




UNIX Interface Reference Manual



Credits and Trademarks

Sun Workstation® is a registered trademark of Sun Microsystems, Inc.

SunStation®, Sun Microsystems®, SunCore®, SunWindows®, DVMA®, and the combination of Sun with a numeric suffix are trademarks of Sun Microsystems, Inc.

UNIX, UNIX/32V, UNIX System III, and UNIX System V are trademarks of AT&T Bell Laboratories.

Intel® and Multibus® are registered trademarks of Intel Corporation.

DEC®, PDP®, VT®, and VAX® are registered trademarks of Digital Equipment Corporation.

Copyright © 1986 by Sun Microsystems.

This publication is protected by Federal Copyright Law, with all rights reserved. No part of this publication may be reproduced, stored in a retrieval system, translated, transcribed, or transmitted, in any form, or by any means manual, electric, electronic, electro-magnetic, mechanical, chemical, optical, or otherwise, without prior explicit written permission from Sun Microsystems.

NAME

intro – introduction to system calls and error numbers

SYNOPSIS

```
#include <errno.h>
```

DESCRIPTION

This section describes all of the system calls. A "(2V)" heading indicates that the system call performs differently when called from programs that use the System V libraries (programs compiled using `/usr/5bin/cc`). On these pages, both the regular behavior and the System V behavior is described.

Most of these calls have one or more error returns. An error condition is indicated by an otherwise impossible return value. This is almost always `-1`; the individual descriptions specify the details. Note that a number of system calls overload the meanings of these error numbers, and that the meanings must be interpreted according to the type and circumstances of the call.

As with normal arguments, all return codes and values from functions are of type integer unless otherwise noted. An error number is also made available in the external variable `errno`, which is not cleared on successful calls. Thus `errno` should be tested only after an error has occurred.

Each system call description attempts to list all possible error numbers. The following is a complete list of the errors and their names as given in `<errno.h>`.

- 0 Error 0
Unused.
- 1 EPERM Not owner
Typically this error indicates an attempt to modify a file in some way forbidden except to its owner or super-user. It is also returned for attempts by ordinary users to do things allowed only to the super-user.
- 2 ENOENT No such file or directory
This error occurs when a filename is specified and the file should exist but doesn't, or when one of the directories in a pathname does not exist.
- 3 ESRCH No such process
The process or process group whose number was given does not exist, or any such process is already dead.
- 4 EINTR Interrupted system call
An asynchronous signal (such as interrupt or quit), which the user has elected to catch, occurred during a system call. If execution is resumed after processing the signal, and the system call is not restarted, it will appear as if the interrupted system call returned this error condition.
- 5 EIO I/O error
Some physical I/O error occurred. This error may in some cases occur on a call following the one to which it actually applies.
- 6 ENXIO No such device or address
I/O on a special file refers to a subdevice which does not exist, or beyond the limits of the device. It may also occur when, for example, a tape drive is not on-line or a disk pack is not loaded on a drive.
- 7 E2BIG Arg list too long
An argument list longer than 10240 bytes is presented to `execve`.
- 8 ENOEXEC Exec format error
A request is made to execute a file which, although it has the appropriate permissions, does not start with a valid magic number (see `a.out(5)`).
- 9 EBADF Bad file number
Either a file descriptor refers to no open file, or a read (respectively, write) request is made to a file which is open only for writing (respectively, reading).

- 10 ECHILD No children
A *wait* was executed by a process that had no existing or unwaited-for child processes.
- 11 EAGAIN No more processes
A *fork* failed because the system's process table is full or the user is not allowed to create any more processes.
- 12 ENOMEM Not enough memory
During an *execve*, *brk*, or *sbrk*, a program asks for more address space or swap space than the system is able to supply, or a process size limit would be exceeded. A lack of swap space is normally a temporary condition; however, a lack of address space is not a temporary condition. The maximum size of the text, data, and stack segments is a system parameter. Soft limits may be increased to their corresponding hard limits.
- 13 EACCES Permission denied
An attempt was made to access a file in a way forbidden by the protection system.
- 14 EFAULT Bad address
The system encountered a hardware fault in attempting to access the arguments of a system call.
- 15 ENOTBLK Block device required
A file which is not a block device was mentioned where a block device was required, for example, in *mount*.
- 16 EBUSY Device busy
An attempt to mount a file system that was already mounted or an attempt was made to dismount a file system on which there is an active file (open file, current directory, mounted-on file, or active text segment).
- 17 EEXIST File exists
An existing file was mentioned in an inappropriate context, for example, *link*.
- 18 EXDEV Cross-device link
A hard link to a file on another file system was attempted.
- 19 ENODEV No such device
An attempt was made to apply an inappropriate system call to a device (for example, an attempt to read a write-only device) or an attempt was made to use a device not configured by the system.
- 20 ENOTDIR Not a directory
A non-directory was specified where a directory is required, for example, in a pathname or as an argument to *chdir*.
- 21 EISDIR Is a directory
An attempt was made to write on a directory.
- 22 EINVAL Invalid argument
A system call was made with an invalid argument; for example, dismounting a non-mounted file system, mentioning an unknown signal in *sigvec* or *kill*, reading or writing a file for which *lseek* has generated a negative pointer, or some other argument inappropriate for the call. Also set by math functions, see *intro(3)*.
- 23 ENFILE File table overflow
The system's table of open files is full, and temporarily no more *opens* can be accepted.
- 24 EMFILE Too many open files
A process tried to have more open files than the system allows a process to have. The customary configuration limit is 30 per process.
- 25 ENOTTY Inappropriate ioctl for device
The code used in an *ioctl* call is not supported by the object that the file descriptor in the call refers to.

- 26 ETXTBSY Text file busy
An attempt to execute a pure-procedure program which is currently open for writing. Also an attempt to open for writing a pure-procedure program that is being executed.
- 27 EFBIG File too large
The size of a file exceeded the maximum file size (1,082,201,088 bytes).
- 28 ENOSPC No space left on device
A *write* to an ordinary file, the creation of a directory or symbolic link, or the creation of a directory entry failed because no more disk blocks are available on the file system, or the allocation of an inode for a newly created file failed because no more inodes are available on the file system.
- 29 ESPIPE Illegal seek
An *lseek* was issued to a socket or pipe. This error may also be issued for other non-seekable devices.
- 30 EROFS Read-only file system
An attempt to modify a file or directory was made on a file system mounted read-only.
- 31 EMLINK Too many links
An attempt to make more than 32767 hard links to a file.
- 32 EPIPE Broken pipe
An attempt was made to write on a pipe or socket for which there is no process to read the data. This condition normally generates a signal; the error is returned if the signal is caught or ignored.
- 33 EDOM Math argument
The argument of a function in the math library (as described in section 3M) is out of the domain of the function.
- 34 ERANGE Result too large
The value of a function in the math library (as described in section 3M) is unrepresentable within machine precision.
- 35 EWOULDBLOCK Operation would block
An operation which would cause a process to block was attempted on an object in non-blocking mode (see *ioctl(2)*).
- 36 EINPROGRESS Operation now in progress
An operation which takes a long time to complete (such as a *connect(2)*) was attempted on a non-blocking object (see *ioctl(2)*).
- 37 EALREADY Operation already in progress
An operation was attempted on a non-blocking object which already had an operation in progress.
- 38 ENOTSOCK Socket operation on non-socket
Self-explanatory.
- 39 EDESTADDRREQ Destination address required
A required address was omitted from an operation on a socket.
- 40 EMSGSIZE Message too long
A message sent on a socket was larger than the internal message buffer.
- 41 EPROTOTYPE Protocol wrong type for socket
A protocol was specified which does not support the semantics of the socket type requested. For example, you cannot use the ARPA Internet UDP protocol with type `SOCK_STREAM`.
- 42 ENOPROTOPT Option not supported by protocol
A bad option was specified in a *getsockopt(2)* or *setsockopt(2)* call.
- 43 EPROTONOSUPPORT Protocol not supported
The protocol has not been configured into the system or no implementation for it exists.

- 44 ESOCKETNOSUPPORT Socket type not supported
The support for the socket type has not been configured into the system or no implementation for it exists.
- 45 EOPNOTSUPP Operation not supported on socket
For example, trying to *accept* a connection on a datagram socket.
- 46 EPFNOSUPPORT Protocol family not supported
The protocol family has not been configured into the system or no implementation for it exists.
- 47 EAFNOSUPPORT Address family not supported by protocol family
An address incompatible with the requested protocol was used. For example, you shouldn't necessarily expect to be able to use PUP Internet addresses with ARPA Internet protocols.
- 48 EADDRINUSE Address already in use
Only one usage of each address is normally permitted.
- 49 EADDRNOTAVAIL Can't assign requested address
Normally results from an attempt to create a socket with an address not on this machine.
- 50 ENETDOWN Network is down
A socket operation encountered a dead network.
- 51 ENETUNREACH Network is unreachable
A socket operation was attempted to an unreachable network.
- 52 ENETRESET Network dropped connection on reset
The host you were connected to crashed and rebooted.
- 53 ECONNABORTED Software caused connection abort
A connection abort was caused internal to your host machine.
- 54 ECONNRESET Connection reset by peer
A connection was forcibly closed by a peer. This normally results from the peer executing a *shutdown(2)* call.
- 55 ENOBUFS No buffer space available
An operation on a socket or pipe was not performed because the system lacked sufficient buffer space.
- 56 EISCONN Socket is already connected
A *connect* request was made on an already connected socket; or, a *sendto* or *sendmsg* request on a connected socket specified a destination other than the connected party.
- 57 ENOTCONN Socket is not connected
An request to send or receive data was disallowed because the socket is not connected.
- 58 ESHUTDOWN Can't send after socket shutdown
A request to send data was disallowed because the socket had already been shut down with a previous *shutdown(2)* call.
- 59 *unused*
- 60 ETIMEDOUT Connection timed out
A *connect* request failed because the connected party did not properly respond after a period of time. (The timeout period is dependent on the communication protocol.)
- 61 ECONNREFUSED Connection refused
No connection could be made because the target machine actively refused it. This usually results from trying to connect to a service which is inactive on the foreign host.
- 62 ELOOP Too many levels of symbolic links
A pathname lookup involved more than 8 symbolic links.

- 63 ENAMETOOLONG File name too long
A component of a pathname exceeded 255 characters, or an entire pathname exceeded 1023 characters.
- 64 EHOSTDOWN Host is down
A socket operation failed because the destination host was down.
- 65 EHOSTUNREACH Host is unreachable
A socket operation was attempted to an unreachable host.
- 66 ENOTEMPTY Directory not empty
An attempt was made to remove a directory with entries other than . and .. by performing a *rmdir* system call or a *rename* system call with that directory specified as the target directory.
- 67 *unused*
- 68 *unused*
- 69 EDQUOT Disc quota exceeded
A *write* to an ordinary file, the creation of a directory or symbolic link, or the creation of a directory entry failed because the user's quota of disk blocks was exhausted, or the allocation of an inode for a newly created file failed because the user's quota of inodes was exhausted.
- 70 ESTALE Stale NFS file handle
A client referenced an open file, when the file has been deleted.
- 71 EREMOTE Too many levels of remote in path
An attempt was made to remotely mount a file system into a path which already has a remotely mounted component.
- 72 *unused*
- 73 *unused*
- 74 *unused*
- 75 ENOMSG No message of desired type
An attempt was made to receive a message of a type that does not exist on the specified message queue; see *msgop(2)*.
- 76 *unused*
- 77 EIDRM Identifier removed
This error is returned to processes that resume execution due to the removal of an identifier from the IPC system's name space (see *msgctl(2)*, *semctl(2)*, and *shmctl(2)*).

DEFINITIONS

Descriptor

An integer assigned by the system when a file is referenced by *open(2V)*, *dup(2)*, or *pipe(2)* or a socket is referenced by *socket(2)* or *socketpair(2)* which uniquely identifies an access path to that file or socket from a given process or any of its children.

Directory

A directory is a special type of file which contains entries which are references to other files. Directory entries are called links. By convention, a directory contains at least two links, . and .., referred to as *dot* and *dot-dot* respectively. Dot refers to the directory itself and dot-dot refers to its parent directory.

Effective User ID, Effective Group ID, and Access Groups

Access to system resources is governed by three values: the effective user ID, the effective group ID, and the group access list.

The effective user ID and effective group ID are initially the process's real user ID and real group ID respectively. Either may be modified through execution of a set-user-ID or set-group-ID file (possibly by one of its ancestors) (see *execve(2)*).

The group access list is an additional set of group ID's used only in determining resource accessibility. Access checks are performed as described below in "File Access Permissions".

File Access Permissions

Every file in the file system has a set of access permissions. These permissions are used in determining whether a process may perform a requested operation on the file (such as opening a file for writing). Access permissions are established at the time a file is created. They may be changed at some later time through the *chmod*(2) call.

File access is broken down according to whether a file may be: read, written, or executed. Directory files use the execute permission to control if the directory may be searched.

File access permissions are interpreted by the system as they apply to three different classes of users: the owner of the file, those users in the file's group, anyone else. Every file has an independent set of access permissions for each of these classes. When an access check is made, the system decides if permission should be granted by checking the access information applicable to the caller.

Read, write, and execute/search permissions on a file are granted to a process if:

The process's effective user ID is that of the super-user.

The process's effective user ID matches the user ID of the owner of the file and the owner permissions allow the access.

The process's effective user ID does not match the user ID of the owner of the file, and either the process's effective group ID matches the group ID of the file, or the group ID of the file is in the process's group access list, and the group permissions allow the access.

Neither the effective user ID nor effective group ID and group access list of the process match the corresponding user ID and group ID of the file, but the permissions for "other users" allow access.

Otherwise, permission is denied.

File Name

Names consisting of up to 255 characters may be used to name an ordinary file, special file, or directory.

These characters may be selected from the set of all ASCII character excluding \0 (null) and the ASCII code for / (slash). (The parity bit, bit 8, must be 0.)

Note that it is generally unwise to use *, ?, [, or] as part of filenames because of the special meaning attached to these characters by the shell. See *sh*(1). Although permitted, it is advisable to avoid the use of unprintable characters in filenames.

Message Queue Identifier

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

```

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

```

msg_perm is an *ipc_perm* structure that specifies the message operation permission (see below). This structure includes the following members:


```

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

```

`msg_qnum` is the number of messages currently on the queue. `msg_qbytes` is the maximum number of bytes allowed on the queue. `msg_lspid` is the process id of the last process that performed a `msgsnd` operation. `msg_lrpid` is the process id of the last process that performed a `msgrcv` operation. `msg_stime` is the time of the last `msgsnd` operation, `msg_rtime` is the time of the last `msgrcv` operation, and `msg_ctime` is the time of the last `msgctl(2)` operation that changed a member of the above structure.

Message Operation Permissions

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

```

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

```

Read and Write permissions on a `msqid` are granted to a process if one or more of the following are true:

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

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

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

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

Otherwise, the corresponding permissions are denied.

Parent Process ID

A new process is created by a currently active process (see `fork(2)`). The parent process ID of a process is the process ID of its creator.

Path Name and Path Prefix

A pathname is a null-terminated character string starting with an optional slash (/), followed by zero or more directory names separated by slashes, optionally followed by a filename. The total length of a pathname must be less than {MAXPATHLEN} (1024) characters.

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

```

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

```

where <file-name> is a string of 1 to 255 characters other than the ASCII slash and null, and <dirname> is a string of 1 to 255 characters (other than the ASCII slash and null) that names a directory.

If a pathname begins with a slash, the search begins at the *root* directory. Otherwise, the search begins at the current working directory.

A slash, by itself, names the root directory. A dot (.) names the current working directory.

A null pathname also refers to the current directory. However, this is not true of all UNIX systems. (On such systems, accidental use of a null pathname in routines that don't check for it may corrupt the current working directory.) For portable code, specify the current directory explicitly using ".", rather than "".

Process Group ID

Each active process is a member of a process group that is identified by a positive integer called the process group ID. This is the process ID of the group leader. This grouping permits the signaling of related processes (see *killpg(2)*) and the job control mechanisms of *cs(1)*.

Process ID

Each active process in the system is uniquely identified by a positive integer called a process ID. The range of this ID is from 0 to 30000.

Real User ID and Real Group ID

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

Each user is also a member of one or more groups. One of these groups is distinguished from others and used in implementing accounting facilities. The positive integer corresponding to this distinguished group is termed the real group ID.

All processes have a real user ID and real group ID. These are initialized from the equivalent attributes of the process which created it.

Root Directory and Current Working Directory

Each process has associated with it a concept of a root directory and a current working directory for the purpose of resolving path name searches. A process's root directory need not be the root directory of the root file system.

Semaphore Identifier

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

```

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

```

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

```

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

```

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

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

```

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

```

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

Semaphore Operation Permissions

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

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

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

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

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

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

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

Otherwise, the corresponding permissions are denied.

Shared Memory Identifier

A shared memory identifier (*shmid*) is a unique positive integer created by a *shmget(2)* system call. Each *shmid* has a segment of memory (referred to as a shared memory segment) and a data structure associated with it. The data structure is referred to as *shmid_ds* and contains the following members:

```

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

```

`shm_perm` is an *ipc_perm* structure that specifies the shared memory operation permission (see below). This structure includes the following members:

```

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

```

`shm_segsz` specifies the size of the shared memory segment. `shm_cpid` is the process id of the process that created the shared memory identifier. `shm_lpid` is the process id of the last process that performed a *shmop(2)* operation. `shm_nattch` is the number of processes that currently have this segment attached. `shm_atime` is the time of the last *shmat* operation, `shm_dtime` is the time of the last *shmdt* operation, and `shm_ctime` is the time of the last *shmctl(2)* operation that changed one of the members of the above structure.

Shared Memory Operation Permissions

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

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

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

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

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

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

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

Otherwise, the corresponding permissions are denied.

Sockets and Address Families

A socket is an endpoint for communication between processes. Each socket has queues for sending and receiving data.

Sockets are typed according to their communications properties. These properties include whether messages sent and received at a socket require the name of the partner, whether communication is reliable, the format used in naming message recipients, etc.

Each instance of the system supports some collection of socket types; consult *socket(2)* for more information about the types available and their properties.

Each instance of the system supports some number of sets of communications protocols. Each protocol set supports addresses of a certain format. An Address Family is the set of addresses for a specific group of protocols. Each socket has an address chosen from the address family in which the socket was created.

Special Processes

The processes with a process ID's of 0, 1, and 2 are special. Process 0 is the scheduler. Process 1 is the initialization process *init*, and is the ancestor of every other process in the system. It is used to control the process structure. Process 2 is the paging daemon.

Super-user

A process is recognized as a *super-user* process and is granted special privileges if its effective user ID is 0.

Tty Group ID

Each active process can be a member of a terminal group that is identified by a positive integer called the tty group ID. This grouping is used to arbitrate between multiple jobs contending for the same terminal (see *csh(1)*, and *ty(4)*).

SEE ALSO

intro(3), *perror(3)*

LIST OF SYSTEM CALLS

<i>Name</i>	<i>Appears on Page</i>	<i>Description</i>
<i>_exit</i>	<i>exit(2)</i>	terminate a process
<i>accept</i>	<i>accept(2)</i>	accept a connection on a socket
<i>access</i>	<i>access(2)</i>	determine accessibility of file
<i>acct</i>	<i>acct(2)</i>	turn accounting on or off
<i>adjtime</i>	<i>adjtime(2)</i>	correct the time to allow synchronization of the system
<i>async_daemon</i>	<i>nfssvc(2)</i>	NFS daemons

bind	bind(2)	bind a name to a socket
brk	brk(2)	change data segment size
chdir	chdir(2)	change current working directory
chmod	chmod(2)	change mode of file
chown	chown(2)	change owner and group of a file
chroot	chroot(2)	change root directory
close	close(2)	delete a descriptor
connect	connect(2)	initiate a connection on a socket
creat	creat(2)	create a new file
dup	dup(2)	duplicate a descriptor
dup2	dup(2)	duplicate a descriptor
execve	execve(2)	execute a file
fchmod	chmod(2)	change mode of file
fchown	chown(2)	change owner and group of a file
fcntl	fcntl(2)	file control
flock	flock(2)	apply or remove an advisory lock on an open file
fork	fork(2)	create a new process
fstat	stat(2)	get file status
fsync	fsync(2)	synchronize a file's in-core state with that on disk
ftruncate	truncate(2)	truncate a file to a specified length
getdirenties	getdirenties(2)	gets directory entries in a filesystem independent format
getdomainname	getdomainname(2)	get name of current domain
getdtablesize	getdtablesize(2)	get descriptor table size
getegid	getgid(2)	get group identity
geteuid	getuid(2)	get effective user identity
getgid	getgid(2)	get group identity
getgroups	getgroups(2)	get group access list
gethostid	gethostid(2)	get unique identifier of current host
gethostname	gethostname(2)	get name of current host
getitimer	getitimer(2)	get value of interval timer
getpagesize	getpagesize(2)	get system page size
getpeername	getpeername(2)	get name of connected peer
getpgrp	setpgrp(2V)	set and/or return the process group of a process
getpid	getpid(2)	get parent process identification
getppid	getpid(2)	get process identification
getpriority	getpriority(2)	get program scheduling priority
getrlimit	getrlimit(2)	control maximum system resource consumption
getrusage	getrusage(2)	get information about resource utilization
getsockname	getsockname(2)	get socket name
getsockopt	getsockopt(2)	get options on sockets
gettimeofday	gettimeofday(2)	get date and time
getuid	getuid(2)	get user identity
ioctl	ioctl(2)	control device
kill	kill(2)	send signal to a process
killpg	killpg(2)	send signal to a process group
link	link(2)	make a hard link to a file
listen	listen(2)	listen for connections on a socket
lseek	lseek(2)	move read/write pointer
lstat	stat(2)	get file status
mkdir	mkdir(2)	make a directory file
mknod	mknod(2)	make a special file
mmap	mmap(2)	map or unmap pages of memory
mount	mount(2)	mount file system

msgctl	msgctl(2)	message control operations
msgget	msgget(2)	get message queue
msgop	msgop(2)	message operations
msgrcv	msgop(2)	message operations
msgsnd	msgop(2)	message operations
munmap	munmap(2)	map or unmap pages of memory
nfssvc	nfssvc(2)	NFS daemons
open	open(2V)	open or create a file for reading or writing
pipe	pipe(2)	create an interprocess communication channel
profil	profil(2)	execution time profile
ptrace	ptrace(2)	process trace
quotactl	quotactl(2)	manipulate disk quotas
read	read(2V)	read input
readlink	readlink(2)	read value of a symbolic link
readv	read(2V)	read input
reboot	reboot(2)	reboot system or halt processor
recv	recv(2)	receive a message from a socket
recvfrom	recv(2)	receive a message from a socket
recvmsg	recv(2)	receive a message from a socket
rename	rename(2)	change the name of a file
rmdir	rmdir(2)	remove a directory file
sbrk	brk(2)	change data segment size
select	select(2)	synchronous I/O multiplexing
semctl	semctl(2)	semaphore control operations
semget	semget(2)	get set of semaphores
semop	semop(2)	semaphore operations
send	send(2)	send a message from a socket
sendmsg	send(2)	send a message from a socket
sendto	send(2)	send a message from a socket
setdomainname	getdomainname(2)	set name of current domain
setgroups	getgroups(2)	set group access list
sethostname	gethostname(2)	set name of current host
setitimer	getitimer(2)	set value of interval timer
setpgrp	setpgrp(2V)	set and/or return the process group of a process
setpriority	getpriority(2)	set program scheduling priority
setregid	setregid(2)	set real and effective group IDs
setreuid	setreuid(2)	set real and effective user IDs
setrlimit	getrlimit(2)	control maximum system resource consumption
setsockopt	getsockopt(2)	set options on sockets
settimeofday	gettimeofday(2)	set date and time
shmat	shmop(2)	shared memory operations
shmctl	shmctl(2)	shared memory control operations
shmdt	shmop(2)	shared memory operations
shmget	shmget(2)	get shared memory segment
shmop	shmop(2)	shared memory operations
shutdown	shutdown(2)	shut down part of a full-duplex connection
sigblock	sigblock(2)	block signals
sigpause	sigpause(2)	atomically release blocked signals and wait for interrupt
sigsetmask	sigsetmask(2)	set current signal mask
sigstack	sigstack(2)	set and/or get signal stack context
sigvec	sigvec(2)	software signal facilities
socket	socket(2)	create an endpoint for communication
socketpair	socketpair(2)	create a pair of connected sockets

stat	stat(2)	get file status
statfs	statfs(2)	get file system statistics
swapon	swapon(2)	add a swap device for interleaved paging/swapping
symlink	symlink(2)	make symbolic link to a file
sync	sync(2)	update super-block
syscall	syscall(2)	indirect system call
tell	lseek(2)	locate read/write pointer
truncate	truncate(2)	truncate a file to a specified length
umask	umask(2)	set file creation mode mask
uname	uname(2V)	get name of current UNIX system
unlink	unlink(2)	remove directory entry
unmount	umount(2)	remove a file system
utimes	utimes(2)	set file times
vadvise	vadvise(2)	give advice to paging system
vfork	vfork(2)	spawn new process in a virtual memory efficient way
vhangup	vhangup(2)	virtually "hangup" the current control terminal
wait	wait(2)	wait for process to terminate or stop
wait3	wait(2)	wait for process to terminate or stop
write	write(2V)	write output
writev	write(2V)	write output

NAME

`accept` – accept a connection on a socket

SYNOPSIS

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

ns = accept(s, addr, addrlen)
int ns, s;
struct sockaddr *addr;
int *addrlen;
```

DESCRIPTION

The argument *s* is a socket that has been created with *socket(2)*, bound to an address with *bind(2)*, and is listening for connections after a *listen(2)*. *Accept* extracts the first connection on the queue of pending connections, creates a new socket with the same properties of *s* and allocates a new file descriptor, *ns*, for the socket. If no pending connections are present on the queue, and the socket is not marked as non-blocking, *accept* blocks the caller until a connection is present. If the socket is marked non-blocking and no pending connections are present on the queue, *accept* returns an error as described below. The accepted socket, *ns*, is used to read and write data to and from the socket which connected to this one; it is not used to accept more connections. The original socket *s* remains open for accepting further connections.

The argument *addr* is a result parameter that is filled in with the address of the connecting entity, as known to the communications layer. The exact format of the *addr* parameter is determined by the domain in which the communication is occurring. The *addrlen* is a value-result parameter; it should initially contain the amount of space pointed to by *addr*; on return it will contain the actual length (in bytes) of the address returned. This call is used with connection-based socket types, currently with `SOCK_STREAM`.

It is possible to *select(2)* a socket for the purposes of doing an *accept* by selecting it for read.

RETURN VALUE

The call returns `-1` on error. If it succeeds, it returns a non-negative integer that is a descriptor for the accepted socket.

ERRORS

The *accept* will fail if:

<code>EBADF</code>	The descriptor is invalid.
<code>ENOTSOCK</code>	The descriptor references a file, not a socket.
<code>EOPNOTSUPP</code>	The referenced socket is not of type <code>SOCK_STREAM</code> .
<code>EFAULT</code>	The <i>addr</i> parameter is not in a writable part of the user address space.
<code>EWOULDBLOCK</code>	The socket is marked non-blocking and no connections are present to be accepted.

SEE ALSO

bind(2), *connect(2)*, *listen(2)*, *select(2)*, *socket(2)*

NAME

`access` – determine accessibility of file

SYNOPSIS

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

#define R_OK      4  /* test for read permission */
#define W_OK      2  /* test for write permission */
#define X_OK      1  /* test for execute (search) permission */
#define F_OK      0  /* test for presence of file */

accessible = access(path, mode)
int accessible;
char *path;
int mode;
```

DESCRIPTION

path points to a path name naming a file. *access* checks the named file for accessibility according to *mode*, which is an inclusive or of the bits `R_OK`, `W_OK` and `X_OK`. Specifying *mode* as `F_OK` (that is, 0) tests whether the directories leading to the file can be searched and the file exists.

