports(5)`, `mountd(8C)`

BUGS

If a client crashes, its entry will not be removed from the list until it reboots and executes `'umount -a'`.

NAME

shutdown – close down the system at a given time

SYNOPSIS

`/usr/etc/shutdown [-fhknr] [time [warning-message ...]`

DESCRIPTION

shutdown provides an automated procedure to notify users when the system is to be shut down. *time* specifies when **shutdown** will bring the system down; it may be the word **now** (indicating an immediate shutdown), or it may specify a future time in one of two formats: *+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 in 24-hour notation.

At intervals that get closer as the apocalypse approaches, warning messages are displayed at terminals of all logged-in users, and of users who have remote mounts on that machine. 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 to the system log daemon, `syslogd(8)`, containing the time of shutdown, the instigator of the shutdown, and the reason. Then a terminate signal is sent to `init`, which brings the system down to single-user mode.

The time of the shutdown and the warning message are placed in `/etc/nologin`, which should be used to inform the users as to when the system will be back up, and why it is going down (or anything else).

OPTIONS

As an alternative to the above procedure, these options can be specified:

- f** Arrange, in the manner of `fastboot(8)`, that when the system is rebooted, the file systems will not be checked.
- h** Execute `halt(8)`.
- k** Simulate shutdown of the system. Do not actually shut down the system.
- n** Prevent the normal `sync(2)` before stopping.
- r** Execute `reboot(8)`.

FILES

<code>/etc/nologin</code>	tells login not to let anyone log in
<code>/etc/xtab</code>	list of remote hosts that have mounted this host

SEE ALSO

`login(1)`, `sync(2)`, `fastboot(8)`, `halt(8)`, `reboot(8)`, `syslogd(8)`

BUGS

Only allows you to bring the system down between “now” and 23:59 if you use the absolute time for shutdown.

NAME

spray – spray packets

SYNOPSIS

/usr/etc/spray host [-c count] [-d delay] [-i delay] [-l length] host

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

spray sends a one-way stream of packets to *host* using RPC, and reports how many were received, as well as the the transfer rate. The *host* argument can be either a name or an internet address.

OPTIONS**-c count**

Specify how many packets to send. The default value of *count* is the numbers of packets required to make the total stream size 100000 bytes.

-d delay

Specify how may microseconds to pause between sending each packet. The default is 0.

-i delay Use ICMP echo packets rather than RPC. Since ICMP automatically echos, this creates a two way stream.

-l length

The *length* parameter is the numbers of bytes in the ethernet packet that holds the RPC call message. Since the data is encoded using XDR, and XDR only deals with 32 bit quantities, not all values of *length* are possible, and spray rounds up to the nearest possible value. When *length* is greater than 1514, then the RPC call can no longer be encapsulated in one Ethernet packet, so the *length* field no longer has a simple correspondence to Ethernet packet size. The default value of *length* is 86 bytes (the size of the RPC and UDP headers)

SEE ALSO

icmp(4P), ping(8C), sprayd(8C)

NAME

sprayd – spray server

SYNOPSIS

/usr/etc/rpc.sprayd

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

rpc.sprayd is a server which records the packets sent by **spray(8C)**. The **rpc.sprayd** daemon is normally invoked by **inetd(8C)**.

SEE ALSO

inetd(8C), **spray(8C)**

NAME

statd – network status monitor

SYNOPSIS

/etc/rpc.statd

DESCRIPTION

statd is an intermediate version of the status monitor. It interacts with **lockd(8C)** to provide the crash and recovery functions for the locking services on NFS.

FILES

/etc/sm
/etc/sm.bak
/etc/state

SEE ALSO

statmon(5), **lockd(8C)**

BUGS

The crash of a site is only detected upon its recovery.

NAME

sticky – persistent text and append-only directories

DESCRIPTION

The *sticky bit* (file mode bit 01000, see `chmod(2)`) is used to indicate special treatment for certain executable files and directories.

Sticky Text Executable Files

While the sticky bit is set on a sharable executable file, the text of that file will not be removed from the system swap area. Thus the file does not have to be fetched from the file system upon each execution. As long as a copy remains in the swap area, the original text cannot be overwritten in the file system, nor can the file be deleted. Directory entries can be removed so long as one link remains.

Sharable executable files are made by the `-n` and `-z` options of `ld(1)`.

To replace a sticky file that has been used:

1. Clear the sticky bit with `chmod(1V)`.
2. Execute the old program to flush the swapped copy. This can be done safely even if others are using it.
3. Overwrite the sticky file. If the file is being executed by any process, writing will be prevented; it suffices to simply remove the file and then rewrite it, being careful to reset the owner and mode with `chmod` and `chown(2)`.
4. Set the sticky bit once again, if still needed.

Only the super-user can set the sticky bit on a sharable executable file.

Sticky Directories

A directory for which the sticky bit is set restricts deletion of files it contains. A file in a sticky directory may only be removed or renamed by a user who has write permission on the directory, and either owns the file, owns the directory, or is the super-user. This is useful for directories such as `/tmp`, which must be publicly writable, but should deny users permission to arbitrarily delete or rename the files of others.

Any user may create a sticky directory. See `chmod` for details about modifying file modes.

BUGS

Since the text areas of sticky text executables are stashed in the swap area, abuse of the feature can cause a system to run out of swap.

Neither `open(2V)` nor `mkdir(2)` will create a file with the sticky bit set.

FILES

`/tmp`

SEE ALSO

`chmod(1V)`, `ld(1)`, `chmod(2)`, `chown(2)`, `mkdir(2)`, `open(2V)`

NAME

suninstall – install and upgrade the Sun Operating System

SYNOPSIS

/usr/etc/install/suninstall

DESCRIPTION

suninstall is a forms-based subsystem for installing and upgrading the Sun Operating System on Sun-2, Sun-3 and Sun-4 systems. Unlike previous installation subsystems, **suninstall** does not require you to recapitulate an interrupted procedure; it allows you to pick up from where you left off. A new invocation of **suninstall** displays the saved information, and gives you an opportunity to make any needed alterations, before it proceeds.

To abort the installation procedure, use the interrupt character (typically CTRL-C).

suninstall allows you to install the operating system onto any system configuration, be it standalone, data-less, a homogeneous file server, or a heterogeneous server. It allows you to install from any distribution tape format, to install the various versions of the operating system needed by clients on a heterogeneous file server; you can install as many different system versions as your disk space may allow.

You can use **suninstall** to convert a 4.0 standalone system into a 4.0 server, without taking down or rebuilding the system. After the initial installation, you can use **setup_client(8)**, to add or remove a diskless client while the server is running in multiuser mode. You can use **setup_exec(8)**, to convert a 4.0 standalone system or server into a heterogeneous file server while it is running multiuser.

USAGE

Refer to *Installing the SunOS* for more information on the various menus and selections.

FILES

/usr/etc/install	directory containing installation programs, scripts and files
/usr/etc/install/files	directory containing default data files for clients and hosts
/usr/etc/install/get_*_info	terminal data-entry forms
/usr/etc/install/installation	subsystem utility program
/usr/etc/install/makedir	subsystem utility program
/usr/etc/install/script	subsystem utility scripts
/usr/etc/install/xdrtoc	subsystem utility program

SEE ALSO

extract_unbundled(8), **setup_client(8)**, **setup_exec(8)**

Installing the SunOS

NAME

swapon – specify additional device for paging and swapping

SYNOPSIS

/usr/etc/swapon -a

/usr/etc/swapon name...

DESCRIPTION

swapon specifies additional devices on which paging and swapping are to take place. The system begins by swapping and paging on only a single device so that only one disk is required at bootstrap time. Calls to **swapon** normally occur in the system multi-user initialization file **/etc/rc** making all swap devices available, so that the paging and swapping activity is interleaved across several devices.

The second form gives individual block devices as given in the system swap configuration table. The call makes only this space available to the system for swap allocation.

OPTIONS

-a Make available all devices of type **swap** in **/etc/fstab**. Using **swapon** with the **-a** option is the normal usage.

FILES

/dev/[ru][pk]?b normal paging devices

/etc/fstab

/etc/rc

SEE ALSO

swapon(2), **init(8)**

BUGS

There is no way to stop paging and swapping on a device. It is therefore not possible to make use of devices which may be dismounted during system operation.

NAME

sysdiag – system diagnostics

SYNOPSIS

/usr/diag/sysdiag/sysdiag

AVAILABILITY

This program is available with the *User Diagnostics* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

sysdiag is a general-purpose system diagnostic facility that tests the system and reports its findings. It concentrates on three areas of system functionality; memory, peripherals and disk.

To use **sysdiag**, log on as **sysdiag**, then enter the command **sysdiag**.

sysdiag creates a **sunview(1)** environment with one window each for memory, peripherals, and disk error messages, plus a window for the console. It also creates date/time and performance monitor graphs. It places abbreviated error messages from the memory, disk, and peripherals in the appropriate windows, and sends console messages to the console window.

When called from a terminal **sysdiag** interleaves all its messages on the screen.

With or without the windows, it places long error messages in files named **log.xx.nn** where:

xx is the name of diagnostic

nn is the pass number (increments each pass)

After it completes its test, **sysdiag** displays the error log files by executing the command '**more log***'. These files remain after **sysdiag** exits.

sysdiag consists of a user account with a home directory, a collection of scripts, and executable files containing the actual test code.

To configure or change **sysdiag**, either change the shell commands in **/usr/diag/sysdiag/sysdiag**, or change the **sysdiag** user configuration files **.login**, **.sunview**, and **.cshrc**.

FILES

.login
.sunview
.cshrc

SEE ALSO

sunview(1) See the appropriate diagnostic manual for your Sun system.

NAME

syslogd – log system messages

SYNOPSIS

/usr/etc/syslogd [**-d**] [**-fconfigfile**] [**-m interval**]

DESCRIPTION

syslogd reads and forwards system messages to the appropriate log files and/or users, depending upon the priority of a message and the system facility from which it originates. The configuration file **/etc/syslog.conf** (see **syslog.conf(5)**) controls where messages are forwarded. **syslogd** logs a mark (timestamp) message every *interval* minutes (default 20) at priority **LOG_INFO** to the facility whose name is given as **mark** in the **syslog.conf** file.

A system message consists of a single line of text, which may be prefixed with a priority code number enclosed in angle-brackets (<>); priorities are defined in <sys/syslog.h>.

syslogd reads from the **AF_UNIX** address family socket **/dev/log**, from an Internet address family socket specified in **/etc/services**, and from the special device **/dev/klog** (for kernel messages).

syslogd reads the configuration file when it starts up, and again whenever it receives a HUP signal, at which time it also closes all files it has open, re-reads its configuration file, and then opens only the log files that are listed in that file. **syslogd** exits when it receives a TERM signal.

As it starts up, **syslogd** creates the file **/etc/syslog.pid**, if possible, containing its process ID.

Sun386i DESCRIPTION

syslogd translates messages using the databases specified on an optional line in the **syslog.conf** as indicated with a **translate** entry.

The format of these databases is described in **translate(5)**.

OPTIONS

-d	Turn on debugging.
-fconfigfile	Specify an alternate configuration file.
-m interval	Specify an interval, in minutes, between mark messages.

FILES

/etc/syslog.conf	configuration file
/etc/syslog.pid	process ID
/dev/log	AF_UNIX address family datagram log socket
/dev/klog	kernel log device
/etc/services	network services database

SEE ALSO

logger(1), **syslog(3)**, **syslog.conf(5)**

NAME

talkd – server for talk program

SYNOPSIS

/usr/etc/in.talkd

DESCRIPTION

talkd is a server used by the **talk(1)** program. It listens at the udp port indicated in the “talk” service description; see **services(5)**. The actual conversation takes place on a tcp connection that is established by negotiation between the two machines involved.

SEE ALSO

talk(1), services(5), inetd(8C)

BUGS

The protocol is architecture dependent, and can not be relied upon to work between Sun systems and other machines.

NAME

telnetd – DARPA TELNET protocol server

SYNOPSIS

`/usr/etc/in.telnetd`

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

`telnetd` is a server which supports the DARPA standard TELNET virtual terminal protocol. `telnetd` is invoked by the internet server (see `inetd(8C)`), normally for requests to connect to the TELNET port as indicated by the `/etc/services` file (see `services(5)`).

`telnetd` operates by allocating a pseudo-terminal device (see `pty(4)`) for a client, then creating a login process which has the slave side of the pseudo-terminal as its standard input, output, and error. `telnetd` manipulates the master side of the pseudo-terminal, implementing the TELNET protocol and passing characters between the remote client and the login process.

When a TELNET session is started up, `telnetd` sends TELNET options to the client side indicating a willingness to do *remote echo* of characters, to *suppress go ahead*, and to receive *terminal type information* from the remote client. If the remote client is willing, the remote terminal type is propagated in the environment of the created login process. The pseudo-terminal allocated to the client is configured to operate in “cooked” mode, and with XTABS, ICRNL, and ONLCR enabled (see `termio(4)`).

`telnetd` is willing to do: *echo*, *binary*, *suppress go ahead*, and *timing mark*. `telnetd` is willing to have the remote client do: *binary*, *terminal type*, and *suppress go ahead*.

SEE ALSO

`telnet(1C)`

Postel, Jon, and Joyce Reynolds, “Telnet Protocol Specification,” RFC 854, Network Information Center, SRI International, Menlo Park, Calif., May 1983.

BUGS

Some TELNET commands are only partially implemented.

The TELNET protocol allows for the exchange of the number of lines and columns on the user’s terminal, but `telnetd` doesn’t make use of them.

Because of bugs in the original 4.2 BSD `telnet(1C)`, `telnetd` performs some dubious protocol exchanges to try to discover if the remote client is, in fact, a 4.2 BSD `telnet(1C)`.

Binary mode has no common interpretation except between similar operating systems

The terminal type name received from the remote client is converted to lower case.

The *packet* interface to the pseudo-terminal (see `pty(4)`) should be used for more intelligent flushing of input and output queues.

`telnetd` never sends TELNET *go ahead* commands.

`telnetd` can only support 64 pseudo-terminals.

NAME

tftpd – DARPA Trivial File Transfer Protocol server

SYNOPSIS

/usr/etc/in.tftpd [**-s**] [*homedir*]

Sun386i SYNOPSIS

/usr/etc/in.tftpd [**-s**] [**-p**] [*homedir*]

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

tftpd is a server that supports the DARPA Trivial File Transfer Protocol (TFTP). This server is normally started by **inetd**(8C) and operates at the port indicated in the **tftp** Internet service description in the **/etc/inetd.conf** file; see **inetd.conf**(5) for details.

Before responding to a request, the server attempts to change its current directory to *homedir*; the default value is **/tftpboot**.

Sun386i DESCRIPTION

The **tftpd** daemon acts as described above, except that it will perform certain filename mapping operations unless instructed otherwise by the **-p** command line argument or when operating in a secure environment. This mapping affects only TFTP boot requests and will not affect requests for existing files.

The semantics of the changes are as follows. Only filenames of the format *ip-address* or *ip-address.arch*, where *ip-address* is the IP address in hex, and *arch* is the hosts's architecture (as returned by the **arch**(1) command), that do not correspond to files in **/tftpboot**, are mapped. If the address is known through a YP lookup, any file of the form **/tftpboot/ip-address*** (with or without a suffix) is returned. If there are multiple such files, any one may be returned. If the *ip-address* is unknown (that is if the **ipaloc** (8C) service says the name service does not know the address), the filename is mapped as follows: Names without the *arch* suffix are mapped into the name **pnp.SUN3**, and names with the suffix are mapped into **pnp.Arch**. That file is returned if it exists.

OPTIONS

-s Secure. When specified, the directory change must succeed; and the daemon also changes its root directory to *homedir*.

The use of *tftp* does not require an account or password on the remote system. Due to the lack of authentication information, *tftpd* will allow only publicly readable files to be accessed. Files may be written only if they already exist and are publicly writable. Note that this extends the concept of "public" to include all users on all hosts that can be reached through the network; this may not be appropriate on all systems, and its implications should be considered before enabling this service.

tftpd runs with the user ID and group ID set to **-2**, under the assumption that no files exist with that owner or group. However, nothing checks this assumption or enforces this restriction.

Sun386i OPTIONS

-p Disable pnp entirely. Do not map filenames.

Sun386i FILES

/tftpboot/* filenames are IP addresses

SEE ALSO

ipalocd(8C), **netconfig**(8C), **inetd**(8C), **tftp**(1C)

Sollins, K.R., *The TFTP Protocol (Revision 2)*, RFC 783, Network Information Center, SRI International, Menlo Park, Calif., June 1981.

Sun386i WARNINGS

A request for an *ip-address* from a Sun-4 can be satisfied by a file named *ip-address.386* for compatibility with some early Sun-4 PROM monitors.

NAME

`tic` – terminfo compiler

SYNOPSIS

`tic [-v[n]] [-c] filename`

DESCRIPTION

Note: Optional Software (System V Option). Refer to *Installing the SunOS* for information on how to install this command.

`tic` compiles a `terminfo(5V)` source file into the compiled format. The results are placed in the directory `/usr/share/lib/terminfo`. The compiled format is used by the `curses(3V)` library.

Each entry in the file describes the capabilities of a particular terminal. When a `use=entry` field is given in a terminal entry, `tic` reads in the binary (compiled) description of the indicated `entry` from `/usr/share/lib/terminfo` to duplicate the contents of that entry within the one being compiled. However, if an `entry` by that name is specified in `filename`, the entry in that source file is used first. Also, if a capability is defined in both entries, the definition in the current entry's source file is used.

If the environment variable `TERMINFO` is set, that directory is searched and written to instead of `/usr/share/lib/terminfo`.

OPTIONS

`-v[n]`

Verbose. Display trace information on the standard error. The optional integer argument is a number from 1 to 10, inclusive, indicating the desired level of detail. If `n` is omitted, the default is 1.

`-c`

Only check `filename` for errors. Errors in `use=` links are not detected.

FILES

`/usr/share/lib/terminfo/?/*`

compiled terminal description data base

SEE ALSO

`fork(2)`, `curses(3V)`, `curses(3X)`, `malloc(3)`, `term(5)`, `terminfo(5V)`

BUGS

Total compiled entries cannot exceed 4096 bytes. The name field cannot exceed 1024 bytes.

When the `-c` option is used, duplicate terminal names will not be diagnosed; however, when `-c` is not used, they will be.

For backward compatibility, cancelled capabilities will not be marked as such within the `terminfo` binary unless the entry name has a '+' within it. Such terminal names are only used for inclusion with a `use=` field, and typically aren't used for actual terminal names.

DIAGNOSTICS

Most diagnostic messages produced by `tic` are preceded with the approximate line number and the name of the entry being processed.

`mkdir name` returned bad status

The named directory could not be created.

File does not start with terminal names in column one

The first thing seen in the file, after comments, must be the list of terminal names.

Token after a `seek(2)` not NAMES

Somehow the file being compiled changed during the compilation.

Not enough memory for `use_list` element

Out of memory

Not enough free memory was available (`malloc(3)` failed).

Can't open *filename*

The named file could not be opened or created.

Error in writing *filename*

The named file could not be written to.

Can't *link filename to filename*

A link failed.

Error in re-reading compiled *filename*

The compiled file could not be read back in.

Premature EOF

The current entry ended prematurely.

Backspaced off beginning of line

This error indicates something wrong happened within **tic**.

Unknown Capability – *filename*

The named invalid capability was found within the file.

Wrong type used for capability ...

For example, a string capability was given a numeric value.

Unknown token type

Tokens must be followed by '@' to cancel, ',' for booleans, '#' for numbers, or '=' for strings.

name*: bad term name*Line *n*: Illegal terminal name – *name*****Terminal names must start with a letter or digit**

The given name was invalid. Names must not contain white space or slashes, and must begin with a letter or digit.

***name*: terminal name too long.**

An extremely long terminal name was found.

***name*: terminal name too short.**

A one-letter name was found.

***name* defined in more than one entry. Entry being used is *name* .**

An entry was found more than once.

Terminal name *name* synonym for itself

A name was listed twice in the list of synonyms.

At least one synonym should begin

At least one of the names of the terminal should begin with a letter.

Illegal character – *c*

The given invalid character was found in the input file.

Newline in middle of terminal name

The trailing comma was probably left off of the list of names.

Missing comma

A comma was missing.

Missing numeric value

The number was missing after a numeric capability.

NULL string value

The proper way to say that a string capability does not exist is to cancel it.

Very long string found. Missing comma?

Self-explanatory.

Unknown option. Usage is:

An invalid option was entered.

Too many file names. Usage is:

Self-explanatory.

***name* non-existent or permission denied**

The given directory could not be written into.

***name* is not a directory**

Self-explanatory.

***name*: Permission denied**

Access denied.

***name*: Not a directory**

tic wanted to use the given name as a directory, but it already exists as a file

SYSTEM ERROR!! Fork failed!!!

A fork(2) failed.

Error in following up use-links.

Either there is a loop in the links or they reference non-existent terminals. The following is a list of the entries involved:

A *terminfo(5V)* entry with a *use=name* capability either referenced a non-existent terminal called *filename* or *filename* somehow referred back to the given entry.

NAME

timed – DARPA Time server

SYNOPSIS

/usr/etc/in.timed

DESCRIPTION

timed is a server which supports the DARPA Time Server Protocol. The time server operates at the port indicated in the “time” service description; see **services(5)**, and is invoked by **inetd(8C)** each time there is a connection to the time server.

SEE ALSO

services(5), **rdate(8)**, **inetd(8C)**

BUGS

A more sophisticated facility that can accept broadcasts and synchronize clocks over an internet is needed.

NAME

tnamed – DARPA Trivial name server

SYNOPSIS

/usr/etc/in.tnamed [-v]

DESCRIPTION

tnamed is a server that supports the DARPA Name Server Protocol. The name server operates at the port indicated in the “name” service description (see **services(5)**), and is invoked by **inetd(8C)** when a request is made to the name server.

Two known clients of this service are the MIT PC/IP software the Bridge boxes.

OPTIONS

-v Invoke the daemon in verbose mode.

SEE ALSO

uucp(1C), **services(5)**, **inetd(8C)**

Postel, Jon, *Internet Name Server*, IEN 116, SRI International, Menlo Park, California, August 1979.

BUGS

The protocol implemented by this program is obsolete. Its use should be phased out in favor of the Internet Domain protocol. See **named(8C)**.

NAME

trpt – transliterate protocol trace

SYNOPSIS

/usr/etc/trpt [**-afjst**] [**-p***hex-address*] [*system* [*core*]]

DESCRIPTION

trpt interrogates the buffer of TCP trace records created when a socket is marked for “debugging” (see **getsockopt(2)**), and prints a readable description of these records. When no options are supplied, **trpt** prints all the trace records found in the system grouped according to TCP connection protocol control block (PCB). The following options may be used to alter this behavior.

OPTIONS

- a** In addition to the normal output, print the values of the source and destination addresses for each packet recorded.
- f** Follow the trace as it occurs, waiting a short time for additional records each time the end of the log is reached.
- j** Just give a list of the protocol control block addresses for which there are trace records.
- s** In addition to the normal output, print a detailed description of the packet sequencing information.
- t** In addition to the normal output, print the values for all timers at each point in the trace.
- p** *hex-address*
Show only trace records associated with the protocol control block, the address of which follows.

The recommended use of **trpt** is as follows. Isolate the problem and enable debugging on the socket(s) involved in the connection. Find the address of the protocol control blocks associated with the sockets using the **-A** option to **netstat(8C)**. Then run **trpt** with the **-p** option, supplying the associated protocol control block addresses. The **-f** option can be used to follow the trace log once the trace is located. If there are many sockets using the debugging option, the **-j** option may be useful in checking to see if any trace records are present for the socket in question.

If debugging is being performed on a system or core file other than the default, the last two arguments may be used to supplant the defaults.

FILES

/vmunix
/dev/kmem

SEE ALSO

getsockopt(2), **netstat(8C)**

DIAGNOSTICS**no namelist**

When the system image does not contain the proper symbols to find the trace buffer; others which should be self explanatory.

BUGS

Should also print the data for each input or output, but this is not saved in the trace record.

The output format is inscrutable and should be described here.

NAME

tunefs – tune up an existing file system

SYNOPSIS

/usr/etc/tunefs [-a maxcontig] [-d rotdelay] [-e maxbpg] [-m minfree] special | filesystem

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 OPTIONS given below:

OPTIONS**-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: 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), dumpfs(8), mkfs(8), newfs(8)

System and Network Administration

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.

NAME

tzsetup – set up old-style time zone information in the kernel

SYNOPSIS

`/usr/etc/tzsetup`

DESCRIPTION

tzsetup attempts to find the offset from GMT and old-style Daylight Savings Time correction type (see `gettimeofday(2)`) that most closely matches the default time zone for the machine, and to pass this information to the kernel with a `settimeofday()` call (see `gettimeofday(2)`). This is necessary if programs built under releases of SunOS prior to 4.0 are to be run; those programs get time zone information from the kernel using `gettimeofday`.

If it cannot find the offset from GMT, the offset is set to 0; if it cannot find the Daylight Savings Time correction type, it is set to `DST_NONE`, indicating that no Daylight Savings Time correction is to be performed.

DIAGNOSTICS

tzsetup: Can't open `/usr/share/lib/zoneinfo/localtime`: *reason*

The time zone file for the current time zone could not be opened.

tzsetup: Error reading `/usr/lib/zoneinfo/localtime`: *reason*

The time zone file for the current time zone could not be read.

tzsetup: Two or more time zone types are equally valid — no DST selected

There were two or more Daylight Savings Time correction types that generated results that were equally close to the correct results. None of them was selected. Programs built under versions of SunOS prior to 4.0 may not convert dates correctly.

tzsetup: No old-style time zone type is valid — no DST selected

None of the Daylight Savings Time correction types generated results that were in any way correct; none of them was selected. Programs built under versions of SunOS prior to 4.0 may not convert dates correctly.

tzsetup: Warning: No old-style time zone type is completely valid

None of the Daylight Savings Time correction types generated results that were completely correct; the best of them was selected. Programs built under versions of SunOS prior to 4.0 may not convert dates correctly.

tzsetup: Can't set time zone

tzsetup was run by a user other than the super-user; only the super-user may change the kernel's notion of the current time zone.

SEE ALSO

`gettimeofday(2)`, `tzfile(5)`, `zic(8)`

NAME

unconfigure -- reset the network configuration for a Sun386i system

SYNOPSIS

/usr/etc/unconfigure [-y]

AVAILABILITY

Sun386i systems only.

DESCRIPTION

unconfigure restores most of the system configuration and status files to the state they were in when delivered by Sun Microsystems, Inc. It also deletes all user accounts (including home directories), Yellow Pages information, and any diskless client configurations that were set up.

After running **unconfigure**, a system halts. Rebooting it to multi-user mode at this point will start automatic system installation.

unconfigure is intended for use in the following situations:

- As one of the final steps in Software Manufacturing.
- In systems being set up with temporary configurations, holding no user accounts or diskless clients. These will occur during demonstrations and evaluation trials.
- To allow systems that had been used as standalones to be upgraded to join a network in a role other than as a master server. (See instructions later.)

unconfigure is potentially a dangerous utility; it does not work unless invoked by the super-user. As a warning, unless the -y option is passed, it will require confirmation that all user files and system software configuration information is to be deleted.

This utility is *not* recommended for routine use of any sort.

Resetting Temporary Configurations

If users need to set up and tear down configurations, **unconfigure** can be used to restore the system to an essentially as-manufactured state. The main concern here is that user accounts will be deleted, so this should not be done casually.

To reset a temporary configuration, just become the super-user and invoke **unconfigure**.

Upgrading Standalones to Network Clients

Systems that are going to be networked should be networked from the very first, if at all possible. This eliminates whole classes of compatibility problems, such as pathnames and (in particular) user account clashes.

Automatic system installation directly supports upgrading a single standalone system to a YP master, and joining any number of unused systems (or systems upon which **unconfigure** has been run) into a network.

However, in the situation where standalone systems that have been used extensively are to be joined to a network, **unconfigure** can be used in conjunction with automatic system installation by a knowledgeable super-user to change a system's configuration from standalone to network client. This procedure is not recommended for use by inexperienced administrators.

The following procedure is not needed unless user accounts or other data need to be preserved; it is intended to ensure that every UID and GID is changed so as not to clash with those in use on the network. It must be applied to each system that is being upgraded from a standalone to a network client.

The procedure is as follows:

1. Identify all accounts and files that you'll want to save. If there are none, just run **unconfigure** and install the system on the network. Do not follow the remaining steps.
2. Copy **/etc/passwd** to **/etc/passwd.bak**.
3. Rename all the files (including home directories) so that they aren't deleted. (See FILES below.) These will probably be only in **/export/home**.

4. Run **unconfigure** and install the system on the network.
5. For each account listed in `/etc/passwd.bak` that you want to save, follow this procedure:
 - a. Create a new account on the network; if the UID and GID are the same as in `/etc/passwd.bak` on the standalone, then skip the next step. However, be sure that you do not make two different accounts with the same UID.
 - b. Use the `'chown -R'` command to change the ownership of the home directories.
 - c. You may need to rename the files you just chowned above, for example to ensure that they are the user's home directory. This may involve updating the `auto.home(5)` and `auto.home(5)` YP maps, as well.
6. Delete `/etc/passwd.bak`.

FILES

unconfigure deletes the following files, if they are present, replacing some of them with the distribution version if one is supposed to exist:

<code>/etc/rootkey</code>	<code>/etc/ethers</code>	<code>/etc/localtime</code>	<code>/etc/publickey</code>
<code>/etc/auto.home</code>	<code>/etc/exports</code>	<code>/etc/net.conf</code>	<code>/etc/sendmail.cf</code>
<code>/etc/auto.vol</code>	<code>/etc/fstab</code>	<code>/etc/netmasks</code>	<code>/etc/syslog.conf</code>
<code>/etc/bootparams</code>	<code>/etc/group</code>	<code>/etc/networks</code>	<code>/etc/systems</code>
<code>/etc/bootservers</code>	<code>/etc/hosts</code>	<code>/etc/passwd</code>	<code>/single/ifconfig</code>
<code>/var/sysex/*</code>			

and all files in `/var/yp` except those distributed with the operating system.

unconfigure truncates all files in `/var/adm`. All user home directories in `/export/home` are deleted, except those for the default user account users, which is shipped with the operating system. All diskless client configuration information stored in `/export/roots`, `/export/swaps`, and `/export/dumps` is deleted.

SEE ALSO

find(1), **passwd(5)**, **group(5)**, **adduser(8)**, **chgrp(1)**, **chown(8)**

BUGS

More of the system configuration files should be reset.

This does not yet support taking a workstation off the network temporarily, for example, to take it home over the weekend for use as a standalone, or to move it to another network while travelling. This should be the default behavior.