The real user ID and the group access list (including the real group ID) are used in verifying permission, so this call is useful to set-UID programs.

The owner of a file has permission checked with respect to the *owner* read, write, and execute mode bits, members of the file's group other than the owner have permission checked with respect to the *group* mode bits, and all others have permissions checked with respect to the *other* mode bits.

Notice that only access bits are checked. A directory may be indicated as writable by *access*, but an attempt to open it for writing will fail (although files may be created there); a file may look executable, but *execve* will fail unless it is in proper format.

RETURN VALUE

If *path* cannot be found or if any of the desired access modes would not be granted, then a `-1` value is returned; otherwise a `0` value is returned.

ERRORS

Access to the file is denied if one or more of the following are true:

<code>ENOTDIR</code>	A component of the path prefix of <i>path</i> is not a directory.
<code>EINVAL</code>	<i>path</i> contains a byte with the high-order bit set.
<code>ENAMETOOLONG</code>	The length of a component of <i>path</i> exceeds 255 characters, or the length of <i>path</i> exceeds 1023 characters.
<code>ENOENT</code>	The file named by <i>path</i> does not exist.
<code>EACCES</code>	Search permission is denied for a component of the path prefix of <i>path</i> .
<code>ELOOP</code>	Too many symbolic links were encountered in translating <i>path</i> .
<code>EROFS</code>	The file named by <i>path</i> is on a read-only file system and write access was requested.
<code>ETXTBSY</code>	The file named by <i>path</i> is a pure procedure (shared text) file that is being executed and write access was requested.
<code>EACCES</code>	Permission bits of the file mode do not permit the requested access to the file named by <i>path</i> .
<code>EFAULT</code>	<i>path</i> points outside the process's allocated address space.
<code>EIO</code>	An I/O error occurred while reading from or writing to the file system.

SEE ALSO

chmod(2), stat(2)

NAME

acct – turn accounting on or off

SYNOPSIS

```
acct(file)
char *file;
```

DESCRIPTION

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

name points to a path name naming the accounting file. The accounting file format is given in *acct(5)*.

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

If accounting is already turned on, and a successful *acct* call is made with a non-zero *name*, all subsequent accounting records will be written to the new accounting file.

NOTES

Accounting is automatically disabled when the file system the accounting file resides on runs out of space; it is enabled when space once again becomes available.

RETURN VALUE

The value -1 is returned if an error occurs, and external variable *errno* is set to indicate the cause of the error. Otherwise the value 0 is returned.

ERRORS

acct will fail if one of the following is true:

EPERM	The caller is not the super-user.
ENOTDIR	A component of the path prefix of <i>file</i> is not a directory.
EINVAL	<i>file</i> contains a character with the high-order bit set.
EINVAL	Support for accounting was not configured into the system.
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 named file does not exist.
EACCES	Search permission is denied for a component of the path prefix of <i>file</i> .
EACCES	The file referred to by <i>file</i> is not a regular file.
ELOOP	Too many symbolic links were encountered in translating the path name.
EROFS	The named file resides on a read-only file system.
EFAULT	<i>file</i> points outside the process's allocated address space.
EIO	An I/O error occurred while reading from or writing to the file system.

SEE ALSO

acct(5), *sa(8)*

BUGS

No accounting is produced for programs running when a crash occurs. In particular non-terminating programs are never accounted for.

NAME

`adjtime` – correct the time to allow synchronization of the system clock

SYNOPSIS

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

adjtime(delta, olddelta)
struct timeval *delta;
struct timeval *olddelta;
```

DESCRIPTION

adjtime adjusts the system's notion of the current time, as returned by *gettimeofday(2)*, advancing or retarding it by the amount of time specified in the **struct timeval** **delta*.

The adjustment is effected by speeding up (if **delta* is positive) or slowing down (if **delta* is negative) the system's clock by a fixed percentage, currently 10%. Thus, the time is always a monotonically increasing function. A time correction from an earlier call to *adjtime* may not be finished when *adjtime* is called again. If *olddelta* is non-zero, then the structure pointed to will contain, upon return, the number of microseconds still to be corrected from the earlier call.

The structures pointed to by *delta* and *olddelta* are defined in *<sys/time.h>* as:

```
struct timeval {
    u_long tv_sec;          /* seconds since Jan. 1, 1970 */
    long tv_usec;         /* and microseconds */
};
```

If *olddelta* is a NULL pointer, the corresponding information will not be returned.

This call may be used in time servers that synchronize the clocks of computers in a local area network. Such time servers would slow down the clocks of some machines and speed up the clocks of others to bring them to the average network time.

Only the super-user may adjust the time of day.

The adjustment value will be silently rounded to the resolution of the system clock.

RETURN

A 0 return value indicates that the call succeeded. A -1 return value indicates an error occurred, and in this case an error code is stored into the global variable *errno*.

ERRORS

The following error codes may be set in *errno*:

EFAULT	<i>delta</i> or <i>olddelta</i> points outside the process's allocated address space, or <i>olddelta</i> points to a region of the process' allocated address space which is not writable.
EPERM	The process's effective user ID is not that of the super-user.

SEE ALSO

`settimeofday(2)`, `date(1)`

NAME

`bind` – bind a name to a socket

SYNOPSIS

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

bind(s, name, namelen)
int s;
struct sockaddr *name;
int namelen;
```

DESCRIPTION

`bind` assigns a name to an unnamed socket. When a socket is created with `socket(2)` it exists in a name space (address family) but has no name assigned. `bind` requests that the name pointed to by `name` be assigned to the socket.

NOTES

Binding a name in the UNIX domain creates a socket in the file system that must be deleted by the caller when it is no longer needed (using `unlink(2)`).

The rules used in name binding vary between communication domains. Consult the manual entries in section 4 for detailed information.

RETURN VALUE

If the `bind` is successful, a 0 value is returned. A return value of `-1` indicates an error, which is further specified in the global `errno`.

ERRORS

The `bind` call will fail if:

EBADF	<code>S</code> is not a valid descriptor.
ENOTSOCK	<code>S</code> is not a socket.
EADDRNOTAVAIL	The specified address is not available from the local machine.
EADDRINUSE	The specified address is already in use.
EINVAL	The socket is already bound to an address.
EACCES	The requested address is protected, and the current user has inadequate permission to access it.
EFAULT	The <code>name</code> parameter is not in a valid part of the user address space.

The following errors are specific to binding names in the UNIX domain.

ENOTDIR	A component of the path prefix of the path name in <code>name</code> is not a directory.
EINVAL	The path name in <code>name</code> contains a character with the high-order bit set.
ENAMETOOLONG	The length of a component of the path name in <code>name</code> exceeds 255 characters, or the length of the path name in <code>name</code> exceeds 1023 characters.
ENOENT	A component of the path prefix of the path name in <code>name</code> does not exist.
EACCES	Search permission is denied for a component of the path prefix of the path name in <code>name</code> .
ELOOP	Too many symbolic links were encountered in translating the path name in <code>name</code> .
EIO	An I/O error occurred while making the directory entry or allocating the inode.
EROFS	The inode would reside on a read-only file system.
EISDIR	A null path name was specified.

SEE ALSO

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

NAME

brk, *sbrk* – change data segment size

SYNOPSIS

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

caddr_t brk(addr)
caddr_t addr;

caddr_t sbrk(incr)
int incr;
```

DESCRIPTION**Brk**

brk sets the system's idea of the lowest data segment location not used by the program (called the break) to *addr* (rounded up to the next multiple of the system's page size). Locations greater than *addr* and below the stack pointer are not in the address space and will thus cause a memory violation if accessed.

Sbrk

In the alternate function *sbrk*, *incr* more bytes are added to the program's data space and a pointer to the start of the new area is returned.

When a program begins execution via *execve* the break is set at the highest location defined by the program and data storage areas. Ordinarily, therefore, only programs with growing data areas need to use *sbrk*.

The *getrlimit(2)* system call may be used to determine the maximum permissible size of the *data* segment; it will not be possible to set the break beyond the *rlim_max* value returned from a call to *getrlimit*, e.g. "etext + rlp → rlim_max." (See *end(3)* for the definition of *etext*.)

RETURN VALUE

Zero is returned if the *brk* could be set; -1 if the program requests more memory than the system limit. *Sbrk* normally returns the current value of the break, but -1 if it could not be set.

ERRORS

Sbrk will fail and no additional memory will be allocated if one of the following are true:

- | | |
|--------|--|
| ENOMEM | The limit, as set by <i>setrlimit(2)</i> , was exceeded. |
| ENOMEM | The maximum possible size of a data segment (compiled into the system) was exceeded. |
| ENOMEM | Insufficient space existed in the swap area to support the expansion. |

SEE ALSO

execve(2), *getrlimit(2)*, *malloc(3)*, *end(3)*

BUGS

Setting the break may fail due to a temporary lack of swap space. It is not possible to distinguish this from a failure caused by exceeding the maximum size of the data segment without consulting *getrlimit*.

NAME

`chdir` – change current working directory

SYNOPSIS

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

DESCRIPTION

path points to the path name of a directory. *chdir* causes this directory to become the current working directory, the starting point for path names not beginning with */*.

In order for a directory to become the current directory, a process must have execute (search) access to the directory.

RETURN VALUE

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

ERRORS

chdir will fail and the current working directory will be unchanged if one or more of the following are true:

ENOTDIR A component of the path prefix of *path* is not a directory.

ENOTDIR The file named by *path* is not a directory.

EINVAL *path* contains a byte with the high-order bit set.

ENAMETOOLONG

The length of a component of *path* exceeds 255 characters, or the length of *path* exceeds 1023 characters.

ENOENT The directory referred to by *path* does not exist.

ELOOP Too many symbolic links were encountered in translating *path*.

EACCES Search permission is denied for a component of the path prefix of *path*.

EACCES Search permission is denied for the directory referred to by *path*.

EFAULT *path* points outside the process's allocated address space.

EIO An I/O error occurred while reading from or writing to the file system.

SEE ALSO

`chroot(2)`

NAME

`chmod`, `fchmod` – change mode of file

SYNOPSIS

```
#include <usr/include/sys/stat.h
```

```
chmod(path, mode)
```

```
char *path;
```

```
int mode;
```

```
fchmod(fd, mode)
```

```
int fd, mode;
```

DESCRIPTION

The file whose name is given by *path* or referenced by the descriptor *fd* has its mode changed to *mode*. Modes are constructed by *or*'ing together some combination of the following:

<code>S_ISUID</code>	04000	set user ID on execution
<code>S_ISGID</code>	02000	set group ID on execution
<code>S_ISVTX</code>	01000	save text image after execution (sticky bit)
<code>S_IREAD</code>	00400	read by owner
<code>S_IWRITE</code>	00200	write by owner
<code>S_IXEXEC</code>	00100	execute (search on directory) by owner
	00070	read, write, execute (search) by group
	00007	read, write, execute (search) by others

These bit patterns are defined in `/usr/include/sys/stat.h`.

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

If the effective user ID of the process is not super-user and the process attempts to set the set group ID bit on a file owned by a group which is not in its group access list, mode bit 02000 (set group ID on execution) is cleared.

If an executable file is set up for sharing (this is the default) then mode 01000 (save text image after execution) prevents the system from abandoning the swap-space image of the program-text portion of the file when its last user terminates. If the effective user ID of the process is not super-user, this bit is cleared.

If a user other than the super-user writes to a file, the set user ID and set group ID bits are turned off. This makes the system somewhat more secure by protecting set-user-ID (set-group-ID) files from remaining set-user-ID (set-group-ID) if they are modified, at the expense of a degree of compatibility.

RETURN VALUE

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

ERRORS

chmod will fail and the file mode will be unchanged if:

`ENOTDIR` A component of the path prefix of *path* is not a directory.

`EINVAL` *path* contains a byte with the high-order bit set.

`ENAMETOOLONG`

The length of a component of *path* exceeds 255 characters, or the length of *path* exceeds 1023 characters.

`ENOENT` The file referred to by *path* does not exist.

`EACCES` Search permission is denied for a component of the path prefix of *path*.

`ELOOP` Too many symbolic links were encountered in translating *path*.

`EPERM` The effective user ID does not match the owner of the file and the effective user ID is

not the super-user.

EINVAL *fd* refers to a socket, not to a file.

EROFS The file referred to by *path* resides on a read-only file system.

EFAULT *path* points outside the process's allocated address space.

EIO An I/O error occurred while reading from or writing to the file system.

fchmod will fail if:

EBADF The descriptor is not valid.

EROFS The file referred to by *fd* resides on a read-only file system.

EPERM The effective user ID does not match the owner of the file and the effective user ID is not the super-user.

EIO An I/O error occurred while reading from or writing to the file system.

FILES

/usr/include/sys/stat.h

SEE ALSO

open(2V), *chown(2)*, *stat(2)*, *sticky(8)*

NAME

chown, fchown – change owner and group of a file

SYNOPSIS

chown(path, owner, group)

char *path;

int owner, group;

fchown(fd, owner, group)

int fd, owner, group;

DESCRIPTION

The file that is named by *path* or referenced by *fd* has its *owner* and *group* changed as specified. Only the super-user may change the owner of the file, because if users were able to give files away, they could defeat the file-space accounting procedures. The owner of the file may change the group to a group of which he is a member; the super-user may change the group arbitrarily.

fchown is particularly useful when used in conjunction with the file locking primitives (see *flock(2)*).

If *owner* or *group* is specified as -1 , the corresponding ID of the file is not changed.

If a process whose effective user ID is not super-user successfully changes the group ID of a file, the set-user-ID and set-group-ID bits of the file mode, 04000 and 02000 respectively, will be cleared.

If the final component of *path* is a symbolic link, the ownership and group of the symbolic link is changed, not the ownership and group of the file or directory to which it points.

RETURN VALUE

Zero is returned if the operation was successful; -1 is returned, and a more specific error code is placed in the global variable *errno*, if an error occurs.

ERRORS

chown will fail and the file will be unchanged if:

ENOTDIR A component of the path prefix of *path* is not a directory.

EINVAL *path* contains a byte with the high-order bit set.

ENAMETOOLONG

The length of a component of *path* exceeds 255 characters, or the length of *path* exceeds 1023 characters.

ENOENT The file referred to by *path* does not exist.

EACCES Search permission is denied for a component of the path prefix of *path*.

ELOOP Too many symbolic links were encountered in translating *path*.

EPERM The user ID specified by *owner* is not the current owner ID of the file, or the group ID specified by *group* is not the current group ID of the file and is not in the process' group access list, and the effective user ID is not the super-user.

EROFS The file referred to by *path* resides on a read-only file system.

EFAULT *path* points outside the process's allocated address space.

EIO An I/O error occurred while reading from or writing to the file system.

fchown will fail if:

EBADF *fd* does not refer to a valid descriptor.

EINVAL *fd* refers to a socket, not a file.

EPERM The user ID specified by *owner* is not the current owner ID of the file, or the group ID specified by *group* is not the current group group access list, and the effective user ID is not the super-user.

EROFS The file referred to by *fd* resides on a read-only file system.

EIO An I/O error occurred while reading from or writing to the file system.

SEE ALSO

chmod(2), flock(2)

NAME

chroot – change root directory

SYNOPSIS

```
chroot(dirname)
char *dirname;
```

DESCRIPTION

dirname points to a path name naming a directory. *chroot* causes this directory to become the root directory, the starting point for path names beginning with /. The current working directory is unaffected by this call. This root directory setting is inherited across *execve(2)* and by all children of this process created with *fork(2)* calls.

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

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

In order for a directory to become the root directory a process must have execute (search) access to the directory.

RETURN VALUE

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

ERRORS

chroot will fail and the root directory will be unchanged if one or more of the following are true:

ENOTDIR	A component of the path prefix of <i>dirname</i> is not a directory.
ENOTDIR	The file referred to by <i>dirname</i> is not a directory.
EINVAL	<i>dirname</i> contains a byte with the high-order bit set.
ENAMETOOLONG	The length of a component of <i>dirname</i> exceeds 255 characters, or the length of <i>dirname</i> exceeds 1023 characters.
ENOENT	The directory referred to by <i>dirname</i> does not exist.
EACCES	Search permission is denied for a component of the path prefix of <i>dirname</i> .
EACCES	Search permission is denied for the directory referred to by <i>dirname</i> .
ELOOP	Too many symbolic links were encountered in translating <i>dirname</i> .
EPERM	The effective user ID is not super-user.
EFAULT	<i>dirname</i> points outside the process's allocated address space.
EIO	An I/O error occurred while reading from or writing to the file system.

SEE ALSO

chdir(2)

NAME

close – delete a descriptor

SYNOPSIS

```
close(d)
int d;
```

DESCRIPTION

The *close* call deletes a descriptor from the per-process object reference table. If this is the last reference to the underlying object, then it will be deactivated. For example, on the last close of a file the current *seek* pointer associated with the file is lost; on the last close of a *socket(2)* associated naming information and queued data are discarded; on the last close of a file holding an advisory lock the lock is released (see *flock(2)* for further information).

A close of all of a process's descriptors is automatic on *exit*, but since there is a limit on the number of active descriptors per process, *close* is necessary for programs that deal with many descriptors.

When a process forks (see *fork(2)*), all descriptors for the new child process reference the same objects as they did in the parent before the fork. If a new process is then to be run using *execve(2)*, the process would normally inherit these descriptors. Most of the descriptors can be rearranged with *dup2(2)* or deleted with *close* before the *execve* is attempted, but if some of these descriptors will still be needed if the *execve* fails, it is necessary to arrange for them to be closed if the *execve* succeeds. For this reason, the call “*fcntl(d, F_SETFD, 1)*” is provided, which arranges that a descriptor will be closed after a successful *execve*; the call “*fcntl(d, F_SETFD, 0)*” restores the default, which is to not close the descriptor.

Close unmaps pages mapped through this file descriptor.

RETURN VALUE

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

ERRORS

Close will fail if:

EBADF	<i>D</i> is not an active descriptor.
EINTR	A read from a slow device was interrupted before any data arrived by the delivery of a signal.

SEE ALSO

accept(2), *flock(2)*, *open(2V)*, *pipe(2)*, *socket(2)*, *socketpair(2)*, *execve(2)*, *fcntl(2)*, *mmap(2)*, *munmap(2)*

NAME

connect – initiate a connection on a socket

SYNOPSIS

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

connect(s, name, namelen)
int s;
struct sockaddr *name;
int namelen;
```

DESCRIPTION

The parameter *s* is a socket. If it is of type `SOCK_DGRAM`, then this call permanently specifies the peer to which datagrams are to be sent; if it is of type `SOCK_STREAM`, then this call attempts to make a connection to another socket. The other socket is specified by *name* which is an address in the communications space of the socket. Each communications space interprets the *name* parameter in its own way.

RETURN VALUE

If the connection or binding succeeds, then 0 is returned. Otherwise a `-1` is returned, and a more specific error code is stored in *errno*.

ERRORS

The call fails if:

EBADF	<i>s</i> is not a valid descriptor.
ENOTSOCK	<i>s</i> is a descriptor for a file, not a socket.
EADDRNOTAVAIL	The specified address is not available on this machine.
EAFNOSUPPORT	Addresses in the specified address family cannot be used with this socket.
EISCONN	The socket is already connected.
ETIMEDOUT	Connection establishment timed out without establishing a connection.
ECONNREFUSED	The attempt to connect was forcefully rejected.
ENETUNREACH	The network isn't reachable from this host.
EADDRINUSE	The address is already in use.
EFAULT	The <i>name</i> parameter specifies an area outside the process address space.
EWouldBLOCK	The socket is non-blocking and the connection cannot be completed immediately. It is possible to <i>select</i> (2) the socket while it is connecting by selecting it for writing.
EINTR	A read from a slow device was interrupted before any data arrived by the delivery of a signal.

The following errors are specific to connecting names in the UNIX domain. These errors may not apply in future versions of the UNIX IPC domain.

ENOTDIR	A component of the path prefix of the path name in <i>name</i> is not a directory.
EINVAL	The path name in <i>name</i> contains a character with the high-order bit set.
ENAMETOOLONG	The length of a component of the path name in <i>name</i> exceeds 255 characters, or the length of the entire path name in <i>name</i> exceeds 1023 characters.
ENOENT	A component of the path prefix of the path name in <i>name</i> does not exist.
ENOENT	The socket referred to by the path name in <i>name</i> does not exist.
EACCES	Search permission is denied for a component of the path prefix of the path name in

name.

ELOOP Too many symbolic links were encountered in translating the path name in *name*.

EIO An I/O error occurred while reading from or writing to the file system.

SEE ALSO

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

NAME

`creat` – create a new file

SYNOPSIS

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

DESCRIPTION

This interface is made obsolete by

`creat` creates a new ordinary file or prepares to rewrite an existing file named by the path name pointed to by *name*. If the file did not exist, it is given mode *mode*, as modified by the process's mode mask (see `umask(2)`). Also see `chmod(2)` for the construction of the *mode* argument.

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

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

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

Upon successful completion, the file descriptor is returned and the file is open for writing, even if the mode does not permit writing. The file pointer is set to the beginning of the file. The file descriptor is set to remain open across `execve` system calls. See `fcntl(2)`.

NOTES

The *mode* given is arbitrary; it need not allow writing. This feature has been used in the past by programs to construct a simple exclusive locking mechanism. It is replaced by the `O_EXCL` open mode, or `flock(2)` facility.

RETURN VALUE

The value `-1` is returned if an error occurs. Otherwise, the call returns a non-negative descriptor which only permits writing.

ERRORS

`creat` will fail and the file will not be created or truncated if one of the following occur:

ENOTDIR	A component of the path prefix of <i>name</i> is not a directory.
EINVAL	<i>name</i> contains a byte with the high-order bit set.
ENAMETOOLONG	The length of a component of <i>name</i> exceeds 255 characters, or the length of <i>name</i> exceeds 1023 characters.
ENOENT	A component of the path prefix of <i>name</i> does not exist.
ELOOP	Too many symbolic links were encountered in translating <i>name</i> .
EACCES	Search permission is denied for a component of the path prefix of <i>name</i> .
EACCES	The file referred to by <i>name</i> does not exist and the directory in which it is to be created is not writable.
EACCES	The file referred to by <i>name</i> exists, but it is unwritable.
EISDIR	The file referred to by <i>name</i> is a directory.
EMFILE	There are already too many files open.
ENFILE	The system file table is full.
ENOSPC	The directory in which the entry for the new file is being placed cannot be extended because there is no space left on the file system containing the directory.

ENOSPC	There are no free inodes on the file system on which the file is being created.
EDQUOT	The directory in which the entry for the new file is being placed cannot be extended because the user's quota of disk blocks on the file system containing the directory has been exhausted.
EDQUOT	The user's quota of inodes on the file system on which the file is being created has been exhausted.
EROFS	The file referred to by <i>name</i> resides, or would reside, on a read-only file system.
ENXIO	The file is a character special or block special file, and the associated device does not exist.
ETXTBSY	The file is a pure procedure (shared text) file that is being executed.
EIO	An I/O error occurred while making the directory entry or allocating the inode.
EFAULT	<i>name</i> points outside the process's allocated address space.
EOPNOTSUPP	The file was a socket (not currently implemented).

SEE ALSO

open(2), write(2V), close(2), chmod(2), fcntl(2), umask(2)

NAME

dup, dup2 – duplicate a descriptor

SYNOPSIS

newd = dup(oldd)

int newd, oldd;

dup2(oldd, newd)

int oldd, newd;

DESCRIPTION

dup duplicates an existing object descriptor. The argument *oldd* is a small non-negative integer index in the per-process descriptor table. The value must be less than the size of the table, which is returned by *getdtablesize(2)*. The new descriptor returned by the call, *newd*, is the lowest numbered descriptor that is not currently in use by the process.

In the second form of the call, the value of *newd* desired is specified. If this descriptor is already in use, the descriptor is first deallocated as if a *close(2)* call had been done first.

The new descriptor has the following in common with the original:

It refers to the same object that the old descriptor referred to.

It uses the same file pointer as the old descriptor. (i.e., both file descriptors share one file pointer).

It has the same access mode (read, write or read/write) as the old descriptor.

Thus if *newd* and *oldd* are duplicate references to an open file, *read(2V)*, *write(2V)* and *lseek(2)* calls all move a single pointer into the file, and append mode, non-blocking I/O and asynchronous I/O options are shared between the references. If a separate pointer into the file is desired, a different object reference to the file must be obtained by issuing an additional *open(2V)* call. The close-on-exec flag on the new file descriptor is unset.

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

RETURN VALUE

The value -1 is returned if an error occurs in either call. The external variable *errno* indicates the cause of the error.

ERRORS

dup and *dup2* fail if:

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

EMFILE Too many descriptors are active.

SEE ALSO

accept(2), *open(2)*, *close(2)*, *fcntl(2)*, *pipe(2)*, *socket(2)*, *socketpair(2)*, *getdtablesize(2)*

NAME

`execve` – execute a file

SYNOPSIS

```
execve(name, argv, envp)
char *name, *argv[], *envp[];
```

DESCRIPTION

`execve` transforms the calling process into a new process. The new process is constructed from an ordinary file, whose name is pointed to by *path*, called the *new process file*. This file is either an executable object file, or a file of data for an interpreter. An executable object file consists of an identifying header, followed by pages of data representing the initial program (text) and initialized data pages. Additional pages may be specified by the header to be initialized with zero data. See *a.out(5)*.

An interpreter file begins with a line of the form “#! *interpreter* [*arg*]”. When an interpreter file is `execve`'d, the system `execve`'s the specified *interpreter*. If the optional *arg* is specified, it becomes the first argument to the *interpreter*, and the name of the originally `execve`'d file becomes the second argument; otherwise, the name of the originally `execve`'d file becomes the first argument. The original arguments are shifted over to become the subsequent arguments. The zeroth argument, normally the name of the `execve`'d file, is left unchanged.

There can be no return from a successful `execve` because the calling core image is lost. This is the mechanism whereby different process images become active.

The argument *argv* is a null-terminated array of character pointers to null-terminated character strings. These strings constitute the argument list to be made available to the new process. By convention, at least one argument must be present in this array, and the first element of this array should be the name of the executed program (i.e., the last component of *name*).

The argument *envp* is also a null-terminated array of character pointers to null-terminated strings. These strings pass information to the new process which are not directly arguments to the command (see *environ(5V)*).

Descriptors open in the calling process remain open in the new process, except for those for which the close-on-exec flag is set (see *close(2)* and *fcntl(2)*). Descriptors which remain open are unaffected by `execve`.

Ignored signals remain ignored across an `execve`, but signals that are caught are reset to their default values. Blocked signals remain blocked regardless of changes to the signal action. The signal stack is reset to be undefined (see *sigvec(2)* for more information).

Each process has a *real* user ID and group ID and an *effective* user ID and group ID. The *real* ID identifies the person using the system; the *effective* ID determines their access privileges. `Execve` changes the effective user or group ID to the owner or group of the executed file if the file has the “set-user-ID” or “set-group-ID” modes. The *real* user ID and group ID are not affected.

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

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

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

process ID	see <i>getpid(2)</i>
parent process ID	see <i>getppid(2)</i>
process group ID	see <i>getpgrp(2)</i>
access groups	see <i>getgroups(2)</i>
semadj values	see <i>semop(2)</i>
working directory	see <i>chdir(2)</i>
root directory	see <i>chroot(2)</i>
control terminal	see <i>tty(4)</i>

trace flag	see <i>ptrace(2)</i> request 0)
resource usages	see <i>getrusage(2)</i>
interval timers	see <i>getitimer(2)</i>
resource limits	see <i>getrlimit(2)</i>
file mode mask	see <i>umask(2)</i>
signal mask	see <i>sigvec(2)</i> , <i>sigmask(2)</i>

When the executed program begins, it is called as follows:

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

where *argc* is the number of elements in *argv* (the “arg count”) and *argv* is the array of character pointers to the arguments themselves.

envp is a pointer to an array of strings that constitute the *environment* of the process. A pointer to this array is also stored in the global variable “*environ*”. Each string consists of a name, an “=”, and a null-terminated value. The array of pointers is terminated by a null pointer. The shell *sh(1)* passes an environment entry for each global shell variable defined when the program is called. See *environ(5V)* for some conventionally used names.

RETURN VALUE

If *execve* returns to the calling process an error has occurred; the return value will be -1 and the global variable *errno* will contain an error code.

ERRORS

execve will fail and return to the calling process if one or more of the following are true:

ENOTDIR	A component of the path prefix of the new process file is not a directory.
EINVAL	<i>name</i> contains a character with the high-order bit set.
ENAMETOOLONG	The length of a component of <i>name</i> exceeds 255 characters, or the length of <i>name</i> exceeds 1023 characters.
ENOENT	One or more components of the path prefix of the new process file does not exist.
ENOENT	The new process file does not exist.
ELOOP	Too many symbolic links were encountered in translating <i>name</i> .
EACCES	Search permission is denied for a component of the new process file’s path prefix.
EACCES	The new process file is not an ordinary file.
EACCES	Execute permission is denied for the new process file.
ENOEXEC	The new process file has the appropriate access permission, but has an invalid magic number in its header.
ETXTBSY	The new process file is a pure procedure (shared text) file that is currently open for writing or reading by some process.
ENOMEM	The new process file requires more virtual memory than is allowed by the imposed maximum (<i>getrlimit(2)</i>).
[E2BIG]	The number of bytes in the new process file’s argument list is larger than the system-imposed limit. The limit in the system as released is 10240 bytes (NCARGS in <sys/param.h>).
EFAULT	The new process file is not as long as indicated by the size values in its header.
EFAULT	<i>Name</i> , <i>argv</i> , or <i>envp</i> point to an illegal address.
EIO	An I/O error occurred while reading from the file system.

CAVEATS

If a program is *setuid* to a non-super-user, but is executed when the real user ID is super-user, then the program has some of the powers of a super-user as well.

SEE ALSO

exit(2), fork(2), execl(3), environ(5V)

NAME

`_exit` – terminate a process

SYNOPSIS

```
_exit(status)  
int status;
```

DESCRIPTION

`_exit` terminates a process with the following consequences:

All of the descriptors open in the calling process are closed. This may entail delays, for example, waiting for output to drain; a process in this state may not be killed, as it is already dying.

If the parent process of the calling process is executing a `wait` or is interested in the SIGCHLD signal, then it is notified of the calling process's termination and the low-order eight bits of `status` are made available to it; see `wait(2)`.

The parent process ID of all of the calling process's existing child processes are also set to 1. This means that the initialization process (see `intro(2)`) inherits each of these processes as well. Any stopped children are restarted with a hangup signal (SIGHUP).

Most C programs will call the library routine `exit(3)` which performs cleanup actions in the standard I/O library before calling `_exit`.

RETURN VALUE

This call never returns.

SEE ALSO

`fork(2)`, `wait(2)`, `exit(3)`

NAME

`fcntl` – file control

SYNOPSIS

```
#include <fcntl.h>

res = fcntl(fd, cmd, arg)
int res;
int fd, cmd, arg;
```

DESCRIPTION

Fcntl performs a variety of functions on open descriptors. The argument *fd* is an open descriptor to be operated on by *cmd* as follows:

- F_DUPFD** Return a new descriptor as follows:
- Lowest numbered available descriptor greater than or equal to *arg*.
 - References the same object as the original descriptor.
 - New descriptor shares the same file pointer if the object was a file.
 - Same access mode (read, write or read/write).
 - Same file status flags (i.e., both descriptors share the same file status flags).
 - The close-on-exec flag associated with the new descriptor is set to remain open across *execve(2)* system calls.
- F_GETFD** Get the close-on-exec flag associated with the descriptor *fd*. If the low-order bit is 0, the file will remain open across *exec*, otherwise the file will be closed upon execution of *exec*.
- F_SETFD** Set the close-on-exec flag associated with *fd* to the low order bit of *arg* (0 or 1 as above).
- F_GETFL** Get descriptor status flags, see *fcntl(5)* for their definitions.
- F_SETFL** Set descriptor status flags, see *fcntl(5)* for their definitions.
- F_GETLK** Get a description of the first lock which would block the lock specified in the *flock* structure pointed to by *arg*. The information retrieved overwrites the information in the *flock* structure. If no lock is found that would prevent this lock from being created, then the structure is passed back unchanged except for the lock type which will be set to **F_UNLCK**.
- F_SETLK** Set or clear an advisory record lock according to the *flock* structure pointed to by *arg*. **F_SETLK** is used to establish shared (**F_RDLCK**) and exclusive (**F_WRLCK**) locks, or to remove either type of lock (**F_UNLCK**). If the specified lock cannot be applied, *fcntl* will return with an error value of -1.
- F_SETLKW** This *cmd* is the same as **F_SETLK** except that if a shared or exclusive lock is blocked by other locks, the requesting process will sleep until the lock may be applied.
- F_GETOWN** Get the process ID or process group currently receiving **SIGIO** and **SIGURG** signals; process groups are returned as negative values.
- F_SETOWN** Set the process or process group to receive **SIGIO** and **SIGURG** signals; process groups are specified by supplying *arg* as negative, otherwise *arg* is interpreted as a process ID.

The **SIGIO** facilities are enabled by setting the **FASYNC** flag with **F_SETFL**.

NOTES

Advisory locks allow cooperating processes to perform consistent operations on files, but do not guarantee exclusive access (i.e., processes may still access files without using advisory locks, possibly resulting in inconsistencies).

The record locking mechanism allows two types of locks: shared locks (F_RDLCK) and exclusive locks (F_WRLCK). More than one process may hold a shared lock for a particular segment of a file at any given time, but multiple exclusive, or both shared and exclusive, locks may not exist simultaneously on any segment.

In order to claim a shared lock, the descriptor must have been opened with read access. The descriptor on which an exclusive lock is being placed must have been opened with write access.

A shared lock may be *upgraded* to an exclusive lock, and vice versa, simply by specifying the appropriate lock type with a *cmd* of F_SETLK or F_SETLKW; the previous lock will be released and the new lock applied (possibly after other processes have gained and released the lock).

If the *cmd* is F_SETLKW and the requested lock cannot be claimed immediately (e.g., another process holds an exclusive lock that partially or completely overlaps the current request) then the calling process will block until the lock may be acquired. Processes blocked awaiting a lock may be awakened by signals.

Care should be taken to avoid deadlock situations in applications in which multiple processes perform blocking locks on a set of common records.

The record that is to be locked or unlocked is described by the *flock* structure, which is defined in `<fcntl.h>` as follows:

```

struct flock {
    short   l_type;           /* F_RDLCK, F_WRLCK, or F_UNLCK */
    short   l_whence;        /* flag to choose starting offset */
    long    l_start;         /* relative offset, in bytes */
    long    l_len;           /* length, in bytes; 0 means lock to EOF */
    short   l_pid;           /* returned with F_GETLK */
};

```

The *flock* structure describes the type (*l_type*), starting offset (*l_whence*), relative offset (*l_start*), and size (*l_len*) of the segment of the file to be affected. *l_whence* must be set to 0, 1, or 2 to indicate that the relative offset will be measured from the start of the file, current position, or end-of-file, respectively. The process id field (*l_pid*) is only used with the F_GETLK *cmd* to return the description of a lock held by another process.

Locks may start and extend beyond the current end-of-file, but may not be negative relative to the beginning of the file. A lock may be set to always extend to the end-of-file by setting *l_len* to zero (0). If such a lock also has *l_whence* and *l_start* set to zero (0), the entire file will be locked. Changing or unlocking a segment from the middle of a larger locked segment leaves two smaller segments at either end. Locking a segment that is already locked by the calling process causes the old lock type to be removed and the new lock type to take affect. All locks associated with a file for a given process are removed when the file is closed or the process terminates. Locks are not inherited by the child process in a *fork(2)* system call.

In order to maintain consistency in the network case, data must not be cached on client machines. For this reason, file buffering for an NFS file is turned off when the first lock is attempted on the file. Buffering will remain off as long as the file is open. Programs that do I/O buffering in the user address space, however, may have inconsistent results (the standard I/O package, for instance, is a common source of unexpected buffering).

The advisory record locking capabilities of *fcntl* are implemented throughout the network by the **network lock daemon**; see *lockd(8C)*. If the file server crashes and is rebooted, the lock daemon will attempt to recover all locks that were associated with that server. If a lock cannot be reclaimed, the process that held the lock will be issued a SIGLOST signal.

RETURN VALUE

Upon successful completion, the value returned depends on *cmd* as follows:

F_DUPFD	A new descriptor.
F_GETFD	Value of flag (only the low-order bit is defined).
F_GETFL	Value of flags.

F_GETOWN Value of descriptor owner.
other Value other than -1.

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

ERRORS

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

EBADF *Fd* is not a valid open descriptor.

EMFILE *Cmd* is **F_DUPFD** and the maximum allowed number of descriptors are currently open.

EINVAL *Cmd* is **F_DUPFD** and *arg* is negative or greater than the maximum allowable number (see *getdtablesize(2)*).

EFAULT *Cmd* is **F_GETLK**, **F_SETLK**, or **F_SETLKW** and *arg* points to an invalid address.

EINVAL *Cmd* is **F_GETLK**, **F_SETLK**, or **F_SETLKW** and the data *arg* points to is not valid.

EBADF *Cmd* is **F_SETLK** or **F_SETLKW** and the process does not have the appropriate read or write permissions on the file.

EAGAIN *Cmd* is **F_SETLK**, the lock type (*l_type*) is **F_RDLCK** (shared lock), and the segment of the file to be locked already has an exclusive lock held by another process. This error will also be returned if the lock type is **F_WRLCK** (exclusive lock) and another process already has the segment locked with either a shared or exclusive lock.

EINTR *Cmd* is **F_SETLKW** and a signal interrupted the process while it was waiting for the lock to be granted.

ENOLCK *Cmd* is **F_SETLK** or **F_SETLKW** and there are no more file lock entries available.

SEE ALSO

close(2), *execve(2)*, *getdtablesize(2)*, *open(2V)*, *sigvec(2)*, *lockf(3)*, *lockd(8C)*

BUGS

File locks obtained through the *fcntl* mechanism do not interact in any way with those acquired via *flock(2)*. They do, however, work correctly with the exclusive locks claimed by *lockf(3)*.

F_GETLK returns **F_UNLCK** if the requesting process holds the specified lock. Thus, there is no way for a process to determine if it is still holding a specific lock after catching a **SIGLOST** signal.

In a network environment, the value of *l_pid* returned by **F_GETLK** is next to useless.

NAME

`flock` – apply or remove an advisory lock on an open file

SYNOPSIS

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

#define LOCK_SH    1    /* shared lock */
#define LOCK_EX    2    /* exclusive lock */
#define LOCK_NB    4    /* don't block when locking */
#define LOCK_UN    8    /* unlock */

flock(fd, operation)
int fd, operation;
```

DESCRIPTION

Flock applies or removes an *advisory* lock on the file associated with the file descriptor *fd*. A lock is applied by specifying an *operation* parameter that is the inclusive OR of LOCK_SH or LOCK_EX and, possibly, LOCK_NB. To unlock an existing lock, the *operation* should be LOCK_UN.

Advisory locks allow cooperating processes to perform consistent operations on files, but do not guarantee exclusive access (i.e., processes may still access files without using advisory locks, possibly resulting in inconsistencies).

The locking mechanism allows two types of locks: *shared* locks and *exclusive* locks. More than one process may hold a shared lock for a file at any given time, but multiple exclusive, or both shared and exclusive, locks may not exist simultaneously on a file.

A shared lock may be *upgraded* to an exclusive lock, and vice versa, simply by specifying the appropriate lock type; the previous lock will be released and the new lock applied (possibly after other processes have gained and released the lock).

Requesting a lock on an object that is already locked normally causes the caller to block until the lock may be acquired. If LOCK_NB is included in *operation*, then this will not happen; instead the call will fail and the error EWOULDBLOCK will be returned.

NOTES

Locks are on files, not file descriptors. That is, file descriptors duplicated through *dup(2)* or *fork(2)* do not result in multiple instances of a lock, but rather multiple references to a single lock. If a process holding a lock on a file forks and the child explicitly unlocks the file, the parent will lose its lock.

Processes blocked awaiting a lock may be awakened by signals.

RETURN VALUE

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

ERRORS

The *flock* call fails if:

EWOULDBLOCK	The file is locked and the LOCK_NB option was specified.
EBADF	The argument <i>fd</i> is an invalid descriptor.
EOPNOTSUPP	The argument <i>fd</i> refers to an object other than a file.

SEE ALSO

open(2V), *close(2)*, *dup(2)*, *execve(2)*, *fcntl(2)*, *fork(2)*, *lockf(3)*

BUGS

Locks obtained through the *flock* mechanism are known only within the system on which they were placed. Thus, multiple clients may successfully acquire exclusive locks on the same remote file. If this behavior is not explicitly desired, the *fcntl(2)* or *lockf(3)* system calls should be used instead; these make use of the services of the **network lock manager** (see *lockd(8C)*).

NAME

`fork` – create a new process

SYNOPSIS

```
pid = fork()
int pid;
```

DESCRIPTION

Fork creates a new process. The new process (child process) is an exact copy of the calling process except for the following:

The child process has a unique process ID.

The child process has a different parent process ID (that is, the process ID of the parent process).

The child process has its own copy of the parent's descriptors. These descriptors reference the same underlying objects, so that, for instance, file pointers in file objects are shared between the child and the parent, so that an *lseek(2)* on a descriptor in the child process can affect a subsequent *read* or *write* by the parent. This descriptor copying is also used by the shell to establish standard input and output for newly created processes as well as to set up pipes.

The child processes resource utilizations are set to 0; see *setrlimit(2)*.

RETURN VALUE

Upon successful completion, *fork* returns a value of 0 to the child process and returns the process ID of the child process to the parent process. Otherwise, a value of -1 is returned to the parent process, no child process is created, and the global variable *errno* is set to indicate the error.

ERRORS

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

- | | |
|--------|--|
| EAGAIN | The system-imposed limit on the total number of processes under execution would be exceeded. This limit is determined when the system is generated. |
| EAGAIN | The system-imposed limit on the total number of processes under execution by a single user would be exceeded. This limit is determined when the system is generated. |
| ENOMEM | There is insufficient swap space for the new process. |

SEE ALSO

execve(2), *wait(2)*

NAME

fsync – synchronize a file's in-core state with that on disk

SYNOPSIS

```
fsync(fd)
int fd;
```

DESCRIPTION

fsync moves all modified data and attributes of *fd* to a permanent storage device: all in-core modified copies of buffers for the associated file have been written to a disk when the call returns. Note that this is different than *sync*(2) which schedules disk I/O for all files (as though an *fsync* had been done on all files) but returns before the I/O completes.

fsync should be used by programs which require a file to be in a known state; for example, a program which contains a simple transaction facility might use it to ensure that all modifications to a file or files caused by a transaction were recorded on disk.

RETURN VALUE

A 0 value is returned on success. A -1 value indicates an error.

ERRORS

The *fsync* fails if:

EBADF *fd* is not a valid descriptor.

EINVAL *fd* refers to a socket, not a file.

EIO An I/O error occurred while reading from or writing to the file system.

SEE ALSO

sync(2), *sync*(8), *cron*(8)

BUGS

The current implementation of this call is expensive for large files.

NAME

`getdirentries` – gets directory entries in a filesystem independent format

SYNOPSIS

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

cc = getdirentries(fd, buf, nbytes, basep)
int cc, fd;
char *buf;
int nbytes;
long *basep;
```

DESCRIPTION

`getdirentries` attempts to put directory entries from the directory referenced by the file descriptor `fd` into the buffer pointed to by `buf`, in a filesystem independent format. Up to `nbytes` of data will be transferred. `nbytes` must be greater than or equal to the block size associated with the file, see `stat(2)`. Sizes less than this may cause errors on certain filesystems.

The data in the buffer is a series of *direct* structures each containing the following entries:

```
    unsigned long  d_fileno;
    unsigned short d_reclen;
    unsigned short d_namlen;
    char           d_name[MAXNAMELEN + 1]; /* see below */
```

The `d_fileno` entry is a number which is unique for each distinct file in the filesystem. Files that are linked by hard links (see `link(2)`) have the same `d_fileno`. The `d_reclen` entry is the length, in bytes, of the directory record. The `d_name` entry contains a null terminated file name. The `d_namlen` entry specifies the length of the file name. Thus the actual size of `d_name` may vary from 2 to `MAXNAMELEN + 1`.

The structures are not necessarily tightly packed. The `d_reclen` entry may be used as an offset from the beginning of a *direct* structure to the next structure, if any.

Upon return, the actual number of bytes transferred is returned. The current position pointer associated with `fd` is set to point to the next block of entries. The pointer is not necessarily incremented by the number of bytes returned by `getdirentries`. If the value returned is zero, the end of the directory has been reached. The current position pointer may be set and retrieved by `lseek(2)`. `getdirentries` writes the position of the block read into the location pointed to by `basep`. It is not safe to set the current position pointer to any value other than a value previously returned by `lseek(2)` or a value previously returned in the location pointed to by `basep` or zero.

RETURN VALUE

If successful, the number of bytes actually transferred is returned. Otherwise, a `-1` is returned and the global variable `errno` is set to indicate the error.

ERRORS

`getdirentries` will fail if one or more of the following are true:

<code>EBADF</code>	<code>fd</code> is not a valid file descriptor open for reading.
<code>EFAULT</code>	Either <code>buf</code> or <code>basep</code> point outside the allocated address space.
<code>EIO</code>	An I/O error occurred while reading from or writing to the file system.
<code>EINTR</code>	A read from a slow device was interrupted before any data arrived by the delivery of a signal.

SEE ALSO

`open(2V)`, `lseek(2)`

NAME

`getdomainname`, `setdomainname` – get/set name of current domain

SYNOPSIS

`getdomainname(name, namelen)`

`char *name;`

`int namelen;`

`setdomainname(name, namelen)`

`char *name;`

`int namelen;`

DESCRIPTION

Getdomainname returns the name of the domain for the current processor, as previously set by *setdomainname*. The parameter *namelen* specifies the size of the *name* array. The returned name is null-terminated unless insufficient space is provided.

Setdomainname sets the domain of the host machine to be *name*, which has length *namelen*. This call is restricted to the super-user and is normally used only when the system is bootstrapped.

The purpose of domains is to enable two distinct networks that may have host names in common to merge. Each network would be distinguished by having a different domain name. At the current time, only the yellow pages service makes use of domains.

RETURN VALUE

If the call succeeds a value of 0 is returned. If the call fails, then a value of -1 is returned and an error code is placed in the global location *errno*.

ERRORS

The following errors may be returned by these calls:

EFAULT The *name* parameter gave an invalid address.

EPERM The caller was not the super-user. This error only applies to *setdomainname*.

BUGS

Domain names are limited to 255 characters.

NAME

getdtablesize – get descriptor table size

SYNOPSIS

```
nds = getdtablesize()
int nds;
```

DESCRIPTION

Each process has a fixed size descriptor table, which is guaranteed to have at least 20 slots. The entries in the descriptor table are numbered with small integers starting at 0. The call *getdtablesize* returns the size of this table.

SEE ALSO

close(2), dup(2), open(2)

NAME

`getgid`, `getegid` – get group identity

SYNOPSIS

```
gid = getgid()
int gid;
```

```
egid = getegid()
int egid;
```

DESCRIPTION

Getgid returns the real group ID of the current process, *getegid* the effective group ID.

The real group ID is specified at login time.

The effective group ID is more transient, and determines additional access permission during execution of a "set-group-ID" process, and it is for such processes that *getgid* is most useful.

SEE ALSO

`getuid(2)`, `setregid(2)`, `setgid(3)`

NAME

`getgroups`, `setgroups` – get or set group access list

SYNOPSIS

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

ngroups = getgroups(gidsetlen, gidset)
int ngroups, gidsetlen, *gidset;

setgroups(ngroups, gidset)
int ngroups, *gidset;
```

DESCRIPTION**Getgroups**

getgroups gets the current group access list of the user process and stores it in the array *gidset*. The parameter *gidsetlen* indicates the number of entries that may be placed in *gidset*. *getgroups* returns the actual number of entries placed in the *gidset* array. No more than NGROUPS, as defined in *<sys/param.h>*, will ever be returned.

Setgroups

setgroups sets the group access list of the current user process according to the array *gidset*. The parameter *ngroups* indicates the number of entries in the array and must be no more than NGROUPS, as defined in *<sys/param.h>*.

Only the super-user may set new groups.

RETURN VALUE**Getgroups**

A return value of greater than zero indicates the number of entries placed in the *gidset* array. A return value of -1 indicates that an error occurred, and the error code is stored in the global variable *errno*.

Setgroups

A 0 value is returned on success, -1 on error, with a error code stored in *errno*.

ERRORS

Either call fails if:

EFAULT The address specified for *gidset* is outside the process address space.

getgroup fails if:

EINVAL The argument *gidsetlen* is smaller than the number of groups in the group set.

setgroups fails if:

EPERM The caller is not the super-user.

SEE ALSO

`initgroups(3)`

NAME

`gethostid` – get unique identifier of current host

SYNOPSIS

```
hostid = gethostid()  
long hostid;
```

DESCRIPTION

Gethostid returns the 32-bit identifier for the current host, which should be unique across all hosts. On the Sun, this number is taken from the CPU board's ID PROM.

SEE ALSO

`hostid(1)`

NAME

gethostname, sethostname – get/set name of current host

SYNOPSIS

gethostname(name, namelen)

char *name;

int namelen;

sethostname(name, namelen)

char *name;

int namelen;

DESCRIPTION

Gethostname returns the standard host name for the current processor, as previously set by *sethostname*. The parameter *namelen* specifies the size of the *name* array. The returned name is null-terminated unless insufficient space is provided.

Sethostname sets the name of the host machine to be *name*, which has length *namelen*. This call is restricted to the super-user and is normally used only when the system is bootstrapped.

RETURN VALUE

If the call succeeds a value of 0 is returned. If the call fails, then a value of -1 is returned and an error code is placed in the global location *errno*.

ERRORS

The following errors may be returned by these calls:

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

EPERM The caller was not the super-user. Note that this error only applies to *sethostname*.

SEE ALSO

gethostid(2)

BUGS

Host names are limited to 31 characters.

NAME

gettimer, setitimer – get/set value of interval timer

SYNOPSIS

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

#define ITIMER_REAL      0      /* real time intervals */
#define ITIMER_VIRTUAL  1      /* virtual time intervals */
#define ITIMER_PROF     2      /* user and system virtual time */
```

```
gettimer(which, value)
int which;
struct itimerval *value;

setitimer(which, value, ovalue)
int which;
struct itimerval *value, *ovalue;
```

DESCRIPTION

The system provides each process with three interval timers, defined in *<sys/time.h>*. The *gettimer* call returns the current value for the timer specified in *which*, while the *setitimer* call sets the value of a timer (optionally returning the previous value of the timer).