The procedure for upgrading standalones to network clients should be automated; currently, only upgrading a standalone to a master server is automated.

NAME

update – periodically update the super block

SYNOPSIS

/usr/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), init(8), sync(8)

NAME

uuclean – uucp spool directory clean-up

SYNOPSIS

/usr/lib/uucp/uuclean [**-m**] [**-ntime**] [**-ppre**]

DESCRIPTION

uuclean scans the spool directory for files with the specified prefix and deletes all those which are older than the specified number of hours.

OPTIONS

- m** Send mail to the owner of the file when it is deleted.
- ntime** Files whose age is more than *time* hours are deleted if the prefix test is satisfied (default time is 72 hours).
- ppre** Scan for files with *pre* as the file prefix. Up to 10 **-p** arguments may be specified. A **-p** without any *pre* following deletes all files older than the specified time.

uuclean will typically be started by **cron(8)**.

FILES

/usr/lib/uucp	directory with commands used by uuclean internally
/usr/lib/uucp/spool	spool directory

SEE ALSO

uucp(1C), **uux(1C)**, **cron(8)**

NAME

vipw – edit the password file

SYNOPSIS

/usr/etc/vipw

DESCRIPTION

vipw edits the password file while setting the appropriate locks, and does any necessary processing after the password file is unlocked. If the password file is already being edited, then you will be told to try again later. The **vi(1)** editor will be used unless the environment variable **VISUAL** or **EDITOR** indicates an alternate editor.

vipw performs a number of consistency checks on the password entry for root, and will not allow a password file with a “mangled” root entry to be installed. It also checks the **/etc/shells** file to verify the login shell for root.

FILES

/etc/ptmp
/etc/shells

SEE ALSO

passwd(1), **vi(1)**, **passwd(5)**, **adduser(8)**

NAME

vmstat – report virtual memory statistics

SYNOPSIS

vmstat [*-fisS*] [*interval* [*count*]]

DESCRIPTION

vmstat delves into the system and normally reports certain statistics kept about process, virtual memory, disk, trap and CPU activity.

Without options, vmstat displays a one-line summary of the virtual memory activity since the system has been booted. If *interval* is specified, vmstat summarizes activity over the last *interval* seconds. If a *count* is given, the statistics are repeated *count* times.

For example, the following command displays a summary of what the system is doing every five seconds. This is a good choice of printing interval since this is how often some of the statistics are sampled in the system.

example% vmstat 5

```

procs      memory                page   faults
r b w    avm fre re at pi po fr de sr x0 x1 x2 x3 in sy cs us sy id
2 0 0    918 286 0 0 0 0 0 0 0 1 0 0 0  4 12  5  3  5 91
1 0 0    846 254 0 0 0 0 0 0 0 6 0 1  0 42 153 31  7 40  54
1 0 0    840 268 0 0 0 0 0 0 0 5 0 0  0 27 103 25  8 26  66
1 0 0    620 312 0 0 0 0 0 0 0 6 0 0  0 26  76 25  6 27  67

```

^C

example%

The fields of vmstat's display are:

- procs** Report the number of processes in each of the three following states:
- r** in run queue
 - b** blocked for resources (i/o, paging, etc.)
 - w** runnable or short sleeper (< 20 secs) but swapped
- memory** Report on usage of virtual and real memory. Virtual memory is considered active if it belongs to processes which are running or have run in the last 20 seconds.
- avm** number of active virtual Kbytes
 - fre** size of the free list in Kbytes
- page** Report information about page faults and paging activity. The information on each of the following activities is averaged each five seconds, and given in units per second.
- re** page reclaims — but see the *-S* option for how this field is modified.
 - at** number of attaches — but see the *-S* option for how this field is modified.
 - pi** kilobytes per second paged in
 - po** kilobytes per second paged out
 - fr** kilobytes freed per second
 - de** anticipated short term memory shortfall in Kbytes
 - sr** pages scanned by clock algorithm, per-second
- disk** Report number of disk operations per second (this field is system dependent). For Sun systems, four slots are available for up to four drives: "x0" (or "s0" for SCSI disks), "x1", "x2", and "x3".
- faults** Report trap/interrupt rate averages per second over last 5 seconds.
- in** (non clock) device interrupts per second
 - sy** system calls per second
 - cs** CPU context switch rate (switches/sec)

cpu Give a breakdown of percentage usage of CPU time.
us user time for normal and low priority processes
sy system time
id CPU idle

OPTIONS

- f** Report on the number of forks and vforks since system startup and the number of pages of virtual memory involved in each kind of fork.
- i** Report the number of interrupts per device. Autovectored interrupts (including the clock) are listed first.
- s** Display the contents of the sum structure, giving the total number of several kinds of paging-related events which have occurred since boot.
- S** Report on swapping rather than paging activity. This option will change two fields in vmstat's "paging" display: rather than the "re" and "at" fields, vmstat will report "si" (swap-ins), and "so" (swap-outs).

FILES

/dev/kmem
/vmunix

BUGS

If more than one autovectored device has the same name, interrupts are counted for all like-named devices regardless of unit number. Such devices are listed with a unit number of '?'.

NAME

ypinit - build and install Yellow Pages database

SYNOPSIS

/usr/etc/yp/ypinit -m

/usr/etc/yp/ypinit -s *master_name*

DESCRIPTION

ypinit sets up a Yellow Pages database on a YP server. It can be used to set up a master or a slave server. You must be the super-user to run it. It asks a few, self-explanatory questions, and reports success or failure to the terminal.

It sets up a master server using the simple model in which that server is master to all maps in the data base. This is the way to bootstrap the YP system; later if you want you can change the association of maps to masters. All databases are built from scratch, either from information available to the program at runtime, or from the ASCII data base files in /etc. These files are listed below under FILES. All such files should be in their "traditional" form, rather than the abbreviated form used on client machines.

A YP database on a slave server is set up by copying an existing database from a running server. The *master_name* argument should be the hostname of YP server (either the master server for all the maps, or a server on which the data base is up-to-date and stable).

Read **ypfiles(5)** and **ypserv(8)** for an overview of the Yellow Pages.

OPTIONS

- m** Indicate that the local host is to be the YP master.
- s** Set up a slave database.

FILES

/etc/passwd
/etc/group
/etc/hosts
/etc/networks
/etc/services
/etc/protocols
/etc/ethers

SEE ALSO

makedbm(8), **ypfiles(5)**, **ypmake(8)**, **yppush(8)**, **ypserv(8)**, **ypxfr(8)**

NAME

yppmake – rebuild Yellow Pages database

SYNOPSIS

cd /var/yp ; make [map]

DESCRIPTION

The file called **Makefile** in **/var/yp** is used by **make** to build the Yellow Pages database. With no arguments, **make** creates **dbm** databases for any YP maps that are out-of-date, and then executes **yppush(8)** to notify slave databases that there has been a change.

If you supply a *map* on the command line, **make** will update that map only. Typing **make passwd** will create and **yppush** the password database (assuming it is out of date). Likewise, **make hosts** and **make networks** will create and **yppush** the host and network files, **/etc/hosts** and **/etc/networks**.

There are three special variables used by **make**: **DIR**, which gives the directory of the source files; **NO-PUSH**, which when non-null inhibits doing a **yppush** of the new database files; and **DOM**, used to construct a domain other than the master's default domain. The default for **DIR** is **/etc**, and the default for **NO-PUSH** is the null string.

Refer to **ypfiles(5)** and **ypserv(8)** for an overview of the YP.

FILES

/var/yp
/etc/hosts
/etc/networks

SEE ALSO

make(1), **ypfiles(5)**, **makedbm(8)**, **yppush(8)**, **ypserv(8)**

NAME

yppasswdd – server for modifying Yellow Pages password file

SYNOPSIS

/usr/etc/rpc.yppasswdd filename [adjunct_file] [-m argument1 argument2 ...]

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

yppasswdd is a server that handles password change requests from **yppasswd(1)**. Unless an *adjunct_file* is specified, it changes a password entry in *filename*, which is assumed to be in the format of **passwd(5)**. If an *adjunct_file* is specified or **/etc/security/passwd.adjunct** exists, this file will be changed instead of the *filename*. An entry in *filename* or *adjunct_file* will only be changed if the password presented by **ypasswd(1)** matches the encrypted password of that entry.

If the **-m** option is given, then after *filename* or *adjunct_file* is modified, a **make(1)** will be performed in */var/yp*. Any arguments following the flag will be passed to *make*.

This server is not run by default, nor can it be started up from **inetd(8C)**. If it is desired to enable remote password updating for the Yellow Pages, then an entry for **yppasswdd** should be put in the **/etc/rc** file of the host serving as the master for the Yellow Pages **passwd** file.

EXAMPLE

If the Yellow Pages password file is stored as **/var/yp/src/passwd**, then to have password changes propagated immediately, the server should be invoked as

```
/usr/etc/rpc.yppasswdd /var/yp/src/passwd -m passwd DIR=/var/yp/src
```

In this case, **src** is the YP domain name.

FILES

/var/yp/Makefile
/etc/security/passwd.adjunct
/etc/rc

SEE ALSO

make(1), **yppasswd(1)**, **passwd(5)**, **passwd.adjunct(5)**, **ypfiles(5)**, **inetd(8C)**, **ypmake(8)**

NAME

yppoll - what version of a YP map is at a YP server host

SYNOPSIS

/usr/etc/yp/yppoll [**-h** *host*] [**-d** *domain*] *mapname*

DESCRIPTION

yppoll asks a **ypserv(8)** process what the order number is, and which host is the master YP server for the named map. If the server is a v.1 YP protocol server, **yppoll** uses the older protocol to communicate with it. In this case, it also uses the older diagnostic messages in case of failure.

OPTIONS

-h *host* Ask the **ypserv** process at *host* about the map parameters. If *host* is not specified, the YP server for the local host is used. That is, the default host is the one returned by **ypwhich(8)**.

-d *domain*

Use *domain* instead of the default domain.

SEE ALSO

ypfiles(5), **ypserv(8)**, **ypwhich(8)**

NAME

yppush - force propagation of a changed YP map

SYNOPSIS

/usr/etc/yp/yppush [**-v**] [**-d** *domain*] *mapname*

DESCRIPTION

yppush copies a new version of a Yellow Pages (YP) map from the master YP server to the slave YP servers. It is normally run only on the master YP server by the Makefile in **/var/yp** after the master databases are changed. It first constructs a list of YP server hosts by reading the YP map **ypservers** within the *domain*. Keys within the map **ypservers** are the ASCII names of the machines on which the YP servers run.

A "transfer map" request is sent to the YP server at each host, along with the information needed by the transfer agent (the program which actually moves the map) to call back the **yppush**. When the attempt has completed (successfully or not), and the transfer agent has sent **yppush** a status message, the results may be printed to stdout. Messages are also printed when a transfer is not possible; for instance when the request message is undeliverable, or when the timeout period on responses has expired.

Refer to **ypfiles(5)** and **ypserv(8)** for an overview of the Yellow Pages.

OPTIONS

- d** Specify a *domain*.
- v** Verbose. This causes messages to be printed when each server is called, and for each response. If this flag is omitted, only error messages are printed.

FILES

/var/yp/domain/ypservers.{dir,pag}
/var/yp

SEE ALSO

ypfiles(5), **ypserv(8)**, **ypxfr(8)**, YP protocol specification

BUGS

In the current implementation (version 2 YP protocol), the transfer agent is **ypxfr(8)**, which is started by the **ypserv** program. If **yppush** detects that it is speaking to a version 1 YP protocol server, it uses the older protocol, sending a version 1 YPPROC_GET request and issues a message to that effect. Unfortunately, there is no way of knowing if or when the map transfer is performed for version 1 servers. **yppush** prints a message saying that an "old-style" message has been sent. The system administrator should later check to see that the transfer has actually taken place.

NAME

ypserv, ypbind – Yellow Pages server and binder processes

SYNOPSIS

`/usr/etc/ypserv`

`/usr/etc/ypbind`

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

The Yellow Pages (YP) provides a simple network lookup service consisting of databases and processes. The databases are `dbm(3X)` files in a directory tree rooted at `/var/yp`. These files are described in `ypfiles(5)`. The processes are `/usr/etc/ypserv`, the YP database lookup server, and `/usr/etc/ypbind`, the YP binder. The programmatic interface to YP is described in `ypclnt(3N)`. Administrative tools are described in `yppush(8)`, `ypxfr(8)`, `yppoll(8)`, `ypwhich(8)`, and `ypset(8)`. Tools to see the contents of YP maps are described in `ypcat(1)`, and `ypmatch(1)`. Database generation and maintenance tools are described in `ypinit(8)`, `ypmake(8)`, and `makedbm(8)`.

Both `ypserv` and `ypbind` are daemon processes typically activated at system startup time from `/etc/rc.local`. `ypserv` runs only on YP server machines with a complete YP database. `ypbind` runs on all machines using YP services, both YP servers and clients.

The `ypserv` daemon's primary function is to look up information in its local database of YP maps. The operations performed by `ypserv` are defined for the implementor by the *YP Protocol Specification*, and for the programmer by the header file `<rpcsvc/yp_prot.h>`. Communication to and from `ypserv` is by means of RPC calls. Lookup functions are described in `ypclnt(3N)`, and are supplied as C-callable functions in the C library. There are four lookup functions, all of which are performed on a specified map within some YP domain: *Match*, *Get_first*, *Get_next*, and *Get_all*. The *Match* operation takes a key, and returns the associated value. The *Get_first* operation returns the first key-value pair from the map, and *Get_next* can be used to enumerate the remainder. *Get_all* ships the entire map to the requester as the response to a single RPC request.

Two other functions supply information about the map, rather than map entries: *Get_order_number*, and *Get_master_name*. In fact, both order number and master name exist in the map as key-value pairs, but the server will not return either through the normal lookup functions. (If you examine the map with `makedbm(8)`, however, they will be visible.) Other functions are used within the YP subsystem itself, and are not of general interest to YP clients. They include *Do_you_serve_this_domain?*, *Transfer_map*, and *Reinitialize_internal_state*.

The function of `ypbind` is to remember information that lets client processes on a single node communicate with some `ypserv` process. `ypbind` must run on every machine which has YP client processes; `ypserv` may or may not be running on the same node, but must be running somewhere on the network.

The information `ypbind` remembers is called a *binding* — the association of a domain name with the internet address of the YP server, and the port on that host at which the `ypserv` process is listening for service requests. This information is cached in the directory `/var/yp/binding` using a filename of `domainname.version`.

The process of binding is driven by client requests. As a request for an unbound domain comes in, the `ypbind` process broadcasts on the net trying to find a `ypserv` process that serves maps within that domain. Since the binding is established by broadcasting, there must be at least one `ypserv` process on every net. If the client is running in C2 secure mode, then `ypbind` will only accept bindings to servers where the `ypserv` process is running as root. Once a domain is bound by a particular `ypbind`, that same binding is given to every client process on the node. The `ypbind` process on the local node or a remote node may be queried for the binding of a particular domain by using the `ypwhich(1)` command.

Bindings and rebindings are handled transparently by the C library routines. If `ypbind` is unable to speak to the `ypserv` process it's bound to, it marks the domain as unbound, tells the client process that the domain is unbound, and tries to bind the domain once again. Requests received for an unbound domain will wait until the domain requested is bound. In general, a bound domain is marked as unbound when the node running `ypserv` crashes or gets overloaded. In such a case, `ypbind` will to bind any YP server (typically one that is less-heavily loaded) available on the net.

`ypbind` also accepts requests to set its binding for a particular domain. The request is usually generated by the YP subsystem itself. `ypset(8)` is a command to access the `Set_domain` facility. It is for unsnarling messes. Note that the `Set Domain` procedure only accepts requests from processes running as root.

FILES

If the file `/var/yp/ypserv.log` exists when `ypserv` starts up, log information will be written to this file when error conditions arise.

The file(s) `/var/yp/binding/domainname.version` will be created to speed up the binding process. These files cache the last successful binding created for the given domain, when a binding is requested these files are checked for validity and then used.

`/var/yp`
`/usr/etc/ypbind`

SEE ALSO

`domainname(1)`, `ypcat(1)`, `ypmatch(1)`, `dbm(3X)`, `ypclnt(3N)`, `ypfiles(5)`, `makedbm(8)`, `ypmake(8)`, `ypinit(8)`, `yppoll(8)`, `yppush(8)`, `ypset(8)`, `ypwhich(8)`, `ypxfr(8)`

YP Protocol Specification in Network Programming

NOTES

Both `ypbind` and `ypserv` support multiple domains. The `ypserv` process determines the domains it serves by looking for directories of the same name in the directory `/var/yp`. It will reply to all broadcasts requesting yp service for that domain. Additionally, the `ypbind` process can maintain bindings to several domains and their servers, the default domain is however the one specified by the `domainname(1)` command at startup time.

NAME

`ypset` - point `ypbind` at a particular server

SYNOPSIS

```
/usr/etc/yp/ypset [ -V1|-V2 ] [ -d domain ] [ -h host ] server
```

DESCRIPTION

`ypset` tells `ypbind` to get YP services for the specified *domain* from the `ypserv` process running on *server*. If *server* is down, or isn't running `ypserv`, this is not discovered until a YP client process tries to get a binding for the domain. At this point, the binding set by `ypset` will be tested by `ypbind`. If the binding is invalid, `ypbind` will attempt to rebind for the same domain.

`ypset` is useful for binding a client node which is not on a broadcast net, or is on a broadcast net which isn't running a YP server host. It also is useful for debugging YP client applications, for instance where a YP map only exists at a single YP server host.

In cases where several hosts on the local net are supplying YP services, it is possible for `ypbind` to rebind to another host even while you attempt to find out if the `ypset` operation succeeded. For example, you can type:

```
example% ypset host1
example% ypwhich
host2
```

which can be confusing. This is a function of the YP subsystem's attempt to load-balance among the available YP servers, and occurs when *host1* does not respond to `ypbind` because it is not running `ypserv` (or is overloaded), and *host2*, running `ypserv`, gets the binding.

server indicates the YP server to bind to, and can be specified as a name or an IP address. If specified as a name, `ypset` will attempt to use YP services to resolve the name to an IP address. This will work only if the node has a current valid binding for the domain in question. In most cases, *server* should be specified as an IP address.

Refer to `ypfiles(5)` and `ypserv(8)` for an overview of the Yellow Pages.

OPTIONS

`-V1` Bind *server* for the (old) v.1 YP protocol.

`-V2` Bind *server* for the (current) v.2 YP protocol.

If no version is supplied, `ypset`, first attempts to set the domain for the (current) v.2 protocol. If this attempt fails, `ypset`, then attempts to set the domain for the (old) v.1 protocol.

`-h host` Set `ypbind`'s binding on *host*, instead of locally. *host* can be specified as a name or as an IP address.

`-d domain`

Use *domain*, instead of the default domain.

SEE ALSO

`ypwhich(1)`, `ypfiles(5)`, `ypserv(8)`

NAME

`ypupdated` – server for changing YP information

SYNOPSIS

`rpc.yupdated [-is]`

DESCRIPTION

`ypupdated` is a daemon that updates information in the Yellow Pages, normally started up by `inetd(8C)`. `ypupdated` consults the file `updaters(5)` in the directory `/var/yp` to determine which YP maps should be updated and how to change them.

By default, the daemon requires the most secure method of authentication available to it, either DES (secure) or UNIX (insecure).

OPTIONS

- `-s` accept only calls authenticated using the secure RPC mechanism (AUTH_DES authentication). This disables programmatic updating of YP maps unless the network supports these calls.
- `-i` also accept RPC calls with the insecure AUTH_UNIX credentials. This allows programmatic updating of YP maps in all networks.

FILES

`/var/yp/updaters`

SEE ALSO

`updaters(5)`, `inetd(8C)`, `keyserv(8C)`

System and Network Administration

Network Programming

NAME

`ypwhich` – what machine is the YP server?

SYNOPSIS

`ypwhich [-d domainname] [hostname]`

`ypwhich [-d domainname] [-t] -m [mname]`

AVAILABILITY

This program is available with the *Networking Tools and Programs* software installation option. Refer to *Installing the SunOS* for information on how to install optional software.

DESCRIPTION

`ypwhich` tells which YP server supplies Yellow Pages to a YP client, and which YP server is the master for a map. If invoked without arguments, it gives the YP server for the local machine. If *hostname* is specified, that machine is queried to find out which YP master it is using.

If the `-m` switch is used without *mname*, a list of every map in the domain and the master of each will be printed. If *mname* is specified, only the master YP server for that map is printed. *mname* may be a mapname, or a nickname for a mapname. Mapnames and nicknames are described in `ypcat(1)`.

OPTIONS

- `-d` *Domainname* specifies the name of a YP domain. The default is the default domain for the local machine.
- `-m` Find the master YP server for a map, or for all maps in a domain. No *hostname* may be specified with `-m`.
- `-t` Inhibit nickname translation; useful if there is a mapname identical to a nickname. This is not true of any Sun-supplied map.

SEE ALSO

`ypcat(1)`, `ypfiles(5)`, `rpcinfo(8C)`, `yppush(8)`, `ypserv(8)`

NAME

ypxfr – transfer YP map from a YP server to here

SYNOPSIS

/usr/etc/yp/ypxfr [**-f**] [**-c**] [**-d domain**] [**-h host**] [**-s domain**] [**-C tid prog ipadd port**] **mapname**

DESCRIPTION

ypxfr moves a YP map in the default domain for the local host to the local host by making use of normal YP services. It creates a temporary map in the directory **/var/yp/domain** (this directory must already exist; *domain* is the default domain for the local host), fills it by enumerating the map's entries, fetches the map parameters (master and order number), and loads them. It then deletes any old versions of the map and moves the temporary map to the real *mapname*.

If run interactively, **ypxfr** it writes its output to the terminal. However, if it is invoked without a controlling terminal, and if the log file **/var/yp/ypxfr.log** exists, it will append all its output to that file. Since **ypxfr** is most often run from the super-user's **crontab** file, or by **ypserv**, you can use the log file to retain a record of what was attempted, and what the results were.

If **issecure(3)** is TRUE, **ypxfr** requires that **ypserv** on the *host* be running as root. If the map being transferred is a secure map, **ypxfr** sets the permissions on the map to 0600.

For consistency between servers, **ypxfr** should be run periodically for every map in the YP data base. Different maps change at different rates: the *services.byname* map may not change for months at a time, for instance, and may therefore be checked only once a day (in the wee hours). You may know that *mail.aliases* or *hosts.byname* changes several times per day. In such a case, you may want to check hourly for updates. A **crontab(5)** entry can be used to perform periodic updates automatically. Rather than having a separate **crontab** entry for each map, you can group comands to update several maps in a shell script. Examples (mnemonically named) are in **/usr/etc/yp**: **ypxfr_1perday**, **ypxfr_2perday**, and **ypxfr_1perhour**. They can serve as reasonable first cuts.

Refer to **ypfiles(5)** and **ypserv(8)** for an overview of the Yellow Pages.

OPTIONS

- f** Force the transfer to occur even if the version at the master is not more recent than the local version.
- c** Do not send a "Clear current map" request to the local **ypserv** process. Use this flag if **ypserv** is not running locally at the time you are running **ypxfr**. Otherwise, **ypxfr** will complain that it can't talk to the local **ypserv**, and the transfer will fail.
- ddomain**
Specify a domain other than the default domain.
- hhost** Get the map from *host*, regardless of what the map says the master is. If *host* is not specified, **ypxfr** will ask the YP service for the name of the master, and try to get the map from there. *host* may be a name or an internet address in the form *a.b.c.d*.
- sdomain**
Specify a source domain from which to transfer a map that should be the same across domains (such as the *services.byname* map).
- Ctid prog ipadd port**
This option is **only** for use by **ypserv**. When **ypserv** invokes **ypxfr**, it specifies that **ypxfr** should call back a **yppush** process at the host with IP address *ipaddr*, registered as program number *prog*, listening on port *port*, and waiting for a response to transaction *tid*.

FILES

/var/yp/ypxfr.log log file
/usr/etc/yp/ypxfr_1perday
 script to run one transfer per day, for use with **cron(8)**

/usr/etc/yp/ypxfr_2perday script to run two transfers per day
/usr/etc/yp/ypxfr_1perhour script for hourly transfers of volatile maps
/var/yp/domain YP domain
/var/spool/cron/crontabs/root Super-user's crontab file

SEE ALSO

issecure(3), crontab(5), ypfiles(5), cron(8), ypserv(8), yppush(8),
YP Protocol Specification, in Network Programming

NAME

`zdump` – time zone dumper

SYNOPSIS

`zdump` [*-v*] [*-c cutoffyear*] [*zonename ...*]

DESCRIPTION

`zdump` prints the current time in each *zonename* named on the command line.

OPTIONS

-v For each *zonename* on the command line, print the current time, the time at the lowest possible time value, the time one day after the lowest possible time value, the times both one second before and exactly at each time at which the rules for computing local time change, the time at the highest possible time value, and the time at one day less than the highest possible time value. Each line ends with *isdst=1* if the given time is Daylight Saving Time or *isdst=0* otherwise.

-c cutoffyear

Cut off the verbose output near the start of the year *cutoffyear*.

FILES

`/usr/share/lib/zoneinfo` standard zone information directory

SEE ALSO

`ctime(3)`, `tzfile(5)`, `zic(8)`

NAME

zic – time zone compiler

SYNOPSIS

zic [-v] [-d *directory*] [-l *localtime*] [*filename ...*]

DESCRIPTION

zic reads text from the file(s) named on the command line and creates the time conversion information files specified in this input. If a *filename* is '-', the standard input is read.

Input lines are made up of fields. Fields are separated from one another by any number of white space characters. Leading and trailing white space on input lines is ignored. An '#' (unquoted sharp character) in the input introduces a comment which extends to the end of the line the sharp character appears on. White space characters and sharp characters may be enclosed in '"' (double quotes) if they're to be used as part of a field. Any line that is blank (after comment stripping) is ignored. Non-blank lines are expected to be of one of three types: rule lines, zone lines, and link lines.

A rule line has the form

```
Rule NAME FROM TO TYPE IN ON AT SAVE LETTER/S
```

For example:

```
Rule USA 1969 1973 - Apr lastSun 2:00 1:00 D
```

The fields that make up a rule line are:

- NAME** Gives the (arbitrary) name of the set of rules this rule is part of.
- FROM** Gives the first year in which the rule applies. The word **minimum** (or an abbreviation) means the minimum year with a representable time value. The word **maximum** (or an abbreviation) means the maximum year with a representable time value.
- TO** Gives the final year in which the rule applies. In addition to **minimum** and **maximum** (as above), the word **only** (or an abbreviation) may be used to repeat the value of the **FROM** field.
- TYPE** Gives the type of year in which the rule applies. If **TYPE** is '-' then the rule applies in all years between **FROM** and **TO** inclusive; if **TYPE** is **uspres**, the rule applies in U.S. Presidential election years; if **TYPE** is **nonpres**, the rule applies in years other than U.S. Presidential election years. If **TYPE** is something else, then zic executes the command

```
yearistype year type
```

to check the type of a year: an exit status of zero is taken to mean that the year is of the given type; an exit status of one is taken to mean that the year is not of the given type.

- IN** Names the month in which the rule takes effect. Month names may be abbreviated.
- ON** Gives the day on which the rule takes effect. Recognized forms include:

```
5          the fifth of the month
lastSun    the last Sunday in the month
lastMon     the last Monday in the month
Sun>=8     first Sunday on or after the eighth
Sun<=25    last Sunday on or before the 25th
```

Names of days of the week may be abbreviated or spelled out in full. Note: there must be no spaces within the ON field.

AT Gives the time of day at which the rule takes effect. Recognized forms include:

2 time in hours
2:00 time in hours and minutes
15:00
 24-hour format time (for times after noon)
1:28:14
 time in hours, minutes, and seconds

Any of these forms may be followed by the letter **w** if the given time is local "wall clock" time or **s** if the given time is local "standard" time; in the absence of **w** or **s**, wall clock time is assumed.

SAVE Gives the amount of time to be added to local standard time when the rule is in effect. This field has the same format as the **AT** field (although, of course, the **w** and **s** suffixes are not used).

LETTER/S

Gives the "variable part" (for example, the "S" or "D" in "EST" or "EDT") of time zone abbreviations to be used when this rule is in effect. If this field is '-', the variable part is null.

A zone line has the form

Zone	NAME	GMTOFF	RULES/SAVE	FORMAT	[UNTIL]
------	------	--------	------------	--------	---------

For example:

Zone	Australia/South-west	9:30	Aus	CST	1987 Mar 15 2:00
------	----------------------	------	-----	-----	------------------

The fields that make up a zone line are:

NAME The name of the time zone. This is the name used in creating the time conversion information file for the zone.

GMTOFF

The amount of time to add to GMT to get standard time in this zone. This field has the same format as the **AT** and **SAVE** fields of rule lines; begin the field with a minus sign if time must be subtracted from GMT.

RULES/SAVE

The name of the rule(s) that apply in the time zone or, alternately, an amount of time to add to local standard time. If this field is '-' then standard time always applies in the time zone.

FORMAT

The format for time zone abbreviations in this time zone. The pair of characters **%s** is used to show where the "variable part" of the time zone abbreviation goes. **UNTIL** The time at which the GMT offset or the rule(s) change for a location. It is specified as a year, a month, a day, and a time of day. If this is specified, the time zone information is generated from the given GMT offset and rule change until the time specified.

The next line must be a "continuation" line; this has the same form as a zone line except that the string "Zone" and the name are omitted, as the continuation line will place information starting at the time specified as the **UNTIL** field in the previous line in the file used by the previous line. Continuation lines may contain an **UNTIL** field, just as zone lines do, indicating that the next line is a further continuation.

A link line has the form

Link LINK-FROM LINK-TO

For example:

Link US/Eastern EST5EDT

The **LINK-FROM** field should appear as the **NAME** field in some zone line; the **LINK-TO** field is used as an alternate name for that zone.

Except for continuation lines, lines may appear in any order in the input.

OPTIONS

-v Complain if a year that appears in a data file is outside the range of years representable by system time values (0:00:00 AM GMT, January 1, 1970, to 3:14:07 AM GMT, January 19, 2038).

-d *directory*

Create time conversion information files in the **directory** rather than in the standard directory **/usr/share/lib/zoneinfo**.

-l *timezone*

Use the time zone **timezone** as local time. **zic** will act as if the file contained a link line of the form

Link *timezone* localtime

FILES

/usr/share/lib/zoneinfo standard directory used for created files

SEE ALSO

time(1V), **ctime(3)**, **tzfile(5)**, **zdump(8)**

NOTE

For areas with more than two types of local time, you may need to use local standard time in the **AT** field of the earliest transition time's rule to ensure that the earliest transition time recorded in the compiled file is correct.