A timer value is defined by the *itimerval* structure:

```
struct itimerval {
    struct timeval it_interval;    /* timer interval */
    struct timeval it_value;      /* current value */
};
```

If *it_value* is non-zero, it indicates the time to the next timer expiration. If *it_interval* is non-zero, it specifies a value to be used in reloading *it_value* when the timer expires. Setting *it_value* to 0 disables a timer. Setting *it_interval* to 0 causes a timer to be disabled after its next expiration (assuming *it_value* is non-zero).

Time values smaller than the resolution of the system clock are rounded up to this resolution.

The ITIMER_REAL timer decrements in real time. A SIGALRM signal is delivered when this timer expires.

The ITIMER_VIRTUAL timer decrements in process virtual time. It runs only when the process is executing. A SIGVTALRM signal is delivered when it expires.

The ITIMER_PROF timer decrements both in process virtual time and when the system is running on behalf of the process. It is designed to be used by interpreters in statistically profiling the execution of interpreted programs. Each time the ITIMER_PROF timer expires, the SIGPROF signal is delivered. Because this signal may interrupt in-progress system calls, programs using this timer must be prepared to restart interrupted system calls.

NOTES

Three macros for manipulating time values are defined in *<sys/time.h>*. *Timerclear* sets a time value to zero, *timerisset* tests if a time value is non-zero, and *timercmp* compares two time values (beware that *>=* and *<=* do not work with this macro).

RETURN VALUE

If the calls succeed, a value of 0 is returned. If an error occurs, the value -1 is returned, and a more precise error code is placed in the global variable *errno*.

ERRORS

The possible errors are:

EFAULT The *value* or *ovalue* parameter specified a bad address.

EINVAL A *value* parameter specified a time which was too large to be handled.

SEE ALSO

sigvec(2), gettimeofday(2)

NAME

`getpagesize` – get system page size

SYNOPSIS

```
pagesize = getpagesize()
int pagesize;
```

DESCRIPTION

Getpagesize returns the number of bytes in a page. Page granularity is the granularity of many of the memory management calls.

The page size is a *system* page size and may not be the same as the underlying hardware page size.

SEE ALSO

`sbrk(2)`, `pagesize(1)`

NAME

`getpeername` – get name of connected peer

SYNOPSIS

```
getpeername(s, name, namelen)  
int s;  
struct sockaddr *name;  
int *namelen;
```

DESCRIPTION

Getpeername returns the name of the peer connected to socket *s*. The *namelen* parameter should be initialized to indicate the amount of space pointed to by *name*. On return it contains the actual size of the name returned (in bytes).

DIAGNOSTICS

A 0 is returned if the call succeeds, -1 if it fails.

ERRORS

The call succeeds unless:

EBADF	The argument <i>s</i> is not a valid descriptor.
ENOTSOCK	The argument <i>s</i> is a file, not a socket.
ENOTCONN	The socket is not connected.
ENOBUFS	Insufficient resources were available in the system to perform the operation.
EFAULT	The <i>name</i> parameter points to memory not in a valid part of the process address space.

SEE ALSO

`bind(2)`, `socket(2)`, `getsockname(2)`

BUGS

Names bound to sockets in the UNIX domain are inaccessible; *getpeername* returns a zero length name.

NAME

getpid, getppid – get process identification

SYNOPSIS

```
pid = getpid()
```

```
int pid;
```

```
ppid = getppid()
```

```
int ppid;
```

DESCRIPTION

Getpid returns the process ID of the current process. Most often it is used to generate uniquely-named temporary files.

Getppid returns the process ID of the parent of the current process.

SEE ALSO

gethostid(2)

NAME

`getpriority`, `setpriority` – get/set program scheduling priority

SYNOPSIS

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

```
prio = getpriority(which, who)
```

```
int prio, which, who;
```

```
setpriority(which, who, prio)
```

```
int which, who, prio;
```

DESCRIPTION

The scheduling priority of the process, process group, or user, as indicated by *which* and *who* is obtained with the *getpriority* call and set with the *setpriority* call. *which* is one of `PRIO_PROCESS`, `PRIO_PGRP`, or `PRIO_USER`, and *who* is interpreted relative to *which* (a process identifier for `PRIO_PROCESS`, process group identifier for `PRIO_PGRP`, and a user ID for `PRIO_USER`). A zero value of *who* denotes the current process, process group, or user. *prio* is a value in the range -20 to 20 . The default priority is 0 ; lower priorities cause more favorable scheduling.

The *getpriority* call returns the highest priority (lowest numerical value) enjoyed by any of the specified processes. The *setpriority* call sets the priorities of all of the specified processes to the specified value. If the specified value is less than -20 , a value of -20 is used; if it is greater than 20 , a value of 20 is used. Only the super-user may lower priorities.

RETURN VALUE

Since *getpriority* can legitimately return the value -1 , it is necessary to clear the external variable *errno* prior to the call, then check it afterward to determine if a -1 is an error or a legitimate value. The *setpriority* call returns 0 if there is no error, or -1 if there is.

ERRORS

getpriority and *setpriority* may return one of the following errors:

`ESRCH` No process was located using the *which* and *who* values specified.

`EINVAL` *which* was not one of `PRIO_PROCESS`, `PRIO_PGRP`, or `PRIO_USER`.

In addition to the errors indicated above, *setpriority* may fail with one of the following errors returned:

`EPERM` A process was located, but neither its effective nor real user ID matched the effective user ID of the caller, and neither the effective nor the real user ID of the process executing the *setpriority* was super-user.

`EACCES` The call to *setpriority* would have changed a process' priority to a value lower than its current value, and the effective user ID of the process executing the call was not that of the super-user.

SEE ALSO

`nice(1)`, `fork(2)`, `renice(8)`

BUGS

It is not possible for the process executing *setpriority()* to lower any other process down to its current priority, without requiring super-user privileges.

NAME

getrlimit, setrlimit – control maximum system resource consumption

SYNOPSIS

```
#include <sys/time.h>
#include <sys/resource.h>

getrlimit(resource, rlp)
int resource;
struct rlimit *rlp;

setrlimit(resource, rlp)
int resource;
struct rlimit *rlp;
```

DESCRIPTION

Limits on the consumption of system resources by the current process and each process it creates may be obtained with the *getrlimit* call, and set with the *setrlimit* call.

The *resource* parameter is one of the following:

RLIMIT_CPU	the maximum amount of cpu time (in seconds) to be used by each process.
RLIMIT_FSIZE	the largest size, in bytes, of any single file that may be created.
RLIMIT_DATA	the maximum size, in bytes, of the data segment for a process; this defines how far a program may extend its break with the <i>sbrk(2)</i> system call.
RLIMIT_STACK	the maximum size, in bytes, of the stack segment for a process; this defines how far a program's stack segment may be extended automatically by the system.
RLIMIT_CORE	the largest size, in bytes, of a <i>core</i> file that may be created.
RLIMIT_RSS	the maximum size, in bytes, to which a process's resident set size may grow. This imposes a limit on the amount of physical memory to be given to a process; if memory is tight, the system will prefer to take memory from processes that are exceeding their declared resident set size.

A resource limit is specified as a soft limit and a hard limit. When a soft limit is exceeded a process may receive a signal (for example, if the cpu time is exceeded), but it will be allowed to continue execution until it reaches the hard limit (or modifies its resource limit). The *rlimit* structure is used to specify the hard and soft limits on a resource,

```
struct rlimit {
    int    rlim_cur;    /* current (soft) limit */
    int    rlim_max;    /* hard limit */
};
```

Only the super-user may raise the maximum limits. Other users may only alter *rlim_cur* within the range from 0 to *rlim_max* or (irreversibly) lower *rlim_max*.

An “infinite” value for a limit is defined as RLIM_INFINITY (0x7fffffff).

Because this information is stored in the per-process information, this system call must be executed directly by the shell if it is to affect all future processes created by the shell; *limit* is thus a built-in command to *csh(1)*.

The system refuses to extend the data or stack space when the limits would be exceeded in the normal way: a *brk* or *sbrk* call will fail if the data space limit is reached, or the process will be 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 that is too large will cause a signal SIGXFSZ to be generated; this normally terminates the process, but may be caught. When the soft CPU time limit is exceeded, a signal SIGXCPU is sent to the offending process.

RETURN VALUE

A 0 return value indicates that the call succeeded, changing or returning the resource limit. A return value of -1 indicates that an error occurred, and an error code is stored in the global location *errno*.

ERRORS

The possible errors are:

EFAULT The address specified for *rlp* is invalid.

EPERM The limit specified to *setrlimit* would have raised the maximum limit value, and the caller is not the super-user.

SEE ALSO

cs(1), *quota*(2)

BUGS

There should be *limit* and *unlimit* commands in *sh*(1) as well as in *cs*.

NAME

getrusage – get information about resource utilization

SYNOPSIS

```
#include <sys/time.h>
#include <sys/resource.h>

getrusage(who, rusage)
int who;
struct rusage *rusage;
```

DESCRIPTION

getrusage returns information about the resources utilized by the current process, or all its terminated child processes. The *who* parameter is one of RUSAGE_SELF or RUSAGE_CHILDREN. The buffer to which *rusage* points will be filled in with the following structure:

```
struct rusage {
    struct timeval ru_utime;          /* user time used */
    struct timeval ru_stime;          /* system time used */
    int ru_maxrss;
    int ru_ixrss;                    /* integral shared text memory size */
    int ru_idrss;                    /* integral unshared data size */
    int ru_isrss;                    /* integral unshared stack size */
    int ru_minflt;                   /* page reclaims */
    int ru_majflt;                   /* page faults */
    int ru_nswap;                    /* swaps */
    int ru_inblock;                  /* block input operations */
    int ru_oublock;                  /* block output operations */
    int ru_msgsnd;                   /* messages sent */
    int ru_msrvcv;                   /* messages received */
    int ru_nsignals;                 /* signals received */
    int ru_nvcsw;                    /* voluntary context switches */
    int ru_nivcsw;                   /* involuntary context switches */
};
```

The fields are interpreted as follows:

ru_utime	the total amount of time spent executing in user mode. Time is given in seconds:microseconds.
ru_stime	the total amount of time spent in the system executing on behalf of the process(es). Time is given in seconds:microseconds.
ru_maxrss	the maximum resident set size utilized. Size is given in pages (the size of a page, in bytes, is given by the <i>getpagesize(2)</i> system call).
ru_ixrss	an “integral” value indicating the amount of memory used by the text segment which was also shared among other processes. This value is expressed in units of pages * clock ticks (1 tick = 1/50 second). The value is calculated by summing the number of shared memory pages in use each time the internal system clock ticks, and then averaging over 1 second intervals.
ru_idrss	an integral value of the amount of unshared memory residing in the data segment of a process. The value is given in pages * clock ticks.
ru_isrss	an integral value of the amount of unshared memory residing in the stack segment of a process. The value is given in pages * clock ticks.
ru_minflt	the number of page faults serviced without any I/O activity; here I/O activity is avoided by “reclaiming” a page frame from the list of pages awaiting reallocation.

<code>ru_majflt</code>	the number of page faults serviced which required I/O activity.
<code>ru_nswap</code>	the number of times a process was “swapped” out of main memory.
<code>ru_inblock</code>	the number of times the file system had to perform input.
<code>ru_outblock</code>	the number of times the file system had to perform output.
<code>ru_msgsnd</code>	the number of messages sent over sockets.
<code>ru_msgrcv</code>	the number of messages received from sockets.
<code>ru_nsignals</code>	the number of signals delivered.
<code>ru_nvcsw</code>	the number of times a context switch resulted due to a process voluntarily giving up the processor before its time slice was completed (usually to await availability of a resource).
<code>ru_nivcsw</code>	the number of times a context switch resulted due to a higher priority process becoming runnable or because the current process exceeded its time slice.

NOTES

The numbers `ru_inblock` and `ru_outblock` account only for real I/O; data supplied by the caching mechanism is charged only to the first process to read or write the data.

ERRORS

getrusage will fail if:

<code>EINVAL</code>	The <i>who</i> parameter is not a valid value.
<code>EFAULT</code>	The address specified by the <i>rusage</i> argument is not in a valid portion of the process's address space.

SEE ALSO

`gettimeofday(2)`, `wait(2)`

BUGS

There is no way to obtain information about a child process which has not yet terminated.

NAME

`getsockname` – get socket name

SYNOPSIS

```
getsockname(s, name, namelen)  
int s;  
struct sockaddr *name;  
int *namelen;
```

DESCRIPTION

Getsockname returns the current *name* for the specified socket. The *namelen* parameter should be initialized to indicate the amount of space pointed to by *name*. On return it contains the actual size of the name returned (in bytes).

DIAGNOSTICS

A 0 is returned if the call succeeds, -1 if it fails.

ERRORS

The call succeeds unless:

EBADF	The argument <i>s</i> is not a valid descriptor.
ENOTSOCK	The argument <i>s</i> is a file, not a socket.
ENOBUFS	Insufficient resources were available in the system to perform the operation.
EFAULT	The <i>name</i> parameter points to memory not in a valid part of the process address space.

SEE ALSO

`bind(2)`, `socket(2)`, `getpeername(2)`

BUGS

Names bound to sockets in the UNIX domain are inaccessible; *getsockname* returns a zero length name.

NAME

getsockopt, setsockopt – get and set options on sockets

SYNOPSIS

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

getsockopt(s, level, optname, optval, optlen)
int s, level, optname;
char *optval;
int *optlen;

setsockopt(s, level, optname, optval, optlen)
int s, level, optname;
char *optval;
int optlen;
```

DESCRIPTION

getsockopt and *setsockopt* manipulate *options* associated with a socket. Options may exist at multiple protocol levels; they are always present at the uppermost “socket” level.

When manipulating socket options the level at which the option resides and the name of the option must be specified. To manipulate options at the “socket” level, *level* is specified as SOL_SOCKET. To manipulate options at any other level the protocol number of the appropriate protocol controlling the option is supplied. For example, to indicate an option is to be interpreted by the TCP protocol, *level* should be set to the protocol number of TCP; see *getprotoent*(3N).

The parameters *optval* and *optlen* are used to access option values for *setsockopt*. For *getsockopt* they identify a buffer in which the value for the requested option(s) are to be returned. For *getsockopt*, *optlen* is a value-result parameter, initially containing the size of the buffer pointed to by *optval*, and modified on return to indicate the actual size of the value returned. If no option value is to be supplied or returned, *optval* may be supplied as 0.

optname and any specified options are passed uninterpreted to the appropriate protocol module for interpretation. The include file *<sys/socket.h>* contains definitions for “socket” level options; see *socket*(2). Options at other protocol levels vary in format and name, consult the appropriate entries in (4P).

RETURN VALUE

A 0 is returned if the call succeeds, -1 if it fails.

ERRORS

The call succeeds unless:

EBADF	The argument <i>s</i> is not a valid descriptor.
ENOTSOCK	The argument <i>s</i> is a file, not a socket.
ENOPROTOOPT	The option is unknown.
EFAULT	The address pointed to by <i>optval</i> is not in a valid part of the process address space. For <i>getsockopt</i> , this error may also be returned if <i>optlen</i> is not in a valid part of the process address space.

SEE ALSO

socket(2), *getprotoent*(3N)

NAME

gettimeofday, settimeofday – get/set date and time

SYNOPSIS

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

gettimeofday(tp, tzp)
struct timeval *tp;
struct timezone *tzp;

settimeofday(tp, tzp)
struct timeval *tp;
struct timezone *tzp;
```

DESCRIPTION

The system's notion of the current Greenwich time and the current time zone is obtained with the *gettimeofday* call, and set with the *settimeofday* call. The time is expressed in seconds and microseconds since midnight (0 hour), January 1, 1970. The resolution of the system clock is hardware dependent, and the time may be updated continuously or in "ticks."

The structures pointed to by *tp* and *tzp* are defined in *<sys/time.h>* as:

```
struct timeval {
    long    tv_sec;        /* seconds since Jan. 1, 1970 */
    long    tv_usec;      /* and microseconds */
};

struct timezone {
    int     tz_minuteswest; /* of Greenwich */
    int     tz_dsttime;     /* type of dst correction to apply */
};
```

The *timezone* structure indicates the local time zone (measured in minutes of time westward from Greenwich), and a flag that, if nonzero, indicates that Daylight Saving time applies locally during the appropriate part of the year.

If *tzp* is a zero pointer, the timezone information is not returned or set.

Only the super-user may set the time of day or time zone.

RETURN

A 0 return value indicates that the call succeeded. A -1 return value indicates an error occurred, and in this case an error code is stored into the global variable *errno*.

ERRORS

The following error codes may be set in *errno*:

EFAULT	An argument address referenced invalid memory.
EPERM	A user other than the super-user attempted to set the time.

SEE ALSO

date(1), adjtime(2), ctime(3)

BUGS

Time is never correct enough to believe the microsecond values. There should a mechanism by which, at least, local clusters of systems might synchronize their clocks to millisecond granularity.

NAME

`getuid`, `geteuid` – get user identity

SYNOPSIS

`uid = getuid()`

`int uid;`

`eid = geteuid()`

`int eid;`

DESCRIPTION

Getuid returns the real user ID of the current process, *geteuid* the effective user ID.

The real user ID identifies the person who is logged in. The effective user ID gives the process additional permissions during execution of “set-user-ID” mode processes, which use *getuid* to determine the real-user-id of the process that invoked them.

SEE ALSO

`getgid(2)`, `setreuid(2)`

NAME

`ioctl` – control device

SYNOPSIS

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

ioctl(d, request, argp)
int d, request;
char *argp;
```

DESCRIPTION

ioctl performs a variety of functions on open descriptors. In particular, many operating characteristics of character special files (e.g. terminals) may be controlled with *ioctl* requests. The writeups of various devices in section 4 discuss how *ioctl* applies to them.

An *ioctl request* has encoded in it whether the argument is an “in” parameter or “out” parameter, and the size of the argument *argp* in bytes. Macros and defines used in specifying an *ioctl request* are located in the file `<sys/ioctl.h>`.

RETURN VALUE

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

If no error has occurred (using a STANDARD device driver), a value of `0` is returned.

ERRORS

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

EBADF	<i>D</i> is not a valid descriptor.
ENOTTY	<i>D</i> is not associated with a character special device.
ENOTTY	The specified request does not apply to the kind of object that the descriptor <i>d</i> references.
EINVAL	<i>Request</i> or <i>argp</i> is not valid.

SEE ALSO

`execve(2)`, `fcntl(2)`, `mtio(4)`, `tty(4)`

NAME

kill – send signal to a process

SYNOPSIS

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

DESCRIPTION

Kill sends the signal *sig* to a process, specified by the process number *pid*. *Sig* may be one of the signals specified in *sigvec(2)*, or it may be 0, in which case error checking is performed but no signal is actually sent. This can be used to check the validity of *pid*.

The sending and receiving processes must have the same effective user ID, otherwise this call is restricted to the super-user. A single exception is the signal SIGCONT, which may always be sent to any descendant of the current process.

If the process number is 0, the signal is sent to all processes in the sender's process group; this is a variant of *killpg(2)*.

If the process number is -1 and the user is the super-user, the signal is broadcast universally except to system processes and the process sending the signal.

Processes may send signals to themselves.

RETURN VALUE

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

ERRORS

Kill will fail and no signal will be sent if any of the following occur:

EINVAL *Sig* is not a valid signal number.

ESRCH No process can be found corresponding to that specified by *pid*.

EPERM The sending process is not the super-user and its effective user id does not match the effective user-id of the receiving process.

SEE ALSO

getpid(2), *getpgrp(2V)*, *killpg(2)*, *sigvec(2)*

NAME

killpg – send signal to a process group

SYNOPSIS

```
killpg(pgrp, sig)  
int pgrp, sig;
```

DESCRIPTION

Killpg sends the signal *sig* to the process group *pgrp*. See *sigvec(2)* for a list of signals.

The sending process and members of the process group must have the same effective user ID, or the sender must have an effective user ID of super-user. As a single special case the continue signal SIGCONT may be sent to any process that is a descendant of the current process.

RETURN VALUE

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

ERRORS

Killpg will fail and no signal will be sent if any of the following occur:

EINVAL	<i>Sig</i> is not a valid signal number.
ESRCH	No process were found in the specified process group.
EPERM	The sending process is not the super-user and one or more of the target processes has an effective user ID different from that of the sending process.

SEE ALSO

kill(2), getpgrp(2V), sigvec(2)

NAME

link – make a hard link to a file

SYNOPSIS

```
link(name1, name2)
char *name1, *name2;
```

DESCRIPTION

name1 points to a path name naming an existing file. *name2* points to a path name naming a new directory entry to be created. A hard link to the first file is created; the link has the name pointed to by *name2*. The file named by *name1* must exist.

With hard links, both files must be on the same file system. Unless the caller is the super-user, the file named by *name1* must not be a directory. Both the old and the new *link* share equal access and rights to the underlying object.

RETURN VALUE

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

ERRORS

link will fail and no link will be created if one or more of the following are true:

ENOTDIR	A component of the path prefix of <i>name1</i> or <i>name2</i> is not a directory.
EINVAL	<i>name1</i> or <i>name2</i> contains a byte with the high-order bit set.
ENAMETOOLONG	The length of a component of <i>name1</i> or <i>name2</i> exceeds 255 characters, or the length of <i>name1</i> or <i>name2</i> exceeds 1023 characters.
ENOENT	A component of the path prefix of <i>name1</i> or <i>name2</i> does not exist.
EACCES	Search permission is denied for a component of the path prefix of <i>name1</i> or <i>name2</i> .
EACCES	The requested link requires writing in a directory for which write permission is denied.
ELOOP	Too many symbolic links were encountered in translating <i>name1</i> or <i>name2</i> .
ENOENT	The file referred to by <i>name1</i> does not exist.
EEXIST	The link referred to by <i>name2</i> does exist.
EPERM	The file named by <i>name1</i> is a directory and the effective user ID is not super-user.
EXDEV	The link named by <i>name2</i> and the file named by <i>name1</i> are on different file systems.
ENOSPC	The directory in which the entry for the new link is being placed cannot be extended because there is no space left on the file system containing the directory.
EDQUOT	The directory in which the entry for the new link is being placed cannot be extended because the user's quota of disk blocks on the file system containing the directory has been exhausted.
EIO	An I/O error occurred while reading from or writing to the file system to make the directory entry.
EROFS	The requested link requires writing in a directory on a read-only file system.
EFAULT	One of the path names specified is outside the process's allocated address space.

SEE ALSO

symlink(2), *unlink(2)*

NAME

listen – listen for connections on a socket

SYNOPSIS

listen(s, backlog)
int s, backlog;

DESCRIPTION

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

The *backlog* parameter defines the maximum length the queue of pending connections may grow to. If a connection request arrives with the queue full the client will receive an error with an indication of ECONNREFUSED.

RETURN VALUE

A 0 return value indicates success; -1 indicates an error.

ERRORS

The call fails if:

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

SEE ALSO

accept(2), *connect(2)*, *socket(2)*

BUGS

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

NAME

lseek, *tell* – move read/write pointer

SYNOPSIS

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

pos = lseek(d, offset, whence)
long pos;
int d;
long offset;
int whence;
```

DESCRIPTION

The descriptor *d* refers to a file or device open for reading and/or writing. *lseek* sets the file pointer of *d* as follows:

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

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

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

Upon successful completion, the resulting pointer location as measured in bytes from beginning of the file is returned. Some devices are incapable of seeking. The value of the pointer associated with such a device is undefined.

The obsolete function *tell(fildes)* is identical to *lseek(fildes, 0L, L_INCR)*.

NOTES

Seeking far beyond the end of a file, then writing, creates a gap or “hole”, which occupies no physical space and reads as zeros.

RETURN VALUE

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

ERRORS

lseek will fail and the file pointer will remain unchanged if:

EBADF	<i>Fildes</i> is not an open file descriptor.
ESPIPE	<i>Fildes</i> is associated with a pipe or a socket.
EINVAL	<i>whence</i> is not a proper value.

SEE ALSO

dup(2), *open(2V)*

NAME

`mkdir` – make a directory file

SYNOPSIS

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

DESCRIPTION

`mkdir` creates a new directory file with name *path*. The mode of the new file is initialized from *mode*. The protection part of the mode is modified by the process's mode mask; see `umask(2)`.

The directory's owner ID is set to the process's effective user ID. The directory's group ID is set to that of the parent directory in which it is created.

The low-order 9 bits of mode are modified by the process's file mode creation mask: all bits set in the process's file mode creation mask are cleared. See `umask(2)`.