Index

Special Characters

- !
 - history substitution — *cs*h, 98
 - logical negation operator — *cs*h, 101
- ! mail command, 296
- != — not equal to operator — *cs*h, 101
- !~ globbing pattern mismatch operator — *cs*h, 101
- # mail command, 296
- #! invoke shell to process script, 103
- \$ — variable substitution, 100
- \$# — word count for variable, 100
- \$\$ — process number of shell, 101
- \$< — read value from terminal — *cs*h, 101
- \$? — variable set inquiry — *cs*h, 101
- *
 - job control, reference to current job — *cs*h, 103
 - job to foreground/background — *cs*h, 108
 - modular division operator — *cs*h, 101
- &
 - bitwise AND operator — *cs*h, 101
 - run command in background, 97
- &&
 - execute on success — *cs*h, 97
 - logical AND operator — *cs*h, 101
- ' quote character, 97
- ()
 - command grouping — *cs*h, 97
 - group operators — *cs*h, 101
- *
 - filename wild card, zero or more of any characters, 101
 - integer multiplication operator — *cs*h, 101
- + — integer addition operator — *cs*h, 101
- — integer subtraction operator — *cs*h, 101
- d C shell file inquiry — directory, 102
- e C shell file inquiry — file exists, 102
- f C shell file inquiry — plain file, 102
- o C shell file inquiry — ownership, 102
- r C shell file inquiry — read accessible, 102
- w C shell file inquiry — write accessible, 102
- x C shell file inquiry — execute accessible, 102
- z C shell file inquiry — zero length, 102
- . (dot) command, 457
- / — integer division operator — *cs*h, 101
- : command, 104, 457
- : modifiers — history substitution — *cs*h, 98
- ; — command separation, 97
- <
 - less than operator — *cs*h, 101
 - redirect standard input, 99
- <<
 - bitwise shift left — *cs*h, 101
 - parse and pass input to command, 99
- <= — less than or equal to operator — *cs*h, 101
- = mail command, 296
- == — is equal to operator — *cs*h, 101
- ==~ — globbing pattern match operator — *cs*h, 101
- >
 - greater than operator — *cs*h, 101
 - redirect standard output, 99
- >& — redirect standard output and standard error — *cs*h, 99
- >= — greater than or equal to operator — *cs*h, 101
- >>
 - append standard output, 99
 - bitwise shift right — *cs*h, 101
- >>& — append standard output and standard error — *cs*h, 99
- ? — filename wild card, any single characters, 101
- ? mail command, 296
- @ — arithmetic on variables — *cs*h, 109
- [] — filename substitution, any character in list or range, 101
- " quote character, 97
- \ escape character, 97
- \!* — alias substitution, include command-line arguments — *cs*h, 99
- ^
 - bitwise XOR operator — *cs*h, 101
 - quick substitution — *cs*h, 99
- _toupper — convert character to upper-case, System V, 1134
- ` — command substitution, 101
- { } — filename substitution, successive strings in enclosed list, 101, 102
- |
 - bitwise OR operator — *cs*h, 101
 - pipe standard output, 97
- | mail command, 298
- |& — pipe standard output and standard error — *cs*h, 97
- ||
 - execute on failure — *cs*h, 97
 - logical OR operator — *cs*h, 101
- ~
 - filename substitution, home directory, 101
 - one's complement operator — *cs*h, 101
- ~! — mail tilde escape, 294

~. — mail tilde escape, 294
 ~: — mail tilde escape, 294
 ~< — mail tilde escape, 295
 ~? — mail tilde escape, 294
 ~_ — mail tilde escape, 294
 ~| — mail tilde escape, 295
 ~A — mail tilde escape, 294
 ~b — mail tilde escape, 295
 ~c — mail tilde escape, 295
 ~d — mail tilde escape, 295
 ~e — mail tilde escape, 295
 ~f — mail tilde escape, 295
 ~h — mail tilde escape, 295
 ~i — mail tilde escape, 295
 ~m — mail tilde escape, 295
 ~p — mail tilde escape, 295
 ~q — mail tilde escape, 295
 ~r — mail tilde escape, 295
 ~s — mail tilde escape, 295
 ~t — mail tilde escape, 295
 ~v — mail tilde escape, 295
 ~w — mail tilde escape, 295
 ~x — mail tilde escape, 295

1

1/2-inch tape drive
 tm — tapemaster, 1318
 xt — Xylogics 472, 1331
 1/4-inch tape drive
 ar — Archive 1/4-inch Streaming Tape Drive, 1193
 st — Sysgen SC 4000 (Archive) Tape Drive, 1292 *thru* 1293
 10 Mb/s 3Com Ethernet interface — ec, 1214
 10 Mb/s Sun Ethernet interface — ie, 1224 *thru* 1225
 10 Mb/s Sun-3/50 Ethernet interface — le, 1254 *thru* 1255

3

3Com 10 Mb/s Ethernet interface — ec, 1214

4

450 SMD Disk driver — xy, 1332 *thru* 1333
 451 SMD Disk driver — xy, 1332 *thru* 1333
 472 1/2-inch tape drive — xt, 1331

7

7053 SMD Disk driver — xd, 1329 *thru* 1330

8

8530 SCC serial communications driver — zs, 1335 *thru* 1336

A

a.out — assembler and link editor output, 1339
 a641 — convert long integer to base-64 ASCII, 809
 abort — generate fault, 810
 abort printer — lpc, 1662
 abs — integer absolute value, 811
 absolute value — abs, 811
 ac — login accounting, 1574
 accept — connection on socket, 628

access, 629
 access times of file, change — utime, 1030
 access times of file, change — utimes, 787
 accounting, display login record — ac, 1574
 process accounting, display record — sa, 1755
 process accounting, on or off — accton, 1755
 process accounting, turn on or off — acct, 630
 accounting file — acct, 1365
 acct
 acct — execution accounting file, 1365
 acct — process accounting on or off, 630
 accton — processing accounting on or off, 1755
 acos — trigonometric arccosine, 1106
 acosh — inverse hyperbolic function, 1088
 Adaptec ST-506 disk driver — sd, 1289 *thru* 1290
 adb — debugger, 13
 adb scripts — adbgen, 1575
 adbgen — generate adb script, 1575
 add password file entry — putpwent, 951
 add route ioctl — SIOCADDRT, 1288
 addbib — create bibliography, 18
 addexportent () function, 845
 additional paging/swapping devices, specify — swapon, 1771
 admntent — get filesystem descriptor file entry, 872
 address resolution display and control — arp, 1579
 adduser — add new user account, 1577
 adjacentscreens, 20
 adjtime — adjust time, 631
 admin — administer SCCS, 21
 adventure — exploration game, 1495
 advise paging system — vadvise, 788
 agt_create () function, 1056
 agt_enumerate () function, 1056
 agt_trap () function, 1056
 aint — aint of, 1102
 alarm — schedule signal, 812
 alias command, 104
 alias mail command, 296
 alias substitution — in C shell, 99
 aliases — sendmail aliases file, 1367
 align_equals — textedit selection filter, 529
 allnet mail variable, 301
 alloca — allocate on stack, 926
 allocate aligned memory — memalign, 925
 allocate aligned memory — valloc, 926
 allocate memory — calloc, 925
 allocate memory — malloc, 925
 allocate on stack — alloca, 926
 allow messages — msg, 327
 alphasort — sort directory, 983
 alter process priority — renice, 1728
 alternates mail command, 296
 alwaysignore mail variable, 301
 anint — anint of, 1102
 ANSI standard terminal emulation, 1204 *thru* 1208
 ANSI terminal emulation — console, 1204 *thru* 1209
 append mail variable, 301
 ar — library maintenance, 24

- ar — Archive 1/4-inch Streaming Tape Drive, 1193
 - ar — archive file format, 1370
 - arc — plot arc, 941
 - arch — display Sun architecture, 26
 - archive
 - ar — library maintenance, 24
 - cpio — copy archive, 87
 - archive file format — ar, 1370
 - archive tapes
 - tar, 37, 509
 - argument list processing — in C shell, 96
 - argument lists, varying length — varargs, 1032
 - argv variable, 109
 - arithmetic — drill in number facts, 1496
 - arp — address resolution display and control, 1579
 - arp ioctl
 - SIOCDDARP — delete arp entry, 1194
 - SIOCGARP — get arp entry, 1194
 - SIOCSARP — set arp entry, 1194
 - arp — Address Resolution Protocol, 1194 *thru* 1195
 - as — assembler, 27
 - ASCII
 - string to long integer — strtol, 1008
 - to integer — atoi, 1008
 - to long — atol, 1008
 - ASCII dump file — od, 348
 - ascii — ASCII character set, 1548
 - ASCII string to double — strtod, 1007
 - ASCII to Ethernet address — ether_aton, 841
 - ASCII to float — atof, 1007
 - asctime — date and time conversion, 824
 - asctime — date and time conversion, System V, 1132
 - asin — trigonometric arcsine, 1106
 - asinh — inverse hyperbolic function, 1088
 - askcc mail variable, 301
 - asksub mail variable, 301
 - assembler output — a.out, 1339
 - assert — program verification, 813
 - assert — program verification, System V, 1131
 - assign buffering to stream
 - setbuf — assign buffering, 987
 - setbuf — assign buffering, System V, 1175
 - setbuffer — assign buffering, 987
 - setbuffer — assign buffering, System V, 1175
 - setlinebuf — assign buffering, 987
 - setlinebuf — assign buffering, System V, 1175
 - setvbuf — assign buffering, 987
 - setvbuf — assign buffering, System V, 1175
 - assign to memory characters — memset, 928
 - async_daemon, 718
 - at — do job at specified time, 29
 - atan — trigonometric arctangent, 1106
 - atan2 — trigonometric arctangent, 1106
 - atanh — inverse hyperbolic function, 1088
 - atof — ASCII to float, 1007
 - atoi — ASCII to integer, 1008
 - atol — ASCII to long, 1008
 - atq — display delayed execution queue, 31
 - atrm — remove delayed execution jobs, 32
 - attributes of file `fstat`, 774
 - attributes of file `lstat`, 774
 - attributes of file `stat`, 774
 - audit — maintain audit trail, 1580
 - audit — audit trail file, 1372, 1374, 1376
 - audit () function, 632
 - audit_args () function, 814
 - audit_text () function, 814
 - audit_warn command, 1581
 - auditd daemon, 1582
 - auditon () function, 633
 - auditsvc () function, 634
 - auto.home, 1344
 - auto.vol, 1345
 - autoboot procedures — boot, 1586, 1650, 1727
 - automatic network install, 1115
 - automount command, 1583
 - autoprint mail variable, 301
 - awk — scan and process patterns, 33
- ## B
- backgammon — backgammon game, 1497
 - backquote substitution, 101
 - backspace magnetic tape files — mt, 333
 - backspace magnetic tape records — mt, 333
 - backup dumps — dump, 1612
 - bang mail variable, 301
 - banner — make posters, 36
 - banner — large banner, 1499
 - bar command, 37
 - bar — tape archive file format, 1346
 - basename — deliver portions of path names, 42
 - battlestar game, 1500
 - bc — calculator language, 43
 - bcd — convert to antique media, 1502
 - bcmp — compare byte strings, 819
 - bcopy — copy byte strings, 819
 - Bessel functions
 - j0, 1084
 - j1, 1084
 - jn, 1084
 - y0, 1084
 - y1, 1084
 - yn, 1084
 - bg command, 104
 - bibliography
 - addbib — create or extend, 18
 - indxbib — make inverted index, 235
 - lookbib — find bibliographic references, 276
 - refer — insert literature references, 415
 - roffbib — print literature references, 422
 - sortbib — sort bibliographic database, 473
 - biff — mail notifier, 45
 - binary file transmission
 - uudecode — decode binary file, 570
 - uuencode — encode binary file, 570
 - binary I/O, buffered
 - fread — read from stream, 855
 - fwrite — write to stream, 855
 - binary search of sorted table — bsearch, 816

- binary tree routines, 1021
 - bind, 636
 - bindresvport () function, 815
 - binmail — version 7 mail, 46
 - bioid daemon, 1703
 - bit string functions — *ffs*, 819
 - bj game, 1503
 - bk — machine-machine communication line discipline, 1196
 - bk ioctl's
 - TIOCGETD — get line discipline, 1196
 - TIOCSETD — set line discipline, 1196
 - block signals, 763
 - block size for tape — 512 bytes, 1612
 - blocked signals, release — *sigpause*, 764
 - blocks, count, in file — *sum*, 486
 - boggle — boggle game, 1504
 - boggletool — SunView game of boggle, 1505
 - boot — system startup procedures, 1586, 1650
 - boot parameter database — *bootparams*, 1377
 - bootparam protocol — *bootparam*, 1108
 - bootparamd daemon, 1589
 - bootparams — boot parameter database, 1377
 - bootstrap procedures — *boot*, 1586, 1650, 1727
 - bootstrap PROM monitor program, 1679
 - both real and effective group ID, set — *setgid*, 991
 - both real and effective group ID, set, System V — *setgid*, 1179
 - both real and effective user ID, set — *setuid*, 991
 - both real and effective user ID, set, System V — *setuid*, 1179
 - bouncedemo — bouncing square graphics demo, 1521
 - Bourne shell, *sh*, 452 *thru* 460
 - Bourne shell commands, 457
 - . command, 457
 - : command, 457
 - break command, 457
 - case command, 453
 - cd command, 458
 - continue command, 458
 - do command, 453
 - done command, 453
 - echo command, 458
 - elif command, 453
 - else command, 453
 - esac command, 453
 - eval command, 458
 - exec command, 458
 - exit command, 458
 - export command, 458
 - fi command, 453
 - for command, 453
 - hash command, 458
 - if command, 453
 - login command, 458
 - newgrp command, 458
 - pwd command, 458
 - read command, 458
 - readonly command, 459
 - return command, 459
 - set command, 459
 - shift command, 459
 - test command, 459
 - then command, 453
 - Bourne shell commands, *continued*
 - times command, 459
 - trap command, 459
 - type command, 459
 - umask command, 459
 - unset command, 459
 - until command, 453
 - wait command, 459
 - while command, 453
 - Bourne shell functions, 453
 - Bourne shell variables, 454 *thru* 455
 - CDPATH variable, 454
 - HOME variable, 454
 - IFS variable, 455
 - MAIL variable, 454
 - MAILCHECK variable, 454
 - MAILPATH variable, 454
 - PATH variable, 454
 - PS1 variable, 454
 - PS2 variable, 454
 - SHELL variable, 455
 - branch, C shell control flow, 102
 - break command, 104, 457
 - breaksw command, 104
 - brk — set data segment break, 638
 - broadcast messages to all users on network — *rwall*, 431
 - bsearch — binary search of a sorted table, 816
 - buffered binary I/O
 - fread* — read from stream, 855
 - fwrite* — write to stream, 855
 - buffered I/O library functions, introduction to, 999, 1183
 - buffering
 - assign to stream — *setbuf*, 987
 - assign to stream, System V — *setbuf*, 1175
 - assign to stream — *setbuffer*, 987
 - assign to stream, System V — *setbuffer*, 1175
 - assign to stream — *setlinebuf*, 987
 - assign to stream, System V — *setlinebuf*, 1175
 - assign to stream — *setvbuf*, 987
 - assign to stream, System V — *setvbuf*, 1175
 - build
 - programs — *make*, 355
 - random library — *ranlib*, 406
 - system configuration files — *config*, 1600
 - Yellow Pages database — *ypinit*, 1793
 - build programs — *make*, 311 *thru* 324
 - bwone — Sun-1 black and white frame buffer, 1197
 - bwtwo — Sun-3/Sun-2 black and white frame buffer, 1198
 - byte order, functions to convert between host and network, 820
 - byte string functions
 - bcmp*, 819
 - bcopy*, 819
 - bzero*, 819
 - bzero* — zero byte strings, 819
- ## C
- C compiler, 52
 - C library functions, introduction to, 797
 - C programming language
 - cflow* — code flow graph, 60
 - cpp* — C preprocessor, 89

C programming language, *continued*

ctags — create tags file, 115
 cxref — cross reference C program, 123
 indent — format C source, 231
 lint — C program verifier, 261
 mkstr — create C error messages, 329
 tcov — code coverage tool, 516
 vgrind — make formatted listings, 580
 xstr — extract strings from C code, 605

C shell

alias substitution, 99
 and Bourne shell scripts, 103
 argument list processing, 96
 arguments list — `argv` variable, 109
 branch, 102
 command execution, 103
 command inquiry, 102
 command substitution, 101
 commands, 104 *thru* 109
 conditional execution — `&&`, 97
 conditional execution — `||`, 97
`.cshrc` file, 96
 escape character, quotes and comments, 97
 expressions, 101
 file inquiries, 102
 filename completion, 97
 filename substitution, 101
 history substitution, 98
 I/O redirection, 99
 job control, 103
 lexical structure, 97
`.login` file, 96
`.logout` file, 96
 loop, 102
 operators, 101
 parentheses — command grouping, 97
 pipeline, 97
 quick substitution, 99
 signal handling, 103
 variable substitution, 100

C shell commands

`%` — job to foreground/background, 108
`:` — null command, 104
`@` — arithmetic on variables, 109
 alias — shell macros, 104
`bg` — job to background, 104
`break` — exit loop, 104
`breaksw` — exit switch, 104
`case` — selector in switch, 104
`cd` — change directory, 104
`chdir` — change directory, 104
`continue` — cycle loop, 104
`default` — catchall in switch, 104
`dirs` — print directory stack, 104
`echo` — echo arguments, 104
`else` — alternative commands, 105
`end` — end loop, 105
`endif` — end conditional, 105
`endsw` — end switch, 108
`eval` — re-evaluate shell data, 104
`exec` — execute command, 104
`exit` — exit shell, 105
`fg` — job to foreground, 105
`foreach` — loop on list of names, 105

C shell commands, *continued*

`glob` — filename expand wordlist, 105
`goto` — command transfer, 105
`hashstat` — display hashing statistics, 105
`history` — display history list, 105
`if` — conditional statement, 105
`jobs` — display job list, 105
`kill` — kill jobs and processes, 106
`limit` — alter resource limitations, 106
`login` — login new user, 106
`logout` — end session, 106
`nice` — run low priority process, 106
`nohup` — run command immune to hangups, 106
`notify` — request immediate notification, 106
`onintr` — handle interrupts in scripts, 106
`popd` — pop shell directory stack, 107
`pushd` — push shell directory stack, 107
`rehash` — recompute command hash table, 107
`repeat` — execute command repeatedly, 107
`set` — change value of shell variable, 107
`setenv` — set or display variables in environment, 107
`shift` — shift argument list, 107
`source` — read commands from file, 107
`stop` — halt job or process, 107
`suspend` — suspend shell, 107
`switch` — multi-way branch, 108
`time` — time command, 108
`umask` — change/display file creation mask, 108
`unalias` — remove aliases, 108
`unhash` — discard hash table, 108
`unlimit` — remove resource limitations, 108
`unset` — discard shell variables, 108
`unsetenv` — remove environment variables, 108
`wait` — wait for background process, 108

C shell metacharacters, 97

C shell variables, 109 *thru* 111

`argv`, 109
`cdpath`, 109
`cwd`, 109
`echo`, 109
`figignore`, 109
`filec`, 109
`hardpaths`, 109
`histchars`, 109
`history`, 109
`home`, 109
`ignoreeof`, 109
`mail`, 109
`nobeep`, 109
`noclobber`, 109
`noglob`, 110
`nomatch`, 110
`notify`, 110
`path`, 110
`prompt`, 110
`savehist`, 110
`shell`, 110
`status`, 110
`time`, 110
`verbose`, 110

`C2conv` — convert to C2 security, 1590

`cal` — display calendar, 48

calculator, 137

- calendar — reminder service, 49
- call-graph, display profile data — `gprof`, 214
- calloc — allocate memory, 925
- canfield — solitaire card game, 1507
- canvas_demo — canvas subwindow demo, 1542
- capitalize — `textedit` selection filter, 529
- captain command, 1591
- case command, 104, 453
- cat — concatenate files, 50
- C/A/T interpreter — `pti`, 363
- catman — create cat files for manual pages, 1593
- cb — format filter for C source files, 51
- cbirt — cube root function, 1105
- cc — C compiler, 52
- ccat — extract files compressed with `compact`, 350
- cd — change directory, 57
- cd command, 104, 458
- cd mail command, 296
- cdc — change delta commentary, 58
- cdpath variable, 109, 454
- ceil — ceiling of, 1102
- cflow — generate C flow graph, 60
- cfree — free memory, 925
- cgfour — Sun-3 color memory frame buffer, 1199
- cgone — Sun-1 color graphics interface, 1200
- cgthree — Sun386i color memory frame buffer, 1201
- cgtwo — Sun-3/Sun-2 color graphics interface, 1202
- change
 - audit characteristics, 1580
 - current working directory, 639
 - data segment size — `sbrk`, 638
 - delta commentary, 58
 - directory, 57
 - file access times — `utime`, 1030
 - file access times — `utimes`, 787
 - file mode — `chmod`, 640
 - file name — `rename`, 741
 - group ID of user — `newgrp`, 335
 - group ownership of file — `chgrp`, 63
 - login password — `passwd`, 379
 - login password in Yellow Pages — `yppasswd`, 611
 - mode of file, 65
 - name of file or directory — `mv`, 334
 - owner and group of file — `chown`, 642
 - owner of file — `chown`, 1595
 - permissions of file, 65
 - priority of command — `nice`, 336
 - process priority — `renice`, 1728
 - root directory — `chroot`, 644
 - working directory, 57
- change mapping protections — `mprotect`, 709
- change translation table entry `ioctl` — `KIOCSSETKEY`, 1235
- character
 - get from stdin — `getchar`, 861
 - get from stdin, System V — `getchar`, 1162
 - get from stream — `fgetc`, 861
 - get from stream, System V — `fgetc`, 1162
 - get from stream — `getc`, 861
 - get from stream, System V — `getc`, 1162
 - push back to stream — `ungetc`, 1028
 - put to stdin — `putchar`, 949
- character, *continued*
 - put to stream — `fputc`, 949
 - put to stream — `putc`, 949
- character classification
 - `isalnum`, 826
 - `isalpha`, 826
 - `isascii`, 826
 - `isctrl`, 826
 - `isdigit`, 826
 - `isgraph`, 826
 - `islower`, 826
 - `isprint`, 826
 - `ispunct`, 826
 - `isspace`, 826
 - `isupper`, 826
 - `isxdigit`, 826
- character classification, System V
 - `isalnum`, 1134
 - `isalpha`, 1134
 - `isascii`, 1134
 - `isctrl`, 1134
 - `isdigit`, 1134
 - `isgraph`, 1134
 - `islower`, 1134
 - `isprint`, 1134
 - `ispunct`, 1134
 - `isspace`, 1134
 - `isupper`, 1134
 - `isxdigit`, 1134
- character conversion
 - `toascii`, 826
 - `tolower`, 826
 - `toupper`, 826
- character conversion, System V
 - `_tolower`, 1134
 - `_toupper`, 1134
 - `toascii`, 1134
 - `tolower`, 1134
 - `toupper`, 1134
- character translation — `tr`, 544
- characters for equations — `eqnchar`, 1550
- characters in file, count — `wc`, 591
- chase — escape killer robots, 1508
- chdir, 639
- chdir command, 104
- chdir mail command, 296
- check buffer state `ioctl` — `GPIO_GET_GBUFFER_STATE`, 1221
- check directory — `dcheck`, 1608
- check file system — `fsck`, 1629
- CHECK () function, 1068
- check heap — `malloc_verify`, 926
- check quota consistency — `quotacheck`, 1721
- check spelling — `spell`, 474
- checkeq — check `eqn` constructs, 173
- checknr — check `nroff`/`troff` files, 62
- chess — chess game, 1509
- chesstool — SunView chess game, 1510
- chgrp — change group ID of file, 63
- ching — book of changes, 1511
- chkey command, 64

- chmod — change mode, 65, 640
- chown, 642
- chown — change owner of file, 1595
- chroot — change root directory, 644, 1596
- circle — plot circle, 941
- clean print queue — lpc, 1662
- clean UUCP spool area — uuclean, 1789
- clear — clear screen, 67
- clear inode — clri, 1598
- clear_colormap — make console text visible, 68
- clear_functions — reset SunView selection service, 69
- clearerr — clear error on stream, 849
- clearerr — clear error on stream, System V, 1159
- click — control keyboard click, 70
- client command, 1597
- clock — display time in window, 71, 821
- clone, STREAMS device driver, 1203
- close, 646
- close database — close, 830
- close directory stream — closedir, 834
- close — close database, 830
- close stream — fclose, 847
- closedir — close directory stream, 834
- closelog — close system log file, 1011
- closepl — close plot device, 941
- clri — clear inode, 1598
- cluster command, 72
- cmd mail variable, 301
- cmdtool — shell or program with SunView text facility, 73
- cmp — compare files, 76
- code coverage tool — tcov, 516
- code flow graph — cflow, 60
- code formatter
 - cb — C source format filter, 51
 - vgrind — troff preprocessor for listings, 580
 - indent — format C source, 231
- COFF, Sun386i executable file format, 1348
- col — filter reverse paper motions, 77
- colcrt — document previewer, 78
- color graphics interface
 - cgfour — Sun-3 color memory frame buffer, 1199
 - cgone — Sun-1 color graphics interface, 1200
 - cgthree — Sun386i color memory frame buffer, 1201
 - cgtwo — Sun-3/Sun-2 color graphics interface, 1202
- coloredit — edit icons, 79
- colrm — remove columns from file, 80
- columns
 - print in multiple — pr, 388
 - remove from file, 80, 121
- comb — combine deltas, 81
- combine SCCS deltas, 81
- comm — display common lines, 82
- command
 - change priority of — nice, 336
 - describe — whatis, 593
 - execution in C shell, 103
 - grouping in the C shell — (), 97
 - inquiry, in C shell, 102
 - locate — whereis, 594
- command, *continued*
 - process options in scripts — getopt, 209
 - return stream to remote — rcmd, 958
 - return stream to remote — rexec, 967
 - run immune to hangup — nohup, 342
 - substitution, 101
- commands
 - Bourne shell, 453, 457, 460
 - comm — display common lines, 82
 - help_viewer — get help_viewer, 225
 - logintool — graphic login interface, 1661
 - organizer, 374
- commands, introduction, 3
- communications
 - cu — connect to remote system, 533
 - enroll — enroll for secret mail, 604
 - mail — send and receive mail, 293 thru 304
 - mesg — permit or deny messages, 327
 - talk — talk to another user, 508
 - telnet — TELNET interface, 518
 - tip — connect to remote system, 533
 - uuclean — clean UUCP spool area, 1789
 - uucp — system to system copy, 568
 - uudecode — decode binary file, 570
 - uuencode — encode binary file, 570
 - uulog — UUCP log, 568
 - uuname — UUCP list of names, 568
 - uuse — send file to remote host, 571
 - uux — system to system command execution, 574
 - write — write to another user, 600
 - xget — receive secret mail, 604
 - xsend — send secret mail, 604
- compact — compress files, 350
- compare
 - byte strings — bcmp, 819
 - files, 76
 - files differentially, 150
 - files side-by-side, 445
 - memory characters — memcmp, 928
 - strings — strcmp, 1001
 - strings — strncmp, 1001
 - three-way differential — diff3, 152
 - versions of SCCS file — sccsdiff, 438
- compile regular expression — re_comp, 961
- compiler generator, 351
- compiler generators
 - lex — lexical analyzer generator, 258
 - yacc — parser generator, 607
- compiler preprocessors
 - cpp — C preprocessor, 89
- compilers
 - cc — C compiler, 52
 - rpcgen — generate RPC protocols, C header files, 424
- compress — compress files, 83
- comsat — biff server, 1599
- concatenate files — cat, 50
- concatenate strings
 - strcat, 1001
 - strncat, 1001
- config — build system configuration files, 1600
- configuration file, system log daemon — syslogd, 1439
- configuration files, build — config, 1600

- configure network interface parameters — `ifconfig`, 1642
- connect, 647
- connect to remote system
 - `cu`, 533
 - tip, 533
- connected peer, get name of, 679
- connection
 - accept on socket — `accept`, 628
 - listen for on socket — `listen`, 697
- `console` — console driver/terminal emulator, 1204 *thru* 1209
- console I/O `ioctl`, `TIOCCONS`, 1204
- `cont` — continue line, 941
- continue command, 104, 458
- control devices — `ioctl`, 692
- control flow — in C shell, 102
- control line printer — `lpc`, 1662 *thru* 1663
- control magnetic tape — `mt`, 333
- control resource consumption — `vlimit`, 1034
- control system log
 - close system log — `closelog`, 1011
 - set log priority mask — `setlogmask`, 1011
 - start system log — `openlog`, 1011
 - write to system log — `syslog`, 1011
- control terminal, hangup — `vhangup`, 790
- conv mail variable, 301
- convert
 - functions to between host and network byte order, 820
 - host to network long — `htonl`, 820
 - host to network short — `htons`, 820
 - network to host long — `ntohl`, 820
 - network to host short — `ntohs`, 820
 - spaces to tabs `unexpand`, 179
 - tabs to spaces `expand`, 179
- convert 8-bit rasterfile to 1-bit rasterfile— `rasfilter8to1`, 407
- convert and copy files, 139
- convert base-64 ASCII to long integer — 164a, 809
- convert character
 - to ASCII — `toascii`, 826
 - to ASCII, System V — `toascii`, 1134
 - to lower-case — `tolower`, 826
 - to lower-case, System V — `_tolower`, 1134
 - to lower-case, System V — `tolower`, 1134
 - to upper-case — `toupper`, 826
 - to upper-case, System V — `_toupper`, 1134
 - to upper-case, System V — `toupper`, 1134
- convert foreign font files — `vswap`, 585
- convert long integer to base-64 ASCII — 164a, 809
- convert numbers to strings
 - `econvert`, 838
 - `fconvert`, 838
 - `fprintf`, 944
 - `gconvert`, 838
 - `printf`, 944
 - `seconvert`, 838
 - `sfconvert`, 838
 - `sgconvert`, 838
 - `sprintf`, 944
- convert numbers to strings, System V
 - `sprintf`, 1168
- convert strings to numbers
 - `atof`, 1007
 - `atoi`, 1008
 - `atol`, 1008
 - `sscanf`, 984
 - `strtod`, 1007
 - `strtol`, 1008
- convert strings to numbers, *continued*
- convert strings to numbers, System V
 - `sscanf`, 1172
- convert time and date
 - `asctime`, 824
 - `ctime`, 824
 - `dysize`, 824
 - `gmtime`, 824
 - `localtime`, 824
 - `timegm`, 824
 - `timelocal`, 824
 - `tzset`, 824
 - `tzsetwall`, 825
- convert time and date, System V
 - `asctime`, 1132
 - `ctime`, 1132
 - `gmtime`, 1132
 - `localtime`, 1132
 - `timegm`, 1132
 - `timelocal`, 1132
 - `tzset`, 1132
 - `tzsetwall`, 1133
- convert units — `units`, 562
- copy
 - archives, 87
 - byte strings — `bcopy`, 819
 - files, 85
 - files from remote machine — `rcp`, 409
 - memory character fields — `memcpy`, 928
 - memory character strings — `memccpy`, 928
 - standard output to many files — `tee`, 517
 - strings — `strcpy`, 1001
 - strings — `strncpy`, 1001
- copy mail command, 296
- Copy mail command, 297
- `copysign()` function, 1093
- core — memory image file format, 1378
- core image, get of process — `gcore`, 202
- `cos` — trigonometric cosine, 1106
- `cosh` — hyperbolic cosine, 1088
- count blocks in file — `sum`, 486
- count lines, words, characters in file — `wc`, 591
- `cp` — copy files, 85
- `cpio` — copy archives, 87
- `cpio` — `cpio` archive format, 1380
- `cpp` — C preprocessor, 89
- CPU PROM monitor program, 1679
- craps game, 1512
- crash — crash information, 1604
- `creat`, 649
- create
 - bibliography — `addbib`, 18
 - cat files for manual pages — `catman`, 1593
 - delta — `delta`, 144
 - directory — `mkdir`, 328

create, continued

error log — `dmesg`, 1611
 fifo — `mknod`, 1674
 file — `open`, 719
 file system — `mkfs`, 1673
 font width table — `vwidth`, 587
 hash table — `hcreate`, 891
 interprocess communication channel — `pipe`, 722
 interprocess communication endpoint — `socket`, 771
 mail aliases database — `newaliases`, 1699
 name for temporary file — `tmpnam`, 1020
 named pipe — `mknod`, 1674
 new file system — `newfs`, 1700
 pair of connected sockets — `socketpair`, 773
 permuted index — `ptx`, 403
 prototype file system — `mkproto`, 1675
 random library — `ranlib`, 406
 SCCS data bases, 21
 SCCS delta — `delta`, 144
 script of terminal session — `script`, 444
 special file, 702
 special file — `mknod`, 1674
 symbolic link — `symlink`, 779
 system configuration files — `config`, 1600
 system log entry — `logger`, 271
 system log entry — `syslog`, 368
 tags file, 115
 unique file name — `mktmp`, 929
 Yellow Pages database — `ypinit`, 1793
 Yellow Pages dbm file — `makedbm`, 1667
create directory, 700
create new process, 661
cribbage — cribbage card game, 1514
cron — clock daemon, 1606
crontab command, 94
crontab — periodic jobs table, 1381
cross reference C program — `cxref`, 123
crt mail variable, 301
crypt — encrypt, 95
crypt — encryption, 822
csh, C shell, 96
.cshrc file, 96
csplit — split file into sections, 113
ctags — create tags file, 115
ctermid — generate filename for terminal, 823
ctime — date and time conversion, 824
ctime — date and time conversion, System V, 1132
ctrace — display program trace, 117
cu — connect to remote system, 533
current directory
 change, 639
 get pathname — `getwd`, 890
current domain, set or display name — `domainname`, 157
current host, get identifier of — `gethostid`, 672
current working directory — `getcwd`, 862
curses functions, System V, 1140
curses library routines, 827
cursor_demo — cursor attributes demo, 1542
curve fitting, `spline`, 476
cuserid — get user name, 829
cut — remove columns from file, 121

cv_broadcast() function, 1058
cv_create() function, 1058
cv_destroy() function, 1058
cv_enumerate() function, 1058
cv_notify() function, 1058
cv_send() function, 1058
cv_wait() function, 1058
cv_waiters() function, 1058
cwd variable, 109
cxref — cross reference C program, 123

D**daemons**

biod daemon, 1703
network file system, 718
 nfsd daemon, 1703
rquotad — remote quota server, 1747
sprayd — spray server, 1767

DARPA

Internet directory service — `whois`, 599
Internet file transfer protocol server — `ftpd`, 1632
DARPA Time server — `timed`, 1781
 to RPC mapper — `portmap`, 1712

DARPA Trivial name server — `tnamed`, 1782

Data Encryption Standard — `des`, 147
data segment size, change — `sbrk`, 638
data types — `types`, 1480

database functions — `dbm`

close, 830
dbminit, 830
delete, 830
fetch, 830
firstkey, 830
nextkey, 830
store, 830

database functions — `ndbm`

dbm_clearerr, 934
dbm_close, 934
dbm_delete, 934
dbm_err, 934
dbm_error, 934
dbm_fetch, 934
dbm_firstkey, 934
dbm_nextkey, 934
dbm_open, 934
dbm_store, 934

database library

-ldb option to `cc`, 830
ndbm, 934

database operator — `join`, 245

datafile

help — get help, 1353, 1355

date and time

get — time, 1016
get — `gettimeofday`, 689
get — `ftime`, 1016
set — `settimeofday`, 689

date and time conversion

asctime, 824
ctime, 824
dysize, 824

- date and time conversion, *continued*
 - gmtime, 824
 - localtime, 824
 - timegm, 824
 - timelocal, 824
 - tzset, 824
 - tzsetwall, 825
- date and time conversion, System V
 - asctime, 1132
 - ctime, 1132
 - gmtime, 1132
 - localtime, 1132
 - timegm, 1132
 - timelocal, 1132
 - tzset, 1132
 - tzsetwall, 1133
- date and time display — fdate, 848
- date — date and time, 124
- dbm_clearerr — clear ndbm database error condition, 934
- dbm_close — close ndbm routine, 934
- dbm_delete — remove data from ndbm database, 934
- dbm_err — ndbm database routine, 934
- dbm_error — return ndbm database error condition, 934
- dbm_fetch — fetch ndbm database data, 934
- dbm_firstkey — access ndbm database, 934
- dbm_nextkey — access ndbm database, 934
- dbm_open — open ndbm database, 934
- dbm_store — add data to ndbm database, 934
- dbm_init — open database, 830
- dbx — source debugger, 126
- dbxtool — debugger, 135
- dc — desk calculator, 137
- dcheck — directory consistency check, 1608
- dd — convert and copy, 139
- DEAD mail variable, 301
- debug mail variable, 301
- debug network — ping, 1709
- debug tools
 - adb — debugger, 13
 - adbgen — generate adb script, 1575
 - ctrace — display program trace, 117
 - dbx — source debugger, 126
 - dbxtool — debugger, 135
 - kadb — kernel debugger, 1653
- debugging memory management
 - malloc_debug — set debug level, 926
 - malloc_verify — verify heap, 926
- debugging support — assert, 813
- debugging support, System V — assert, 1131
- decimal dump file — od, 348
- decimal record from double-precision floating —
 - double_to_decimal, 850
- decimal record from single-precision floating —
 - single_to_decimal, 850
- decimal record to double-precision floating —
 - decimal_to_double, 832
- decimal record to extended-precision floating —
 - decimal_to_extended, 832
- decimal record to extended-precision floating —
 - extended_to_decimal, 850
- decimal record to single-precision floating —
 - decimal_to_single, 832
- decimal_to_double — decimal record to double-precision floating, 832
- decimal_to_extended — decimal record to extended-precision floating, 832
- decimal_to_single — decimal record to single-precision floating, 832
- decode binary file — uudecode, 570
- decode files
 - crypt, 95
 - des — Data Encryption Standard, 147
- crypt — decrypt, 95
- default command, 104
- defaults, update kernel from — input_from_defaults, 239
- defaults_from_input — update defaults from kernel, 239
- defaultsed — changing SunView default settings, 141
- delayed execution
 - add job to queue — at, 29
 - display queue — atq, 31
 - remove jobs from queue — atrm, 32
- delete
 - columns from file, 80, 121
 - directory — rmdir, 420, 743
 - directory entry — unlink, 785
 - file — rm, 420
 - filename affixes — basename, 42
 - m/c address ioctl — SIOCDELMULTI, 1227
 - nroff, troff, tbl and eqn constructs — deroff, 146
 - print jobs — lprm, 283
 - repeated lines — uniq, 561
- delete arp entry ioctl — SIOCDELRARP, 1194
- delete datum and key — delete, 830
- delete delayed execution jobs — atrm, 32
- delete descriptor, 646
- delete — delete datum and key, 830
- delete mail command, 297
- delete route ioctl — SIOCDELRT, 1288
- delta
 - change commentary, 58
 - combine, 81
 - make SCCS delta — delta, 144
 - remove — rmdel, 421
- demos
 - bouncedemo — bouncing square graphics demo, 1521
 - canvas_demo — canvas subwindow demo, 1542
 - cursor_demo — cursor attributes demo, 1542
 - framedemo — graphics demo, 1521
 - graphics_demos, 1521
 - introduction, 1493
 - jumpdemo — graphics demo, 1521
 - spheresdemo — graphics demo, 1521
 - SunView demos, 1542
- demount file system — umount, 1687
- demount file system — unmount, 786
- deny messages — msg, 327
- deroff — remove troff constructs, 146
- des — data encryption, 147
- des — DES encryption chip interface, 1210
- DES encryption

- DES encryption, *continued*
 - cbc_crypt, 833
 - des_setparity, 833
- describe command — *what is*, 593
- descriptors
 - close, 646
 - delete, 646
 - dup, 651
 - dup2, 651
 - fcntl, 656
 - flock, 660
 - getdtablesize, 669
 - lockf, 921
 - select, 744
- DESIOCBLOCK — process block, 1210
- DESIOCQUICK — process quickly, 1210
- desk calculator, 137
- destroy hash table — *hdestroy*, 891
- device controls — *ioctl*, 692
- devices
 - paging, specify — *swapon*, 1771
 - swapping, specify — *swapon*, 1771
- devices, introduction to, 1189
- devnm command, 1609
- df — display free space, 149
- diagnostics
 - I/O error mapping page, 479
- diagnostics — *sysdiag*, 1772
- diff — differential compare, 150
- diff3 — three-way differential compare, 152
- diffmk — add change marks to documents, 154
- dir — directory format, 1383
- dircmp — compare directories, 155
- directory
 - change current, 639
 - change name of — *mv*, 334
 - change root — *chroot*, 644
 - change working, 57
 - check consistency — *dcheck*, 1608
 - delete — *rmdir*, 420
 - delete — *rmdir*, 743
 - differential compare, 150
 - display name of working — *pwd*, 404
 - erase — *rmdir*, 743
 - get entries, 664, 666
 - list contents of — *ls*, 285
 - make — *mkdir*, 328, 329, 700
 - make link to — *ln*, 265
 - move — *mv*, 334
 - remove — *rmdir*, 420
 - remove — *rmdir*, 743
 - rename — *mv*, 334
 - scan, 983
- directory operations
 - closedir, 834
 - opendir, 834
 - readdir, 834
 - rewinddir, 834
 - seekdir, 834
 - telldir, 834
- dirs command, 104
- dis command, 156
- disable print queue — *lpc*, 1662
- discard mail command, 297
- disk
 - control operations — *dkio*, 1211
 - diagnostics — *sysdiag*, 1772
 - dkinfo — geometry information, 1610
- disk driver
 - sd — Adaptec ST-506, 1289 *thru* 1290
 - fd — Sun floppy, 1218
 - si — Sun SCSI, 1289 *thru* 1290
 - xd — Xylogics, 1329 *thru* 1330, 1332 *thru* 1333
- disk quotas
 - edquota — edit user quotas, 1616
 - quotacheck — check quota consistency, 1721
 - quotaoff — turn file system quotas off, 1722
 - quotaon — turn file system quotas on, 1722
 - repquota — summarize quotas, 1729
 - rquotad — remote quota server, 1747
- disk quotas — *quotactl*, 732
- display
 - architecture of current Sun host, 26
 - call-graph profile data — *gprof*, 214
 - current domain name — *domainname*, 157
 - current host identifier, 226
 - current host name, 227
 - date, 124
 - date and time, 124
 - delayed execution queue — *atq*, 31
 - disk usage, 162
 - disk usage and limits — *quota*, 405
 - dynamic dependencies — *ldd*, 256
 - effective user name — *whoami*, 598
 - file by screenfuls — *more*, 330
 - file names — *ls*, 285
 - file system quotas — *repquota*, 1729
 - first lines of file, 223
 - free space in file system, 149
 - group membership, 222
 - identifier of current host, 226
 - last commands — *lastcomm*, 249
 - last part of file — *tail*, 507
 - login name — *logname*, 274
 - name list of object file or library — *nm*, 339
 - name of current host, 227
 - page size — *pagesize*, 378
 - printer queue — *lpq*, 278
 - process status — *ps*, 399
 - processor of current Sun host, 291
 - program profile — *prof*, 392
 - program trace — *ctrace*, 117
 - SCCS file editing status — *sact*, 433
 - selected lines from file — *sed*, 446
 - status of network hosts — *rup*, 428
 - system up time — *uptime*, 566
 - time and date, 124
 - time in window, 71
 - user and group IDs — *id*, 230
 - users on system — *users*, 567
 - waiting mail — *prmail*, 362
 - working directory name — *pwd*, 404
- display editor — *vi*, 582
- display status of local hosts — *ruptime*, 429
- dkinfo — disk geometry information, 1610

- dkio — disk control operations, 1211
 - DKIOCGGEOM — get disk geometry, 1211
 - DKIOCGPART — get disk partition info, 1211
 - DKIOCINFO — get disk info, 1211
 - DKIOCSGEOM — set disk geometry, 1211
 - DKIOCSPART — set disk partition info, 1211
 - dmesg — create error log, 1611
 - dn_comp — Internet name server routines, 965
 - dn_expand — Internet name server routines, 965
 - do command, 453
 - document production
 - addbib — create bibliography, 18
 - checknr — check nroff/troff files, 62
 - col — filter reverse paper motions, 77
 - colcrt command, 78
 - deroff — delete troff, tbl and eqn constructs, 146
 - diffmk — add change marks, 154
 - eqn — set mathematical equations, 173
 - eqnchar — special characters for equations, 1550
 - fmt — simple formatter, 188
 - indxbib — make inverted index, 235
 - lookbib — find bibliographic references, 276
 - man — macros to format manual pages, 1560
 - me — macro package, 1563
 - ms — macro package, 1565
 - nroff — document formatter, 344
 - pti — (old) troff interpreter, 363
 - ptx — generate permuted index, 403
 - refer — insert literature references, 415
 - roffbib — print bibliographic database, 422
 - soelim — eliminate .so's from nroff input, 469
 - sortbib — sort bibliographic database, 473
 - spell — check spelling, 474
 - tbl — table formatter, 513
 - troff — typeset documents, 548
 - vfontinfo — examine font files, 579
 - vtroff — format document for raster printer, 586
 - width — make font width table, 587
 - DoD Internet host table, get from host — gettable, 1635
 - domain
 - get name of current — getdomainname, 668
 - set name of current — setdomainname, 668
 - domainname — set/display domain name, 157
 - done command, 453
 - dos — window for IBM PC/AT applications, 158
 - dos2unix — convert text file from DOS format to SunOS format, 161
 - dot mail variable, 301
 - double_to_decimal — decimal record from double-precision floating, 850
 - down, take printer — lpc, 1662
 - dp mail command, 297
 - drand48 — generate uniformly distributed random numbers, 836
 - draw graph, 216
 - drum — paging device, 1213
 - dt mail command, 297
 - du — display disk usage, 162
 - dump — dump file system, 1612
 - dump — incremental dump format, 1385
 - dump frame buffer image — screendump, 440
 - dumpfs — dump file system information, 1615
 - dup, 651
 - dup2, 651
 - duplicate descriptor, 651
 - dysize — date and time conversion, 824
- ## E
- E2BIG error number, 613
 - EACCES error number, 614
 - EADDRINUSE error number, 616
 - EADDRNOTAVAIL error number, 616
 - EAFNOSUPPORT error number, 616
 - EAGAIN error number, 614
 - EALREADY error number, 615
 - EBADF error number, 613
 - EBADMSG error number, 617
 - EBUSY error number, 614
 - ec — 3Com 10 Mb/s Ethernet interface, 1214
 - ECHILD error number, 614
 - echo command, 104, 163, 458
 - echo mail command, 297
 - echo variable, 109
 - ECONNABORTED error number, 616
 - ECONNREFUSED error number, 616
 - ECONNRESET error number, 616
 - econvert — convert number to ASCII, 838
 - ed — line editor, 164
 - edata — end of program data, 840
 - EDESTADDRREQ error number, 615
 - edit
 - fonts — fontedit, 190
 - icons — coloredit, 79
 - icons — iconedit, 228
 - password file — vipw, 1790
 - SunView defaults — defaultsed, 141
 - user quotas — edquota, 1616
 - edit — line editor, 177
 - edit mail command, 297
 - editing text
 - ed — line editor, 164
 - edit — line editor, 177
 - ex — line editor, 177
 - sed — stream editor, 446
 - EDITOR mail variable, 302
 - EDOM error number, 615
 - EDQUOT error number, 617
 - edquota — edit user quotas, 1616
 - EEPROM display and load program — eeprom, 1617
 - EEXIST error number, 614
 - EFAULT error number, 614
 - EFBIG error number, 615
 - effective group ID
 - get, 670
 - set, 754
 - effective group ID, set — setegid, 991
 - effective user ID
 - get, 691
 - set — setreuid, 755
 - effective user ID, set — seteuid, 991
 - egrep — pattern scanner, 218

- EHOSTDOWN error number, 617
- EHOSTUNREACH error number, 617
- EIDRM error number, 617
- EINPROGRESS error number, 615
- EINTR error number, 613
- EINVAL error number, 614
- EIO error number, 613
- EISCONN error number, 616
- EISDIR error number, 614
- elif command, 453
- eliminate `#ifdef`'s from C input — `unifdef`, 560
- eliminate `.so`'s from `nrff` input — `soelim`, 469
- ELOOP error number, 617
- else command, 105, 453
- else mail command, 298
- EMFILE error number, 614
- EMLINK error number, 615
- EMSGSIZE error number, 615
- emulate Tektronix 4014 — `tektool`, 369
- enable print queue — `lpc`, 1662
- ENAMETOOLONG error number, 617
- encode binary file — `uencode`, 570
- encode files
 - crypt, 95
 - des — Data Encryption Standard, 147
- encrypt — encryption, 822
- encrypted mail
 - enroll for — `enroll`, 604
 - receive — `enroll`, 604
 - send — `xsend`, 604
- encryption
 - cbc_crypt, 833
 - crypt, 822
 - des_setparity, 833
 - encrypt, 822
 - setkey, 822
- encryption chip — `des`, 1210
- encryption key, change, `chkey` command, 64
- encryption key, generate — `makekey`, 1669
- end command, 105
- end — end of program, 840
- end locations in program, 840
- endac () function, 859
- endxportent () function, 845
- endfsent — get file system descriptor file entry, 865
- endgraent () function, 866
- endgrent — get group file entry, 867
- endhostent — get network host entry, 869
- endif command, 105
- endif mail command, 298
- endmntent — get filesystem descriptor file entry, 872
- endnetent — get network entry, 874
- endnetgrent — get network group entry, 875
- endprotoent — get protocol entry, 879
- endpwaent () function, 881
- endpwent — get password file entry, 882
- endpwent — get password file entry, System V, 1164
- endrpcent — get RPC entry, 884
- endservent — get service entry, 886
- endsw command, 108
- endttyent () function, 887
- endusershell () function, 889
- ENETDOWN error number, 616
- ENETRESET error number, 616
- ENETUNREACH error number, 616
- ENFILE error number, 614
- ENOBUFS error number, 616
- ENODEV error number, 614
- ENOENT error number, 613
- ENOEXEC error number, 613
- ENOMEM error number, 614
- ENOMSG error number, 617
- ENOPROTOPT error number, 615
- ENOSPC error number, 615
- ENOSR error number, 617
- ENOSTR error number, 617
- ENOTBLK error number, 614
- ENOTCONN error number, 616
- ENOTDIR error number, 614
- ENOTEMPTY error number, 617
- ENOTSOCK error number, 615
- ENOTTY error number, 614
- enquire stream status
 - clearerr — clear error on stream, 849
 - clearerr — clear error on stream, System V, 1159
 - feof — enquire EOF on stream, 849
 - feof — enquire EOF on stream, System V, 1159
 - ferror — inquire error on stream, 849
 - ferror — inquire error on stream, System V, 1159
 - fileno — get stream descriptor number, 849
 - fileno — get stream descriptor number, System V, 1159
- enroll — enroll for secret mail, 604
- env — obtain or alter environment variables, 172
- environ — user environment, 1387
- environ — execute file, 842
- environment
 - display variables — `printenv`, 391
 - get value — `getenv`, 863
 - set value — `putenv`, 950
 - tset — set terminal characteristics for, 551
- environment variables — in C shell, 109
- environment variables in mail, see also *mail environment variables*
- ENXIO error number, 613
- EOPNOTSUPP error number, 616
- EPERM error number, 613
- EPFNOSUPPORT error number, 616
- EPIPE error number, 615
- EPROTONOSUPPORT error number, 615
- EPROTOTYPE error number, 615
- eqn — remove constructs — `deroff`, 146
- eqn — mathematical typesetting, 173
- eqnchar — special characters for equations, 1550
- erand48 — generate uniformly distributed random numbers, 836
- ERANGE error number, 615
- erase
 - directory — `rmdir`, 420, 743
 - directory entry — `unlink`, 785

- erase, *continued*
 file — rm, 420
 erase — start new plot frame, 941
 erase magnetic tape — mt, 333
 EREMOTE error number, 617
 erf — error functions, 1085
 erfc — error functions, 1085
 EROFS error number, 615
 errno — system error messages, 940
 error — analyze error messages, 175
 error messages, 940
 esac command, 453
 escape character, quotes and comments, C shell, 97
 escape mail variable, 302
 ESHUTDOWN error number, 616
 ESOCKETNOSUPPORT error number, 616
 ESPIPE error number, 615
 ESRCH error number, 613
 ESTALE error number, 617
 etext — end of program text, 840
 etherd — Ethernet statistics server daemon, 1618
 etherfind — find packets on the Ethernet, 1619
 Ethernet
 find packets — etherfind, 1619
 statistics server daemon — etherd, 1618
 Ethernet address mapping, 841
 Ethernet address to ASCII — ether_ntoa, 841
 Ethernet address to hostname — ether_ntohost, 841
 Ethernet controller
 ec — 10 Mb/s 3Com Ethernet interface, 1214
 ie — Sun Ethernet interface, 1224 *thru* 1225
 le — 10 Mb/s LANCE Ethernet interface, 1254 *thru* 1255
 ethers file — Ethernet addresses, 1388
 ETIME error number, 617
 ETIMEDOUT error number, 616
 Euclidean distance functions
 hypot, 1089
 eval command, 104, 458
 evaluate expressions, 180
 EWOULDBLOCK error number, 615
 ex — line editor, 177
 exc_bound() function, 1061
 exc_handle() function, 1061
 exc_notify() function, 1061
 exc_on_exit() function, 1061
 exc_raise() function, 1061
 exc_unhandle() function, 1061
 EXDEV error number, 614
 exec command, 104, 458
 execl — execute file, 842
 execlx — execute file, 842
 execlp — execute file, 842
 execute commands at specified times — cron, 1606
 execute file, 652, 842
 environ, 842
 execl, 842
 execlx, 842
 execlp, 842
 execv, 842
 execve, 652
 execvp — execute file, 842
 exit, 655
 exit command, 105, 458
 exit — terminate process, 844
 exit mail command, 297
 exp — exponential function, 1086
 exp10 — exponential function, 1086
 exp2 — exponential function, 1086
 expand assembly-language calls in-line, inline, 236
 expand — expand tabs, 179
 expm1 — exponential function, 1086
 exponent and significand, split into — frexp, 1087
 exponential function — exp, 1086
 export command, 458
 exportable file system table — exports, 1389
 exported file system table — xtab, 1389
 exportent() function, 845
 exportfs command, 1621
 exports — exported file system table, 1389
 expr — evaluate expressions, 180
 expression evaluation, 180
 expressions — in C shell, 101
 extend bibliography — addbib, 18
 extended_to_decimal — decimal record from extended-precision floating, 850
 extract strings from C code — xstr, 605
 extract_unbundled command, 1623
 eyacc — compiler generator, 351
- ## F
- fabs() function, 1093
 factor game, 1516
 fastboot — reboot system, 1624
 fasthalt — halt system, 1624
 fb — Sun console frame buffer driver, 1215
 fbio — frame buffers general properties, 1216 *thru* 1217
 fchmod, 640
 fchown, 642
 fclose — close stream, 847
 fcntl — file control system call, 656
 fcntl — file control options, 1391
 fconvert — convert number to ASCII, 838
 fdate — return date and time in ASCII format, 848
 fdopen — associate descriptor, 854
 fdopen — associate descriptor, System V, 1160
 feof — enquire EOF on stream, 849
 feof — enquire EOF on stream, System V, 1159

- ferret** — inquire error on stream, 849
ferret — inquire error on stream, System V, 1159
fetch — retrieve datum under key, 830
fflush — flush stream, 847
ffs — find first one bit, 819
fg command, 105
fgetc — get character from stream, 861
fgetc — get character from stream, System V, 1162
fgetgraent () function, 866
fgetgrent — get group file entry, 867
fgetpwaent () function, 881
fgetpwent — get password file entry, 882
setpwfile — get password file entry, System V, 1164
fgets — get string from stream, 885
fgrep — pattern scanner, 218
fi command, 453
fifo, make — **mknod**, 1674
ignore variable, 109
file
 ftw — traverse file tree, 858
 browse through text— **more**, 330
 browse through text— **page**, 330
 browse through text— **pg**, 383
 change name of — **mv**, 334
 change ownership — **chown**, 1595
 copy from remote machine — **rcp**, 409
 count lines, words, characters in — **wc**, 591
 create new, 649
 create temporary name — **tmpnam**, 1020
 delete — **rm**, 420
 determine accessibility of, 629
 display last part of — **tail**, 507
 dump — **od**, 348
 execute, 652
 find lines in sorted — **look**, 275
 identify version — **what**, 592
 make hard link to, 696
 make link to — **ln**, 265
 move — **mv**, 334
 print — **lpr**, 280
 remove — **rm**, 420
 rename — **mv**, 334
 reverse lines in — **rev**, 417
 send to remote host — **uucp**, 571
 split into pieces — **split**, 477
 sum — sum and count blocks in file, 486
 synchronize state — **fsync**, 662
 update last modified date of — **touch**, 541
file attributes
 fstat, 774
 lstat, 774
 stat, 774
file — get file type, 182
file control
 options header file — **fcntl**, 1391
 system call — **fcntl**, 656
file formats, 1337
file inquiries — in C shell, 102
file mail command, 297
file position, move — **lseek**, 698
file system
 file system, *continued*
 4.2 format — **fs**, 1392
 access, 629
 cd — change directory, 57
 chdir, 639
 check and repair — **fsck**, 1629
 check consistency — **icheck**, 1641
 check directory — **dcheck**, 1608
 chmod, 640
 chown, 642
 create file — **open**, 719
 create new — **newfs**, 1700
 delete directory entry — **unlink**, 785
 delete directory — **rmdir**, 743
 unmount — demount file system, 786, 1687
 display disk usage and limits — **quota**, 405
 display free space, 149
 dump information — **dumpfs**, 1615
 edquota — edit user quotas, 1616
 erase directory entry — **unlink**, 785
 erase directory — **rmdir**, 743
 exported table — **xtab**, 1389
 exports table — **exports**, 1389
 fchmod, 640
 fchown, 642
 free space display, 149
 fstab — static information, 1396
 ftruncate, 782
 get file descriptor entry, 865
 getdents, 664
 getdirentries, 666
 link, 696
 lseek, 698
 make — **mkfs**, 1673
 make prototype— **mkproto**, 1675
 mkdir, 700
 mknod, 702
 mount, 706, 1687
 mounted table — **mtab**, 1396, 1412
 open, 719
 quotacheck — check quota consistency, 1721
 quotactl — disk quotas, 732
 quotaoff — turn file system quotas off, 1722
 quotaon — turn file system quotas on, 1722
 readlink, 737
 remove directory entry — **unlink**, 785
 remove directory — **rmdir**, 743
 rename file — **rename**, 741
 repquota — summarize quotas, 1729
 rquotad — remote quota server, 1747
 statistics — **fstatfs**, 776
 statistics — **statfs**, 776
 summarize ownership — **quot**, 1720
 symlink, 779
 tell, 698
 truncate, 782
 tune — **tunefs**, 1784
 umask, 783
 unmount — demount file system, 786, 1687
 utime — set file times, 1030
 utimes — set file times, 787
 where am I — **pwd**, 404
 file system dump — **dump**, 1612
 file system restore — **restore**, 1730

- file transfer protocol
 - ftp command, 195
 - server — ftpd, 1632
 - trivial, tftp command, 530
- file_to_decimal — decimal record from character stream, 1004
- filec variable, 109
- filemerge command, 352
- filename completion, C shell, 97
- filename substitution, 101
- filename, change — rename, 741
- fileno — get stream descriptor number, 849
- fileno — get stream descriptor number, System V, 1159
- files
 - csplit — split file into sections, 113
 - basename — strip affixes, 42
 - cat — concatenate, 50
 - ccat — extract files compressed with compact, 350
 - chmod — change mode, 65
 - cmp — compare files, 76
 - colrm — remove columns from, 80
 - compact — compress files, 350
 - compare, 76
 - compare, three-way differential — diff3, 152
 - compress — compress files, 83
 - convert and copy, 139
 - copy, 85
 - copy standard output to many — tee, 517
 - cp — copy files, 85
 - cpio — copy archives, 87
 - crypt — encrypt/decrypt, 95
 - .cshrc and the C shell, 96
 - cut — remove columns from, 121
 - des — encrypt/decrypt, data encryption standard, 147
 - determine type of, 182
 - differential compare, 150
 - display first lines of, 223
 - display names — ls, 285
 - find, 183
 - find differences, 150
 - .login and the C shell, 96
 - .logout and the C shell, 96
 - pack — pack files, 376
 - paste — horizontal merge, 380
 - pcat — pack files, 376
 - prepare files for printing — pr, 388
 - search for patterns in — grep, 218
 - side-by-side compare, 445
 - sort — sort and collate lines, 470
 - transfer, 195, 530
 - uncompact — uncompress files, 350
 - uncompress — uncompress files, 83
 - unpack — unpack files, 376
 - zcat — extract compress files, 83
- file system organization, 1551
- filesystem descriptor, get file entry, 872
- filter reverse paper motions — col, 77
- find
 - first key in dbm database — firstkey, 830
 - first one bit — ffs, 819
 - find lines in sorted file — look, 275
 - find literature references — refer, 415
 - find, *continued*
 - name of terminal — ttyname, 1024
 - next key in dbm database — nextkey, 830
 - find object file size — size, 465
 - find ordering for object library — lorder, 277
 - patterns in file — grep, 218
 - find printable strings in binary file — strings, 478
 - program — whereis, 594
 - find — find files, 183
 - finger — info on users, 186
 - fingerd daemon, 1625
 - finite() function, 1093
 - FIOASYNC — set/clear async I/O, 1219
 - FIOCLEX — set close-on-exec flag for fd, 1219
 - FIOGETOWN — get file owner, 1219
 - FIONBIO — set/clear non-blocking I/O, 1219
 - FIONCLEX — remove close-on-exec flag, 1219
 - FIONREAD — get # bytes to read, 1219
 - FIOSETOWN — set file owner, 1219
 - firstkey — find first key, 830
 - fish — Go Fish game, 1517
 - floating-point, 193
 - reliability tests — fparel, 1627
 - version and tests — fpaversion, 1628
 - floating-point accellerator, fpa, 1220
 - floatingpoint
 - IEEE floating point definitions, 852
 - flock, 660
 - floor — floor of, 1102
 - flush disk activity — sync, 502
 - flush stream — fflush, 847
 - fmod() function, 1093
 - fmt — simple formatter, 188
 - fold — fold long lines, 189
 - folder mail command, 297
 - folder mail variable, 302
 - folders mail command, 297
 - followup mail command, 297
 - Followup mail command, 297
 - font
 - files, convert foreign — vswap, 585
 - vwidth — make font width table, 587
 - fontedit — font editor, 190
 - fopen — open stream, 854
 - fopen — open stream, System V, 1160
 - foption — determine available floating-point code generation options, 193
 - for command, 453
 - foreach command, 105
 - fork a new process — fork, 661
 - format C programs — indent, 231
 - format command, 1626
 - format document for raster printer — vtroff, 586
 - format of memory image file — core, 1378
 - format tables — tbl, 513
 - formatted input conversion
 - fscanf — convert from stream, 984
 - fscanf — convert from stream, System V, 1172
 - scanf — convert from stdin, 984