RETURN VALUE

A 0 return value indicates success. A -1 return value indicates an error, and an error code is stored in *errno*.

ERRORS

`mkdir` will fail and no directory will be created if:

ENOTDIR	A component of the path prefix of <i>path</i> is not a directory.
EINVAL	<i>path</i> contains a byte with the high-order bit set.
ENAMETOOLONG	The length of a component of <i>path</i> exceeds 255 characters, or the length of <i>path</i> exceeds 1023 characters.
ENOENT	A component of the path prefix of <i>path</i> does not exist.
EACCES	Search permission is denied for a component of the path prefix of <i>path</i> .
ELOOP	Too many symbolic links were encountered in translating <i>path</i> .
EROFS	The file referred to by <i>path</i> resides on a read-only file system.
EEXIST	The file referred to by <i>path</i> exists.
ENOSPC	The directory in which the entry for the new file is being placed cannot be extended because there is no space left on the file system containing the directory.
ENOSPC	The new directory cannot be created because there is no space left on the file system which will contain the directory.
ENOSPC	There are no free inodes on the file system on which the file is being created.
EDQUOT	The directory in which the entry for the new file is being placed cannot be extended because the user's quota of disk blocks on the file system containing the directory has been exhausted.
EDQUOT	The new directory cannot be created because the user's quota of disk blocks on the file system which will contain the directory has been exhausted.
EDQUOT	The user's quota of inodes on the file system on which the file is being created has been exhausted.
EIO	An I/O error occurred while reading from or writing to the file system.
EFAULT	<i>Path</i> points outside the process's allocated address space.

SEE ALSO

chmod(2), stat(2), rmdir(2), umask(2)

NAME

mknod – make a special file

SYNOPSIS

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

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

DESCRIPTION

mknod creates a new file named by the path name pointed to by *path*. The mode of the new file (including file type bits) is initialized from *mode*. The values of the file type bits which are permitted are:

```
#define S_IFCHR      0020000    /* character special */
#define S_IFBLK      0060000    /* block special */
#define S_IFREG      0100000    /* regular */
#define S_IFIFO      0010000    /* FIFO special */
```

Values of *mode* other than those above are undefined and should not be used.

The protection part of the mode is modified by the process's mode mask (see *umask(2)*).

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

If mode indicates a block or character special file, *dev* is a configuration dependent specification of a character or block I/O device. If *mode* does not indicate a block special or character special device, *dev* is ignored.

mknod may be invoked only by the super-user for file types other than FIFO special.

RETURN VALUE

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

ERRORS

mknod fails and the file mode remains unchanged if:

ENOTDIR	A component of the path prefix of <i>path</i> is not a directory.
EINVAL	<i>path</i> contains a character with the high-order bit set.
ENAMETOOLONG	The length of a component of <i>path</i> exceeds 255 characters, or the length of <i>path</i> exceeds 1023 characters.
ENOENT	A component of the path prefix of <i>path</i> does not exist.
EACCES	Search permission is denied for a component of the path prefix of <i>path</i> .
ELOOP	Too many symbolic links were encountered in translating <i>path</i> .
EPERM	An attempt was made to create a file of type other than FIFO special and the process's effective user ID is not super-user.
EIO	An I/O error occurred while reading from or writing to the file system.
EISDIR	The specified <i>mode</i> would have created a directory.
ENOSPC	The directory in which the entry for the new file is being placed cannot be extended because there is no space left on the file system containing the directory.
ENOSPC	There are no free inodes on the file system on which the file is being created.
EDQUOT	The directory in which the entry for the new file is being placed cannot be extended

because the user's quota of disk blocks on the file system containing the directory has been exhausted.

- EDQUOT The user's quota of inodes on the file system on which the node is being created has been exhausted.
- EROFS The file referred to by *path* resides on a read-only file system.
- EEXIST The file referred to by *path* exists.
- EFAULT *path* points outside the process's allocated address space.

SEE ALSO

chmod(2), stat(2), umask(2)

NAME

`mmap`, `munmap` – map or unmap pages of memory

SYNOPSIS

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

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

```
mmap(addr, len, prot, share, fd, off)
```

```
caddr_t addr; int len, prot, share, fd; off_t off;
```

```
munmap(addr, len)
```

```
caddr_t addr; int len;
```

DESCRIPTION**Mmap**

mmap maps pages of memory from the memory device associated with the file *fd* into the address space of the calling process, one page at a time. Pages are mapped from the memory device, beginning at *off*, and into the caller's address space, beginning at *addr*, and continuing for *len* bytes. *fd* is a file descriptor obtained by opening the device from which to map pages. Only character-special devices are currently supported.

share specifies whether modifications made to mapped-in copies of pages are to be kept "private" or are to be "shared" with other references. Currently, it must be set to `MAP_SHARED`.

The parameter *prot* specifies the read/write accessibility of the mapped pages. The *addr* and *len* parameters, and the sum of the current position in *fd* and *off* parameters, must be multiples of the page size (found using the *getpagesize(2)* call). For this reason, local memory space beginning at *addr* should be allocated using *valloc(2)*, which supplies a buffer with proper page alignment.

When mapping an area of 128K or more, the kernel releases the swap area associated with it. Consequently, when the pages are unmapped, they are marked invalid; the next call to *valloc(2)* returns the invalid pages, and any attempt to refer to those pages results in a segmentation violation. To avoid this, do not *free(2)* such large areas; instead, call *valloc(2)* again without calling *free(2)*.

All pages are automatically unmapped when *fd* is closed. Specific pages can be unmapped explicitly using *munmap*.

mmap can sometimes be used to install memory-mapped devices without writing a device driver. However, this does not always work. In particular, devices that are *mmap*'ed into user space and then accessed by user programs will see those accesses in user mode. If the device contains registers that must be accessed in supervisor mode, *mmap* cannot be used to drive it. (See *Writing Device Drivers for the Sun Workstation* for more information.)

Munmap

munmap unmaps previously mapped pages starting at *addr* and continuing for *len* bytes. Unmapped pages refer, once again, to private pages within the caller's address space. Pages are initialized to zero, unless *len* is greater than or equal to 128K, in which case the pages are marked invalid.

RETURN VALUE

Each call returns 0 on success, -1 on failure.

ERRORS

Both calls fail when:

EINVAL The argument address or length is not a multiple of the page size as returned by *getpagesize(2)*, or the length is negative.

EINVAL The entire range of pages specified in the call is not part of data space.

In addition *mmap* fails when:

EINVAL The specified *fd* does not refer to a character special device which supports mapping (e.g. a frame buffer).

EINVAL The specified *fd* is not open for reading and read access is requested, or not open for writing when write access is requested.

EINVAL The sharing mode was not specified as `MAP_SHARED`.

SEE ALSO

`getpagesize(2)`, `munmap(2)`, `close(2)`

BUGS

The kernel may panic when more than 128k of memory has been unmapped with `munmap(1)` and `mmap` is subsequently called with an incorrect *length* value.

If 128K of memory, or more, is unmapped as a result of closing *fd*, the resulting invalid pages cannot be reclaimed within the life of the calling process.

NAME

mount – mount file system

SYNOPSIS

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

mount(type, dir, flags, data)
int type;
char *dir;
int flags;
caddr_t data;
```

DESCRIPTION

mount attaches a file system to a directory. After a successful return, references to directory *dir* will refer to the root directory on the newly mounted file system. *dir* is a pointer to a null-terminated string containing a path name. *dir* must exist already, and must be a directory. Its old contents are inaccessible while the file system is mounted.

mount may be invoked only by the super-user.

The *flags* argument determines whether the file system can be written on, and if set-uid execution is allowed. Physically write-protected and magnetic tape file systems must be mounted read-only or errors will occur when access times are updated, whether or not any explicit write is attempted.

type indicates the type of the filesystem. It must be one of the types defined in *mount.h*. *data* is a pointer to a structure which contains the type specific arguments to mount. Below is a list of the filesystem types supported and the type specific arguments to each:

MOUNT_UFS

```
struct ufs_args {
    char *fspec;      /* Block special file to mount */
};
```

MOUNT_NFS

```
#include <nfs/nfs.h>
#include <netinet/in.h>
struct nfs_args {
    struct sockaddr_in *addr; /* file server address */
    fhandle_t *fh;          /* File handle to be mounted */
    int flags;              /* flags */
    int wsize;              /* write size in bytes */
    int rsize;              /* read size in bytes */
    int timeo;              /* initial timeout in .1 secs */
    int retrans;            /* times to retry send */
};
```

RETURN VALUE

mount returns 0 if the action occurred, and -1 if *fspec* is inaccessible or not an appropriate file, if *name* does not exist, if *fspec* is already mounted, if *dir* is in use, or if there are already too many file systems mounted.

ERRORS

mount fails when one of the following occurs:

EPERM	The caller is not the super-user.
ENOTBLK	<i>fspec</i> is not a block device.
ENXIO	The major device number of <i>fspec</i> is out of range (this indicates no device driver exists for the associated hardware).
EBUSY	<i>dir</i> is not a directory, or another process currently holds a reference to it.

EBUSY	No space remains in the mount table.
EBUSY	The super block for the file system had a bad magic number or an out of range block size.
EBUSY	Not enough memory was available to read the cylinder group information for the file system.
EIO	An I/O error occurred while reading the super block or cylinder group information.
ENOTDIR	A component of the path prefix in <i>fspec</i> or <i>dir</i> is not a directory.
EINVAL	The path name of <i>fspec</i> or <i>dir</i> contains a character with the high-order bit set.
ENAMETOOLONG	The length of a component of the path name of <i>fspec</i> or <i>dir</i> exceeds 255 characters, or the length of the entire path name of <i>fspec</i> or <i>dir</i> exceeds 1023 characters.
ENOENT	<i>fspec</i> or <i>dir</i> does not exist.
ENOTDIR	The file named by <i>dir</i> is not a directory.
EACCES	Search permission is denied for a component of the path prefix of <i>fspec</i> or <i>dir</i> .
EFAULT	<i>fspec</i> or <i>dir</i> points outside the process's allocated address space.
ELOOP	Too many symbolic links were encountered in translating the path name of <i>fspec</i> or <i>dir</i> .
EIO	An I/O error occurred while reading from or writing to the file system.

SEE ALSO

umount(2), mount(8)

BUGS

The error codes are in a state of disarray; too many errors appear to the caller as one value.

NAME

`msgctl` – message control operations

SYNOPSIS

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

int msgctl (msqid, cmd, buf)
int msqid, cmd;
struct msqid_ds *buf;
```

DESCRIPTION

`msgctl` provides a variety of message control operations as specified by `cmd`. The following `cmds` are available:

IPC_STAT Place the current value of each member of the data structure associated with `msqid` into the structure pointed to by `buf`. The contents of this structure are defined in `intro(2)`. {READ}

IPC_SET Set the value of the following members of the data structure associated with `msqid` to the corresponding value found in the structure pointed to by `buf`:

```
msg_perm.uid
msg_perm.gid
msg_perm.mode /* only low 9 bits */
msg_qbytes
```

This `cmd` can only be executed by a process that has an effective user ID equal to either that of super user or to the value of `msg_perm.uid` in the data structure associated with `msqid`. Only super user can raise the value of `msg_qbytes`.

IPC_RMID Remove the message queue identifier specified by `msqid` from the system and destroy the message queue and data structure associated with it. This `cmd` can only be executed by a process that has an effective user ID equal to either that of super user or to the value of `msg_perm.uid` in the data structure associated with `msqid`.

ERRORS

`msgctl` will fail if:

EINVAL	<code>msqid</code> is not a valid message queue identifier.
EINVAL	<code>cmd</code> is not a valid command.
EACCES	<code>cmd</code> is equal to <code>IPC_STAT</code> and {READ} operation permission is denied to the calling process (see <code>intro(2)</code>).
EPERM	<code>cmd</code> is equal to <code>IPC_RMID</code> or <code>IPC_SET</code> . The effective user ID of the calling process is not equal to that of super user and it is not equal to the value of <code>msg_perm.uid</code> in the data structure associated with <code>msqid</code> .
EPERM	<code>cmd</code> is equal to <code>IPC_SET</code> , an attempt is being made to increase to the value of <code>msg_qbytes</code> , and the effective user ID of the calling process is not equal to that of super user.
EFAULT	<code>Buf</code> points to an illegal address.

RETURN VALUE

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

SEE ALSO

`intro(2)`, `msgget(2)`, `msgop(2)`

NAME

`msgget` – get message queue

SYNOPSIS

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

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

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

```
int msgget (key, msgflg)
```

```
key_t key;
```

```
int msgflg;
```

DESCRIPTION

`msgget` returns the message queue identifier associated with *key*.

A message queue identifier and associated message queue and data structure (see *intro(2)*) are created for *key* if one of the following are true:

- *key* is equal to `IPC_PRIVATE`.
- *key* does not already have a message queue identifier associated with it, and (`msgflg & IPC_CREAT`) is “true”.

Upon creation, the data structure associated with the new message queue identifier is initialized as follows:

`msg_perm.cuid`, `msg_perm.uid`, `msg_perm.cgid`, and `msg_perm.gid` are set equal to the effective user ID and effective group ID, respectively, of the calling process.

The low-order 9 bits of `msg_perm.mode` are set equal to the low-order 9 bits of `msgflg`.

`msg_qnum`, `msg_lspid`, `msg_lrpid`, `msg_stime`, and `msg_rtime` are set equal to 0.

`msg_ctime` is set equal to the current time.

`msg_qbytes` is set equal to the system limit.

ERRORS

`msgget` will fail if one or more of the following are true:

- | | |
|---------------|--|
| EACCES | A message queue identifier exists for <i>key</i> , but operation permission (see <i>intro(2)</i>) as specified by the low-order 9 bits of <code>msgflg</code> would not be granted. |
| ENOENT | A message queue identifier does not exist for <i>key</i> and (<code>msgflg & IPC_CREAT</code>) is “false”. |
| ENOSPC | A message queue identifier is to be created but the system-imposed limit on the maximum number of allowed message queue identifiers system wide would be exceeded. |
| EEXIST | A message queue identifier exists for <i>key</i> but (<code>msgflg & IPC_CREAT</code>) & (<code>msgflg & IPC_EXCL</code>) is “true”. |

RETURN VALUE

Upon successful completion, a non-negative integer, namely a message queue identifier, is returned. Otherwise, a value of `-1` is returned and `errno` is set to indicate the error.

SEE ALSO

intro(2), *msgctl(2)*, *msgop(2)*

NAME

msgop, msgsnd, msgrcv – message operations

SYNOPSIS

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

int msgsnd (msqid, msgp, msgsz, msgflg)
int msqid;
struct msgbuf *msgp;
int msgsz, msgflg;

int msgrcv (msqid, msgp, msgsz, msgtyp, msgflg)
int msqid;
struct msgbuf *msgp;
int msgsz;
long msgtyp;
int msgflg;
```

DESCRIPTION

msgsnd is used to send a message to the queue associated with the message queue identifier specified by *msqid*. {WRITE} *msgp* points to a structure containing the message. This structure is composed of the following members:

```
long    mtype;    /* message type */
char    mtext[]; /* message text */
```

mtype is a positive integer that can be used by the receiving process for message selection (see *msgrcv* below). *mtext* is any text of length *msgsz* bytes. *msgsz* can range from 0 to a system-imposed maximum.

msgflg specifies the action to be taken if one or more of the following are true:

The number of bytes already on the queue is equal to *msg_qbytes* (see *intro(2)*).

The total number of messages on all queues system-wide is equal to the system-imposed limit.

These actions are as follows:

If (*msgflg* & IPC_NOWAIT) is “true”, the message will not be sent and the calling process will return immediately.

If (*msgflg* & IPC_NOWAIT) is “false”, the calling process will suspend execution until one of the following occurs:

The condition responsible for the suspension no longer exists, in which case the message is sent.

msqid is removed from the system (see *msgctl(2)*). When this occurs, *errno* is set equal to EIDRM, and a value of -1 is returned.

The calling process receives a signal that is to be caught. In this case the message is not sent and the calling process resumes execution in the manner prescribed in *signal(2)*.

msgsnd will fail and no message will be sent if one or more of the following are true:

EINVAL	<i>msqid</i> is not a valid message queue identifier.
EACCES	Operation permission is denied to the calling process (see <i>intro(2)</i>).
EINVAL	<i>mtype</i> is less than 1.
EAGAIN	The message cannot be sent for one of the reasons cited above and (<i>msgflg</i> & IPC_NOWAIT) is “true”.
EINVAL	<i>msgsz</i> is less than zero or greater than the system-imposed limit.

EFAULT *msgp* points to an illegal address.

Upon successful completion, the following actions are taken with respect to the data structure associated with *msqid* (see intro (2)).

msg_qnum is incremented by 1.

msg_lspid is set equal to the process ID of the calling process.

msg_stime is set equal to the current time.

msgrcv reads a message from the queue associated with the message queue identifier specified by *msqid* and places it in the structure pointed to by *msgp*. {READ} This structure is composed of the following members:

```
long    mtype;      /* message type */
char    mtext[];   /* message text */
```

mtype is the received message's type as specified by the sending process. *mtext* is the text of the message. *msgsz* specifies the size in bytes of *mtext*. The received message is truncated to *msgsz* bytes if it is larger than *msgsz* and (*msgflg* & MSG_NOERROR) is "true". The truncated part of the message is lost and no indication of the truncation is given to the calling process.

msgtyp specifies the type of message requested as follows:

If *msgtyp* is equal to 0, the first message on the queue is received.

If *msgtyp* is greater than 0, the first message of type *msgtyp* is received.

If *msgtyp* is less than 0, the first message of the lowest type that is less than or equal to the absolute value of *msgtyp* is received.

msgflg specifies the action to be taken if a message of the desired type is not on the queue. These are as follows:

If (*msgflg* & IPC_NOWAIT) is "true", the calling process will return immediately with a return value of -1 and *errno* set to ENOMSG.

If (*msgflg* & IPC_NOWAIT) is "false", the calling process will suspend execution until one of the following occurs:

A message of the desired type is placed on the queue.

msqid is removed from the system. When this occurs, *errno* is set equal to EIDRM, and a value of -1 is returned.

The calling process receives a signal that is to be caught. In this case a message is not received and the calling process resumes execution in the manner prescribed in *signal(2)*.

msgrcv will fail and no message will be received if one or more of the following are true:

EINVAL *msqid* is not a valid message queue identifier.

EACCES Operation permission is denied to the calling process.

EINVAL *msgsz* is less than 0.

[E2BIG] *mtext* is greater than *msgsz* and (*msgflg* & MSG_NOERROR) is "false".

ENOMSG The queue does not contain a message of the desired type and (*msgtyp* & IPC_NOWAIT) is "true".

EFAULT *msgp* points to an illegal address.

Upon successful completion, the following actions are taken with respect to the data structure associated with *msqid* (see intro (2)).

msg_qnum is decremented by 1.

msg_lrpid is set equal to the process ID of the calling process.

msg_rtime is set equal to the current time.

RETURN VALUES

If *msgsnd* or *msgrcv* return due to the receipt of a signal, a value of -1 is returned to the calling process and *errno* is set to EINTR. If they return due to removal of *msgid* from the system, a value of -1 is returned and *errno* is set to EIDRM.

Upon successful completion, the return value is as follows:

msgsnd returns a value of 0.

msgrcv returns a value equal to the number of bytes actually placed into *mtext*.

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

SEE ALSO

intro(2), msgctl(2), msgget(2), signal(2).

NAME

nfssvc, *async_daemon* – NFS daemons

SYNOPSIS

***nfssvc*(sock)**

int sock;

***async_daemon*()**

DESCRIPTION

Nfssvc starts an NFS daemon listening on socket *sock*. The socket must be AF_INET, and SOCK_DGRAM (protocol UDP/IP). The system call will return only if the process is killed.

Async_daemon implements the NFS daemon that handles asynchronous I/O for an NFS client. The system call never returns.

BUGS

These two system calls allow kernel processes to have user context.

SEE ALSO

mountd(8)

NAME

`open` – open or create a file for reading or writing

SYNOPSIS

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

```
int open(path, flags [ , mode ] )
```

```
char *path;
```

```
int flags, mode;
```

DESCRIPTION

path points to the pathname of a file. *open* opens the named file for reading and/or writing, as specified by the *flags* argument, and returns a descriptor for that file. The *flags* argument may indicate the file is to be created if it does not already exist (by specifying the `O_CREAT` flag), in which case the file is created with mode *mode* as described in *chmod(2)* and modified by the process' umask value (see *umask(2)*). If the path is a null string, the kernel maps this null pathname to `.`, the current directory. *flags* values are constructed by ORing flags from the following list (only one of the first three flags below may be used):

`O_RDONLY` Open for reading only.

`O_WRONLY` Open for writing only.

`O_RDWR` Open for reading and writing.

`O_NDELAY` When opening a FIFO with `O_RDONLY` or `O_WRONLY` set:

If `O_NDELAY` is set:

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

If `O_NDELAY` is clear:

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

When opening a file associated with a communication line:

If `O_NDELAY` is set:

The *open* will return without waiting for carrier. The first time the process attempts to perform I/O on the open file it will block (not currently implemented).

If `O_NDELAY` is clear:

The *open* will block until carrier is present.

`O_APPEND` If set, the file pointer will be set to the end of the file prior to each write.

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

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

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

`O_TRUNC` If the file exists, its length is truncated to 0 and the mode and owner are unchanged.

`O_EXCL` If `O_EXCL` and `O_CREAT` are set, *open* will fail if the file exists. This can be used to implement a simple exclusive access locking mechanism. If `O_EXCL` is set and the last component of the pathname is a symbolic link, the *open* will fail even if the symbolic link points to a non-existent name.

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

The new descriptor is set to remain open across *execve* system calls; see *close(2)* and *fcntl(2)*.

There is a system enforced limit on the number of open file descriptors per process, whose value is returned by the *getdtablesize(2)* call.

SYSTEM V DESCRIPTION

If the *O_NDELAY* flag is set on an *open*, that flag is set for that file descriptor (see *fcntl*) and may affect subsequent reads and writes. See *read(2V)* and *write(2V)*.

RETURN VALUE

The value *-1* is returned if an error occurs, and external variable *errno* is set to indicate the cause of the error. Otherwise a non-negative numbered file descriptor for the new open file is returned.

ERRORS

Open fails if:

ENOTDIR	A component of the path prefix of <i>path</i> is not a directory.
EINVAL	<i>path</i> contains a character with the high-order bit set.
ENAMETOOLONG	The length of a component of <i>path</i> exceeds 255 characters, or the length of <i>path</i> exceeds 1023 characters.
ENOENT	<i>O_CREAT</i> is not set and the named file does not exist.
ENOENT	A component of the path prefix of <i>path</i> does not exist.
ELOOP	Too many symbolic links were encountered in translating <i>path</i> .
EACCES	Search permission is denied for a component of the path prefix of <i>path</i> .
EACCES	The required permissions (for reading and/or writing) are denied for the file named by <i>path</i> .
EACCES	The file referred to by <i>path</i> does not exist, <i>O_CREAT</i> is specified, and the directory in which it is to be created does not permit writing.
EISDIR	The named file is a directory, and the arguments specify it is to be opened for writing.
ENXIO	<i>O_NDELAY</i> is set, the named file is a FIFO, <i>O_WRONLY</i> is set, and no process has the file open for reading.
EMFILE	The system limit for open file descriptors per process has already been reached.
ENFILE	The system file table is full.
ENOSPC	The file does not exist, <i>O_CREAT</i> is specified, and the directory in which the entry for the new file is being placed cannot be extended because there is no space left on the file system containing the directory.
ENOSPC	The file does not exist, <i>O_CREAT</i> is specified, and there are no free inodes on the file system on which the file is being created.
EDQUOT	The file does not exist, <i>O_CREAT</i> is specified, and the directory in which the entry for the new file is being placed cannot be extended because the user's quota of disk blocks on the file system containing the directory has been exhausted.
EDQUOT	The file does not exist, <i>O_CREAT</i> is specified, and the user's quota of inodes on the file system on which the file is being created has been exhausted.
EROFS	The named file does not exist, <i>O_CREAT</i> is specified, and the file system on which it is to be created is a read-only file system.
EROFS	The named file resides on a read-only file system, and the file is to be opened for writing.
ENXIO	The file is a character special or block special file, and the associated device does not

	exist.
EINTR	A signal was caught during the <i>open</i> system call.
ETXTBSY	The file is a pure procedure (shared text) file that is being executed and the <i>open</i> call requests write access.
EIO	An I/O error occurred while reading from or writing to the file system.
EFAULT	<i>path</i> points outside the process's allocated address space.
EEXIST	O_EXCL and O_CREAT were both specified and the file exists.
EOPNOTSUPP	An attempt was made to open a socket (not currently implemented).

SEE ALSO

chmod(2), close(2), dup(2), fcntl(2), lseek(2), read(2V), write(2V), umask(2)

NAME

pipe – create an interprocess communication channel

SYNOPSIS

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

DESCRIPTION

The *pipe* system call creates an I/O mechanism called a pipe and returns two file descriptors, *fildes*[0] and *fildes*[1]. *fildes*[0] is opened for reading and *fildes*[1] is opened for writing. When the pipe is written using the descriptor *fildes*[1] up to 4096 bytes of data are buffered before the writing process is blocked. A read only file descriptor *fildes*[0] accesses the data written to *fildes*[1] on a first-in-first-out (FIFO) basis.

It is assumed that after the pipe has been set up, two (or more) cooperating processes (created by subsequent *fork* calls) will pass data through the pipe with *read* and *write* calls.

The shell has a syntax to set up a linear array of processes connected by pipes.

Read calls on an empty pipe (no buffered data) with only one end (all write file descriptors closed) returns an end-of-file.

Pipes are really a special case of the *socketpair*(2) call and, in fact, are implemented as such in the system.

A signal is generated if a write on a pipe with only one end is attempted.

RETURN VALUE

The function value zero is returned if the pipe was created; -1 if an error occurred.

ERRORS

The *pipe* call will fail if:

EMFILE	Too many descriptors are active.
ENFILE	The system file table is full.
EFAULT	The <i>fildes</i> buffer is in an invalid area of the process's address space.

SEE ALSO

sh(1), read(2V), write(2V), fork(2), socketpair(2)

BUGS

Should more than 4096 bytes be necessary in any pipe among a loop of processes, deadlock will occur.

NAME

profil – execution time profile

SYNOPSIS