- formatted input conversion, *continued*
 scanf — convert from stdin, System V, 1172
 sscanf — convert from string, 984
 sscanf — convert from string, System V, 1172
- FORTRAN**
 tags file — ctags, 115
- fortune — get fortune, 1518
 .forward file, 300
 .forward — mail forwarding file, 1367
- forwarding mail, 300
- fp_class () function, 1093
- fpa, floating-point accelerator, 1220
- fparel — floating-point reliability tests, 1627
- fpaversion — floating-point version and tests, 1628
- fprintf — formatted output conversion, 944
- fprintf — format to stream, System V, 1168
- fputc — put character on stream, 949
- fputs — put string to stream, 952
- frame buffer
 bwone — Sun-1 black and white frame buffer, 1197
 bwtwo — Sun-3/Sun-2 black and white frame buffer, 1198
- framedemo — graphics demo, 1521
- fread — read from stream, 855
- free — free memory, 925
- free memory — cfree, 925
- free memory — free, 925
- free static block ioctl — GP1IO_FREE_STATIC_BLOCK, 1221
- freopen — reopen stream, 854
- freopen — reopen stream, System V, 1160
- frexp — split into significand and exponent, 1087
- from — who is mail from, 194
- from mail command, 297
- fs — 4.2 file system format, 1392
- fscanf — convert from stream, 984
- fscanf — convert from stream, System V, 1172
- fsck — check and repair file system, 1629
- fseek — seek on stream, 856
- fsirand — install random inode generation numbers, 1631
- fspec text file tabstop specifications, 1394
- fstab — file mountable information, 1396
- fstat — obtain file attributes, 774
- fstatfs — obtain file system statistics, 776
- fsync — synchronize disk file with core image, 662
- ftell — get stream position, 856
- ftime — get date and time, 1016
- ftok — interprocess communication routine, 857
- ftp — file transfer, 195
- ftp — remote login data — .netrc file, 1415
 .netrc — ftp remote login data file, 1415
- ftpd — file transfer protocol server, 1632
- ftpusers — ftp prohibited users list, 1398
- fttruncate, 782
- ftw — traverse file tree, 858
- full-duplex connection, shut down — shutdown, 762
- file_to_decimal — decimal record from character function, 1004
- kvm_open () function, 901, 910, 1070
- functions, Bourne shell, 453
- fwrite — write to stream, 855
- G**
- games
 boggletool — SunView game of boggle, 1505
 canfield — solitaire card game, 1507
 chess — chess game, 1509
 chesstool — SunView chess game, 1510
 gammontool — SunView backgammon game, 1519
 introduction, 1493
 life — SunView game of life, 1527
- gamma — log gamma, 1098
- gammontool — SunView backgammon game, 1519
- getaid () function, 663
- gather write — writev, 794
- gcd — multiple precision GCD, 932
- gconvert — convert number to ASCII, 838
- gcore — core image of process, 202
- general magnetic tape interface — mtio, 1268
- generate
 adb script — adbgen, 1575
 encryption key — makekey, 1669
 fault — abort, 810
 lexical analyzer — lex, 258
 permuted index — ptx, 403
- generate random numbers
 initstate, 956
 rand, 955
 random, 956
 setstate, 956
 srand, 955
 srandom, 956
 drand48, 836
 erand48, 836
 jrand48, 836
 lcong48, 836
 lrand48, 836
 mrand48, 836
 nrand48, 836
 seed48, 836
 srand48, 836
- generate random numbers, System V
 rand, 1171
 srand, 1171
- generic disk control operations — dkio, 1211
- generic operations
 gather write — writev, 794
 ioctl, 692
 read, 734
 scatter read — readv, 734
 write, 794
- get
 arp entry ioctl — SIOCGARP, 1194
 character from stream — fgetc, 861
 character from stream —getc, 861
 console I/O ioctl — TIOCCONS, 1204
 count of bytes to read ioctl — FIONREAD, 1219
 current working directory pathname — getwd, 890
 date and time — ftime, 1016
 date and time — time, 1016
 disk geometry ioctl — DKIOCGGEOM, 1211
 disk info ioctl — DKIOCINFO, 1211

get, continued

- disk partition info `ioctl` — `DKIOCGPART`, 1211
- entries from kernel symbol table — `kvm_nlist`, 900
- entries from symbol table — `nlist`, 937, 1167
- environment value — `getenv`, 863
- file owner `ioctl` — `FIOGETOWN`, 1219
- file system descriptor file entry, 865
- high water mark `ioctl` — `SIOCGLHIWAT`, 1304
- ifnet address `ioctl` — `SIOCGIFADDR`, 1226
- ifnet flags `ioctl` — `SIOCGIFFLAGS`, 1226
- ifnet list `ioctl` — `SIOCGIFCONF`, 1226
- info on resource usage — `vtimes`, 1037
- line discipline `ioctl` — `TIOCGETD`, 1196
- login name — `getlogin`, 871
- low water mark `ioctl` — `SIOCGLOWAT`, 1304
- magnetic tape unit status — `mt`, 333
- network entry — `getnetent`, 874
- network group entry — `getnetgrent`, 875
- network host entry — `gethostent`, 869
- network service entry — `getservent`, 886
- options on sockets — `getsockopt`, 687
- p-p address `ioctl` — `SIOCGIFDSTADDR`, 1226
- parent process identification — `getppid`, 680
- pathname of current working directory — `getcwd`, 862
- position of stream — `ftell`, 856
- process domain name — `getdomainname`, 668
- process identification — `getpid`, 680
- process times — `times`, 1017
- protocol entry — `getprotoent`, 879
- requested minor device `ioctl` — `GP1IO_GET_REQDEV`, 1221
- restart count `ioctl` — `GP1IO_GET_RESTART_COUNT`, 1221
- RPC program entry — `getrpcnt`, 884
- scheduling priority — `getpriority`, 681
- signal stack context — `sigstack`, 766
- static block `ioctl` — `GP1IO_GET_STATIC_BLOCK`, 1221
- string from `stdin` — `gets`, 885
- string from stream — `fgets`, 885
- tape status `ioctl` — `MTIOCGET`, 1269
- terminal name — `tty`, 556
- terminal state — `gtty`, 1009
- true minor device `ioctl` — `GP1IO_GET_TRUMINORDEV`, 1222
- user limits — `ulimit`, 1027
- word from stream — `getw`, 861
- get character from stream, System V — `fgetc`, 1162
- get character from stream, System V — `getc`, 1162
- get — get SCCS file, 203
- get date and time, 689
- get filesystem descriptor file entry
 - `addmntent`, 872
 - `endmntent`, 872
 - `getmntent`, 872
 - `hasmntopt`, 872
 - `setmntent`, 872
- get group file entry
 - `endgrent`, 867
 - `fgetgrent`, 867
 - `getgrent`, 867
 - `getgrgid`, 867
 - `getgrnam`, 867
- get group file entry, *continued*
 - `setgrent`, 867
- get high water mark `ioctl` — `SIOCGLHIWAT`, 1325
- get keyboard “direct input” state `ioctl` — `KIOCGDIRECT`, 1236
- get keyboard translation `ioctl` — `KIOCGTRANS`, 1235
- get keyboard type `ioctl` — `KIOCTYPE`, 1236
- get low water mark `ioctl` — `SIOCGLOWAT`, 1325
- get password file entry
 - `endpwent`, 882
 - `fgetpwent`, 882
 - `getpwent`, 882
 - `getpwnam`, 882
 - `getpwuid`, 882
 - `setpwent`, 882
 - `fgetpwent`, 882
- get password file entry, System V
 - `endpwent`, 1164
 - `fgetpwent`, 1164
 - `getpwent`, 1164
 - `getpwnam`, 1164
 - `getpwuid`, 1164
 - `setpwent`, 1164
 - `fgetpwent`, 1164
- get process times, System V — `times`, 1185
- get time zone name
 - `timezone`, 1018
- get translation table entry `ioctl` — `KIOCGGETKEY`, 1236
- get user name — `cuserid`, 829
- get word from stream, System V — `getw`, 1162
- `get_selection` — copy a SunView selection to standard output, 208
- `getacdir()` function, 859
- `getacflg()` function, 859
- `getacinfo()` function, 859
- `getacmin()` function, 859
- `getauditflags()` function, 864
- `getauditflagsbin()` function, 860
- `getauditflagschar()` function, 860
- `getc` — get character from stream, 861
- `getc` — get character from stream, System V, 1162
- `getchar` — get character from `stdin`, 861
- `getchar` — get character from `stdin`, System V, 1162
- `getcwd` — get pathname of current directory, 862
- `getdents`, 664
- `getdirent`, 666
- `getdomainname` — get process domain, 668
- `getdtablesize`, 669
- `getegid` — get effective group ID, 670
- `getenv` — get value from environment, 863
- `geteuid` — get effective user ID, 691
- `getexportent()` function, 845
- `getexportopt()` function, 845
- `getfsent` — get file system descriptor file entry, 865
- `getfsfile` — get file system descriptor file entry, 865
- `getfsspec` — get file system descriptor file entry, 865
- `getfstype` — get file system descriptor file entry, 865
- `getgid` — get group ID, 670
- `getgraent()` function, 866
- `getgranam()` function, 866

- getgrent — get group file entry, 867
- getgrgid — get group file entry, 867
- getgrnam — get group file entry, 867
- getgroups, 671
- gethostbyaddr — get network host entry, 869
- gethostbyname — get network host entry, 869
- gethostent — get network host entry, 869
- gethostid, 672
- gethostname, 673
- getitimer, 674
- getlogin — get login name, 871
- getmntent — get filesystem descriptor file entry, 872
- getmsg () function, 676
- getnetbyaddr — get network entry, 874
- getnetbyname — get network entry, 874
- getnetent — get network entry, 874
- getnetgrent — get network group entry, 875
- getopt — process options in scripts, 209
- getopt () function, 876
- getopts command, 211
- getpagesize, 678
- getpass — read password, 878
- getpass — read password, System V, 1163
- getpeername, 679
- getpgrp, 753
- getpid, 680
- getppid, 680
- getpriority, 681
- getprotobyname — get protocol entry, 879
- getprotoent — get protocol entry, 879
- getpublickey () function, 1116
- getpw — get name from uid, 880
- getpwaent () function, 881
- getpwanam () function, 881
- getpwent — get password file entry, 882
- getpwent — get password file entry, System V, 1164
- getpwnam — get password file entry, 882
- getpwnam — get password file entry, System V, 1164
- getpwuid — get password file entry, 882
- getpwuid — get password file entry, System V, 1164
- getrlimit, 682
- getrpcbyname — get RPC entry, 884
- getrpcbynumber — get RPC entry, 884
- getrpcent — get RPC entry, 884
- getrpcport — get RPC port number, 1110
- getrusage, 684
- gets — get string from stdin, 885
- getsecretkey () function, 1116
- getservbyname — get service entry, 886
- getservbyport — get service entry, 886
- getservent — get service entry, 886
- getsockname, 686
- getsockopt, 687
- gettable — get DoD Internet host table, 1635
- gettimeofday, 689
- getttyent () function, 887
- getttynam () function, 887
- getty — set terminal mode, 1636
- gettytab — terminal configuration data base, 1399
- getuid — get user ID, 691
- getusershell () function, 889
- getw — get word from stream, 861
- getw — get word from stream, System V, 1162
- getwd — get current working directory pathname, 890
- gfxtool — SunWindows graphics tool, 213
- glob command, 105
- gmtime — date and time conversion, 824
- gmtime — date and time conversion, System V, 1132
- goto command, 105
- GP, initialize graphics processor — gpconfig, 1637
- GP1IO_CHK_GP — restart GP, 1221
- GP1IO_FREE_STATIC_BLOCK — free static block, 1221
- GP1IO_GET_GBUFFER_STATE — check buffer state, 1221
- GP1IO_GET_REQDEV — get requested minor device, 1221
- GP1IO_GET_RESTART_COUNT — get restart count, 1221
- GP1IO_GET_STATIC_BLOCK — get static block, 1221
- GP1IO_GET_TRUMINORDEV — get true minor device, 1222
- GP1IO_PUT_INFO — pass framebuffer info, 1221
- GP1IO_REDIRECT_DEVFB — reconfigure fb, 1221
- gpconfig — bind cgtwo frame buffers to GP, 1637
- gpone — graphics processor interface, 1221 thru 1222
- gprof — call-graph profile, 214
- graph — draw graph, 216
- graphics
 - spline — interpolate smooth curve, 476
 - vplot — plot on Versatec, 584
- graphics filters — plot, 386
- graphics interface
 - arc, 941
 - circle, 941
 - closepl, 941
 - cont, 941
 - erase, 941
 - label, 941
 - line, 941
 - linemod, 941
 - move, 941
 - openpl, 941
 - point, 941
 - space, 941
- graphics interface files — plot, 1421
- graphics processor interface — gpone, 1221 thru 1222
- graphics tool — gfxtool, 213
- grep — pattern scanner, 218
- group access list
 - initialize — initgroups, 895
- group entry, network — getnetgrent, 875
- group — group file format, 1402
- group file entry — getgrent, 867
- group ID
 - chgrp — change group ID of file, 63
 - id — display user and group IDs, 230
 - newgrp — change group ID of user, 335
 - get, 670
 - get effective, 670
 - set real and effective, 754
- group mail command, 296
- group.adjunct — password file, 1404

grouping commands in the C shell, 97
 groups access list, get — `getgroups`, 671
 groups access list, set — `setgroups`, 671
 groups — display group membership, 222
`grpauth()` function, 953
`grpck` — check group database entries, 1638
`gtty` — get terminal state, 1009

H

`hack game`, 1522
`halt` — stop processor, 1639
`halt processor`, 738
`halt system` — `fasthalt`, 1624
`hangman` — hangman game, 1523
`hangup`, control terminal — `vhangup`, 790
 hard link to file — `link`, 696
 hard link, make — `ln`, 265
`hardpaths` variable, 109
 hardware support, introduction to, 1189
`hash command`, 458
 hash table search routine — `hsearch`, 891
`hashstat command`, 105
`hasmntopt` — get filesystem descriptor file entry, 872
`havedisk` — disk inquiry of remote kernel, 1120
`hcreate` — create hash table, 891
`hdestroy` — destroy hash table, 891
`head` — display head of file, 223
`header mail variable`, 302
`headers mail command`, 297
`help` — get SCCS help, 224
`help` — get help, 1353, 1355
`help mail command`, 297
`help_viewer` — get `help_viewer`, 225
`hexadecimal dump file` — `od`, 348
`hier`, file system hierarchy, 1555
`histchars` variable, 109
`history command`, 105
 history substitution — in C shell, 98
 history substitution modifiers, 98
`history variable`, 109
`hold mail command`, 298
`hold mail variable`, 302
`HOME mail environment variable`, 300
`home variable`, 109, 454
`host`
 functions to convert to network byte order, 820
 get identifier of, 672
 get network entry — `gethostent`, 869
 get/set name — `gethostname`, 673
 phone numbers file — `phones`, 1420
`hostid` — display host ID, 226
`hostname` — display host name, 227
`hostname to Ethernet address` — `ether_hostton`, 841
`hosts` — host name data base, 1405
`hosts.equiv` — trusted hosts list, 1406
`hsearch` — hash table search routine, 891
`htable` — convert DoD Internet format host table, 1640
`htonl` — convert network to host long, 820

`htons` — convert host to network short, 820
`HUGE()` function, 1097
`HUGE_VAL()` function, 1097
`hunt game`, 1524
 hyperbolic functions
 `cosh`, 1088
 `sinh`, 1088
 `tanh`, 1088
`hypot` — Euclidean distance, 1089

I

I/O
 socket, see `sockio(4)`, 1291
 STREAMS, see `streamio(4)`, 1294
 terminals, see `termio(4)`, 1305
 tty, see `termio(4)`, 1305
 I/O redirection in the C shell, 99
 I/O statistics report — `iostat`, 1651
 I/O, buffered binary
 `fread` — read from stream, 855
 `frwrite` — write to stream, 855
`i386` — machine type indication, 292
`iAPX286` — machine type indication, 292
`icheck` — file system consistency check, 1641
`icmp` — Internet Control Message Protocol, 1223
`iconedit` — edit icons, 228
`id` — display user and group IDs, 230
 identifier of current host, get — `gethostid`, 672
 identify file version — `what`, 592
`ie` — Sun 10 Mb/s Ethernet interface, 1224 thru 1225
`ieee_flags()` function, 1090
`ieee_handler()` function, 1094
`ieeefp`
 IEEE floating point definitions, 852
`if command`, 105, 453
`if` — network interface general properties, 1226 thru 1227
`if mail command`, 298
`ifconfig` — configure network interface parameters, 1642
`IFS variable` — `sh`, 455
`ignore mail command`, 297
`ignore mail variable`, 302
`ignoreeof mail variable`, 302
`ignoreeof variable`, 109
`lkon 10071-5 printer interface` — `vp`, 1326
`ilogb()` function, 1093
`inc mail command`, 298
`incremental dump format` — `dump`, 1385
`incremental file system dump` — `dump`, 1612
`incremental file system restore` — `restore`, 1730
`indent` — format C source, 231, 1407
`indentprefix mail variable`, 302
`index` — find character in string, 1001
`index memory characters` — `memchr`, 928
`index strings` — `index`, 1001
`index strings` — `rindex`, 1001
`indexing`, generate permuted index — `ptx`, 403
`indirect system call`, 781
`indxbib` — make inverted index, 235
`inet` — Internet protocol family, 1228 thru 1229

- inet_addr — Internet address manipulation, 893
- inet_lnaof — Internet address manipulation, 893
- inet_makeaddr — Internet address manipulation, 893
- inet_netof — Internet address manipulation, 893
- inet_network — Internet address manipulation, 893
- inet_ntoa — Internet address manipulation, 893
- inetd — Internet server daemon, 1644
- inetd.conf — Internet server database, 1408, 1436
- infinity () function, 1097
- infocmp command, 1645
- information
 - miscellaneous, 1547
 - non-, miscellaneous, 1547
- inhibit messages — `msg`, 327
- init — process control initialization, 1648
- initgroups — initialize group access list, 895
- initialize group access list — `initgroups`, 895
- initiate
 - connection on socket — `connect`, 647
 - I/O to or from process — `popen`, 943
- initstate — random number routines, 956
- inline command, 236
- innetgr — get network group entry, 875
- inode, clear — `clri`, 1598
- input conversion
 - `fscanf` — convert from stream, 984
 - `fscanf` — convert from stream, System V, 1172
 - `scanf` — convert from stdin, 984
 - `scanf` — convert from stdin, System V, 1172
 - `sscanf` — convert from string, 984
 - `sscanf` — convert from string, System V, 1172
- input stream, push character back to — `ungetc`, 1028
- input_from_defaults — update kernel from defaults database, 239
- inquire stream status
 - `clearerr` — clear error on stream, 849
 - `clearerr` — clear error on stream, System V, 1159
 - `feof` — enquire EOF on stream, 849
 - `feof` — enquire EOF on stream, System V, 1159
 - `ferror` — inquire error on stream, 849
 - `ferror` — inquire error on stream, System V, 1159
 - `fileno` — get stream descriptor number, 849
 - `fileno` — get stream descriptor number, System V, 1159
- insert element in queue — `insque`, 896
- insert literature references — `refer`, 415
- insert_brackets — `textedit` selection filter, 529
- insque — insert element in queue, 896
- install — install files, 240
- install Yellow Pages database — `ypinit`, 1793
- installboot procedures — `boot`, 1650
- integer absolute value — `abs`, 811
- internat — key mapping table for internationalization, 1356
- Internet
 - control message protocol — `icmp`, 1223
 - directory service — `whois`, 599
 - file transfer protocol server — `ftpd`, 1632
 - protocol family — `inet`, 1228 *thru* 1229
 - Protocol — `ip`, 1230 *thru* 1232
 - to Ethernet address resolution — `arp`, 1194 *thru* 1195
 - Transmission Control Protocol — `tcp`, 1303, 1304
- Internet, *continued*
 - User Datagram Protocol — `udp`, 1324, 1325
- Internet address manipulation functions, 893
- Internet name server routines, 965
- Internet servers database — `servers`, 1408, 1436
- interpolate smooth curve — `spline`, 476
- interpret (old) `troff` output — `pti`, 363
- interprocess communication
 - accept connection — `accept`, 628
 - `bind`, 636
 - `connect`, 647
 - `ftok`, 857
 - `getsockname`, 686
 - `getsockopt`, 687
 - `ipcrm`, 241
 - `ipcs`, 242
 - `listen`, 697
 - `pipe`, 722
 - `recv`, 739
 - `recvfrom`, 739
 - `recvmsg`, 739
 - `send`, 751
 - `sendmsg`, 751
 - `sendto`, 751
 - `setsockopt`, 687
 - `shutdown`, 762
 - `socket`, 771
 - `socketpair`, 773
- interrupts, release blocked signals — `sigpause`, 764
- interval timers
 - `clock`, 821
 - `get`, 674
 - `set`, 674
 - `timerclear` — macro, 674
 - `timercmp` — macro, 674
 - `timerisset` — macro, 674
- introduction
 - C library functions, 797
 - commands, 3
 - devices, 1189
 - file formats, 1337
 - games and demos, 1493
 - hardware support, 1189
 - mathematical library functions, 1081
 - miscellaneous information, 1547
 - miscellaneous noninformation, 1547
 - network interface, 1189
 - protocols, 1189
 - RPC library functions, 1107
 - standard I/O library functions, 999, 1183
 - system calls, 613 *thru* 624
 - system error numbers, 613 *thru* 617
 - system maintenance and operation, 1569
 - System V library functions, 1127
- `ioctl`, 692
- `ioctl`'s for des chip
 - `DESIOCBLOCK` — process block, 1210
 - `DESIOCQUICK` — process quickly, 1210
- `ioctl`s for disks
 - `DKIOCGGEO` — get disk geometry, 1211
 - `DKIOCGPART` — get disk partition info, 1211
 - `DKIOCINFO` — get disk info, 1211
 - `DKIOCSGEO` — set disk geometry, 1211

- ioctl's for disks, *continued*
 - DKIOCSPART — set disk partition info, 1211
 - ioctl's for files
 - FIOASYNC — set/clear async I/O, 1219
 - FIOCLEX — set close-on-exec for fd, 1219
 - FIOGETOWN — get owner, 1219
 - FIONBIO — set/clear non-blocking I/O, 1219
 - FIONCLEX — remove close-on-exec flag, 1219
 - FIONREAD — get # bytes to read, 1219
 - FIOSETOWN — set owner, 1219
 - ioctl's for graphics processor
 - GP1IO_CHK_GP — restart GP, 1221
 - GP1IO_FREE_STATIC_BLOCK — free static block, 1221
 - GP1IO_GET_GBUFFER_STATE — check buffer state, 1221
 - GP1IO_GET_REQDEV — get requested minor device, 1221
 - GP1IO_GET_RESTART_COUNT — get restart count, 1221
 - GP1IO_GET_STATIC_BLOCK — get static block, 1221
 - GP1IO_GET_TRUMINORDEV — get true minor device, 1222
 - GP1IO_PUT_INFO — pass framebuffer info, 1221
 - GP1IO_REDIRECT_DEVFB — reconfigure fb, 1221
 - ioctl's for keyboards
 - KIOCCMD — send a keyboard command, 1236
 - KIOCGDIRECT — get keyboard “direct input” state, 1236
 - KIOCGTKEY — get translation table entry, 1236
 - KIOCGTRANS — get keyboard translation, 1235
 - KIOCSDIRECT — set keyboard “direct input” state, 1236
 - KIOCSKEY — change translation table entry, 1235
 - KIOCTRANS — set keyboard translation, 1235
 - KIOCTYPE — get keyboard type, 1236
 - ioctl's for sockets
 - SIOCADMULTI — set m/c address, 1227
 - SIOCADDRT — add route, 1288
 - SIOCDAEP — delete arp entry, 1194
 - SIOCDELMULTI — delete m/c address, 1227
 - SIOCDELRT — delete route, 1288
 - SIOCGARP — get arp entry, 1194
 - SIOCGHIWAT — get high water mark, 1304, 1325
 - SIOCGIFADDR — get ifnet address, 1226
 - SIOCGIFCONF — get ifnet list, 1226
 - SIOCGIFDSTADDR — get p-p address, 1226
 - SIOCGIFFLAGS — get ifnet flags, 1226
 - SIOCGLOWAT — get low water mark, 1304, 1325
 - SIOCSARP — set arp entry, 1194
 - SIOCSHIWAT — set high water mark, 1304, 1325
 - SIOCSIFADDR — set ifnet address, 1226
 - SIOCSIFDSTADDR — set p-p address, 1226
 - SIOCSIFFLAGS — set ifnet flags, 1226
 - SIOCSLOWAT — set low water mark, 1304, 1325
 - SIOCSMISC — toggle promiscuous mode, 1227
 - ioctl's for tapes
 - MTIOCGET — get tape status, 1269
 - MTIOCTOP — tape operation, 1268
 - ioctl's for terminals
 - TIOCCONS — get console I/O, 1204
 - TIOCGTD — get line discipline, 1196
 - TIOCPKT — set/clear packet mode (pty), 1285
 - TIOCREMOTE — remote input editing, 1285
 - TIOCSETD — set line discipline, 1196
 - TIOCPSTART — start output (like control-Q), 1285
 - TIOCPSTOP — stop output (like control-S), 1285
 - iostat — report I/O statistics, 1651
 - IP address allocation, 1113
 - IP address mapping, 1113
 - ip — Internet Protocol, 1230 *thru* 1232
 - ipalloc — IP address mapper, 1113
 - ipalloc.net range file, 1358
 - ipallocald — Ethernet-to-IP address mapper, 1652
 - ipcrm — remove interprocess communication identifiers, 241
 - iipcs — display interprocess communication status, 242
 - irint — irint of, 1102
 - isalnum — is character alphanumeric, 826
 - isalnum — is character alphanumeric, System V, 1134
 - isalpha — is character letter, 826
 - isalpha — is character letter, System V, 1134
 - isascii — is character ASCII, 826
 - isascii — is character ASCII, System V, 1134
 - isatty — test if device is terminal, 1024
 - isctrl — is character control, 826
 - isctrl — is character control, System V, 1134
 - isdigit — is character digit, 826
 - isdigit — is character digit, System V, 1134
 - isgraph — is character graphic, 826
 - isgraph — is character graphic, System V, 1134
 - isinf () function, 1093
 - islower — is character lower-case, 826
 - islower — is character lower-case, System V, 1134
 - isnan () function, 1093
 - isnormal () function, 1093
 - isprint — is character printable, 826
 - isprint — is character printable, System V, 1134
 - ispunct — is character punctuation, 826
 - ispunct — is character punctuation, System V, 1134
 - issecure () function, 897
 - isspace — is character whitespace, 826
 - isspace — is character whitespace, System V, 1134
 - isubnormal () function, 1093
 - issue shell command — system, 1013
 - isupper — is character upper-case, 826
 - isupper — is character upper-case, System V, 1134
 - isxdigit — is character hex digit, 826
 - isxdigit — is character hex digit, System V, 1134
 - iszero () function, 1093
 - itom — integer to multiple precision, 932
- ## J
- j0 — Bessel function, 1084
 - j1 — Bessel function, 1084
 - jn — Bessel function, 1084
 - job control \m csh, 103
 - jobs command, 105
 - join — relational database operator, 245
 - jrand48 — generate uniformly distributed random numbers, 836
 - jumpdemo — graphics demo, 1521
- ## K
- kadb — kernel debugger, 1653
 - kb — Sun keyboard
 - kb — Sun keyboard STREAMS module, 1233
 - keep mail variable, 302
 - keepsave mail variable, 302

- kernel and local lock manager protocol, 1111
 - kernel symbol table, get entries from — `kvm_nlist`, 900
 - keyboard click, control with — `click`, 70
 - kbd — Sun keyboard, 1251
 - keyenvoy command, 1655
 - keyenvoy server, 1656
 - keylogin command, 246
 - kgmon — dump profile buffers, 1657
 - kill — send signal to process, 693
 - kill command, 106, 247
 - killpg — send signal to process group, 695
 - KIOCCMD — send a keyboard command, 1236
 - KIOCGDIRECT — get keyboard “direct input” state, 1236
 - KIOCGTKEY — get translation table entry, 1236
 - KIOCGTRANS — get keyboard translation, 1235
 - KIOCSDIRECT — set keyboard “direct input” state, 1236
 - KIOCSKEY — change translation table entry, 1235
 - KIOCTRANS — set keyboard translation, 1235
 - KIOCTYPE — get keyboard type, 1236
 - kmem — kernel memory space, 1261 *thru* 1262
 - `kvm_close()` function, 901
 - `kvm_getcmd()` function, 898
 - `kvm_getproc()` function, 899
 - `kvm_getu()` function, 898
 - `kvm_nextproc()` function, 899
 - `nlist` — get entries from kernel symbol table, 900
 - `kvm_read()` function, 903
 - `kvm_setproc()` function, 899
 - `kvm_write()` function, 903
- L**
- 164a — convert base-64 ASCII to long integer, 809
 - label — plot label, 941
 - LANCE 10 Mb/s Ethernet interface — `le`, 1254 *thru* 1255
 - languages
 - `cb` — format filter for C sources, 51
 - `cc` — C compiler, 52
 - `tcflow` — code flow graph, 60
 - `cpp` — C preprocessor, 89
 - `cxref` — cross reference C program, 123
 - `indent` — format C source, 231
 - `lex` — generate lexical analyzer, 258
 - `lint` — C program verifier, 261
 - `mkstr` — create C error messages, 329
 - `tcov` — code coverage tool, 516
 - `xstr` — extract strings from C code, 605
 - `last` — list last logins, 248
 - last locations in program, 840
 - `lastcomm` — display last commands, 249
 - `lastlog` — login records, 1485
 - `lcong48` — generate uniformly distributed random numbers, 836
 - `ld` — link editor, 250
 - `ldaclose()` function, 905
 - `ldaopen()` function, 913
 - `ldclose()` function, 905
 - `ldconfig` — configure link-editor, 1658
 - `ldd` — list dynamic dependencies, 256
 - `ldfcn()` function, 906
 - `ldfhread()` function, 908
 - `ldgetname()` function, 909
 - `ldlitem()` function, 910
 - `ldlread()` function, 910
 - `ldlseek()` function, 911
 - `ldnlseek()` function, 911
 - `ldnrseek()` function, 915
 - `ldnshread()` function, 916
 - `ldnsseek()` function, 917
 - `ldohseek()` function, 912
 - `ldopen()` function, 913
 - `ldrseek()` function, 915
 - `ldshread()` function, 916
 - `ldsseek()` function, 917
 - `ldtbindx()` function, 918
 - `ldtbread()` function, 919
 - `ldtbseek()` function, 920
 - `ldterm`, terminal STREAMS module, 1252
 - `le` — Sun-3/50 10 Mb/s Ethernet interface, 1254 *thru* 1255
 - `leave` — remind you of leaving time, 257
 - `lex` — generate lexical analyzer, 258
 - LEX language tags file — `ctags`, 115
 - lexical analysis, C shell, 97
 - `lfind` — linear search routine, 923
 - library
 - find ordering for object — `lorder`, 277
 - make random — `ranlib`, 406
 - library file format — `ar`, 1370
 - library functions
 - introduction to C, 797
 - introduction to mathematical, 1081
 - introduction to RPC, 1107
 - introduction to standard I/O, 999, 1183
 - introduction to System V, 1127
 - library management
 - `ar` — library maintenance, 24
 - `life` — SunView game of life, 1527
 - lightweight processes library, 1051
 - `limit` command, 106
 - limits
 - disk space — `quota`, 405
 - get for user — `ulimit`, 1027
 - set for user — `ulimit`, 1027
 - `line` — read one line, 260
 - line discipline — `bk`, 1196
 - line discipline `ioctl`'s
 - `TIOCGETD` — get line discipline, 1196
 - `TIOCSETD` — set line discipline, 1196
 - `line` — plot line, 941
 - line numbering — `nl`, 337
 - line printer control — `lpc`, 1662 *thru* 1663
 - line printer daemon — `lpd`, 1664
 - line to Ethernet address — `ether_line`, 841
 - linear search and update routine — `lsearch`, 923
 - linear search routine — `lfind`, 923
 - `linemod` — set line style, 941
 - lines
 - count — `wc`, 591
 - find, in sorted file — `look`, 275
 - `link`, 696, 1659

- link, *continued*
 - make symbolic, 779
 - read value of symbolic, 737
- link editor — ld, 250, 1409
- link editor output — a.out, 1339
- lint — C program verifier, 261
- list mail command, 298
- listen, 697
- LISTER mail variable, 302
- literature references, find and insert — refer, 415
- lo — software loopback network interface, 1256
- load command, 267
- load frame buffer image — screenload, 442
- load mail command, 298
- loadc command, 267
- localtime — date and time conversion, 824
- localtime — date and time conversion, System V, 1132
- locate program — whereis, 594
- lock
 - file — flock, 660
 - record — fcntl, 656, 921
- lockd — network lock daemon, 1660
- lockf, 921
- lockscreen — save window context, 269
- log files and system log daemon — syslogd, 1773
- log — natural logarithm, 1086
- log gamma function — gamma, 1098
- log10 — logarithm, base 10, 1086
- log1p — natural logarithm, 1086
- log2 — logarithm, base 2, 1086
- logarithm, base 10 — log10, 1086
- logarithm, base 2 — log2, 1086
- logarithm, natural — log, 1086
- logb () function, 1096
- logger — make system log entry, 271
- login
 - change password — passwd, 379
 - display effective user name — whoami, 598
 - display login name — logname, 274
 - display user and group IDs — id, 230
 - info on users — finger, 186
 - list last — last, 248
 - make script of session — script, 444
 - rusers — who is on local network, 430
 - rwho — who is on local network, 432
 - save window context — lockscreen, 269
 - to local machine — login, 272
 - to remote machine — rlogin, 418
 - what are users doing — w, 588
 - who — who is logged in, 597
- login accounting, display login record — ac, 1574
- login command, 106, 458
- login environment
 - display variables — printenv, 391
 - tset — set terminal characteristics, 551
 - tty — get terminal name, 556
- login environment — environ, 1387
- .login file, 96
- login name, get — getlogin, 871
- login password
 - continued*
 - change password — passwd, 379
 - change in Yellow Pages — yppasswd, 611
- login records
 - lastlog file, 1485
 - utmp file, 1485
 - wtmp file, 1485
- logintool — graphic login interface, 1661
- logname — display login name, 274
- logout command, 106
 - .logout file, 96
- longjmp — non-local goto, 989, 1177
- look — find lines in a sorted file, 275
- look for pattern in file — grep, 218
- lookbib — find bibliographic references, 276
- loop, C shell control flow, 102
- loopback filesystem, 1257
- lorder — find ordering for object library, 277
- lpc — line printer control, 1662
- lpd — line printer daemon, 1664
- lpq — display printer queue, 278
- lpr — print files, 280
- lprm — remove print jobs, 283
- lptest command, 284
- lrand48 — generate uniformly distributed random numbers, 836
- ls — list files, 285
- lsearch — linear search and update routine, 923
- lseek — move file position, 698
- lstat — obtain file attributes, 774
- lwp_checkstkset () function, 1068
- lwp_create () function, 1064
- lwp_ctxinit () function, 1066
- lwp_ctxmemget () function, 1066
- lwp_ctxmemset () function, 1066
- lwp_ctxremove () function, 1066
- lwp_ctxset () function, 1066
- lwp_datastk () function, 1068
- lwp_destroy () function, 1064
- lwp_enumerate () function, 1071
- lwp_errstr () function, 1070
- lwp_fpset () function, 1066
- lwp_getregs () function, 1071
- lwp_getstate () function, 1071
- lwp_join () function, 1072
- lwp_libcset () function, 1066
- lwp_newstk () function, 1068
- lwp_perror () function, 1070
- lwp_ping () function, 1071
- lwp_resched () function, 1072
- lwp_resume () function, 1072
- lwp_self () function, 1071
- lwp_setpri () function, 1072
- lwp_setregs () function, 1071
- lwp_setstkcache () function, 1068
- lwp_sleep () function, 1072
- lwp_stkcswwset () function, 1068
- lwp_suspend () function, 1072
- lwp_yield () function, 1072

M

m4 — macro processor, 288
 m68k — machine type truth value, 292
 mach — display Sun processor, 291
 machine-dependent values — values, 1031
 machine-machine communication line discipline — bk, 1196
 macro processor — m4, 288
 madd — multiple precision add, 932
 magic file — file command's magic numbers table, 1410
 magnetic tape
 backspace files — mt, 333
 backspace records — mt, 333
 copy — tcopy, 515
 erase — mt, 333
 forward space files — mt, 333
 forward space records — mt, 333
 get unit status — mt, 333
 interface — mtio, 1268
 manipulate — mt, 333
 place unit off-line — mt, 333
 retension — mt, 333
 rewind — mt, 333
 scan — tcopy, 515
 skip backward files — mt, 333
 skip backward records — mt, 333
 skip forward files — mt, 333
 skip forward records — mt, 333
 write EOF mark on — mt, 333
 magnetic tape ioctl's
 MTIOCGET — get tape status, 1269
 MTIOCTOP — tape operation, 1268
 magnify raster image — rastrepl, 408
 mail
 enroll for secret — enroll, 604
 print waiting — prmail, 362
 receive secret mail — enroll, 604
 send secret mail — xsend, 604
 mail — send and receive mail, 293 thru 304
 mail commands, 296 thru 300
 !, 296
 #, 296
 =, 296
 ?, 296
 |, 298
 alias, 296
 alternates, 296
 cd, 296
 chdir, 296
 copy, 296
 Copy, 297
 delete, 297
 discard, 297
 dp, 297
 dt, 297
 echo, 297
 edit, 297
 else, 298
 endif, 298
 exit, 297
 file, 297
 folder, 297
 folders, 297

mail commands, *continued*

 followup, 297
 Followup, 297
 from, 297
 group, 296
 headers, 297
 help, 297
 hold, 298
 if, 298
 ignore, 297
 inc, 298
 list, 298
 load, 298
 mail, 298
 mbox, 298
 new, 298
 next, 298
 pipe, 298
 preserve, 298
 print, 298
 Print, 299
 quit, 299
 reply, 299
 Reply, 299
 Replyall, 299
 repliesender, 299
 respond, 299
 Respond, 299
 save, 299
 Save, 299
 set, 299
 shell, 299
 size, 299
 source, 299
 top, 299
 touch, 300
 type, 298
 Type, 299
 undelete, 300
 unread, 298
 unset, 300
 version, 300
 visual, 300
 write, 300
 xit, 297
 z, 300
 mail delivery server — sendmail, 1758
 mail environment variables
 HOME, 300
 MAIL, 300
 MAILRC, 300
 mail, forwarding messages, 300
 mail mail command, 298
 MAIL mail environment variable, 300
 mail services
 biff — mail notifier, 45
 binmail — version 7 mail, 46
 who is mail from — from, 194
 mail tilde escapes, 294 thru 295
 ~! , 294
 ~. , 294
 ~: , 294
 ~< , 295

mail tilde escapes, *continued*

~? , 294
 ~_ , 294
 ~| , 295
 ~A , 294
 ~b , 295
 ~c , 295
 ~d , 295
 ~e , 295
 ~f , 295
 ~h , 295
 ~i , 295
 ~m , 295
 ~p , 295
 ~q , 295
 ~r , 295
 ~s , 295
 ~t , 295
 ~v , 295
 ~w , 295
 ~x , 295

mail utilities

comsat — biff server, 1599
 create aliases database — newaliases, 1699
 statistics — mailstats, 1666

mail variable, 109, 454

mail variables, 300 *thru* 303

allnet, 301
 alwaysignore, 301
 append, 301
 askcc, 301
 asksub, 301
 autoprint, 301
 bang, 301
 cmd, 301
 conv, 301
 crt, 301
 DEAD, 301
 debug, 301
 dot, 301
 EDITOR, 302
 escape, 302
 folder, 302
 header, 302
 hold, 302
 ignore, 302
 ignoreeof, 302
 indentprefix, 302
 keep, 302
 keepsave, 302
 LISTER, 302
 MBOX, 302
 metoo, 302
 no, 302
 onehop, 302
 out folder, 302
 page, 302
 PAGER, 302
 prompt, 302
 quiet, 303
 record, 303
 replyall, 303
 save, 303
 sendmail, 303

mail variables, *continued*

sendwait, 303
 SHELL, 303
 showto, 303
 sign, 303
 toplines, 303
 verbose, 303
 VISUAL, 303

MAILCHECK variable — sh, 454

MAILPATH variable — sh, 454

MAILRC mail environment variable, 300

mailstats — mail delivery statistics, 1666

mailtool — SunView mail interface, 305

maintain programs — make, 311 *thru* 324, 355 *thru* 361

maintenance and operation, 1569

make

delta, SCCS — delta, 144

directory — mkdir, 328

fifo — mknod, 1674

file system — mkfs, 1673

hard link to file — ln, 265

implicit rules, list of — /usr/include/make/default.mk, 319

named pipe — mknod, 1674

new file system — newfs, 1700

SCCS delta — delta, 144

special file, 702

special file — mknod, 1674

symbolic link to file — ln, 265

system log entry — logger, 271

system log entry — syslog, 368

system special files — makedev, 1668

make — build programs, 311 *thru* 324, 355 *thru* 361

make directory, 700

make hard link to file, 696

makedbm — make Yellow Pages dbm file, 1667

makedev — make system special files, 1668

makekey — generate encryption key, 1669

malloc — allocate memory, 925

malloc_debug — set debug level, 926

malloc_verify — verify heap, 926

man — online display of reference pages, 325

-man — macros to format manual pages, 1560

manipulate Internet addresses, 893

manipulate magnetic tape — mt, 333

manual pages

create cat files for — catman, 1593

describe command — whatis, 593

map memory pages — mmap, 704

mask, set current signal — sigsetmask, 765

mathematical functions

acos, 1106

aint — convert to integral floating, 1102

anint — convert to integral floating, 1102

asin, 1106

atan, 1106

atan2, 1106

ceil — convert to integral floating, 1102

cos, 1106

cosh, 1088

exp — exponential, 1086

floor — convert to integral floating, 1102

- mathematical functions, *continued*
- gamma, 1098
 - hypot, 1089
 - rint — convert to integer, 1102
 - j0, 1084
 - j1, 1084
 - jn, 1084
 - log — natural logarithm, 1086
 - log10 — logarithm, base 10, 1086
 - log2 — logarithm, base 2, 1086
 - nint — convert to integer, 1102
 - pow — raise to power, 1086
 - rint — convert to integral floating, 1102
 - sin, 1106
 - sinh, 1088
 - tan, 1106
 - tanh, 1088
 - y0, 1084
 - y1, 1084
 - yn, 1084
- mathematical library functions, introduction to, 1081
- matherr — math library exception-handline function, 1099
- max_normal () function, 1097
- max_subnormal () function, 1097
- mbio — Multibus I/O space, 1261 *thru* 1262
- mbmem — Multibus memory space, 1261 *thru* 1262
- mbox mail command, 298
- MBOX mail variable, 302
- mc68881version — display MC68881 version, 1670
- mconnect — open connection to remote mail server, 1671
- mcp — Sun MCP Multiprotocol Communications Processor, 1258
- mdiv — multiple precision divide, 932
- me — macro package, 1563
- mem — main memory space, 1261 *thru* 1262
- memalign — allocate aligned memory, 925
- memcpy — copy memory character strings, 928
- memchr — index memory characters, 928
- memcmp compare memory characters, 928
- memcpy copy memory character fields, 928
- memory allocation debugging
- memory diagnostics — sysdiag, 1772
- memory image file format — core, 1378
- memory images
- kmem — kernel memory space, 1261 *thru* 1262
 - mbio — Multibus I/O space, 1261 *thru* 1262
 - mbmem — Multibus memory space, 1261 *thru* 1262
 - mem — main memory space, 1261 *thru* 1262
 - virtual — virtual address space, 1261 *thru* 1262
 - vme16 — VMEbus 16-bit space, 1261 *thru* 1262
 - vme16d16 — VMEbus address space, 1261 *thru* 1262
 - vme16d32 — VMEbus address space, 1261 *thru* 1262
 - vme24 — VMEbus 24-bit space, 1261 *thru* 1262
 - vme24d16 — VMEbus address space, 1261 *thru* 1262
 - vme24d32 — VMEbus address space, 1261 *thru* 1262
 - vme32d16 — VMEbus address space, 1261 *thru* 1262
 - vme32d32 — VMEbus address space, 1261 *thru* 1262
- memory management, 925 *thru* 926
- alloca — allocate on stack, 926
 - brk — set data segment break, 638
 - calloc — allocate memory, 925
 - cfree — free memory, 925
- memory management, *continued*
- free — free memory, 925
 - getpagesize, 678
 - malloc — allocate memory, 925
 - malloc_debug — set debug level, 926
 - malloc_verify — verify heap, 926
 - memalign — allocate aligned memory, 925
 - mmap, 704
 - mprotect, 709
 - munmap, 717
 - realloc — reallocate memory, 925
 - sbrk — change data segment size, 638
 - valloc — allocate aligned memory, 926
- memory management debugging
- memory operations, 928
- memset assign to memory characters, 928
- sort and collate lines — sort, 470
- merge files — paste, 380
- msg — permit or deny messages, 327
- message
- receive from socket — recv, 739
 - send from socket — send, 751
- message control operations
- msgctl, 710
 - msgget, 712
 - msgsnd, 713
- messages
- permit or deny — msg, 327
 - system error, 940
 - system signal, 948
- metacharacters in C shell, 97
- metoo mail variable, 302
- mfree — release multiple precision storage, 932
- mille — Mille Bornes game, 1528
- min — multiple precision decimal input, 932
- min_normal () function, 1097
- min_subnormal () function, 1097
- mincore () function, 699
- MINSTACKSZ () function, 1068
- miscellaneous information, 1547
- miscellaneous noninformation, 1547
- mkdir, 700
- mkdir — make directory, 328
- mkfile command, 1672
- mkfs — make file system, 1673
- mknod, 702
- mknod — make special file, 1674
- mkproto — make prototype file system, 1675
- mkstr — create C error messages, 329
- mktemp — make unique file name, 929
- mmap, 704
- modes, change permission — chmod, 65
- modf — split into integer part and fraction part, 1087
- modload command, 1676
- modstat command, 1677
- modunload command, 1678
- mon_break () function, 1074
- mon_cond_enter () function, 1074
- mon_create () function, 1074
- mon_destroy () function, 1074

mon_enter() function, 1074
 mon_enumerate() function, 1074
 mon_exit() function, 1074
 mon_waiters() function, 1074
 moncontrol — make execution profile, 930
 monitor — make execution profile, 930, 1074
 monitor program, 1679
 monitor traffic on the Ethernet, 1109
 monochrome frame buffer — bwone, 1197
 monochrome frame buffer — bwtwo, 1198
 monop — Monopoly game, 1531
 monstartup — make execution profile, 930
 moo game, 1533
 more — browse text file, 330
 mount, 706
 mount — mount filesystem, 1687
 mount file system — mount, 1687
 mount() function, 1112
 mountd — NFS mount server, 1690
 mounted file system table — mtab, 1396, 1412
 mouse — Sun mouse, 1263
 mouse — Sun mouse, 1264
 mout — multiple precision decimal output, 932
 move directory — mv, 334
 move file — mv, 334
 move file position — lseek, 698
 move — move current point, 941
 move print jobs — lpc, 1663
 mprotect, 709
 srand48 — generate uniformly distributed random numbers, 836
 -ms — macro package, 1565
 msg_enumrecv() function, 1077
 msg_enumsend() function, 1077
 msg_recv() function, 1077
 MSG_RECVALL() function, 1077
 msg_reply() function, 1077
 msg_send() function, 1077
 msgctl, 710
 msgget, 712
 msgsnd, 713
 msqrt — multiple precision exponential, 932
 msub — multiple precision subtract, 932
 msync function, 716
 mt — manipulate magnetic tape, 333
 mtab — mounted file system table, 1396, 1412
 mti — Systech MTI-800/1600 multi-terminal interface, 1266 *thru* 1267
 mtio — general magnetic tape interface, 1268
 MTIOCGET — get tape status, 1269
 MTIOCTOP — tape operation, 1268
 mtox — multiple precision to hexadecimal string, 932
 mult — multiple precision multiply, 932
 multiple columns, print in — pr, 388
 multiple precision integer arithmetic
 gcd, 932
 itom, 932
 madd, 932
 mdiv, 932

multiple precision integer arithmetic, *continued*

mfree, 932
 min, 932
 mout, 932
 msqrt, 932
 msub, 932
 mtox, 932
 mult, 932
 pow, 932
 rpow, 932
 sdiv, 932
 xtom, 932
 munmap, 717
 mv — move or rename files or directory, 334

N

name of terminal, find — ttyname, 1024
 name server routines, Internet, 965
 name termination handler — on_exit, 938
 named — internet domain name server daemon, 1691
 named pipe, make — mknod, 1674
 natural logarithm — log, 1086
 ncheck — convert i-numbers to filenames, 1693
 ndbootd daemon, 1694
 neqn — mathematical typesetting, 173
 netconfig — pnp diskful boot service, 1695
 netgroup — network groups list, 1413
 netmasks — netmask data base, 1414
 netstat — display network status, 1696
 network
 copy files across — rcp, 409
 rusers — who is logged in on local network, 430
 rwall — write to all users, 431
 rwho — who is logged in on local network, 432
 network byte order
 function to convert to host, 820
 network debugging — ping, 1709
 network entry, get — getnetent, 874
 network file system
 biod daemon, 1703
 nfsd daemon, 1703
 network file system daemons, 718
 network group entry
 get, 875
 network host entry, get — gethostent, 869
 network interface ioctl's
 SIOCADMULTI — set m/c address, 1227
 SIOCDELMULTI — delete m/c address, 1227
 SIOCGIFADDR — get ifnet address, 1226
 SIOCGIFCONF — get ifnet list, 1226
 SIOCGIFDSTADDR — get p-p address, 1226
 SIOCGIFFLAGS — get ifnet flags, 1226
 SIOCSIFADDR — set ifnet address, 1226
 SIOCSIFDSTADDR — set p-p address, 1226
 SIOCSIFFLAGS — set ifnet flags, 1226
 SIOCSPPROMISC toggle promiscuous mode, 1227
 network interface parameters, configure — ifconfig, 1642
 network interface, introduction to, 1189
 network lock manager protocol, 1114
 network loopback interface — lo, 1256
 network packet routing device — routing, 1288

- network routing daemon — `routed`, 1743
 - network rwall server — `rwalld`, 1752
 - network service entry, get — `getservent`, 886
 - network services status monitor files, 1438
 - network status, display — `netstat`, 1696
 - networks — network name data base, 1416
 - new mail command, 298
 - newaliases — make mail aliases database, 1699
 - newfs — make new file system, 1700
 - newgrp — change group ID of user, 335, 458
 - newkey command, 1702
 - next mail command, 298
 - nextafter() function, 1093
 - nextkey — find next key, 830
 - NFS directories to export— `exports`, 1389
 - NFS exported directories— `xtab`, 1389
 - NFS mount server — `mountd`, 1690
 - NFS statistics, display — `nfstat`, 1704
 - NFS, network file system protocol, 1271
 - nsd daemon, 1703
 - nfstat — display network statistics, 1704
 - nfssvc, 718
 - nice command, 106, 336
 - nice — change priority of a process, 936
 - nice — change priority of a process, System V, 1166
 - nint — nint of, 1102
 - NIT, Network Interface Tap, 1272
 - nit_buf, NIT buffering module, 1275
 - nit_if, NIT device interface, 1277
 - nit_pf, NIT packet filtering module, 1279
 - nl — number lines, 337
 - nlist — get entries from symbol table, 937, 1167
 - nm — display name list, 339
 - no mail variable, 302
 - nobeep variable, 109
 - noclobber variable, 109
 - noglob variable, 110
 - nohup command, 106, 342
 - non-local goto
 - non-local goto — `longjmp`, 989, 1177
 - non-local goto — `setjmp`, 989, 1177
 - noninformation miscellaneous, 1547
 - information miscellaneous, 1547
 - nonomatch variable, 110
 - notify command, 106
 - notify variable, 110
 - nrnd48 — generate uniformly distributed random numbers, 836
 - nroff — document formatter, 344
 - nroff utilities
 - checknr — check nroff/troff files, 62
 - col — filter reverse paper motions, 77
 - colcrt — filter nroff output for CRT, 78
 - nroff utilities, 146
 - soelim — eliminate `.so`'s, incorporate sourced-in files, 469
 - nslookup command, 1705
 - ntohl — convert network to host long, 820
 - ntohs — convert host to network short, 820
 - null — null device, 1282
 - null-terminated strings
 - compare — `strcmp`, 1001
 - compare — `strncmp`, 1001
 - concatenate — `strcat`, 1001
 - concatenate — `strncat`, 1001
 - copy — `strcpy`, 1001
 - copy — `strncpy`, 1001
 - index — `index`, 1001
 - index — `rindex`, 1001
 - reverse index — `rindex`, 1001
 - number — convert Arabic numerals to English, 1534
 - numbers
 - convert to strings, 838
- ## O
- objdump command, 346
 - object code management
 - ar — library maintenance, 24
 - ranlib — make random library, 406
 - object file
 - find printable strings in — `strings`, 478
 - size — find object file size, 465
 - strip — strip symbols and relocation bits, 479
 - object library
 - find ordering for — `lorder`, 277
 - octal dump file — `od`, 348
 - od — dump file, 348
 - on — remote command execution, 373
 - on_exit — name termination handler, 938
 - onehop mail variable, 302
 - onintr command, 106
 - online reference — `man`, 325
 - open, 719
 - open database — `dbmopen`, 830
 - open directory stream — `opendir`, 834
 - open stream — `fopen`, 854
 - open stream, System V — `fopen`, 1160
 - opendir — open directory stream, 834
 - openlog — initialize system log file, 1011
 - openpl — open plot device, 941
 - optarg() function, 876
 - optind() function, 876
 - options on sockets
 - get, 687
 - set, 687
 - organizer — get organizer, 374
 - outfolder mail variable, 302
 - output conversion
 - fprintf — convert to stream, 944
 - fprintf — convert to stream, System V, 1168
 - printf — convert to stdout, 944
 - printf — convert to stdout, System V, 1168
 - sprintf — convert to string, 944
 - sprintf — convert to string, System V, 1168
 - overview — take over screen w/ graphics, 375
 - owner of file, change — `chown`, 1595

P

- pac — printer/plotter accounting, 1708
- pack — pack files, 376
- packet routing device — routing, 1288
- packet routing `ioctl`'s
 - `SIOCADDRT` — add route, 1288
 - `SIOCDELRT` — delete route, 1288
- page — browse text file, 330
- page mail variable, 302
- page size, display — `pagesize`, 378
- page size, get — `getpagesize`, 678
- PAGER mail variable, 302
- `pagesize` — display page size, 378
- paging device — `swapon`, 778, 1213
- paging devices, specify — `swapon`, 1771
- paging system, advise — `vadvise`, 788
- parent process identification, get — `getppid`, 680
- parentheses, C shell command grouping, 97
- parser generator — `yacc`, 607
- Pascal language
 - tags file — `ctags`, 115
- pass framebuffer info `ioctl` — `GPIO_PUT_INFO`, 1221
- `passwd` — change login password, 379
- `passwd` — password file, 1417
- `passwd.adjunct` — password file, 1419
- `passwd2des` () function, 1124
- password
 - change in Yellow Pages — `yppasswd`, 611
 - change login — `passwd`, 379
 - read — `getpass`, 878
 - read, System V — `getpass`, 1163
- password file
 - add entry — `putpwent`, 951
 - edit — `vipw`, 1790
 - get entry — `endpwnam`, 882
 - get entry, System V — `endpwnam`, 1164
 - get entry — `fgetpwent`, 882
 - get entry, System V — `fgetpwent`, 1164
 - get entry — `getpwent`, 882
 - get entry, System V — `getpwent`, 1164
 - get entry — `getpwnam`, 882
 - get entry, System V — `getpwnam`, 1164
 - get entry — `getpwuid`, 882
 - get entry, System V — `getpwuid`, 1164
 - get entry — `setpwent`, 882
 - get entry, System V — `setpwent`, 1164
 - get entry — `fsetpwfile`, 882
 - get entry, System V — `fgetpwent`, 1164
- paste — horizontal merge, 380
- path variable, 110, 454
- patterns, search in file for — `grep`, 218
- pause — stop until signal, 939
- `pcat` — pack files, 376
- `pclose` — close stream to process, 943
- `pdpl1` — machine type truth value, 292
- peer name, get — `getpeername`, 679
- `perfmeter` — display performance statistics, 381
- performance monitoring — `perfmeter`, 381
 - display call-graph profile data — `gprof`, 214
 - `prof` — display program profile, 392
 - performance monitoring — `perfmeter`, *continued*
 - time — time command, 532
- periodic jobs table — `crontab`, 1381
- peripheral diagnostics — `sysdiag`, 1772
- permissions, change mode — `chmod`, 65
- permit messages — `msg`, 327
- permuted index, generate — `ptx`, 403
- `perror` — system error messages, 940
- `pg` — browse text file, 383
- phones — remote host phone numbers, 1420
- `ping` — debug network, 1709
- pipe, 722
- pipe mail command, 298
- pipeline, C shell, 97
- place magnetic tape unit off-line — `mt`, 333
- plot — graphics filters, 386
- plot — graphics interface files, 1421
- plot on Versatec — `vplot`, 584
- `pnp` — automatic network installation, 1115
- PNP
 - `pnpd` — PNP daemon, 1711
- `pnplib` — `pnplib` diskless boot service, 1710
- `pnplib` — `pnplib` diskless boot service, 1710
- `pnplib` — `pnplib` diskless boot service, 1710
- `pnplib` — PNP daemon, 1711
- `pod_exit` () function, 1064
- `pod_getexit` () function, 1064
- `pod_getmaxpri` () function, 1079
- `pod_getmaxsize` () function, 1079
- `pod_setexit` () function, 1064
- `pod_setmaxpri` () function, 1079
- point — plot point, 941
- policies file, 1359
- `poll` function, 723
- `popd` command, 107
- `popen` — open stream to process, 943
- `portmap` — DARPA to RPC mapper, 1712
- position of directory stream — `telldir`, 834
- `pow` — multiple precision exponential, 932, 1086
- power function — `pow`, 1086
- `pp`, Sun386i parallel printer port, 1283
 - `bcd` — convert to antique media, 1502
- `pr` — prepare files for printing, 388
- `praudit` — display audit trail, 1713
- predefined variables, in C shell, 109
- prepare execution profile
 - `moncontrol` — make execution profile, 930
 - `monitor` — make execution profile, 930
 - `monstartup` — make execution profile, 930
- prepare files for printing — `pr`, 388
- preserve mail command, 298
- pretty printer
 - `indent` — format C source, 231
 - `vgrind` — make formatted listings, 580
- `colcrt` command, 78
- primes game, 1535
- primitive system data types — `types`, 1480
- print
 - print waiting mail — `prmail`, 362