```
profil(buff, bufsiz, offset, scale)
char *buff;
int bufsiz, offset, scale;
```

DESCRIPTION

Buff points to an area of core whose length (in bytes) is given by *bufsiz*. After this call, the user's program counter (*pc*) is examined each clock tick (20 milliseconds); *offset* is subtracted from it, and the result multiplied by *scale*. If the resulting number corresponds to a word inside *buff*, that word is incremented.

The scale is interpreted as an unsigned, fixed-point fraction with binary point at the left: 0x10000 gives a 1-1 mapping of *pc*'s to words in *buff*; 0x8000 maps each pair of instruction words together. 0x2 maps all instructions onto the beginning of *buff* (producing a non-interrupting core clock).

Profiling is turned off by giving a *scale* of 0 or 1. It is rendered ineffective by giving a *bufsiz* of 0. Profiling is turned off when an *execve* is executed, but remains on in child and parent both after a *fork*. Profiling is turned off if an update in *buff* would cause a memory fault.

RETURN VALUE

A 0, indicating success, is always returned.

SEE ALSO

gprof(1), setitimer(2), monitor(3)

NAME

`ptrace` – process trace

SYNOPSIS

```
#include <signal.h>
#include <sys/ptrace.h>
#include <sys/wait.h>

ptrace(request, pid, addr, data [ , addr2 ] )
enum ptracereq request;
int pid;
char *addr;
int data;
char *addr2;
```

DESCRIPTION

`ptrace` provides a means by which a process may control the execution of another process, and examine and change its core image. Its primary use is for the implementation of breakpoint debugging. There are five arguments whose interpretation depends on the *request* argument. Generally, *pid* is the process ID of the traced process. A process being traced behaves normally until it encounters some signal whether internally generated like ‘illegal instruction’ or externally generated like ‘interrupt’. See `sigvec(2)` for the list. Then the traced process enters a stopped state and the tracing process is notified via `wait(2)`. When the traced process is in the stopped state, its core image can be examined and modified using `ptrace`. If desired, another `ptrace` request can then cause the traced process either to terminate or to continue, possibly ignoring the signal.

Note that several different values of the *request* argument can make `ptrace` return data values — since `-1` is a possibly legitimate value, to differentiate between `-1` as a legitimate value and `-1` as an error code, you should clear the *errno* global error code before doing a `ptrace` call, and then check the value of *errno* afterwards.

The value of the *request* argument determines the precise action of the call:

PTRACE_TRACEME

This request is the only one used by the traced process; it declares that the process is to be traced by its parent. All the other arguments are ignored. Peculiar results will ensue if the parent does not expect to trace the child.

PTRACE_PEEKTEXT, PTRACE_PEEKDATA

The word in the traced process’s address space at *addr* is returned. If the instruction and data spaces are separate (for example, historically on a PDP-11), request `PTRACE_PEEKTEXT` indicates instruction space while `PTRACE_PEEKDATA` indicates data space. Otherwise, either request may be used, with equal results. *addr* must be even, the child must be stopped and the input *data* and *addr2* are ignored.

PTRACE_PEEKUSER

The word of the system’s per-process data area corresponding to *addr* is returned. *addr* must be a valid offset within the kernel’s per-process data pages. This space contains the registers and other information about the process; its layout corresponds to the *user* structure in the system (see `<sys/user.h>`).

PTRACE_POKETEXT, PTRACE_POKEDATA

The given *data* is written at the word in the process’s address space corresponding to *addr*, which must be even. No useful value is returned. If the instruction and data spaces are separate, request `PTRACE_PEEKTEXT` indicates instruction space while `PTRACE_PEEKDATA` indicates data space. The `PTRACE_POKETEXT` request must be used to write into a process’s text space even if the instruction and data spaces are not separate. Attempts to write in a pure text space fail if another process is executing the same file.

PTRACE_POKEUSER

The process's system data is written, as it is read with request `PTRACE_PEEKUSER`. Only a few locations can be written in this way: the general registers, the floating point status and registers, and certain bits of the processor status word.

`PTRACE_CONT`

The *data* argument is taken as a signal number and the child's execution continues at location *addr* as if it had incurred that signal. Normally the signal number will be either 0 to indicate that the signal that caused the stop should be ignored, or that value fetched out of the process's image indicating which signal caused the stop. If *addr* is `(int *)1` then execution continues from where it stopped.

`PTRACE_KILL`

The traced process terminates, with the same consequences as `exit(2)`.

`PTRACE_SINGLESTEP`

Execution continues as in request `PTRACE_CONT`; however, as soon as possible after execution of at least one instruction, execution stops again. The signal number from the stop is `SIGTRAP`. On the Sun, the T-bit is used and just one instruction is executed. This is part of the mechanism for implementing breakpoints.

`PTRACE_ATTACH`

Attach to the process identified by the *pid* argument and begin tracing it. Process *pid* does not have to be a child of the requestor, but the requestor must have permission to send process *pid* a signal and the effective user IDs of the requesting process and process *pid* must match.

`PTRACE_DETACH`

Detach the process being traced. Process *pid* is no longer being traced and continues its execution. The *data* argument is taken as a signal number and the process continues at location *addr* as if it had incurred that signal.

`PTRACE_GETREGS`

The traced process's registers are returned in a structure pointed to by the *addr* argument. The registers include the general purpose registers, the program counter and the program status word. The 'regs' structure defined in `<machine/reg.h>` describes the data that is returned.

`PTRACE_SETREGS`

The traced process's registers are written from a structure pointed to by the *addr* argument. The registers include the general purpose registers, the program counter and the program status word. The 'regs' structure defined in `<machine/reg.h>` describes the data that is set.

`PTRACE_GETFPREGS`

(Sun-3 only) The traced process's FPP status is returned in a structure pointed to by the *addr* argument. The status includes the 68881 floating point registers and the control, status, and instruction address registers. The 'fp_status' structure defined in `<machine/reg.h>` describes the data that is returned.

`PTRACE_SETFPREGS`

(Sun-3 only) The traced process's FPP status is written from a structure pointed to by the *addr* argument. The status includes the 68881 floating point registers and the control, status, and instruction address registers. The 'fp_status' structure defined in `<machine/reg.h>` describes the data that is set.

`PTRACE_GETFPAREGS`

(Sun-3 with FPA only) The traced process's FPA registers are returned in a structure pointed to by the *addr* argument. The 'fpa_regs' structure defined in `<machine/reg.h>` describes the data that is returned.

`PTRACE_SETFPAREGS`

(Sun-3 with FPA only) The traced process's FPA registers are written from a structure pointed to by the *addr* argument. The 'fpa_regs' structure defined in `<machine/reg.h>` describes the data that is set.

`PTRACE_READTEXT, PTRACE_READDATA`

Read data from the address space of the traced process. If the instruction and data spaces are separate, request `PTRACE_READTEXT` indicates instruction space while `PTRACE_READDATA` indicates data space. The *addr* argument is the address within the traced process from where the data is read, the *data* argument is the number of bytes to read, and the *addr2* argument is the address within the requesting process where the data is written.

`PTRACE_WRITETEXT`, `PTRACE_WRITEDATA`

Write data into the address space of the traced process. If the instruction and data spaces are separate, request `PTRACE_READTEXT` indicates instruction space while `PTRACE_READDATA` indicates data space. The *addr* argument is the address within the traced process where the data is written, the *data* argument is the number of bytes to write, and the *addr2* argument is the address within the requesting process from where the data is read.

As indicated, these calls (except for requests `PTRACE_TRACEME` and `PTRACE_ATTACH`) can be used only when the subject process has stopped. The *wait* call is used to determine when a process stops; in such a case the 'termination' status returned by *wait* has the value `WSTOPPED` to indicate a stop rather than genuine termination.

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

On the Sun, 'word' also means a 32-bit integer.

RETURN VALUE

In general, a 0 value is returned if the call succeeds. Note that this is not always true because requests such as `PTRACE_PEEKTEXT` and `PTRACE_PEEKDATA` return legitimate values. If the call fails then a -1 is returned and the global variable *errno* is set to indicate the error.

ERRORS

<code>EIO</code>	The request code is invalid.
<code>ESRCH</code>	The specified process does not exist.
<code>ESRCH</code>	The request requires the process to be one which is traced by the current process and stopped, but it is not stopped or it is not being traced by the current process.
<code>EIO</code>	The given signal number is invalid.
<code>EIO</code>	The specified address is out of bounds.
<code>EPERM</code>	The specified process cannot be traced.

SEE ALSO

wait(2), *sigvec(2)*, *adb(1)*

BUGS

ptrace is unique and arcane; it should be replaced with a special file which can be opened and read and written. The control functions could then be implemented with *ioctl(2)* calls on this file. This would be simpler to understand and have much higher performance.

The requests `PTRACE_TRACEME` thru `PTRACE_SINGLESTEP` are standard UNIX *ptrace* requests. The requests `PTRACE_ATTACH` thru `PTRACE_WRITEDATA` and the fifth argument, *addr2*, are unique to Sun UNIX.

The request `PTRACE_TRACEME` should be able to specify signals which are to be treated normally and not cause a stop. In this way, for example, programs with simulated floating point (which use 'illegal instruction' signals at a very high rate) could be efficiently debugged.

The error indication, -1, is a legitimate function value; *errno*, (see *intro(2)*), can be used to clarify.

It should be possible to stop a process on occurrence of a system call; in this way a completely controlled environment could be provided.

NAME

quotactl – manipulate disk quotas

SYNOPSIS

```
#include <ufs/quota.h>

quotactl(cmd, special, uid, addr)
int cmd;
char *special;
int uid;
caddr_t addr;
```

DESCRIPTION

The *quotactl* call manipulates disk quotas. The *cmd* parameter indicates a command to be applied to the user ID *uid*. *Special* is a pointer to a null-terminated string containing the path name of the block special device for the file system being manipulated. The block special device must be mounted. *Addr* is the address of an optional, command specific, data structure which is copied in or out of the system. The interpretation of *addr* is given with each command below.

Q_QUOTAON

Turn on quotas for a file system. *Addr* is a pointer to a null terminated string containing the path name of file containing the quotas for the file system. The quota file must exist; it is normally created with the *quotacheck*(8) program. This call is restricted to the super-user.

Q_QUOTAOFF

Turn off quotas for a file system. This call is restricted to the super-user.

Q_GETQUOTA

Get disk quota limits and current usage for user *uid*. *Addr* is a pointer to a struct *dqblk* structure (defined in *<ufs/quota.h>*). Only the super-user may get the quotas of a user other than himself.

Q_SETQUOTA

Set disk quota limits and current usage for user *uid*. *Addr* is a pointer to a struct *dqblk* structure (defined in *<ufs/quota.h>*). This call is restricted to the super-user.

Q_SETQLIM

Set disk quota limits for user *uid*. *Addr* is a pointer to a struct *dqblk* structure (defined in *<ufs/quota.h>*). This call is restricted to the super-user.

Q_SYNC

Update the on-disk copy of quota usages. This call is restricted to the super-user.

RETURN VALUE

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

ERRORS

A *quotactl* call will fail when one of the following occurs:

EINVAL	<i>Cmd</i> is invalid.
EPERM	The call is privileged and the caller was not the super-user.
EINVAL	The <i>special</i> parameter is not a mounted file system or is a mounted file system without quotas enabled.
ENOTBLK	The <i>special</i> parameter is not a block device.
EFAULT	An invalid <i>addr</i> is supplied; the associated structure could not be copied in or out of the kernel.
EINVAL	The <i>addr</i> parameter is being interpreted as the path of a quota file which exists but is either not a regular file or is not on the file system pointed to by the <i>special</i> parameter.

EUSERS The quota table is full.

SEE ALSO

quotaon(8), quotacheck(8)

BUGS

There should be some way to integrate this call with the resource limit interface provided by *setrlimit(2)* and *getrlimit(2)*. Incompatible with Melbourne quotas.

NAME

`read`, `readv` – read input

SYNOPSIS

```
cc = read(d, buf, nbytes)
int cc, d;
char *buf;
int nbytes;

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

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

DESCRIPTION

`read` attempts to read *nbytes* of data from the object referenced by the descriptor *d* into the buffer pointed to by *buf*. `readv` performs the same action, but scatters the input data into the *iovcnt* buffers specified by the members of the *iov* array: *iov*[0], *iov*[1], ..., *iov*[*iovcnt* - 1].

For `readv`, the *iovec* structure is defined as

```
struct iovec {
    caddr_t iov_base;
    int     iov_len;
};
```

Each *iovec* entry specifies the base address and length of an area in memory where data should be placed. `readv` will always fill an area completely before proceeding to the next.

On objects capable of seeking, the `read` starts at a position given by the pointer associated with *d* (see `lseek(2)`). Upon return from `read`, the pointer is incremented by the number of bytes actually read.

Objects that are not capable of seeking always read from the current position. The value of the pointer associated with such an object is undefined.

Note: For read access to a directory, use `readdir(3)` function. (Directory access using `readdir(3)` is no longer optional.)

Upon successful completion, `read` and `readv` return the number of bytes actually read and placed in the buffer. The system guarantees to read the number of bytes requested if the descriptor references a normal file which has that many bytes left before the end-of-file, but in no other case.

If the returned value is 0, then end-of-file has been reached.

When attempting to read from a descriptor associated with an empty pipe, socket, or FIFO:

If `O_NDELAY` is set, the read will return a -1 and *errno* will be set to `EWOULDBLOCK`.

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

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

If `O_NDELAY` is set, the read will return a -1 and *errno* will be set to `EWOULDBLOCK`.

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

If `O_NDELAY` is set, and less data are available than are requested by the `read` or `readv`, only the data that are available are returned, and the count indicates how many bytes of data were actually read.

SYSTEM V DESCRIPTION

When an attempt is made to read a descriptor which is in no-delay mode, and there is no data currently available, `read` will return a 0 instead of returning a -1 and setting *errno* to `EWOULDBLOCK`. Note that this

is indistinguishable

RETURN VALUE

If successful, the number of bytes actually read is returned. Otherwise, a -1 is returned and the global variable *errno* is set to indicate the error.

ERRORS

read and *readv* will fail if one or more of the following are true:

- EBADF *d* is not a valid file descriptor open for reading.
- EISDIR *d* refers to a directory which is on a file system mounted using the NFS.
- EFAULT *buf* points outside the allocated address space.
- EIO An I/O error occurred while reading from or writing to the file system.
- EINTR A read from a slow device was interrupted before any data arrived by the delivery of a signal.
- EINVAL The pointer associated with *d* was negative.

EWOULDBLOCK

The file was marked for non-blocking I/O, and no data were ready to be read. In addition, *readv* may return one of the following errors:

- EINVAL *iovcnt* was less than or equal to 0, or greater than 16.
- EINVAL One of the *iov_len* values in the *iov* array was negative.
- EINVAL The sum of the *iov_len* values in the *iov* array overflowed a 32-bit integer.
- EFAULT Part of *iov* points outside the process's allocated address space.

SEE ALSO

dup(2), *fcntl(2)*, *open(2)*, *pipe(2)*, *select(2)*, *socket(2)*, *socketpair(2)*

NAME

`readlink` – read value of a symbolic link

SYNOPSIS

```
cc = readlink(path, buf, bufsiz)  
int cc;  
char *path, *buf;  
int bufsiz;
```

DESCRIPTION

readlink places the contents of the symbolic link *name* in the buffer *buf* which has size *bufsiz*. The contents of the link are not null terminated when returned.

RETURN VALUE

The call returns the count of characters placed in the buffer if it succeeds, or a `-1` if an error occurs, placing the error code in the global variable *errno*.

ERRORS

readlink will fail and the buffer will be unchanged if:

EINVAL *path* contained a byte with the high-order bit set.

ENAMETOOLONG

The length of a component of *path* exceeds 255 characters, or the length of *path* exceeds 1023 characters.

ENOENT The named file does not exist.

EACCES Search permission is denied for a component of the path prefix of *path*.

ELOOP Too many symbolic links were encountered in translating *path*.

EINVAL The named file is not a symbolic link.

EIO An I/O error occurred while reading from or writing to the file system.

EFAULT *path* or *buf* extends outside the process's allocated address space.

SEE ALSO

`stat(2)`, `lstat(2)`, `symlink(2)`

NAME

reboot – reboot system or halt processor

SYNOPSIS

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

```
reboot(howto)
```

```
int howto;
```

DESCRIPTION

Reboot reboots the system, and is invoked automatically in the event of unrecoverable system failures. *Howto* is a mask of options passed to the bootstrap program. The system call interface permits only RB_HALT or RB_AUTOBOOT to be passed to the reboot program; the other flags are used in scripts stored on the console storage media, or used in manual bootstrap procedures. When none of these options (e.g. RB_AUTOBOOT) is given, the system is rebooted from file “vmunix” in the root file system of unit 0 of a disk chosen in a processor specific way. An automatic consistency check of the disks is then normally performed.

The bits of *howto* are:

RB_HALT

the processor is simply halted; no reboot takes place. RB_HALT should be used with caution.

RB_ASKNAME

Interpreted by the bootstrap program itself, causing it to inquire as to what file should be booted. Normally, the system is booted from the file “vmunix” without asking.

RB_SINGLE

Normally, the reboot procedure involves an automatic disk consistency check and then multi-user operations. RB_SINGLE prevents the consistency check, rather simply booting the system with a single-user shell on the console. RB_SINGLE is interpreted by the *init*(8) program in the newly booted system.

Only the super-user may *reboot* a machine.

RETURN VALUES

If successful, this call never returns. Otherwise, a -1 is returned and an error is returned in the global variable *errno*.

ERRORS

EPERM The caller is not the super-user.

SEE ALSO

crash(8S), halt(8), init(8), reboot(8)

NAME

recv, *recvfrom*, *recvmsg* – receive a message from a socket

SYNOPSIS

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

cc = recv(s, buf, len, flags)
int cc, s;
char *buf;
int len, flags;

cc = recvfrom(s, buf, len, flags, from, fromlen)
int cc, s;
char *buf;
int len, flags;
struct sockaddr *from;
int *fromlen;

cc = recvmsg(s, msg, flags)
int cc, s;
struct msghdr msg[];
int flags;
```

DESCRIPTION

recv, *recvfrom*, and *recvmsg* are used to receive messages from a socket.

The *recv* call may be used only on a *connected* socket (see *connect(2)*), while *recvfrom* and *recvmsg* may be used to receive data on a socket whether it is in a connected state or not.

If *from* is non-zero, the source address of the message is filled in. *fromlen* is a value-result parameter, initialized to the size of the buffer associated with *from*, and modified on return to indicate the actual size of the address stored there. The length of the message is returned in *cc*. If a message is too long to fit in the supplied buffer, excess bytes may be discarded depending on the type of socket the message is received from (see *socket(2)*).

If no messages are available at the socket, the receive call waits for a message to arrive, unless the socket is nonblocking (see *ioctl(2)*) in which case a *cc* of -1 is returned with the external variable *errno* set to *EWOULDBLOCK*.

The *select(2)* call may be used to determine when more data arrives.

The *flags* argument to a *recv* call is formed by *or*'ing one or more of the values,

```
#define MSG_OOB      0x1    /* process out-of-band data */
#define MSG_PEEK     0x2    /* peek at incoming message */
```

The *recvmsg* call uses a *msghdr* structure to minimize the number of directly supplied parameters. This structure has the following form, as defined in *<sys/socket.h>*:

```
struct msghdr {
    caddr_t msg_name;           /* optional address */
    int msg_namelen;           /* size of address */
    struct iovec *msg_iov;      /* scatter/gather array */
    int msg_iovlen;            /* # elements in msg_iov */
    caddr_t msg_accrightrights; /* access rights sent/received */
    int msg_accrightrightslen;
};
```

Here *msg_name* and *msg_namelen* specify the destination address if the socket is unconnected; *msg_name* may be given as a null pointer if no names are desired or required. The *msg_iov* and *msg_iovlen* describe the scatter gather locations, as described in *read(2V)*. A buffer to receive any access rights sent along with

the message is specified in *msg_accrights*, which has length *msg_accrightslen*.

RETURN VALUE

These calls return the number of bytes received, or -1 if an error occurred.

ERRORS

The calls fail if:

EBADF	The argument <i>s</i> is an invalid descriptor.
ENOTSOCK	The argument <i>s</i> is not a socket.
EWOULDBLOCK	The socket is marked non-blocking and the receive operation would block.
EINTR	The receive was interrupted by delivery of a signal before any data was available for the receive.
EFAULT	The data was specified to be received into a non-existent or protected part of the process address space.

SEE ALSO

fcntl(2), *read(2V)*, *send(2)*, *select(2)*, *getsockopt(2)*, *socket(2)*

NAME

`rename` – change the name of a file

SYNOPSIS

```
rename(from, to)
char *from, *to;
```

DESCRIPTION

`rename` renames the link named *from* as *to*. If *to* exists, then it is first removed. Both *from* and *to* must be of the same type (that is, both directories or both non-directories), and must reside on the same file system.

`Rename` guarantees that an instance of *to* will always exist, even if the system should crash in the middle of the operation.

If the final component of *from* is a symbolic link, the symbolic link is renamed, not the file or directory to which it points.

CAVEAT

The system can deadlock if a loop in the file system graph is present. This loop takes the form of an entry in directory “a”, say “a/foo”, being a hard link to directory “b”, and an entry in directory “b”, say “b/bar”, being a hard link to directory “a”. When such a loop exists and two separate processes attempt to perform “rename a/foo b/bar” and “rename b/bar a/foo”, respectively, the system may deadlock attempting to lock both directories for modification. Hard links to directories should be replaced by symbolic links by the system administrator.

RETURN VALUE

A 0 value is returned if the operation succeeds, otherwise `rename` returns `-1` and the global variable `errno` indicates the reason for the failure.

ERRORS

`rename` will fail and neither of the argument files will be affected if any of the following are true:

ENOTDIR	A component of the path prefix of either <i>from</i> or <i>to</i> is not a directory.
EINVAL	Either <i>from</i> or <i>to</i> contains a byte with the high-order bit set.
ENAMETOOLONG	The length of a component of either <i>from</i> or <i>to</i> exceeds 255 characters, or the length of either <i>from</i> or <i>to</i> exceeds 1023 characters.
ENOENT	A component of the path prefix of either <i>from</i> or <i>to</i> does not exist.
ENOENT	The file named by <i>from</i> does not exist.
EACCES	A component of the path prefix of either <i>from</i> or <i>to</i> denies search permission.
EACCES	The requested rename requires writing in a directory with a mode that denies write permission.
ELOOP	Too many symbolic links were encountered while translating either <i>from</i> or <i>to</i> .
EXDEV	The link named by <i>to</i> and the file named by <i>from</i> are on different logical devices (file systems).
ENOSPC	The directory in which the entry for the new name is being placed cannot be extended because there is no space left on the file system containing the directory.
EDQUOT	The directory in which the entry for the new name is being placed cannot be extended because the user’s quota of disk blocks on the file system containing the directory has been exhausted.
EIO	An I/O error occurred while reading from or writing to the file system.
EROFS	The requested rename requires writing in a directory on a read-only file system.
EFAULT	Either or both of <i>from</i> or <i>to</i> point outside the process’s allocated address space.

EINVAL *from* is a parent directory of *to*, or an attempt is made to rename “.” or “..”.

ENOTEMPTY *to* is a directory and is not empty.

EBUSY *to* is a directory and is the mount point for a mounted file system.

SEE ALSO

open(2V)

NAME

rmdir – remove a directory file

SYNOPSIS

```
rmdir(path)
char *path;
```

DESCRIPTION

rmdir removes a directory file whose name is given by *path*. The directory must not have any entries other than “.” and “..”.

RETURN VALUE

A 0 is returned if the remove succeeds; otherwise a -1 is returned and an error code is stored in the global location *errno*.

ERRORS

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

- | | |
|--------------|---|
| ENOTDIR | A component of the path prefix of <i>path</i> is not a directory. |
| ENOTDIR | The file referred to by <i>path</i> is not a directory. |
| EINVAL | <i>path</i> contains a character with the high-order bit set. |
| ENAMETOOLONG | The length of a component of <i>path</i> exceeds 255 characters, or the length of <i>path</i> exceeds 1023 characters. |
| ENOENT | The directory referred to by <i>path</i> does not exist. |
| ELOOP | Too many symbolic links were encountered in translating <i>path</i> . |
| ENOTEMPTY | The directory referred to by <i>path</i> contains files other than “.” and “..” in it. |
| EACCES | Search permission is denied for a component of the path prefix of <i>path</i> . |
| EACCES | Write permission is denied for the directory containing the link to be removed. |
| EBUSY | The directory to be removed is the mount point for a mounted file system. EIO An I/O error occurred while reading from or writing to the file system. |
| EROFS | The directory to be removed resides on a read-only file system. |
| EFAULT | <i>path</i> points outside the process’s allocated address space. |
- mkdir(2), unlink(2)**

NAME

`select` – synchronous I/O multiplexing

SYNOPSIS

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

nfd = select(width, readfds, writefds, exceptfds, timeout)
int width, *readfds, *writefds, *exceptfds;
struct timeval timeout;
```

DESCRIPTION

`select` examines the I/O descriptors specified by the bit masks `readfds`, `writefds`, and `exceptfds` to see if they are ready for reading, writing, or have an exceptional condition pending, respectively. `width` is the number of significant bits in each bit mask that represent a file descriptor. Typically `width` has the value returned by `getdtablesize(2)` for the maximum number of file descriptors or is the constant 32 (number of bits in an int). File descriptor `f` is represented by the bit “1<<f” in the mask. `select` returns, in place, a mask of those descriptors which are ready. The total number of ready descriptors is returned in `nfd`.