- print, *continued*
 values from YP database — `ypcat`, 609
 working directory name — `pwd`, 404
- print bibliographic database — `roffbib`, 422
- print files
`lpr` — print files, 280
- print mail command, 298
- Print mail command, 299
- printcap — printer capability data base, 1422
- printenv — display environment, 391
- printer
 abort — `lpc`, 1662
 clean queue — `lpc`, 1662
 control — `lpc`, 1662
 daemon — `lpd`, 1664
 disable queue — `lpc`, 1662
`lpq` — display queue, 278
 enable queue — `lpc`, 1662
 move jobs — `lpc`, 1663
 remove jobs from queue — `lprm`, 283
 restart — `lpc`, 1662
 start — `lpc`, 1662
 status of — `lpc`, 1663
 stop — `lpc`, 1663
 take printer down — `lpc`, 1662
- printer interface
`vp` — Ikon 10071-5 Versatec parallel printer interface, 1326
`vpc` — Systech VPC-2200 Versatec/Centronics interface, 1327
- printer/plotter accounting, 1708
- printf — formatted output conversion, 944
- printf — format to stdout, System V, 1168
- priority
`get`, 681
`set`, 681
- priority of process — `nice`, 936
- priority of process, System V — `nice`, 1166
- prmail — print waiting mail, 362
- procedure calls, assembler, expand in-line, `inline`, 236
- process
 and child process times, System V — `times`, 1185
 change priority — `renice`, 1728
 create, 661
 display status — `ps`, 399
 get core image of, 202
 get identification — `getpid`, 680
 get times — `times`, 1017
 initiate I/O to or from, 943
 priority — `nice`, 936
 priority, System V — `nice`, 1166
 send signal to — `kill`, 693
 software signals — `sigvec`, 767 thru 770
 terminate — `kill`, 247, 655
 terminate and cleanup — `exit`, 844
 tracing — `ptrace`, 726
 wait — wait process completion, 589
- process block `ioctl` — `DESIOCBLOCK`, 1210
- process group
`get` — `getpgrp`, 753
 send signal to — `killpg`, 695
`set` — `setpgrp`, 753
- process quickly `ioctl` — `DESIQCQUICK`, 1210
- processes and protection
`execve`, 652
`exit`, 655
`fork`, 661
`getdomainname`, 668
`getegid`, 670
`geteuid`, 691
`getgid`, 670
`getgroups`, 671
`gethostid`, 672
`gethostname`, 673
`getpgrp`, 753
`getpid`, 680
`getppid`, 680
`getuid`, 691
`ptrace`, 726
`setdomainname`, 668
`setgroups`, 671
`sethostname`, 673
`setpgrp`, 753
`setregid`, 754
`setreuid`, 755
`vfork`, 789
`vhangup`, 790
`wait`, 791
`wait3`, 791
`wait4`, 791
- `prof` — profile within a function, 947
- `prof` — display program profile, 392
- `profil`, 725
- profile
 display call-graph — `gprof`, 214
- profile, execution — `monitor`, 930
- profiling
`prof` — display program profile, 392
- `prof`, 947
- program verification — `assert`, 813
- program verification, System V — `assert`, 1131
- programming languages
 analyze and disperse compiler error messages, 175
 assembler, 27
`cc` — C compiler, 52
`cpp` — C preprocessor, 89
`cxref` — cross reference C program, 123
`lex` — generate lexical analyzer, 258
`lint` — C program verifier, 261
`vgrind` — make formatted listings, 580
`xstr` — extract strings from C code, 605
- programming tools
`adb` — debug tool, 13
`bc` — calculator language, 43
`cflow` — code flow graph, 60
 compiler generator, 351
`ctags` — create tags file, 115
`ctrace` — display program trace, 117
`dbx` — source debugger, 126
`dbxtool` — debugger, 135
 display call-graph profile data — `gprof`, 214
`indent` — format C source, 231
`install` — install files, 240
`ld` — link editor, 250
`lex` — generate lexical analyzer, 258
`lint` — C program verifier, 261

programming tools, *continued*

lorder — find ordering for object library, 277
m4 — macro processor, 288
 maintain object libraries, 24
make — build programs, 311 *thru* 324, 355 *thru* 361
mkstr — create C error messages, 329
nm — display name list, 339
prof — display program profile, 392
ranlib — make random library, 406
sccs — source code control system, 434
size — find object file size, 465
strings — find printable strings in binary file, 478
strip — strip symbols and relocation bits, 479
tcov — code coverage tool, 516
time — time command, 532
touch — update last modified date of file, 541
unifdef — eliminate `#ifdef`'s from C input, 560
yacc — parser generator, 607
 programs, introduction, 3
 PROM monitor program, 1679
 PROM monitor program, display and load program — `eeprom`, 1617
 prompt mail variable, 302
 prompt variable, 110
 protocol entry
 get, 879
 protocol specifications, 1107
protocols — protocol name data base, 1425
 protocols, introduction to, 1189
 provide truth values — `true`, 550
prs — display SCCS history, 394
prt — display SCCS history, 397
ps — display process status, 399
 PS1 variable — `sh`, 454
 PS2 variable — `sh`, 454
psignal — system signal messages, 948
pstat — display system statistics, 1714
pti — (old) `troff` interpreter, 363
ptrace, 726
ptx — generate permuted index, 403
pty — pseudo-terminal driver, 1284 *thru* 1286
publickey file, 1426
 push character back to input stream — `ungetc`, 1028
pushd command, 107
 put character to stdout — `putchar`, 949
 put character to stream — `fputc`, 949
 put character to stream — `putc`, 949
 put string to stdout — `puts`, 952
 put string to stream — `fputs`, 952
 put word to stream — `putw`, 949
putc — put character on stream, 949
putchar — put character on stdout, 949
putenv — set environment value, 950
putmsg () function, 730
putpwent — add password file entry, 951
puts — put string to stdout, 952
putw — put word on stream, 949
pwck — check password database entries, 1718
pwd — print working directory name, 404, 458
pwdauth () function, 953

pwdauthd daemon, 1719

Q

qsort — quicker sort, 954
queue
 atq — display delayed execution, 31
 lpq — display printer, 278
 insert element in — `insque`, 896
 remove element from — `remque`, 896
 remove jobs from delayed execution — `atrm`, 32
 remove jobs from printer — `lprm`, 283
queuedefs file, 1427
 quick substitution — in C shell, 99
 quicker sort — `qsort`, 954
 quiet mail variable, 303
quiet_nan () function, 1097
 quit mail command, 299
 quiz — test knowledge, 1536
quot — summarize file system ownership, 1720
quota — display disk usage and limits, 405
quotacheck — check quota consistency, 1721
quotactl — disk quotas, 732
quotaoff — turn file system quotas off, 1722
quotaon — turn file system quotas on, 1722
quotas
 edquota — edit user quotas, 1616
 quotacheck — check quota consistency, 1721
 quotaoff — turn file system quotas off, 1722
 quotaon — turn file system quotas on, 1722
 repquota — summarize quotas, 1729
 rquotad — remote quota server, 1747

R

rain — display raindrops, 1537
rand — generate random numbers, 955
rand — generate random numbers, System V, 1171
random — generate random number, 956
random game, 1538
random number generator
 drand48, 836
 erand48, 836
 initstate, 956
 jrand48, 836
 lcg48, 836
 lrand48, 836
 mrnd48, 836
 nrnd48, 836
 rand, 955
 random, 956
 seed48, 836
 setstate, 956
 srand, 955
 srand48, 836
 srandom, 956
random number generator, System V
 rand, 1171
 srand, 1171
ranlib — make random library, 406
rarpd — reverse Address Resolution Protocol daemon, 1723
rasfilter8to1 — convert 8-bit rasterfile to 1-bit rasterfile, 407

- rasterfile, 1428
- `rast repl` — magnify raster image, 408
- `rc` — startup commands, 1724
- `rcmd` — execute command remotely, 958
- `rcp` — remote file copy, 409
- `rdate` — remote date, 1726
- `rdist` — remote file distribution, 411
- `re_comp` — compile regular expression, 961
- `re_exec` — execute regular expression, 961
- `read`, 734
- `read` command, 458
- `read` directory stream — `readdir`, 834
- `read` formatted
 - `fscanf` — convert from stream, 984
 - `fscanf` — convert from stream, System V, 1172
 - `scanf` — convert from stdin, 984
 - `scanf` — convert from stdin, System V, 1172
 - `sscanf` — convert from string, 984
 - `sscanf` — convert from string, System V, 1172
- `read` from stream — `fread`, 855
- `read` mail — `mail`, 293 *thru* 304
- `read` password — `getpass`, 878
- `read` password, System V — `getpass`, 1163
- `read` scattered — `readv`, 734
- `read/write` pointer, move — `lseek`, 698
- `readdir` — `read` directory stream, 834
- `readlink`, 737
- `readonly` command, 459
- `real` group ID
 - set, 754
- `real` group ID, set — `setrgid`, 991
- `real` user ID
 - get — `getuid`, 691
 - set — `setreuid`, 755
- `real` user ID, set — `setruid`, 991
- `realloc` — reallocate memory, 925
- `reallocate` memory — `realloc`, 925
- `realpath` () function, 960
- `reboot` — halt processor, 738
- `reboot` — system startup procedures, 1727
- `reboot` system — `fastboot`, 1624
- `rebuild` Yellow Pages database — `ypmake`, 1794
- `receive` message from socket, 739
- `receive` secret mail — `enroll`, 604
- `reconfigure` fb `ioctl` — `GPIO_REDIRECT_DEVFB`, 1221
- `record` mail variable, 303
- `recv` — `receive` message from socket, 739
- `recvfrom`, 739
- `recvmsg`, 739
- `refer` — insert literature references, 415
- `regenerate` programs — `make`, 311 *thru* 324, 355 *thru* 361
- `regexp` — regular expression compile and match routines, 962
- regular expressions
 - compile — `re_comp`, 961
 - execute — `re_exec`, 961
- `rehash` command, 107
- relational database operator — `join`, 245
- `release` blocked signals — `sigpause`, 764
- `remainder` () function, 1093
- `remexportent` () function, 845
- reminder services
 - `biff` — mail notifier, 45
 - `calendar` — reminder service, 49
 - `leave` — remind you of leaving time, 257
- remote command execution — on, 373
- remote command, return stream to — `rcmd`, 958
- remote command, return stream to — `rexec`, 967
- remote execution protocol — `rex`, 1117
- remote execution server — `rexecd`, 1735
- remote — remote host descriptions, 1429
- remote file copy — `rcp`, 409
- remote host
 - number of users — `rusers`, 1118
 - phone numbers — `phones`, 1420
 - send file to — `uusend`, 571
- remote input editing `ioctl` — `TIOCREMOTE`, 1285
- remote kernel performance, 1120
- remote login
 - `rlogin`, 418
 - server — `rlogind`, 1737
- remote magtape protocol server — `rmt`, 1739
- remote procedure call services
 - `rquotad` — remote quota server, 1747
 - `sprayd` — spray server, 1767
- remote procedure calls, 968, 1044
- remote shell — `rsh`, 426
- remote shell server — `rshd`, 1748
- remote system
 - connect to — `cu`, 533
 - connect to — `tip`, 533
- remote users, number of — `rnusers`, 1118
- remove
 - close-on-exec flag `ioctl` — `FIONCLEX`, 1219
 - columns from file, 80, 121
 - delayed execution jobs — `atrm`, 32
 - remove delta from SCCS file — `rmdel`, 421
 - directory — `rmdir`, 420, 743
 - directory entry — `unlink`, 785
 - element from queue — `remque`, 896
 - file — `rm`, 420
 - file system — `unmount`, 786
 - filename affixes — `basename`, 42
 - `nroff`, `troff`, `tbl` and `eqn` constructs — `deroff`, 146
 - print jobs — `lprm`, 283
 - repeated lines — `uniq`, 561
- `remque` — remove element from queue, 896
- rename directory — `mv`, 334
- rename file — `mv`, 334, 741
- `renice` — change process priority, 1728
- reopen stream — `freopen`, 854
- reopen stream, System V — `freopen`, 1160
- repeat command, 107
- reply mail command, 299
- Reply mail command, 299
- `replyall` mail variable, 303
- Replyall mail command, 299
- `repliesender` mail command, 299
- report file system quotas — `repquota`, 1729

- reposition stream
 - fseek, 856
 - ftell, 856
 - rewind, 856
 - repquota — summarize quotas, 1729
 - res_init — Internet name server routines, 965
 - res_mkquery — Internet name servers, 965
 - res_send — Internet name server routines, 965
 - reset — reset terminal bits, 551
 - reset terminal bits — reset, 551
 - resolve.conf file — name server initialization info, 1431
 - resource consumption, control — vlimit, 1034
 - resource control
 - getrlimit, 682
 - getrusage, 684
 - setrlimit, 682
 - resource controls
 - getpriority, 681
 - setpriority, 681
 - resource usage, get information about — vtimes, 1037
 - resource utilization, get information about — getrusage, 684
 - respond mail command, 299
 - Respond mail command, 299
 - restart GP ioctl — GP1IO_CHK_GP, 1221
 - restart printer — lpc, 1662
 - restore — restore file system, 1730
 - restore file system — restore, 1730
 - restore frame buffer image — screenload, 442
 - retension magnetic tape — mt, 333
 - retrieve datum under key — fetch, 830
 - return command, 459
 - return stream to remote command — rcmd, 958
 - return stream to remote command — rexec, 967
 - return to saved environment — longjmp, 989, 1177
 - rev — reverse lines in file, 417
 - reverse index strings — rindex, 1001
 - reverse lines in file — rev, 417
 - rewind directory stream — rewinddir, 834
 - rewind — rewind stream, 856
 - rewind magnetic tape — mt, 333
 - rewind stream — rewind, 856
 - rewinddir — rewind directory stream, 834
 - rexcd — remote execution daemon, 1734
 - rexec — return stream to remote command, 967
 - rexecd — remote execution server, 1735
 - .rgb file, 1360
 - rindex — find character in string, 1001
 - rint — rint of, 1102
 - rlogin — remote login, 418
 - rlogind — remote login server, 1737
 - rm — remove file or directory, 420
 - rmail — process remote mail, 1738
 - rmdel — remove delta from SCCS file, 421
 - rmdir — remove directory, 743
 - rmdir — remove directory, 420
 - rmt — remote magtape protocol server, 1739
 - robots game, 1539
 - roffbib — print bibliographic database, 422
 - root directory, change — chroot, 644
 - root directory, change for a command — chroot, 1596
 - root, Sun386i root disk device, 1287
 - route — manipulate routing tables, 1741
 - routed — network routing daemon, 1743
 - routing — local network packet routing, 1288
 - routing ioctl's
 - SIOCADDRT — add route, 1288
 - SIOCDELRT — delete route, 1288
 - RPC routines, 968, 1044
 - RPC
 - generate protocols — rpcgen, 424
 - report RPC information — rpcinfo, 1745
 - RPC library functions, introduction to, 1107
 - RPC program entry, get — getrpcent, 884
 - rpc — rpc name data base, 1432
 - RPC protocol specifications, 1107
 - rpcgen — generate RPC protocol, C header files, and server skeleton, 424
 - rpcinfo — report RPC information, 1745
 - rpow — multiple precision exponential, 932
 - rquota () function, 1119
 - rquotad — remote quota server, 1747
 - rresvport — get privileged socket, 958
 - rsh — remote shell, 426
 - rshd — remote shell server, 1748
 - rstat — performance data from remote kernel, 1120
 - rstatd — kernel statistics server, 1750, 1751
 - rtime () function, 982
 - rup — display status of network hosts, 428
 - ruptime — display status of local hosts, 429
 - ruserok — authenticate user, 958
 - rusers — who is logged in on local network, 430
 - rwall — write to specified remote machines, 1121
 - rwall.d — network rwall server, 1752
 - rwho — who is logged in on local network, 432
 - rwhod — system status server, 1753
- ## S
- sa — process accounting summary, 1755
 - SAMECV () function, 1058
 - SAMEMON () function, 1074
 - SAMETHREAD () function, 1064
 - save mail command, 299
 - save mail variable, 303
 - save stack environment — setjmp, 989, 1177
 - savecore — save OS core dump, 1757
 - savehist variable, 110
 - sbrk — change data segment size, 638
 - scalb () function, 1096
 - scalbn () function, 1093
 - scan directory — alphasort, 983
 - scan directory — scandir, 983
 - scandir — scan directory, 983
 - scanf — convert from stdin, 984
 - scanf — convert from stdin, System V, 1172
 - scatter read — readv, 734
 - sccs — source code control system, 434
 - SCCS commands

SCCS commands, *continued*

- admin — administer SCCS, 21
- cdc — change delta commentary, 58
- comb — combine deltas, 81
- get — get SCCS file, 203
- help — get SCCS help, 224
- cdc — display SCCS history, 394
- prt — display SCCS history, 397
- rmDEL — remove delta, 421
- sact — display SCCS file editing status, 433
- sccsdiff — compare versions of SCCS file, 438
- unget — unget SCCS file, 559
- val — validate SCCS file, 578

SCCS delta

- change commentary, 58
- combine, 81
- create — delta, 144
- remove — rmDEL, 421

sccsdiff — compare versions of SCCS file, 438

sccsfile — SCCS file format, 1433

schedule signal — alarm, 812, 1026

scheduling priority

- get, 681
- set, 681

screen fonts, edit — fontedit, 190

screen-oriented editor — vi, 582

screenblank — turn of idle screen, 439

screenDump — dump frame buffer image, 440

screenload — load frame buffer image, 442

script — make script of terminal session, 444

sd — Adaptec ST-506 Disk driver, 1289 *thru* 1290

sdiff — side-by-side compare, 445

sdiv — multiple precision divide, 932

search for files, 183

search for pattern in file — grep, 218

search functions

- bsearch binary search, 816
- hsearch — hash table search, 891
- lsearch — linear search and update, 923

seconvert — convert number to ASCII, 838

secret mail

- enroll for — enroll, 604
- receive — enroll, 604
- send — xsend, 604

sed — stream editor, 446

seed48 — generate uniformly distributed random numbers, 836

seek in directory stream — seekdir, 834

seek on stream — fseek, 856

seekdir — seek in directory stream, 834

select, 744

selection, copy to standard output — get_selection, 208

selection_svc, 451

semaphore

- control — semctl, 746
- get set of — semget, 748
- operations — semop, 749

semctl — semaphore controls, 746

semget — get semaphore set, 748

semop — semaphore operations, 749

send

send, *continued*

- file to remote host — uuse, 571
- message from socket — send, 751
- secret mail — xsend, 604
- signal to process — kill, 247, 693
- signal to process group — killpg, 695

send a keyboard command ioctl — KIOCCMD, 1236

send and receive mail — mail, 293 *thru* 304

sendmail aliases file — aliases, 1367

sendmail — mail delivery system, 1758

sendmail aliases file — .forward, 1367

sendmail mail variable, 303

sendmsg — send message over socket, 751

sendto — send message to socket, 751

sendwait mail variable, 303

serial communications driver — zs, 1335 *thru* 1336

servers

- comsat — biff server, 1599
- ftpd — Internet File Transfer Protocol, 1632
- inetd — Internet server daemon, 1644
- lockd — network lock daemon, 1660
- mountd — mount request server, 1690
- named — internet domain name server daemon, 1691
- npd — PNP daemon, 1711
- rexeed — remote execution server, 1735
- rlogind — remote login server, 1737
- rshd — remote shell server, 1748
- rstatd — kernel statistics server, 1750, 1751
- rwalld — network rwall server, 1752
- rwhod — system status server, 1753
- statd — network status monitor, 1768
- talkd — talk program server, 1774
- timed — DARPA Time server, 1781
- tname — DARPA Trivial name server, 1782
- yppasswdd — Yellow Pages password server, 1795

service entry, get — getservent, 886

set

- arp entry ioctl — SIOCSARP, 1194
- close-on-exec for fd ioctl — FIOCLEX, 1219
- current domain name — domainname, 157
- current host name, 227
- current signal mask — sigsetmask, 765
- date and time — gettimeofday, 689
- disk geometry ioctl — DKIOCSGEOM, 1211
- disk partition info ioctl — DKIOCSPART, 1211
- environment value — putenv, 950
- file creation mode mask — umask, 783
- file owner ioctl — FIOSETOWN, 1219
- high water mark ioctl — SIOCShiwat, 1304
- ifnet address ioctl — SIOCSIFADDR, 1226
- ifnet flags ioctl — SIOCSIFFLAGS, 1226
- line discipline ioctl — TIOCSETD, 1196
- low water mark ioctl — SIOCSLOWAT, 1304
- m/c address ioctl — SIOCADDMULTI, 1227
- memory management debug level — malloc_debug, 926
- name of current host, 227
- network group entry — setnetgrent, 875
- network service entry — getservent, 886
- p-p address ioctl — SIOCSIFDSTADDR, 1226
- process domain name — setdomainname, 668
- RPC program entry — setrpcnt, 884
- scheduling priority — setpriority, 681

set, *continued*

- signal stack context — sigstack, 766
- terminal characteristics — stty, 480
- terminal characteristics — tset, 551
- terminal state — stty, 1009
- user limits — ulimit, 1027
- user mask — umask, 783
- set command, 107, 459
- set high water mark `ioctl` — SIOCShiwat, 1325
- set keyboard “direct input” state `ioctl` — KIOCSDIRECT, 1236
- set keyboard translation `ioctl` — KIOCTRANS, 1235
- set low water mark `ioctl` — SIOCSLOWAT, 1325
- set mail command, 299
- set options sockets, 687
- set/clear
 - async I/O `ioctl` — FIOASYNC, 1219
 - non-blocking I/O `ioctl` — FIONBIO, 1219
 - packet mode (pty) `ioctl` — TIOCPKT, 1285
- setac() function, 859
- setuseraudit() function, 756
- setbuf — assign buffering, 987
- setbuf — assign buffering, System V, 1175
- setbuffer — assign buffering, 987
- setbuffer — assign buffering, System V, 1175
- setdomainname — set process domain, 668
- setegid — set effective group ID, 991
- setenv command, 107
- seteuid — set effective user ID, 991
- setexportent() function, 845
- setfsent — get file system descriptor file entry, 865
- setgid — set group ID, 991
- setgid — set group ID, System V, 1179
- setgraent() function, 866
- setgrent — get group file entry, 867
- setgroups, 671
- sethostent — get network host entry, 869
- sethostname, 673
- setitimer, 674
- setjmp — save stack environment, 989, 1177
- setjmp — non-local goto, 989, 1177
- setkey — encryption, 822
- setkeys — change keyboard layout, 364
- setlinebuf — assign buffering, 987
- setlinebuf — assign buffering, System V, 1175
- closelog — set log priority mask, 1011
- setmntent — get filesystem descriptor file entry, 872
- setnetent — get network entry, 874
- setnetgrent — get network group entry, 875
- setpgrp, 753
- setpriority, 681
- setprotoent — get protocol entry, 879
- setpwaent() function, 881
- setpwent — get password file entry, 882
- setpwent — get password file entry, System V, 1164
- fgetpwent — get password file entry, 882
- setregid, 754
- setreuid, 755
- setrgid — set real group ID, 991
- setrlimit, 682
- setrpcnt — get RPC entry, 884
- setruid — set real user ID, 991
- setservent — get service entry, 886
- setsockopt, 687
- setstate — random number routines, 956
- settimeofday, 689
- setttyent() function, 887
- setuid — set user ID, 991
- setuid — set user ID, System V, 1179
- setup_client command, 1761
- setup_exec command, 1763
- setuseraudit() function, 756
- setusershell() function, 889
- setvbuf — assign buffering, 987
- setvbuf — assign buffering, System V, 1175
- sfconvert — convert number to ASCII, 838
- sgconvert — convert number to ASCII, 838
- sh command, Bourne shell, 452 *thru* 460
- shared libraries
 - display users of — ldd, 256
- shared memory
 - control — shmctl, 757
 - get segment — shmget, 758
 - operation — shmop, 760
- shell
 - remote — rsh, 426
- shell command, issuing — system, 1013
- shell functions, Bourne, 453
- shell mail command, 299
- SHELL mail variable, 303
- shell variable, 110, 455
- shell variables, in Bourne shell, 454 *thru* 455
- shell window
 - cmdtool, 73
 - shelltool, 461
- shelltool — shell terminal window, 461
- shift command, 107, 459
- shift_lines — textedit selection filter, 529
- shmctl — shared memory control, 757
- shmget — get shared memory segment, 758
- shmop — get shared memory operations, 760
- showmount — display remote mounts, 1764
- showto mail variable, 303
- shutdown, 762
- shutdown — shut down multiuser operation, 1765
- si — Sun SCSI Disk driver, 1289 *thru* 1290
- sigblock, 763
- sigfpe function
 - sigfpe() function, 992
- siginterrupt — interrupt system calls with software signal, 994
- sign mail variable, 303
- login — sign on, 272
 - to remote machine — rlogin, 418
- sign-on last — last, 248
- signal
 - schedule — alarm, 812, 1026
 - stop until — pause, 939

- signal — software signals, 995, 998
- signal — software signals, System V, 1180
- signal handling, in C shell, 103
- signal messages
 - psignal, 948
 - sys_siglist, 948
- signaling_nan () function, 1097
- signals
 - kill, 693
 - killpg — send to process group, 695
 - sigblock, 763
 - sigpause, 764
 - sigsetmask, 765
 - sigstack — signal stack context, 766
- signbit () function, 1093
- significant and exponent, split into — frexp, 1087
- significant () function, 1096
- sigpause, 764
- sigsetmask, 765
- sigstack — signal stack context, 766
- sigvec — software signals, 767 thru 770
- sin — trigonometric sine, 1106
- single-precision versions of math functions, 1104
- single_to_decimal — decimal record from single-precision floating, 850
- sinh — hyperbolic sine, 1088
- SIOCADMULTI — set m/c address, 1227
- SIOCADDRT — add route, 1288
- SIOCDDARP — delete arp entry, 1194
- SIOCDELMULTI — delete m/c address, 1227
- SIOCDELRT — delete route, 1288
- SIOCGARP — get arp entry, 1194
- SIOCGHIWAT — get high water mark, 1304, 1325
- SIOCGIFADDR — get ifnet address, 1226
- SIOCGIFCONF — get ifnet list, 1226
- SIOCGIFDSTADDR — get p-p address, 1226
- SIOCGIFFLAGS — get ifnet flags, 1226
- SIOCGLOWAT — get low water mark, 1304, 1325
- SIOCSARP — set arp entry, 1194
- SIOCSHIWAT — set high water mark, 1304, 1325
- SIOCSIFADDR — set ifnet address, 1226
- SIOCSIFDSTADDR — set p-p address, 1226
- SIOCSIFFLAGS — set ifnet flags, 1226
- SIOCSLOWAT — set low water mark, 1304, 1325
- SIOCSMISC — toggle promiscuous mode, 1227
- size — find object file size, 465
- size mail command, 299
- skip backward magnetic tape files — mt, 333
- skip backward magnetic tape records — mt, 333
- skip forward magnetic tape files — mt, 333
- skip forward magnetic tape records — mt, 333
- sleep — suspend execution, 466
- sleep — suspend execution, 997, 1182
- sm, file, 1437
- SMD disk controller
 - xy — Xylogics 450, 1332 thru 1333
 - xy — Xylogics 451, 1332 thru 1333
 - xd — Xylogics 7053, 1329 thru 1330
- smoothing, interpolate curve — spline, 476
- snake — display chase game, 1541
- snap command, 467
- socket, 771
- socket I/O, see sockio(4), 1291
- socket operations
 - async_daemon, 718
 - bind, 636
 - connect, 647
 - getpeername, 679
 - getsockname, 686
 - getsockopt, 687
 - listen, 697
 - nfssvc, 718
 - recv, 739
 - recvfrom, 739
 - recvmsg, 739
 - send, 751
 - sendmsg, 751
 - sendto, 751
 - setsockopt, 687
 - shutdown, 762
 - socket, 771
 - socketpair, 773
- socket operations, accept connection
 - accept, 628
- socket options
 - get, 687
 - set, 687
- socketpair create connected socket pair, 773
- soelim — eliminate .so's from nroff input, 469
- interrupt system calls with software signal — siginterrupt, 994, 995, 998
- software signal — signal, System V, 1180
- sort bibliographic database — sortbib, 473
- sort — sort and collate lines, 470
- sort and collate lines — sort, 470
- sort quicker — qsort, 954
- sort topologically — tsort, 555
- sortbib — sort bibliographic database, 473
- sorted file
 - find lines in — look, 275
 - remove repeated lines — uniq, 561
- source code control system — sccs, 434
- source command, 107
- source mail command, 299
- space — specify plot space, 941
- spaces, to tabs — unexpand, 179
- sparc — machine type truth value, 292
- spawn process, 789
- special characters for equations — eqnchar, 1550
- special file
 - make, 702
 - make — mknod, 1674
- special files — makedev, 1668
- specify paging/swapping device — swapon, 778
- spell — check spelling, 474
- spellin — check spelling, 474
- spellout — check spelling, 474
- spheresdemo — graphics demo, 1521
- spline — interpolate smooth curve, 476