If `timeout` is a non-zero pointer, it specifies a maximum interval to wait for the selection to complete. If `timeout` is a zero pointer, the `select` blocks indefinitely. To effect a poll, the `timeout` argument should be non-zero, pointing to a zero-valued `timeval` structure.

Any of `readfds`, `writefds`, and `exceptfds` may be given as NULL pointers if no descriptors are of interest.

RETURN VALUE

`select` returns the number of ready descriptors that are contained in the bit masks, or `-1` if an error occurred. If the time limit expires then `select` returns 0.

ERRORS

An error return from `select` indicates:

EBADF	One of the bit masks specified an invalid descriptor.
EINTR	A signal was delivered before any of the selected events occurred or the time limit expired.
EINVAL	The specified time limit is unacceptable. One of its components is negative or too large.
EFAULT	One of the pointers given in the call referred to a non-existent portion of the process' address space.

SEE ALSO

`accept(2)`, `connect(2)`, `gettimeofday(2)`, `read(2V)`, `write(2V)`, `recv(2)`, `send(2)`, `getdtablesize(2)`

BUGS

The descriptor masks are always modified on return, even if the call returns as the result of the timeout.

NAME

`semctl` – semaphore control operations

SYNOPSIS

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

int semctl (semid, semnum, cmd, arg)
int semid, cmd;
int semnum;
union semun {
    val;
    struct semid_ds *buf;
    ushort *array;
} arg;
```

DESCRIPTION

`semctl` provides a variety of semaphore control operations as specified by `cmd`.

The following `cmds` are executed with respect to the semaphore specified by `semid` and `semnum`:

- GETVAL** Return the value of `semval` (see `intro(2)`). {READ}
- SETVAL** Set the value of `semval` to `arg.val`. {ALTER} When this `cmd` is successfully executed, the `semadj` value corresponding to the specified semaphore in all processes is cleared.
- GETPID** Return the value of `sempid`. {READ}
- GETNCNT** Return the value of `semncnt`. {READ}
- GETZCNT** Return the value of `semzcnt`. {READ}

The following `cmds` return and set, respectively, every `semval` in the set of semaphores.

- GETALL** Place `semvals` into array pointed to by `arg.array`. {READ}
- SETALL** Set `semvals` according to the array pointed to by `arg.array`. {ALTER} When this `cmd` is successfully executed the `semadj` values corresponding to each specified semaphore in all processes are cleared.

The following `cmds` are also available:

- IPC_STAT** Place the current value of each member of the data structure associated with `semid` into the structure pointed to by `arg.buf`. The contents of this structure are defined in `intro(2)`. {READ}
- IPC_SET** Set the value of the following members of the data structure associated with `semid` to the corresponding value found in the structure pointed to by `arg.buf`:
sem_perm.uid
sem_perm.gid
sem_perm.mode /* only low 9 bits */
 This `cmd` can only be executed by a process that has an effective user ID equal to either that of super-user or to the value of `sem_perm.uid` in the data structure associated with `semid`.
- IPC_RMID** Remove the semaphore identifier specified by `semid` from the system and destroy the set of semaphores and data structure associated with it. This `cmd` can only be executed by a process that has an effective user ID equal to either that of super-user or to the value of `sem_perm.uid` in the data structure associated with `semid`.

ERRORS

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

EINVAL	<i>semid</i> is not a valid semaphore identifier.
EINVAL	<i>semnum</i> is less than zero or greater than <i>sem_nsems</i> .
EINVAL	<i>cmd</i> is not a valid command.
EACCES	Operation permission is denied to the calling process (see <i>intro(2)</i>).
ERANGE	<i>cmd</i> is SETVAL or SETALL and the value to which <i>semval</i> is to be set is greater than the system imposed maximum.
EPERM	<i>cmd</i> is equal to IPC_RMID or IPC_SET and the effective user ID of the calling process is not equal to that of super-user and it is not equal to the value of <i>sem_perm.uid</i> in the data structure associated with <i>semid</i> .
EFAULT	<i>arg.buf</i> points to an illegal address.

RETURN VALUE

Upon successful completion, the value returned depends on *cmd* as follows:

GETVAL	The value of <i>semval</i> .
GETPID	The value of <i>sempid</i> .
GETNCNT	The value of <i>semmcnt</i> .
GETZCNT	The value of <i>semzcnt</i> .
All others	A value of 0.

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

SEE ALSO

intro(2), *semget(2)*, *semop(2)*.

NAME

`semget` – get set of semaphores

SYNOPSIS

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

int semget (key, nsems, semflg)
key_t key;
int nsems, semflg;
```

DESCRIPTION

`semget` returns the semaphore identifier associated with *key*.

A semaphore identifier and associated data structure and set containing *nsems* semaphores (see *intro(2)*) are created for *key* if one of the following are true:

key is equal to `IPC_PRIVATE`.

key does not already have a semaphore identifier associated with it, and $(semflg \& IPC_CREAT)$ is “true”.

Upon creation, the data structure associated with the new semaphore identifier is initialized as follows:

`Sem_perm.cuid`, `sem_perm.uid`, `sem_perm.cgid`, and `sem_perm.gid` are set equal to the effective user ID and effective group ID, respectively, of the calling process.

The low-order 9 bits of `sem_perm.mode` are set equal to the low-order 9 bits of *semflg*.

`sem_nsems` is set equal to the value of *nsems*.

`sem_otime` is set equal to 0 and `sem_ctime` is set equal to the current time.

ERRORS

`semget` will fail if one or more of the following are true:

<code>EINVAL</code>	<i>nsems</i> is either less than or equal to zero or greater than the system-imposed limit.
<code>EACCES</code>	A semaphore identifier exists for <i>key</i> , but operation permission (see <i>intro(2)</i>) as specified by the low-order 9 bits of <i>semflg</i> would not be granted.
<code>EINVAL</code>	A semaphore identifier exists for <i>key</i> , but the number of semaphores in the set associated with it is less than <i>nsems</i> and <i>nsems</i> is not equal to zero.
<code>ENOENT</code>	A semaphore identifier does not exist for <i>key</i> and $(semflg \& IPC_CREAT)$ is “false”.
<code>ENOSPC</code>	A semaphore identifier is to be created but the system-imposed limit on the maximum number of allowed semaphore identifiers system wide would be exceeded.
<code>ENOSPC</code>	A semaphore identifier is to be created but the system-imposed limit on the maximum number of allowed semaphores system wide would be exceeded.
<code>EEXIST</code>	A semaphore identifier exists for <i>key</i> but $(semflg \& IPC_CREAT)$ and $(semflg \& IPC_EXCL)$ is “true”.

RETURN VALUE

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

SEE ALSO

intro(2), *semctl(2)*, *semop(2)*.

NAME

semop – semaphore operations

SYNOPSIS

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

int semop (semid, sops, nsops)
int semid;
struct sembuf **sops;
int nsops;
```

DESCRIPTION

semop is used to automatically perform an array of semaphore operations on the set of semaphores associated with the semaphore identifier specified by *semid*. *sops* is a pointer to the array of semaphore-operation structures. *nsops* is the number of such structures in the array. The contents of each structure includes the following members:

```
short    sem_num;    /* semaphore number */
short    sem_op;     /* semaphore operation */
short    sem_flg;    /* operation flags */
```

Each semaphore operation specified by *sem_op* is performed on the corresponding semaphore specified by *semid* and *sem_num*.

sem_op specifies one of three semaphore operations as follows:

If *sem_op* is a negative integer, one of the following will occur: {ALTER}

If *semval* (see *intro(2)*) is greater than or equal to the absolute value of *sem_op*, the absolute value of *sem_op* is subtracted from *semval*. Also, if (*sem_flg* & SEM_UNDO) is “true”, the absolute value of *sem_op* is added to the calling process’s *semadj* value (see *exit(2)*) for the specified semaphore.

If *semval* is less than the absolute value of *sem_op* and (*sem_flg* & IPC_NOWAIT) is “true”, *semop* will return immediately.

If *semval* is less than the absolute value of *sem_op* and (*sem_flg* & IPC_NOWAIT) is “false”, *semop* will increment the *semncnt* associated with the specified semaphore and suspend execution of the calling process until one of the following conditions occur.

semval becomes greater than or equal to the absolute value of *sem_op*. When this occurs, the value of *semncnt* associated with the specified semaphore is decremented, the absolute value of *sem_op* is subtracted from *semval* and, if (*sem_flg* & SEM_UNDO) is “true”, the absolute value of *sem_op* is added to the calling process’s *semadj* value for the specified semaphore.

The *semid* for which the calling process is awaiting action is removed from the system (see *semctl(2)*). When this occurs, *errno* is set equal to EIDRM, and a value of -1 is returned.

The calling process receives a signal that is to be caught. When this occurs, the value of *semcnt* associated with the specified semaphore is decremented, and the calling process resumes execution in the manner prescribed in *signal(2)*.

If *sem_op* is a positive integer, the value of *sem_op* is added to *semval* and, if (*sem_flg* & SEM_UNDO) is “true”, the value of *sem_op* is subtracted from the calling process’s *semadj* value for the specified semaphore. {ALTER}

If *sem_op* is zero, one of the following will occur: {READ}

If *semval* is zero, *semop* will return immediately.

If *semval* is not equal to zero and (*sem_flg* & IPC_NOWAIT) is “true”, *semop* will return immediately.

If *semval* is not equal to zero and (*sem_flg* & IPC_NOWAIT) is “false”, *semop* will increment the *semcnt* associated with the specified semaphore and suspend execution of the calling process until one of the following occurs:

semval becomes zero, at which time the value of *semcnt* associated with the specified semaphore is decremented.

The *semid* for which the calling process is awaiting action is removed from the system. When this occurs, *errno* is set equal to EIDRM, and a value of -1 is returned.

The calling process receives a signal that is to be caught. When this occurs, the value of *semcnt* associated with the specified semaphore is decremented, and the calling process resumes execution in the manner prescribed in *signal(2)*.

ERRORS

semop will fail if one or more of the following are true for any of the semaphore operations specified by *sops*:

EINVAL	<i>semid</i> is not a valid semaphore identifier.
EFBIG	<i>sem_num</i> is less than zero or greater than or equal to the number of semaphores in the set associated with <i>semid</i> .
[E2BIG]	<i>nsops</i> is greater than the system-imposed maximum.
EACCESS	Operation permission is denied to the calling process (see <i>intro(2)</i>).
EAGAIN	The operation would result in suspension of the calling process but (<i>sem_flg</i> & IPC_NOWAIT) is “true”.
ENOSPC	The limit on the number of individual processes requesting an SEM_UNDO would be exceeded.
EINVAL	The number of individual semaphores for which the calling process requests a SEM_UNDO would exceed the limit.
ERANGE	An operation would cause a <i>semval</i> or <i>semadj</i> value to overflow the system-imposed limit.
EFAULT	<i>sops</i> points to an illegal address.

Upon successful completion, the value of *sempid* for each semaphore specified in the array pointed to by *sops* is set equal to the process ID of the calling process.

RETURN VALUE

If *semop* returns due to the receipt of a signal, a value of -1 is returned to the calling process and *errno* is set to EINTR. If it returns due to the removal of a *semid* from the system, a value of -1 is returned and *errno* is set to EIDRM.

Upon successful completion, the value of *semval* at the time of the call for the last operation in the array pointed to by *sops* is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

SEE ALSO

exec(2), exit(2), fork(2), intro(2), semctl(2), semget(2).

NAME

`send`, `sendto`, `sendmsg` – send a message from a socket

SYNOPSIS

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

cc = send(s, msg, len, flags)
int cc, s;
char *msg;
int len, flags;

cc = sendto(s, msg, len, flags, to, tolen)
int cc, s;
char *msg;
int len, flags;
struct sockaddr *to;
int tolen;

cc = sendmsg(s, msg, flags)
int cc, s;
struct msghdr msg[];
int flags;
```

DESCRIPTION

S is a socket created with `socket(2)`. `Send`, `sendto`, and `sendmsg` are used to transmit a message to another socket. `Send` may be used only when the socket is in a *connected* state, while `sendto` and `sendmsg` may be used at any time.

The address of the target is given by *to* with *tolen* specifying its size. The length of the message is given by *len*. If the message is too long to pass atomically through the underlying protocol, then the error `EMSGSIZE` is returned, and the message is not transmitted.

No indication of failure to deliver is implicit in a `send`. Return values of `-1` indicate some locally detected errors.

If no messages space is available at the socket to hold the message to be transmitted, then `send` normally blocks, unless the socket has been placed in non-blocking *i/o* mode. The `select(2)` call may be used to determine when it is possible to send more data.

The *flags* parameter may be set to `MSG_OOB` to send “out-of-band” data on sockets which support this notion (e.g. `SOCK_STREAM`).

See `recv(2)` for a description of the `msghdr` structure.

RETURN VALUE

The call returns the number of characters sent, or `-1` if an error occurred.

ERRORS

<code>EBADF</code>	An invalid descriptor was specified.
<code>ENOTSOCK</code>	The argument <i>s</i> is not a socket.
<code>EFAULT</code>	An invalid user space address was specified for a parameter.
<code>EMSGSIZE</code>	The socket requires that message be sent atomically, and the size of the message to be sent made this impossible.
<code>EWOULDBLOCK</code>	The socket is marked non-blocking and the requested operation would block.

SEE ALSO

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

NAME

setpgrp, *getpgrp* – set and/or return the process group of a process

SYNOPSIS

setpgrp(pid, pgrp)

pgrp = getpgrp(pid)

int pgrp;

int pid;

int pid, pgrp;

SYSTEM V SYNOPSIS

int setpgrp ()

DESCRIPTION**Setpgrp**

setpgrp sets the process group of the specified process, (*pid*) to the specified *pgrp*. If *pid* is zero, then the call applies to the current (calling) process.

If the effective user ID is not that of the super-user, then the process to be affected must have the same effective user ID as that of the caller or be a descendant of that process.

Getpgrp

getpgrp returns the process group of the indicated process. If *pid* is zero, then the call applies to the calling process.

Process groups are used for distribution of signals, and by terminals to arbitrate requests for their input. Processes that have the same process group as the terminal run in the foreground and may read from the terminal, while others block with a signal when they attempt to read.

This call is thus used by programs such as *cs*(1) to create process groups in implementing job control. The TIOCGPRG and TIOCSPGRP calls described in *ty*(4) are used to get/set the process group of the control terminal.

RETURN VALUE

setpgrp returns 0 when the operation was successful. If the request failed, -1 is returned and the global variable *errno* indicates the reason.

ERRORS

setpgrp fails, and the process group is not altered when one of the following occurs:

ESRCH The requested process does not exist.

EPERM The effective user ID of the requested process is different from that of the caller and the process is not a descendent of the calling process.

SYSTEM V DESCRIPTION

In the System V implementation, *setpgrp* takes no parameters. It sets the process group of the calling process to match its process ID, and returns the new process group ID.

SEE ALSO

exec(2), *fork*(2), *getpid*(2), *getuid*(2), *intro*(2), *kill*(2), *signal*(2), *ty*(4)

NAME

setregid – set real and effective group IDs

SYNOPSIS

```
int setregid(rgid, egid)
int rgid, egid;
```

DESCRIPTION

setregid is used to set the real and effective group IDs of the calling process. If *rgid* is -1 , the real group ID is not changed; if *egid* is -1 , the effective group ID is not changed. The real and effective group IDs may be set to different values in the same call.

If the effective user ID of the calling process is super-user, the real group ID and the effective group ID can be set to any legal value.

If the effective user ID of the calling process is not super-user, either the real group ID can be set to the saved set-group ID from *execve*(2), or the effective group ID can either be set to the saved set-group ID or the real group ID. Note that if a set-GID process sets its effective group ID to its real group ID, it can still set its effective group ID back to the saved set-group ID.

In either case, if the real group ID is changed to a particular value (i.e., if *rgid* is not -1), the saved set-group ID is set to that same value.

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

Setregid will fail and neither of the group IDs will be changed if:

EPERM The calling process' effective user ID is not the super-user and a change other than changing the real group ID to the saved set-group ID, or changing the effective group ID to the real group-id or the saved set-group ID, was specified.

SEE ALSO

getgid(2), *execve*(2), *setreuid*(2), *setgid*(3)

NAME

setreuid – set real and effective user IDs

SYNOPSIS

```
int setreuid(ruid, euid)
int ruid, euid;
```

DESCRIPTION

setreuid is used to set the real and effective user IDs of the calling process. If *ruid* is -1 , the real user ID is not changed; if *euid* is -1 , the effective user ID is not changed. The real and effective user IDs may be set to different values in the same call.

If the effective user ID of the calling process is super-user, the real user ID and the effective user ID can be set to any legal value.

If the effective user ID of the calling process is not super-user, either the real user ID can be set to the effective user ID, or the effective user ID can either be set to the saved set-user ID from *execve*(2) or the real user ID. Note that if a set-UID process sets its effective user ID to its real user ID, it can still set its effective user ID back to the saved set-user ID.

In either case, if the real user ID is changed to a particular value (i.e., if *ruid* is not -1), the saved set-user ID is set to that same value.

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

Setreuid will fail and neither of the user IDs will be changed if:

EPERM	The calling process' effective user ID is not the super-user and a change other than changing the real user ID to the effective user ID, or changing the effective user ID to the real user-id or the saved set-user ID, was specified.
-------	---

SEE ALSO

getuid(2), *execve*(2), *setregid*(2), *setuid*(3)

NAME

shmctl – shared memory control operations

SYNOPSIS

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

int shmctl (shmids, cmd, buf)
int shmids, cmd;
struct shmids_ds *buf;
```

DESCRIPTION

shmctl provides a variety of shared memory control operations as specified by *cmd*. The following *cmds* are available:

- IPC_STAT** Place the current value of each member of the data structure associated with *shmids* into the structure pointed to by *buf*. The contents of this structure are defined in *intro(2)*. {READ}
- IPC_SET** Set the value of the following members of the data structure associated with *shmids* to the corresponding value found in the structure pointed to by *buf*:
 - shm_perm.uid
 - shm_perm.gid
 - shm_perm.mode /* only low 9 bits */
 This *cmd* can only be executed by a process that has an effective user ID equal to either that of super-user or to the value of *shm_perm.uid* in the data structure associated with *shmids*.
- IPC_RMID** Remove the shared memory identifier specified by *shmids* from the system and destroy the shared memory segment and data structure associated with it. This *cmd* can only be executed by a process that has an effective user ID equal to either that of super-user or to the value of *shm_perm.uid* in the data structure associated with *shmids*.

ERRORS

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

- EINVAL** *Shmids* is not a valid shared memory identifier.
- EINVAL** *cmd* is not a valid command.
- EACCES** *cmd* is equal to **IPC_STAT** and {READ} operation permission is denied to the calling process (see *intro(2)*).
- EPERM** *cmd* is equal to **IPC_RMID** or **IPC_SET** and the effective user ID of the calling process is not equal to that of super-user and it is not equal to the value of *shm_perm.uid* in the data structure associated with *shmids*.
- EFAULT** *buf* points to an illegal address.

RETURN VALUE

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

SEE ALSO

intro(2), *shmget(2)*, *shmop(2)*.

NAME

shmget – get shared memory segment

SYNOPSIS

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

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

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

```
int shmget (key, size, shmflg)
```

```
key_t key;
```

```
int size, shmflg;
```

DESCRIPTION

shmget returns the shared memory identifier associated with *key*.

A shared memory identifier and associated data structure and shared memory segment of size *size* bytes (see *intro(2)*) are created for *key* if one of the following are true:

key is equal to `IPC_PRIVATE`.

key does not already have a shared memory identifier associated with it, and (*shmflg* & `IPC_CREAT`) is “true”.

Upon creation, the data structure associated with the new shared memory identifier is initialized as follows:

`shm_perm.cuid`, `shm_perm.uid`, `shm_perm.cgid`, and `shm_perm.gid` are set equal to the effective user ID and effective group ID, respectively, of the calling process.

The low-order 9 bits of `shm_perm.mode` are set equal to the low-order 9 bits of *shmflg*. `shm_segsz` is set equal to the value of *size*.

`shm_lpid`, `shm_nattch`, `shm_atime`, and `shm_dtime` are set equal to 0.

`shm_ctime` is set equal to the current time.

ERRORS

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

EINVAL	<i>size</i> is less than the system-imposed minimum or greater than the system-imposed maximum.
EACCES	A shared memory identifier exists for <i>key</i> but operation permission (see <i>intro(2)</i>) as specified by the low-order 9 bits of <i>shmflg</i> would not be granted.
EINVAL	A shared memory identifier exists for <i>key</i> but the size of the segment associated with it is less than <i>size</i> and <i>size</i> is not equal to zero.
ENOENT	A shared memory identifier does not exist for <i>key</i> and (<i>shmflg</i> & <code>IPC_CREAT</code>) is “false”.
ENOSPC	A shared memory identifier is to be created but the system-imposed limit on the maximum number of allowed shared memory identifiers system wide would be exceeded.
ENOMEM	A shared memory identifier and associated shared memory segment are to be created but the amount of available physical memory is not sufficient to fill the request.
EEXIST	A shared memory identifier exists for <i>key</i> but ((<i>shmflg</i> & <code>IPC_CREAT</code>) and (<i>shmflg</i> & <code>IPC_EXCL</code>)) is “true”.

RETURN VALUE

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

SEE ALSO

intro(2), shmctl(2), shmop(2)

NAME

shmop, shmat, shmdt – shared memory operations

SYNOPSIS

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

char *shmat (shmid, shmaddr, shmflg)
int shmid;
char *shmaddr
int shmflg;

int shmdt (shmaddr)
char *shmaddr
```

DESCRIPTION

shmat attaches the shared memory segment associated with the shared memory identifier specified by *shmid* to the data segment of the calling process. The segment is attached at the address specified by one of the following criteria:

If *shmaddr* is equal to zero, the segment is attached at the first available address as selected by the system.

If *shmaddr* is not equal to zero and (*shmflg* & SHM_RND) is “true”, the segment is attached at the address given by (*shmaddr* - (*shmaddr* modulus SHMLBA)).

If *shmaddr* is not equal to zero and (*shmflg* & SHM_RND) is “false”, the segment is attached at the address given by *shmaddr*.

The segment is attached for reading if (*shmflg* & SHM_RDONLY) is “true” {READ}, otherwise it is attached for reading and writing {READ/WRITE}.

shmat will fail and not attach the shared memory segment if one or more of the following are true:

EINVAL	<i>Shmid</i> is not a valid shared memory identifier.
EACCES	Operation permission is denied to the calling process (see <i>intro(2)</i>).
ENOMEM	The available data space is not large enough to accommodate the shared memory segment.
EINVAL	<i>shmaddr</i> is not equal to zero, and the value of (<i>shmaddr</i> - (<i>shmaddr</i> modulus SHMLBA)) is an illegal address.
EINVAL	<i>shmaddr</i> is not equal to zero, (<i>shmflg</i> & SHM_RND) is “false”, and the value of <i>shmaddr</i> is an illegal address.
EMFILE	The number of shared memory segments attached to the calling process would exceed the system-imposed limit.
EINVAL	<i>shmdt</i> detaches from the calling process's data segment the shared memory segment located at the address specified by <i>shmaddr</i> .
EINVAL	<i>shmdt</i> will fail and not detach the shared memory segment if <i>shmaddr</i> is not the data segment start address of a shared memory segment.

RETURN VALUES

Upon successful completion, the return value is as follows:

shmat returns the data segment start address of the attached shared memory segment.

shmdt returns a value of 0.

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

SEE ALSO

exec(2), *exit(2)*, *fork(2)*, *intro(2)*, *shmctl(2)*, *shmget(2)*.

NAME

shutdown – shut down part of a full-duplex connection

SYNOPSIS

```
shutdown(s, how)
int s, how;
```

DESCRIPTION

The *shutdown* call causes all or part of a full-duplex connection on the socket associated with *s* to be shut down. If *how* is 0, then further receives will be disallowed. If *how* is 1, then further sends will be disallowed. If *how* is 2, then further sends and receives will be disallowed.

DIAGNOSTICS

A 0 is returned if the call succeeds, -1 if it fails.

ERRORS

The call succeeds unless:

EBADF	<i>S</i> is not a valid descriptor.
ENOTSOCK	<i>S</i> is a file, not a socket.
ENOTCONN	The specified socket is not connected.

SEE ALSO

connect(2), socket(2)

BUGS

The *how* values should be defined constants.

NAME

sigblock – block signals

SYNOPSIS

```
#include <signal.h>

oldmask = sigblock(mask);
int mask;

mask = sigmask(signum)
```

DESCRIPTION

Sigblock adds the signals specified in *mask* to the set of signals currently being blocked from delivery. Signals are blocked if the corresponding bit in *mask* is a 1; the macro *sigmask* is provided to construct the mask for a given *signum*. The previous mask is returned, and may be restored using *sigsetmask(2)*.

It is not possible to block SIGKILL, SIGSTOP, or SIGCONT; this restriction is silently imposed by the system.

RETURN VALUE

The previous set of masked signals is returned.

SEE ALSO

kill(2), sigvec(2), sigsetmask(2), signal(3)

NAME

sigpause – atomically release blocked signals and wait for interrupt

SYNOPSIS

```
sigpause(sigmask)  
int sigmask;
```

DESCRIPTION

Sigpause assigns *sigmask* to the set of masked signals and then waits for a signal to arrive; on return the set of masked signals is restored. *Sigmask* is usually 0 to indicate that no signals are now to be blocked. *Sigpause* always terminates by being interrupted, returning EINTR.