- `split` — split file into pieces, 477
- split into significand and exponent — `frexp`, 1087
- `spray` — spray packets, 1766
- `spray()` function, 1123
- `sprayd` — spray server, 1767
- `sprintf` — formatted output conversion, 944
- `sprintf` — format to string, System V, 1168
- `sqrt` — square root function, 1105
- `srand` — generate random numbers, 955
- `srand` — generate random numbers, System V, 1171
- `srandom` — generate random number, 956
- `sscanf` — convert from string, 984
- `sscanf` — convert from string, System V, 1172
- `st` — Sysgen SC 4000 (Archive) Tape Driver, 1292 *thru* 1293
- stand-alone utilities
 - `kadb` — kernel debugger, 1653
- standard I/O library functions, introduction to, 999, 1183
- standard output
 - copy to many files — `tee`, 517
- start output (like control-Q) `ioctl` — `TIOCSTART`, 1285
- start printer — `lpc`, 1662
- startup procedures — `boot`, 1586, 1650, 1727
- `stat` — obtain file attributes, 774
- `statd` — network status monitor, 1768
- state of terminal
 - `get` — `gtty`, 1009
 - `set` — `stty`, 1009
- `statfs` — obtain file system statistics, 776
- static file system information — `fstab`, 1396
- statistics
 - I/O — `iostat`, 1651
 - of file system — `fstatfs`, 776
 - of file system — `statfs`, 776
 - `profil`, 725
 - `rstatd` — kernel statistics server, 1750, 1751
- statistics of NFS, display — `nfsstat`, 1704
- status monitor files for network services, 1438
- status monitor protocol, 1122
- status of network
 - display — `netstat`, 1696
- status of printer — `lpc`, 1663
- status variable, 110
- `stdin`
 - get character — `getchar`, 861
 - get character, System V — `getchar`, 1162
 - get string from — `gets`, 885
 - input conversion — `scanf`, 984
 - input conversion, System V — `scanf`, 1172
- `stdout`
 - output conversion, System V — `printf`, 1168
 - put character to — `putchar`, 949
- sticky bit — `chmod`, 640
- sticky directory, 1769
- `STKTOP()` function, 1068
- stop command, 107
- stop output (like control-S) `ioctl` — `TIOCSTOP`, 1285
- stop printer — `lpc`, 1663
- stop processor, 738
- stop processor — `halt`, 1639
- stop until signal — `pause`, 939
- storage allocation, 925 *thru* 926
 - `alloca` — allocate on stack, 926
 - `calloc` — allocate memory, 925
 - `cfree` — free memory, 925
 - `free` — free memory, 925
 - `malloc` — allocate memory, 925
 - `malloc_debug` — set debug level, 926
 - `malloc_verify` — verify heap, 926
 - `memalign` — allocate aligned memory, 925
 - `realloc` — reallocate memory, 925
 - `valloc` — allocate aligned memory, 926
- storage management, 925 *thru* 926
- storage management debugging
- store datum under key — `store`, 830
- `store` — store datum under key, 830
- `strcat` — concatenate strings, 1001
- index — find character in string, 1001
- `strcmp` — compare strings, 1001
- `strcpy` — copy strings, 1001
- `strcat` — duplicate string, 1001
- stream
 - `fopen` — open stream, System V, 1160
 - assign buffering — `setbuf`, 987
 - assign buffering, System V — `setbuf`, 1175
 - assign buffering — `setbuffer`, 987
 - assign buffering, System V — `setbuffer`, 1175
 - assign buffering — `setlinebuf`, 987
 - assign buffering, System V — `setlinebuf`, 1175
 - assign buffering — `setvbuf`, 987
 - assign buffering, System V — `setvbuf`, 1175
 - associate descriptor — `fdopen`, 854
 - associate descriptor, System V — `fdopen`, 1160
 - `close` — `fclose`, 847
 - `flush` — `fflush`, 847
 - `fprintf` — format to stream, System V, 1168
 - get character — `fgetc`, 861
 - get character, System V — `fgetc`, 1162
 - get character — `getc`, 861
 - get character, System V — `getc`, 1162
 - get character — `getchar`, 861
 - get character, System V — `getchar`, 1162
 - get position of — `ftell`, 856
 - get string from — `fgets`, 885
 - get word — `getw`, 861
 - get word, System V — `getw`, 1162
 - input conversion — `scanf`, 984
 - input conversion, System V — `scanf`, 1172
 - `open` — `fopen`, 854
 - output conversion, System V — `printf`, 1168
 - `printf` — format to stdout, System V, 1168
 - push character back to — `ungetc`, 1028
 - put character to — `fputc`, 949
 - put character to — `putc`, 949
 - put string to — `puts`, 952
 - put string to — `fputs`, 952
 - put word to — `putw`, 949
 - read from stream — `fread`, 855
 - `reopen` — `freopen`, 854
 - `reopen`, System V — `freopen`, 1160
 - reposition — `rewind`, 856
 - return to remote command — `rcmd`, 958
 - return to remote command — `rexec`, 967
 - `rewind` — `rewind`, 856

- stream, *continued*
 - write to stream — `fwrite`, 855
 - seek — `fseek`, 856
 - `sprintf` — format to string, System V, 1168
- stream editor — `sed`, 446
- stream status enquiries
 - `clearerr` — clear error on stream, 849
 - `clearerr` — clear error on stream, System V, 1159
 - `feof` — enquire EOF on stream, 849
 - `feof` — enquire EOF on stream, System V, 1159
 - `ferror` — inquire error on stream, 849
 - `ferror` — inquire error on stream, System V, 1159
 - `fileno` — get stream descriptor number, 849
 - `fileno` — get stream descriptor number, System V, 1159
- stream, formatted output
 - `fprintf` — format to stream, System V, 1168
 - `printf` — format to stdout, System V, 1168
 - `sprintf` — format to string, System V, 1168
- streaming 1/4-inch tape drive — `ar`, 1193
- STREAMS
 - clone device driver, 1203
 - I/O, see `streamio(4)`, 1294
 - `ldterm` terminal module, 1252
 - NIT, Network Interface Tap, 1272
 - `nit_buf`, NIT buffering module, 1275
 - `nit_if`, NIT device interface, 1277
 - `nit_pf`, NIT packet filtering module, 1279
 - `ttcompat`, V7, BSD compatibility module, 1319
- string
 - number conversion — `printf`, 944, 984
 - number conversion, System V — `printf`, 1168, 1172
- string operations
 - compare — `strcmp`, 1001
 - compare — `strncmp`, 1001
 - concatenate — `strcat`, 1001
 - concatenate — `strncat`, 1001
 - copy — `strcpy`, 1001
 - copy — `strncpy`, 1001
 - get from stdin — `gets`, 885
 - get from stream — `fgets`, 885
 - index — `ndex`, 1001
 - put to stdout — `puts`, 952
 - put to stream — `fputs`, 952
 - reverse index — `rindex`, 1001
 - reverse index — `rindex`, 1001
- `string_to_decimal` — decimal record from character string, 1004
- strings
 - convert from numbers, 838
- `strings` — find printable strings in binary file, 478
- `strip` — strip symbols and relocation bits, 479
- strip filename affixes — `basename`, 42
- `strlen` — get length of string, 1001
- `strncat` — concatenate strings, 1001
- `strncmp` — compare strings, 1001
- `strncpy` — copy strings, 1001
- `rindex` — find character in string, 1001
- `strtod` — ASCII string to double, 1007
- `strtol` — ASCII string to long integer, 1008
- `stty` command, 480
- `stty` — set terminal state, 1009
- `stty_from_defaults` — set terminal from SunView defaults, 484
- `su` — substitute user id, 485
- substitute user id — `su`, 485
- sum — sum and count blocks in file, 486
- summarize file system quotas — `repquota`, 1729
- sun — machine type truth value, 292
- Sun 10 Mb/s Ethernet interface — `ie`, 1224 *thru* 1225
- Sun floppy disk driver — `fd`, 1218
- Sun keyboard device — `kbd`, 1251
- Sun mouse device — `mouse`, 1263
- Sun mouse streams module — `mouse`, 1264
- Sun SCSI disk driver — `si`, 1289 *thru* 1290
- Sun-3/50 10 Mb/s Ethernet interface — `le`, 1254 *thru* 1255
- `sun3cvt` — convert large Sun-2 executables to Sun-3, 367
- `suninstall` command, 1770
- SunView
 - `coloredit`, 79
 - `iconedit`, 228
 - start up environment, 487
- `sunview` — Suntools window environment, 487
- SunView environment, changing default settings — `defaultsedit`, 141
- SunWindows, graphics tool — `gfxtool`, 213
- super block, update — `sync`, 780
- super-user command — `su`, 485
- suspend command, 107
- suspend execution — `sleep`, 466
- suspend execution — `sleep`, 997, 1182
- suspend execution for interval in microseconds — `usleep`, 1029
- `swab` — swap bytes, 1010
- swap bytes — `swab`, 1010
- `swapon` — specify paging device, 778
- `swapon` — specify paging device, 1771
- swapping device — `swapon`, 778
- swapping devices, specify — `swapon`, 1771
- `swin` — set window input behavior, 496
- switch command, 108
- `switcher`, 499
- symbol table, get entries from — `nlist`, 937, 1167
- symbolic link
 - create, 779
 - read value of, 737
- symbolic link, make — `ln`, 265
- `symlink`, 779
- `symorder` — update symbol table ordering, 501
- `sync` — update super block, 780
- `sync` — update super block, 502
- synchronize file state — `fsync`, 662
- synchronous I/O multiplexing, 744
- `sys_errlist` — system error messages, 940
- `sys_nerr` — system error messages, 940
- `sys_siglist` — system signal messages, 948
- `syscall`, 781
- `sysdiag` — system diagnostics, 1772
- `sysex` command, 503
- Sysgen SC 4000 (Archive) Tape Driver — `st`, 1292 *thru* 1293
- `syslog` — make system log entry, 368
- `syslog` — write message to system log, 1011
- `syslogd.conf` — system log daemon configuration file, 1439

syslogd — system log message daemon, 1773
 Systech VPC-2200 interface — `vpc`, 1327
 system administration
 adduser — add new user account, 1577
 catman — create cat files for manual pages, 1593
 install — install files, 240
 system calls, introduction to, 613 *thru* 624
 system configuration files, build — `config`, 1600
 system data types — `types`, 1480
 system EEPROM display and load program, 1617
 system error messages
 errno — system error messages, 940
 perror — system error messages, 940
 sys_errlist — system error messages, 940
 sys_nerr — system error messages, 940
 system error numbers, introduction to, 613 *thru* 617
 system — issue shell command, 1013
 system log configuration file — `syslogd.conf`, 1439
 system log daemon — `syslog`, 1773
 system log, control — `syslog`, 1011
 system maintenance and operation, 1569
 system operation support
 mount, 706
 process accounting — `acct`, 630
 reboot, 738
 swapon — specify paging device, 778
 sync, 780
 vadvise, 788
 system page size, get — `getpagesize`, 678
 system PROM monitor program, 1679
 system resource consumption
 control — `vlimit`, 1034
 system signal messages
 psignal, 948
 sys_siglist, 948
 system special files — `makedev`, 1668
 system status server — `rwhod`, 1753
 system to system command execution — `uux`, 574
 system to system copy — `uucp`, 568
 System V commands
 banner, 36
 cat, 50
 cc, 52
 chmod, 65
 col, 77
 date, 124
 diff3, 152
 dircmp, 155
 du, 162
 echo, 163
 expr, 180
 grep, 218
 grpck, 1638
 lint, 261
 ls, 285
 m4, 288
 nohup, 342
 od, 348
 pg, 383
 pr, 388
 pwck, 1718

System V commands, *continued*

 sed, 446
 sort, 470
 sum, 486
 test, 522
 time, 532
 touch, 541
 tr, 544
 uname, 558
 System V library functions, introduction to, 1127
 System V library, system call versions
 getpgrp, 753
 open, 719
 setpgrp, 753
 uname, 784
 write, 794
 syswait — execute a command, 505

T

taac device, 1302
 tabs command, 505
 tabs, expand to spaces — `expand`, 179
 tabstop specifications in text files — `fspec`, 1394
 tail — display last part of file, 507
 talk — talk to another user, 508
 talkd — talk server, 1774
 tan — trigonometric tangent, 1106
 tanh — hyperbolic tangent, 1088
 tape
 backspace files — `mt`, 333
 backspace records — `mt`, 333
 copy, blocking preserved — `tcopy`, 515
 erase — `mt`, 333
 forward space files — `mt`, 333
 forward space records — `mt`, 333
 get unit status — `mt`, 333
 manipulate magnetic — `mt`, 333
 place unit off-line — `mt`, 333
 retension — `mt`, 333
 rewind — `mt`, 333
 scan — `tcopy`, 515
 skip backward files — `mt`, 333
 skip backward records — `mt`, 333
 skip forward files — `mt`, 333
 skip forward records — `mt`, 333
 write EOF mark on — `mt`, 333
 tape archives — `tar`, 509
 bar command, 37
 tape block size — 512 bytes, 1612
 tape drive, 1/2-inch
 tm — tapemaster, 1318
 xt — Xylogics 472, 1331
 tape drive, 1/4-inch
 ar — Archive 1/4-inch Streaming Tape Drive, 1193
 Sysgen SC 4000 (Archive) Tape Driver — `st`, 1292 *thru* 1293
 tape interface — `mtio`, 1268
 tape operation ioctl — `MTIOCTOP`, 1268
 tapemaster 1/2-inch tape drive — `tm`, 1318
 tar — tape archiver, 509
 tar — tape archive file format, 1442

- tbl — remove construction — deroff, 146
 - table formatter, 513
- tcov — code coverage tool, 516
- TCP ioctl's
 - SIOCGHIWAT — get high water mark, 1304
 - SIOCLOWAT — get low water mark, 1304
 - SIOCSHIWAT — set high water mark, 1304
 - SIOCSLOWAT — set low water mark, 1304
- tcp — Internet Transmission Control Protocol, 1303 *thru* 1304
- tdelete — delete binary tree node, 1021
- tee — copy standard output to many files, 517
- tektool — emulate Tektronix 4014, 369
- Tektronix 4014, emulate — tektool, 369
- tell, 698
- tellmdir — position of directory stream, 834
- telnet — TELNET interface, 518
- telnetd daemon, 1775
- temporary file
 - create name for — tmpnam, 1020
- term — terminal driving tables, 1444, 1450
- termcap — terminal capability data base, 1452
- terminal
 - configuration data base — gettytab, 1399
 - find name of — ttyname, 1024
 - get name of — tty, 556
 - I/O, see `termio(4)`, 1305
 - make script of session — `script`, 444
 - reset bits — `reset`, 551
 - set characteristics — `stty`, 480, 551
- terminal emulation, ANSI, 1204 *thru* 1208
- terminal emulator — `console`, 1204 *thru* 1209
- terminal independent operations
 - `tgetent`, 1014
 - `tgetflag`, 1014
 - `tgetnum`, 1014
 - `tgetstr`, 1014
 - `tgoto`, 1014
 - `tputs`, 1014
- alm — Sun ALM-2 Asynchronous Line Multiplexer, 1258
 - alm — Sun ALM-2 Asynchronous Line Multiplexer, 1259
- terminal state
 - `get` — `gtty`, 1009
 - `set` — `stty`, 1009
- terminal types — `ttytype`, 1479
- terminate process, 655, 844
- terminate program — `abort`, 810
- termination handler, name — `on_exit`, 938
- terminfo — System V terminal capability data base, 1460
- test command, 459, 522
- text editing
 - `ed` — line editor, 164
 - `edit` — line editor, 177
 - `ex` — line editor, 177
 - `sed` — stream editor, 446
 - `vi` — visual editor, 582
- text file
 - browse through — `pg`, 383
- text file, browse through
 - `more`, 330
 - page, 330
- text processing utilities
 - text processing utilities, *continued*
 - `awk` — scan and process patterns, 33
 - `cat` — concatenate files, 50
 - reverse lines in file — `rev`, 417
 - search for patterns — `grep`, 218
 - `sort` — sort and collate lines, 470
 - `spell` — check spelling, 474
 - `split` — split file into pieces, 477
 - `tail` — display last part of file, 507
 - `tr` — translate characters, 544
 - `tsort` — topological sort, 555
 - `ul` — underline text, 557
 - `uniq` — remove repeated lines, 561
 - `textedit` — SunView text editor, 524
 - `tfind` — search binary tree, 1021
 - `tftp` command, 530
 - `tftpd` daemon, 1776
 - `tgetent` — get entry for terminal, 1014
 - `tgetflag` — get Boolean capability, 1014
 - `tgetnum` — get numeric capability, 1014
 - `tgetstr` — get string capability, 1014
 - `tgoto` — go to position, 1014
 - `then` command, 453
 - `tic` command, 1778
 - tilde escapes in `mail`, 294 *thru* 295
 - time
 - `adjust` — `adjtime`, 631
 - display date and, 124
 - display in window, 71
 - time and date
 - `get` — `time`, 1016
 - `get` — `gettimeofday`, 689
 - `get` — `ftime`, 1016
 - `set` — `settimeofday`, 689
 - time and date conversion
 - `asctime`, 824
 - `ctime`, 824
 - `dysize`, 824
 - `gmtime`, 824
 - `localtime`, 824
 - `timegm`, 824
 - `timelocal`, 824
 - `tzset`, 824
 - `tzsetwall`, 825
 - time and date conversion, System V
 - `asctime`, 1132
 - `ctime`, 1132
 - `gmtime`, 1132
 - `localtime`, 1132
 - `timegm`, 1132
 - `timelocal`, 1132
 - `tzset`, 1132
 - `tzsetwall`, 1133
 - `time` command, 108, 532
 - `time` — get date and time, 1016
 - time variable, 110
 - `timed` — time server, 1781
 - timed event jobs table — `crontab`, 1381
 - timed event services
 - `at` — do job at specified time, 29
 - `calendar` — reminder service, 49

- timed event services, *continued*
 leave — remind you of leaving time, 257
 timed events — cron, 1606
 timegm — date and time conversion, 824
 timegm — date and time conversion, System V, 1132
 timelocal — date and time conversion, 824
 timelocal — date and time conversion, System V, 1132
 timerclear — macro, 674
 timercmp — macro, 674
 timerisset — macro, 674
 times command, 459
 times — get process times, 1017
 times — get process and child process times, System V, 1185
 timezone — get time zone name, 1018
 timing and statistics
 clock, 821
 getitimer, 674
 gettimeofday, 689
 profil, 725
 setitimer, 674
 settimeofday, 689
 timerclear — macro, 674
 timercmp — macro, 674
 timerisset — macro, 674
 TIOCCONS — get console I/O, 1204
 TIOCGTD — get line discipline, 1196
 TIOCPKT — set/clear packet mode (pty), 1285
 TIOCREMOTE — remote input editing, 1285
 TIOCSETD — set line discipline, 1196
 TIOCSTART — start output (like control-Q), 1285
 TIOCSTOP — stop output (like control-S), 1285
 tip — connect to remote system, 533
 tm — tapemaster 1/2-inch tape drive, 1318
 tmpfile — create temporary file, 1019
 tmpnam — make temporary file name, 1020
 tnamed — name server, 1782
 toascii — convert character to ASCII, System V, 1134
 toascii — convert character to ASCII, 826
 toc file, 1361
 toggle promiscuous mode ioctl — SIOCPROMISC, 1227
 tolower — convert character to lower-case, System V, 1134
 tolower — convert character to lower-case, 826
 _tolower — convert character to lower-case, System V, 1134
 toolplaces — show current window info, 539
 tools
 mailtool, 305
 textedit, 524
 top mail command, 299
 toplines mail variable, 303
 topological sort — tsort, 555
 touch — update last modified date of file, 541
 touch mail command, 300
 toupper — convert character to upper-case, System V, 1134
 toupper — convert character to upper-case, 826
 tput command, 542
 tputs — decode padding information, 1014
 tr — translate characters, 544
 trace command, 545
 trace process — ptrace, 726
 traffic — show Ethernet traffic, 547
 translate — input and output files for system message translation, 1363
 translate characters — tr, 544
 transliterate protocol trace — trpt, 1783
 trap command, 459
 trek — Star Trek game, 1543
 trigonometric functions, 1106
 acos, 1106
 asin, 1106
 atan, 1106
 atan2, 1106
 cos, 1106
 sin, 1106
 tan, 1106
 troff — typeset documents, 548
 troff utilities
 checknr — check nroff/troff files, 62
 col — filter reverse paper motions, 77
 troff utilities, 146
 soelim — eliminate .so's, incorporate sourced-in files, 469
 trpt — transliterate protocol trace, 1783
 true — provide truth values, 550
 truncate, 782
 trusted hosts list — hosts.equiv, 1406, 1413
 tsearch — build and search binary tree, 1021
 tset — set terminal characteristics, 551
 tsort — topological sort, 555
 ttcompat STREAMS module, 1319
 tty
 I/O, see termio(4), 1305
 tty — get terminal name, 556
 tty terminal interface, 1323
 tty, set characteristics — stty, 480
 tty, set characteristics — tset, 551
 ttyname — find terminal name, 1024
 ttyslot — get utmp slot number, 1025
 ttyslot — get utmp slot number, System V, 1186
 ttytab file, 1478
 ttytype — connected terminal types, 1479
 tunefs — tune file system, 1784
 twalk — traverse binary tree, 1021
 type command, 459
 type mail command, 298
 Type mail command, 299
 types — primitive system data types, 1480
 typeset documents — troff, 548
 tzfile file, 1483
 tzset — date and time conversion, 824
 tzset — date and time conversion, System V, 1132
 tzsetup command, 1785
 tzsetwall — date and time conversion, 825
 tzsetwall — date and time conversion, System V, 1133
- ## U
- u3b — machine type truth value, 292
 u3b15 — machine type truth value, 292
 u3b2 — machine type truth value, 292

- u3b5 — machine type truth value, 292
 - ualarm — schedule signal in microsecond precision, 1026
 - UDP *ioctl*'s
 - SIOCGHIWAT — get high water mark, 1325
 - SIOCGLOWAT — get low water mark, 1325
 - SIOCSEHIWAT — set high water mark, 1325
 - SIOCSLOWAT — set low water mark, 1325
 - udp — Internet User Datagram Protocol, 1324 *thru* 1325
 - ul — underline text, 557
 - ulimit — get and set user limits, 1027
 - umask, 783
 - umask command, 108, 459
 - umount — unmount file system, 1687
 - unalias command, 108
 - uname — print hostname, 558
 - uname — get system name, 784
 - uncompact — uncompress files, 350
 - uncompress — uncompress files, 83
 - unconfigure command, 1786
 - undeletemail command, 300
 - underline text — *ul*, 557
 - unexpand — spaces to tabs, 179
 - unget — unget SCCS file, 559
 - ungetc — push character back to stream, 1028
 - unhash command, 108
 - unifdef — eliminate *#ifdef*'s from C input, 560
 - uniq — remove repeated lines, 561
 - unique file name
 - create — *mktemp*, 929
 - units — convert units, 562
 - unix2dos — convert text file from DOS format to SunOS format, 563
 - unlimit command, 108
 - unlink — remove directory entry, 785, 1659
 - unload command, 564
 - unmap memory pages — *mmap*, 717
 - unmount — demount file system, 786
 - zero — source of zeroed unnamed memory, 1334
 - unpack — unpack files, 376
 - unreadmail command, 298
 - unset command, 108, 459
 - unsetmail command, 300
 - unsetenv command, 108
 - until command, 453
 - update — update super block, 1788
 - update last modified date of file — *touch*, 541
 - update programs — *make*, 311 *thru* 324, 355 *thru* 361
 - update super block — *sync*, 502
 - update super block — *sync*, 780
 - updaters file, 1484
 - uptime — display system up time, 566
 - user
 - display effective name — *logname*, 274, 598
 - talk to another — *talk*, 508
 - write to another — *write*, 600
 - user ID
 - chown* — change user ID of file, 1595
 - id* — display user and group IDs, 230
 - get*, 691
 - user ID, *continued*
 - set real and effective — *setreuid*, 755
 - substitute — *su*, 485
 - user limits
 - get* — *ulimit*, 1027
 - set* — *ulimit*, 1027
 - user mask, set — *umask*, 783
 - user name, get — *cuserid*, 829
 - user quotas
 - edquota* — edit user quotas, 1616
 - quotacheck* — check quota consistency, 1721
 - quotaoff* — turn file system quotas off, 1722
 - quotaon* — turn file system quotas on, 1722
 - repquota* — summarize quotas, 1729
 - rquotad* — remote quota server, 1747
 - users
 - info on users — *finger*, 186
 - list last logins — *last*, 248
 - what are they doing — *w*, 588
 - who — who is logged in, 597
 - write to all — *wall*, 590
 - users — display users on system, 567
 - usleep — suspend execution for interval in microseconds, 1029
 - utilities, introduction, 3
 - utime — set file times, 1030
 - utimes — set file times, 787
 - utmp — login records, 1485
 - uuclean — clean UUCP spool area, 1789
 - uucp — system to system copy, 568
 - UUCP log — *uulog*, 568
 - uudecode — decode binary file, 570
 - uuencode — encode binary file, 570
 - uuencode — UUCP encoded file format, 1487
 - uulog* — UUCP log, 568
 - uuname* — UUCP list of names, 568
 - uusend* — send file to remote host, 571
 - uustat* command, 572
 - uux* — system to system command execution, 574
- ## V
- va_arg* — next argument in variable list, 1032
 - va_dcl* — variable argument declarations, 1032
 - va_end* — finish variable argument list, 1032
 - va_list* — variable argument declarations, 1032
 - va_start* — initialize *varargs*, 1032
 - vacation* — automatic mail replies, 576
 - vadvise* — advise paging system, 788
 - val* — validate SCCS file, 578
 - validate SCCS file — *val*, 578
 - valloc* — allocate aligned memory, 926
 - values — machine-dependent values, 1031
 - varargs* — variable argument list, 1032
 - variable argument list, — *varargs*, 1032
 - variable substitution, in C shell, 100
 - variables
 - in Bourne shell, 454, 455
 - in C shell, 109
 - environment variables in *mail*
 - vax* — machine type truth value, 292
 - vc* command, 371

verbose mail variable, 303
 verbose variable, 110
 verifier, C programs — `lint`, 261
 verify heap — `malloc_verify`, 926
 plot graphics on — `vplot`, 584
 version mail command, 300
 version of file — `what`, 592
`vfont` — font formats, 1488
`vfontinfo` — examine font files, 579
`vfork`, 789
`vfprintf` — format and print variable argument list, 1035
`vfprintf` — format and print variable argument list, System V, 1187
`vgrind` — make formatted listings, 580
`vgrindefs` — `vgrind` language definitions, 1489
`vhangup`, 790
`vi` — visual editor, 582
`vipw` — edit password file, 1790
`virtual` — virtual address space, 1261 *thru* 1262
 visual editor — `vi`, 582
 visual mail command, 300
 VISUAL mail variable, 303
`vlimit` — control consumption, 1034
`vme16` — VMEbus 16-bit space, 1261 *thru* 1262
`vme16d16` — VMEbus address space, 1261 *thru* 1262
`vme16d32` — VMEbus address space, 1261 *thru* 1262
`vme24` — VMEbus 24-bit space, 1261 *thru* 1262
`vme24d16` — VMEbus address space, 1261 *thru* 1262
`vme24d32` — VMEbus address space, 1261 *thru* 1262
`vme32d16` — VMEbus address space, 1261 *thru* 1262
`vme32d32` — VMEbus address space, 1261 *thru* 1262
`vmstat` — display virtual memory statistics, 1791
`vp` — Ikon 10071-5 Versatec parallel printer interface, 1326
`vpc` — Systech VPC-2200 Versatec/Centronics interface, 1327
`vplot` — plot on Versatec, 584
`vprintf` — format and print variable argument list, 1035
`vprintf` — format and print variable argument list, System V, 1187
`vsprintf` — format and print variable argument list, 1035
`vsprintf` — format and print variable argument list, System V, 1187
`vswap` — convert foreign font files, 585
`vprintf` — log message with variable argument list, 1036
`vtimes` — resource use information, 1037
`vtroff` — format document for raster printer, 586
`vwidth` — make font width table, 587

W

`w` — what are users doing, 588
`wait`, 791
`wait` command, 108, 459, 589
`wait3`, 791
`wait4`, 791
`wall` — write to all users, 590
`wc` — count lines, words, characters in file, 591
 what are users doing — `w`, 588
`what` — identify file version, 592
`whatis` — describe command, 593

`whereis` — find program, 594
`which` — find program file, 596
`while` command, 108, 453
`while` — repeat commands — `cs`, 108
`who` — who is logged in, 597
 who is logged in on local network — `rusers`, 430, 432
`whoami` — display effective user name, 598
`whois` — Internet directory service, 599
`win` — Sun window system, 1328
 window environment — `sunview`, 487
 window management
 `adjacentscreens` command, 20
 `switcher` utility, 499
 window, save context — `lockscreen`, 269
 word
 get from stream — `getw`, 861
 get from stream, System V — `getw`, 1162
 put to stream — `putw`, 949
 words in file, count — `wc`, 591
 working directory
 `cd` — change directory, 57
 change, 639
 display name of — `pwd`, 404
 get pathname — `getwd`, 890
 worm — growing worm game, 1544
 worms — animate worms on display, 1545
`write`, 794
`write` — write to another user, 600
 write EOF mark on magnetic tape — `mt`, 333
 write formatted
 `fprintf` — convert to stream, System V, 1168
 `printf` — convert to stdout, System V, 1168
 `sprintf` — convert to string, System V, 1168
 write gathered — `writev`, 794
 write mail command, 300
 write to all users — `wall`, 590
 write to all users on network — `rwall`, 431
 write to stream — `fwrite`, 855
`wtmp` — login records, 1485
`wump` — hunt the Wumpus game, 1546

X

`xargs` — construct and use initial arguments lists, 602
`xcrypt` () function, 1124
`xd` — Xylogics SMD Disk driver, 1329 *thru* 1330
`xdecrypt` () function, 1124
`xdr` networking functions, 1038
`xget` — receive secret mail, 604
`xit` mail command, 297
`xsend` — send secret mail, 604
`xstr` — extract strings from C code, 605
`xt` — Xylogics 472 1/2-inch tape drive, 1331
`xtab` — exported file system table, 1389
`xtom` — hexadecimal string to multiple precision, 932
`xy` — Xylogics SMD Disk driver, 1332 *thru* 1333
 Xylogics 472 1/2-inch tape drive — `xt`, 1331
 Xylogics SMD Disk driver — `xd`, 1329 *thru* 1330, 1332 *thru* 1333

Y

- y0 — Bessel function, 1084
- y1 — Bessel function, 1084
- yacc — parser generator, 607
- YACC language tags file — `ctags`, 115
- Yellow Pages
 - change login password in — `yppasswd`, 611
 - make database — `ypinit`, 1793
 - make dbm file — `makedbm`, 1667
 - print values from database — `ypcat`, 609
 - rebuild database — `ypmake`, 1794
- Yellow Pages client interface, 1044
- yes — be repetitively affirmative, 608
- yn — Bessel function, 1084
- yp () function, 1125
- yp_all — Yellow Pages client interface, 1044
- yp_bind — Yellow Pages client interface, 1044
- yp_first — Yellow Pages client interface, 1044
- yp_get_default_domain — Yellow Pages client interface, 1044
- yp_master — Yellow Pages client interface, 1044
- yp_match — Yellow Pages client interface, 1044
- yp_next — Yellow Pages client interface, 1044
- yp_order — Yellow Pages client interface, 1044
- yp_unbind — Yellow Pages client interface, 1044
- yp_update () function, 1049
- ypcat — print values from YP database, 609
- yperr_string — Yellow Pages client interface, 1044
- ypfiles — Yellow Pages database and directory, 1491
- ypinit — make Yellow Pages database, 1793
- ypmake — rebuild Yellow Pages database, 1794
- ypmatch — match YP keys, 610
- yppasswd — update YP password entry, 1126
- yppasswd — change login password in Yellow Pages, 611
- yppoll — Yellow Pages version inquiry, 1796
- ypprot_err — Yellow Pages client interface, 1044
- yppush — force propagation of changed Yellow Pages map, 1797
- ypserv — Yellow Pages server process, 1798
- ypset — direct `ypbind` to a server, 1800
- ypupdated daemon, 1801
- ypwhich — who is Yellow Pages server, 612, 1802
- ypxfr — move remote Yellow Pages map to local host, 1803
- yppasswdd — Yellow Pages password server, 1795

Z

- z mail command, 300
- zcat — extract compressed files, 83
- zdump command, 1805
- zero byte strings — `bzero`, 819
- zic command, 1806
- zs — zilog 8530 SCC serial communications driver, 1335 *thru* 1336

Notes

Notes

Notes