In normal usage, a signal is blocked using *sigblock(2)*, to begin a critical section, variables modified on the occurrence of the signal are examined to determine that there is no work to be done, and the process pauses awaiting work by using *sigpause* with the mask returned by *sigblock*.

SEE ALSO

sigblock(2), *sigvec(2)*, *signal(3)*

NAME

sigsetmask – set current signal mask

SYNOPSIS

```
#include <signal.h>

sigsetmask(mask);
int mask;

mask = sigmask(signum)
```

DESCRIPTION

sigsetmask sets the current signal mask (those signals that are blocked from delivery). Signals are blocked if the corresponding bit in *mask* is a 1; the macro *sigmask* is provided to construct the mask for a given *signum*.

The system quietly disallows SIGKILL, SIGSTOP, or SIGCONT from being blocked.

RETURN VALUE

The previous set of masked signals is returned.

SEE ALSO

kill(2), sigvec(2), sigblock(2), sigpause(2), signal(3)

NAME

sigstack – set and/or get signal stack context

SYNOPSIS

```
#include <signal.h>

struct sigstack {
    caddr_t ss_sp;
    int     ss_onstack;
};

sigstack(ss, oss)
struct sigstack *ss, *oss;
```

DESCRIPTION

Sigstack allows users to define an alternate stack on which signals are to be processed. If *ss* is non-zero, it specifies a *signal stack* on which to deliver signals and tells the system if the process is currently executing on that stack. When a signal's action indicates its handler should execute on the signal stack (specified with a *sigvec(2)* call), the system checks to see if the process is currently executing on that stack. If the process is not currently executing on the signal stack, the system arranges a switch to the signal stack for the duration of the signal handler's execution. If *oss* is non-zero, the current signal stack state is returned.

NOTES

Signal stacks are not "grown" automatically, as is done for the normal stack. If the stack overflows unpredictable results may occur.

RETURN VALUE

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

ERRORS

Sigstack will fail and the signal stack context will remain unchanged if one of the following occurs.

EFAULT Either *ss* or *oss* points to memory that is not a valid part of the process address space.

SEE ALSO

sigvec(2), setjmp(3), signal(3)

NAME

sigvec – software signal facilities

SYNOPSIS

```
#include <signal.h>

struct sigvec {
    int      (*sv_handler)();
    int      sv_mask;
    int      sv_flags;
};

sigvec(sig, vec, ovec)
int sig;
struct sigvec *vec, *ovec;
```

DESCRIPTION

The system defines a set of signals that may be delivered to a process. Signal delivery resembles the occurrence of a hardware interrupt: the signal is blocked from further occurrence, the current process context is saved, and a new one is built. A process may specify a *handler* to which a signal is delivered, or specify that a signal is to be *blocked* or *ignored*. A process may also specify that a default action is to be taken by the system when a signal occurs. Normally, signal handlers execute on the current stack of the process. This may be changed, on a per-handler basis, so that signals are taken on a special *signal stack*.

All signals have the same *priority*. Signal routines execute with the signal that caused their invocation *blocked*, but other signals may yet occur. A global *signal mask* defines the set of signals currently blocked from delivery to a process. The signal mask for a process is initialized from that of its parent (normally 0). It may be changed with a *sigblock(2)* or *sigsetmask(2)* call, or when a signal is delivered to the process.

When a signal condition arises for a process, the signal is added to a set of signals pending for the process. If the signal is not currently *blocked* by the process then it is delivered to the process. When a signal is delivered, the current state of the process is saved, a new signal mask is calculated (as described below), and the signal handler is invoked. The call to the handler is arranged so that if the signal handling routine returns normally the process will resume execution in the context from before the signal's delivery. If the process wishes to resume in a different context, then it must arrange to restore the previous context itself.

When a signal is delivered to a process a new signal mask is installed for the duration of the process' signal handler (or until a *sigblock* or *sigsetmask* call is made). This mask is formed by taking the current signal mask, adding the signal to be delivered, and *or*'ing in the signal mask associated with the handler to be invoked.

Sigvec assigns a handler for a specific signal. If *vec* is non-zero, it specifies a handler routine and mask to be used when delivering the specified signal. Further, if the SV_ONSTACK bit is set in *sv_flags*, the system will deliver the signal to the process on a *signal stack*, specified with *sigstack(2)*. If *ovec* is non-zero, the previous handling information for the signal is returned to the user.

The mask specified in *vec* is not allowed to block SIGKILL, SIGSTOP, or SIGCONT. The system enforces this restriction silently.

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 (other than A-line or F-line op code)
SIGTRAP	5*	trace trap
SIGIOT	6*	IOT trap (not generated on Suns)
SIGEMT	7*	EMT trap (A-line or F-line op code)
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 <i>fcntl(2)</i>)
SIGXCPU	24	cpu time limit exceeded (see <i>setrlimit(2)</i>)
SIGXFSZ	25	file size limit exceeded (see <i>setrlimit(2)</i>)
SIGVTALRM	26	virtual time alarm (see <i>setitimer(2)</i>)
SIGPROF	27	profiling timer alarm (see <i>setitimer(2)</i>)
SIGWINCH	28•	window changed (see <i>win(4S)</i>)
SIGLOST	29*	resource lost (see <i>lockd(8C)</i>)
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.

Once a signal handler is installed, it remains installed until another *sigvec* call is made, or an *execve(2)* is performed, except that if the `SV_RESETHAND` bit is set in *sv_flags*, the value of *sv_handler* for the caught signal will be set to `SIG_DFL` before entering the signal-catching function, unless the signal is `SIGILL` or `SIGTRAP`. If this bit is set, the bit for that signal in the signal mask will not be set; unless the signal mask associated with that signal blocks that signal, further occurrences of that signal will not be blocked. The `SV_RESETHAND` flag is not available in 4.2BSD, hence it should not be used if backward compatibility is needed.

The default action for a signal may be reinstated by setting *sv_handler* to `SIG_DFL`; this default is termination 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 the process is terminated, a “core image” will be made in the current working directory of the receiving process if the signal is one for which an asterisk appears in the above list *and* the following conditions are met:

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

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

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

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

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

a file group ID that is the same as the file group ID of the current directory

If *sv_handler* is `SIG_IGN` the signal is subsequently ignored, and pending instances of the signal are discarded.

Note: the signals `SIGKILL`, `SIGSTOP`, and `SIGCONT` cannot be ignored.

If a caught signal occurs during certain system calls, the call is normally restarted. The call can be forced to terminate prematurely with an `EINTR` error return by setting the `SV_INTERRUPT` bit in *sv_flags*. The `SV_INTERRUPT` flag is not available in 4.2BSD, hence it should not be used if backward compatibility is needed. The affected system calls are *read(2V)* or *write(2V)* on a

slow device (such as a terminal or pipe or other socket, but not a file) and during a *wait(2)*.

After a *fork(2)* or *vfork(2)* the child inherits all signals, the signal mask, the signal stack, and the restart/interrupt and reset-signal-handler flags.

The *execve(2)* call resets all caught signals to default action and resets all signals to be caught on the user stack. Ignored signals remain ignored; the signal mask remains the same; signals that interrupt system calls continue to do so.

NOTES

The handler routine can be declared:

```
handler(sig, code, scp)
int sig, code;
struct sigcontext *scp;
```

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.

Programs that must be portable to UNIX systems other than 4.2 BSD should use the *signal(3)* interface instead.

CODES

The following defines the codes for signals which produce them. All of these symbols are defined in `<signal.h>`:

Hardware condition	Signal	Code
Illegal instruction	SIGILL	ILL_INSTR_FAULT
Privilege violation	SIGILL	ILL_PRIVVIO_FAULT
Coprocessor protocol error	SIGILL	ILL_INSTR_FAULT
Trap # <i>n</i> (1 <= <i>n</i> <= 14)	SIGILL	ILL_TRAP <i>n</i> _FAULT
A-line op code	SIGEMT	EMT_EMU1010
F-line op code	SIGEMT	EMT_EMU1111
Integer division by zero	SIGFPE	FPE_INTDIV_TRAP
CHK or CHK2 instruction	SIGFPE	FPE_CHKINST_TRAP
TRAPV or TRAPcc or cpTRAPcc	SIGFPE	FPE_TRAPV_TRAP
IEEE floating point compare unordered	SIGFPE	FPE_FLTBSUN_TRAP
IEEE floating point inexact	SIGFPE	FPE_FLTINEX_TRAP
IEEE floating point division by zero	SIGFPE	FPE_FLTDIV_TRAP
IEEE floating point underflow	SIGFPE	FPE_FLTUND_TRAP
IEEE floating point operand error	SIGFPE	FPE_FLTOPERR_TRAP
IEEE floating point overflow	SIGFPE	FPE_FLTOVF_FAULT
IEEE floating point signaling NaN	SIGFPE	FPE_FLTNAN_TRAP

RETURN VALUE

A 0 value indicated that the call succeeded. A -1 return value indicates an error occurred and *errno* is set to indicate the reason.

ERRORS

Sigvec will fail and no new signal handler will be installed if one of the following occurs:

EFAULT	Either <i>vec</i> or <i>ovec</i> points to memory that is not a valid part of the process address space.
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), ptrace(2), kill(2), sigblock(2), sigsetmask(2), sigpause(2), sigstack(2), setjmp(3), signal(3), tty(4)

NAME

`socket` – create an endpoint for communication

SYNOPSIS

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

s = socket(af, type, protocol)
int s, af, type, protocol;
```

DESCRIPTION

Socket creates an endpoint for communication and returns a descriptor.

The *af* parameter specifies an address format with which addresses specified in later operations using the `socket` should be interpreted. These formats are defined in the include file `<sys/socket.h>`. The currently understood formats are

AF_UNIX	(UNIX path names),
AF_INET	(ARPA Internet addresses),
AF_PUP	(Xerox PUP-I Internet addresses), and
AF_IMPLINK	(IMP “host at IMP” addresses).

The `socket` has the indicated *type* which specifies the semantics of communication. Currently defined types are:

```
SOCK_STREAM
SOCK_DGRAM
SOCK_RAW
SOCK_SEQPACKET
SOCK_RDM
```

A `SOCK_STREAM` type provides sequenced, reliable, two-way connection based byte streams with an out-of-band data transmission mechanism. A `SOCK_DGRAM` socket supports datagrams (connectionless, unreliable messages of a fixed (typically small) maximum length). `SOCK_RAW` sockets provide access to internal network interfaces. The types `SOCK_RAW`, which is available only to the super-user, and `SOCK_SEQPACKET` and `SOCK_RDM`, which are planned, but not yet implemented, are not described here.

The *protocol* specifies a particular protocol to be used with the socket. Normally only a single protocol exists to support a particular socket type using a given address format. However, it is possible that many protocols may exist in which case a particular protocol must be specified in this manner. The protocol number to use is particular to the “communication domain” in which communication is to take place; see *services(5)* and *protocols(5)*.

Sockets of type `SOCK_STREAM` are full-duplex byte streams, similar to pipes. A stream socket must be in a *connected* state before any data may be sent or received on it. A connection to another socket is created with a *connect(2)* call. Once connected, data may be transferred using *read(2V)* and *write(2V)* calls or some variant of the *send(2)* and *recv(2)* calls. When a session has been completed a *close(2)* may be performed. Out-of-band data may also be transmitted as described in *send(2)* and received as described in *recv(2)*.

The communications protocols used to implement a `SOCK_STREAM` insure that data is not lost or duplicated. If a piece of data for which the peer protocol has buffer space cannot be successfully transmitted within a reasonable length of time, then the connection is considered broken and calls will indicate an error with `-1` returns and with `ETIMEDOUT` as the specific code in the global variable `errno`. The protocols optionally keep sockets “warm” by forcing transmissions roughly every minute in the absence of other activity. An error is then indicated if no response can be elicited on an otherwise idle connection for an extended period (e.g. 5 minutes). A `SIGPIPE` signal is raised if a process sends on a broken stream; this causes naive processes, which do not handle the signal, to exit.

SOCK_DGRAM and SOCK_RAW sockets allow sending of datagrams to correspondents named in *send(2)* calls. It is also possible to receive datagrams at such a socket with *recv(2)*.

An *fcntl(2)* call can be used to specify a process group to receive a SIGURG signal when the out-of-band data arrives.

The operation of sockets is controlled by socket level *options*. These options are defined in the file *<sys/socket.h>* and explained below. *setsockopt* and *getsockopt(2)* are used to set and get options, respectively.

SO_DEBUG	turn on recording of debugging information
SO_REUSEADDR	allow local address reuse
SO_KEEPALIVE	keep connections alive
SO_DONTROUTE	do not apply routing on outgoing messages
SO_LINGER	linger on close if data present
SO_DONTLINGER	do not linger on close

SO_DEBUG enables debugging in the underlying protocol modules. SO_REUSEADDR indicates the rules used in validating addresses supplied in a *bind(2)* call should allow reuse of local addresses. SO_KEEPALIVE enables the periodic transmission of messages on a connected socket. Should the connected party fail to respond to these messages, the connection is considered broken and processes using the socket are notified via a SIGPIPE signal. SO_DONTROUTE indicates that outgoing messages should bypass the standard routing facilities. Instead, messages are directed to the appropriate network interface according to the network portion of the destination address. SO_LINGER and SO_DONTLINGER control the actions taken when unsent messages are queued on socket and a *close(2)* is performed. If the socket promises reliable delivery of data and SO_LINGER is set, the system will block the process on the *close* attempt until it is able to transmit the data or until it decides it is unable to deliver the information (a timeout period, termed the linger interval, is specified in the *setsockopt* call when SO_LINGER is requested). If SO_DONTLINGER is specified and a *close* is issued, the system will process the close in a manner which allows the process to continue as quickly as possible.

RETURN VALUE

A -1 is returned if an error occurs, otherwise the return value is a descriptor referencing the socket.

ERRORS

The *socket* call fails if:

EAFNOSUPPORT	The specified address family is not supported in this version of the system.
ESOCKTNOSUPPORT	The specified socket type is not supported in this address family.
EPROTONOSUPPORT	The specified protocol is not supported.
EMFILE	The per-process descriptor table is full.
ENOBUFS	No buffer space is available. The socket cannot be created.

SEE ALSO

accept(2), *bind(2)*, *connect(2)*, *getsockname(2)*, *getsockopt(2)*, *ioctl(2)*, *listen(2)*, *recv(2)*, *select(2)*, *send(2)*, *shutdown(2)*, *socketpair(2)*

Inter-Process Communication Primer in *Networking on the Sun Workstation*

BUGS

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

NAME

socketpair – create a pair of connected sockets

SYNOPSIS

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

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

```
socketpair(d, type, protocol, sv)
```

```
int d, type, protocol;
```

```
int sv[2];
```

DESCRIPTION

The *socketpair* system call creates an unnamed pair of connected sockets in the specified domain *d*, of the specified *type* and using the optionally specified *protocol*. The descriptors used in referencing the new sockets are returned in *sv*[0] and *sv*[1]. The two sockets are indistinguishable.

DIAGNOSTICS

A 0 is returned if the call succeeds, -1 if it fails.

ERRORS

The call succeeds unless:

EMFILE Too many descriptors are in use by this process.

EAFNOSUPPORT The specified address family is not supported on this machine.

EPROTONOSUPPORT
The specified protocol is not supported on this machine.

EOPNOSUPPORT The specified protocol does not support creation of socket pairs.

EFAULT The address *sv* does not specify a valid part of the process address space.

SEE ALSO

read(2V), write(2V), pipe(2)

BUGS

This call is currently implemented only for the UNIX domain.

NAME

stat, lstat, fstat – get file status

SYNOPSIS

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

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

```
stat(path, buf)
```

```
char *path;
```

```
struct stat *buf;
```

```
lstat(path, buf)
```

```
char *path;
```

```
struct stat *buf;
```

```
fstat(fd, buf)
```

```
int fd;
```

```
struct stat *buf;
```

DESCRIPTION

stat obtains information about the file named by *path*. Read, write or execute permission of the named file is not required, but all directories listed in the path name leading to the file must be searchable.

lstat is like *stat* except in the case where the named file is a symbolic link, in which case *lstat* returns information about the link, while *stat* returns information about the file the link references.

fstat obtains the same information about an open file referenced by the argument descriptor, such as would be obtained by an *open* call.

buf is a pointer to a *stat* structure into which information is placed concerning the file. The contents of the structure pointed to by *buf* include the following members:

```

dev_t      st_dev;      /* device inode resides on */
ino_t      st_ino;     /* this inode's number */
u_short    st_mode;    /* protection */
short      st_nlink;   /* number of hard links to the file */
short      st_uid;     /* user ID of owner */
short      st_gid;     /* group ID of owner */
dev_t      st_rdev;    /* the device type, for inode that is device */
off_t      st_size;    /* total size of file, in bytes */
time_t     st_atime;   /* file last access time */
time_t     st_mtime;   /* file last modify time */
time_t     st_ctime;   /* file last status change time */
long       st_blksize; /* optimal blocksize for file system i/o ops */
long       st_blocks;  /* actual number of blocks allocated */

```

st_atime Time when file data was last read or modified. Changed by the following system calls: *mknod(2)*, *utimes(2)*, *read(2V)*, *write(2V)*, and *truncate(2)*. For reasons of efficiency, *st_atime* is not set when a directory is searched, although this would be more logical.

st_mtime Time when data was last modified. It is not set by changes of owner, group, link count, or mode. Changed by the following system calls: *mknod(2)*, *utimes(2)*, *write(2V)*.

st_ctime Time when file status was last changed. It is set both both by writing and changing the i-node. Changed by the following system calls: *chmod(2)*, *chown(2)*, *link(2)*, *mknod(2)*, *rename(2)*, *unlink(2)*, *utimes(2)*, *write(2V)*, *truncate(2)*.

The status information word *st_mode* has bits:

```

#define S_IFMT      0170000 /* type of file */
#define S_IFIFO     0010000 /* fifo special */
#define S_IFCHR     0020000 /* character special */

```

```

#define S_IFDIR      0040000 /* directory */
#define S_IFBLK      0060000 /* block special */
#define S_IFREG      0100000 /* regular file */
#define S_IFLNK      0120000 /* symbolic link */
#define S_IFSOCK     0140000 /* socket */
#define S_ISUID      0004000 /* set user id on execution */
#define S_ISGID      0002000 /* set group id on execution */
#define S_ISVTX      0001000 /* save swapped text even after use */
#define S_IRREAD     0000400 /* read permission, owner */
#define S_IWRITE     0000200 /* write permission, owner */
#define S_IXEXEC     0000100 /* execute/search permission, owner */

```

The mode bits 0000070 and 0000007 encode group and others permissions (see *chmod(2)*).

RETURN VALUE

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

ERRORS

stat and *lstat* will fail if one or more of the following are true:

ENOTDIR A component of the path prefix of *path* is not a directory.

EINVAL *path* contains a character with the high-order bit set.

ENAMETOOLONG

The length of a component of *path* exceeds 255 characters, or the length of *path* exceeds 1023 characters.

ENOENT The file referred to by *path* does not exist.

EACCES Search permission is denied for a component of the path prefix of *path*.

ELOOP Too many symbolic links were encountered in translating *path*.

EFAULT *buf* or *path* points to an invalid address.

EIO An I/O error occurred while reading from or writing to the file system.

fstat will fail if one or both of the following are true:

EBADF *fd* is not a valid open file descriptor.

EFAULT *buf* points to an invalid address.

EIO An I/O error occurred while reading from or writing to the file system.

CAVEAT

The fields in the *stat* structure currently marked *st_spare1*, *st_spare2*, and *st_spare3* are present in preparation for inode time stamps expanding to 64 bits. This, however, can break certain programs which depend on the time stamps being contiguous (in calls to *utimes(2)*).

SEE ALSO

chmod(2), *chown(2)*, *readlink(2)*, *utimes(2)*

NAME

`statfs` – get file system statistics

SYNOPSIS

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

```
statfs(path, buf)
```

```
char *path;
```

```
struct statfs *buf;
```

```
fstatfs(fd, buf)
```

```
int fd;
```

```
struct statfs *buf;
```

DESCRIPTION

`statfs` returns information about a mounted file system. *path* is the path name of any file within the mounted filesystem. *Buf* is a pointer to a `statfs` structure defined as follows:

```
typedef struct {
    long    val[2];
} fsid_t;

struct statfs {
    long    f_type;      /* type of info, zero for now */
    long    f_bsize;    /* fundamental file system block size */
    long    f_blocks;   /* total blocks in file system */
    long    f_bfree;    /* free blocks */
    long    f_bavail;   /* free blocks available to non-superuser */
    long    f_files;    /* total file nodes in file system */
    long    f_ffree;    /* free file nodes in fs */
    fsid_t  f_fsid;     /* file system id */
    long    f_spare[7]; /* spare for later */
};
```

Fields that are undefined for a particular file system are set to `-1`. `fstatfs` returns the same information about an open file referenced by descriptor *fd*.

RETURN VALUE

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

ERRORS

`statfs` fails if one or more of the following are true:

ENOTDIR A component of the path prefix of *path* is not a directory.

EINVAL *path* contains a character with the high-order bit set.

ENAMETOOLONG

The length of a component of *path* exceeds 255 characters, or the length of *path* exceeds 1023 characters.

ENOENT The file referred to by *path* does not exist.

EACCES Search permission is denied for a component of the path prefix of *path*.

ELOOP Too many symbolic links were encountered in translating *path*.

EFAULT *buf* or *path* points to an invalid address.

EIO An I/O error occurred while reading from or writing to the file system.

fstatfs fails if one or both of the following are true:

- EBADF *fd* is not a valid open file descriptor.
- EFAULT *buf* points to an invalid address.
- EIO An I/O error occurred while reading from or writing to the file system.

NAME

`swapon` – add a swap device for interleaved paging/swapping

SYNOPSIS

```
swapon(special)
char *special;
```

DESCRIPTION

`swapon` makes the block device *special* available to the system for allocation for paging and swapping. The names of potentially available devices are known to the system and defined at system configuration time. The size of the swap area on *special* is calculated at the time the device is first made available for swapping.

SEE ALSO

`swapon(8)`, `config(8)`

RETURN VALUE

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

ERRORS

ENOTDIR	A component of the path prefix of <i>special</i> is not a directory.
EINVAL	<i>special</i> contains a character with the high-order bit set.
ENAMETOOLONG	The length of a component of <i>special</i> exceeds 255 characters, or the length of <i>special</i> exceeds 1023 characters.
ENOENT	The device referred to by <i>special</i> does not exist.
EACCES	Search permission is denied for a component of the path prefix of <i>special</i> .
ELOOP	Too many symbolic links were encountered in translating <i>special</i> .
EPERM	The caller is not the super-user.
ENOTBLK	The file referred to by <i>special</i> is not a block device.
EBUSY	The device referred to by <i>special</i> has already been made available for swapping.
ENODEV	The device referred to by <i>special</i> was not configured into the system as a swap device.
ENXIO	The major device number of the device referred to by <i>special</i> is out of range (this indicates no device driver exists for the associated hardware).
EIO	An I/O error occurred while reading from or writing to the file system or opening the swap device.
EFAULT	<i>special</i> points outside the process's address space.

BUGS

There is no way to stop swapping on a disk so that the pack may be dismounted.
This call will be upgraded in future versions of the system.

NAME

symlink – make symbolic link to a file

SYNOPSIS

```
symlink(name1, name2)
char *name1, *name2;
```

DESCRIPTION

A symbolic link *name2* is created to *name1* (*name2* is the name of the file created, *name1* is the string used in creating the symbolic link). Either name may be an arbitrary path name; the files need not be on the same file system.

RETURN VALUE

Upon successful completion, a zero value is returned. If an error occurs, the error code is stored in *errno* and a -1 value is returned.

ERRORS

The symbolic link is made unless one or more of the following are true:

ENOTDIR	A component of the path prefix of <i>name2</i> is not a directory.
EINVAL	<i>name2</i> contains a character with the high-order bit set.
ENAMETOOLONG	The length of a component of either <i>name1</i> or <i>name2</i> exceeds 255 characters, or the length of either <i>name1</i> or <i>name2</i> exceeds 1023 characters.
ENOENT	A component of the path prefix of <i>name2</i> does not exist.
EACCES	Search permission is denied for a component of the path prefix of <i>name2</i> .
ELOOP	Too many symbolic links were encountered in translating <i>name2</i> .
EEXIST	The file referred to by <i>name2</i> already exists.
EIO	An I/O error occurred while reading from or writing to the file system.
EROFS	The file <i>name2</i> would reside on a read-only file system.
ENOSPC	The directory in which the entry for the new symbolic link is being placed cannot be extended because there is no space left on the file system containing the directory.
ENOSPC	The new symbolic link cannot be created because there is no space left on the file system which will contain the link.
ENOSPC	There are no free inodes on the file system on which the file is being created.
EDQUOT	The directory in which the entry for the new symbolic link is being placed cannot be extended because the user's quota of disk blocks on the file system containing the directory has been exhausted.
EDQUOT	The new symbolic link cannot be created because the user's quota of disk blocks on the file system which will contain the link has been exhausted.
EDQUOT	The user's quota of inodes on the file system on which the file is being created has been exhausted.
EFAULT	<i>name1</i> or <i>name2</i> points outside the process's allocated address space.

SEE ALSO

link(2), ln(1), readlink(2), unlink(2)

NAME

`sync` – update super-block

SYNOPSIS

`sync()`

DESCRIPTION

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

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

SEE ALSO

`fsync(2)`, `sync(8)`, `cron(8)`

BUGS

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

NAME

syscall – indirect system call

SYNOPSIS

```
#include <syscall.h>
```

```
syscall(number, arg, ...)
```

DESCRIPTION

syscall performs the system call whose assembly language interface has the specified *number*, and arguments *arg* Symbolic constants for system calls can be found in the header file *<syscall.h>*.

The register d0 value of the system call is returned.

DIAGNOSTICS

When the C-bit is set, *syscall* returns -1 and sets the external variable *errno* (see *intro(2)*).

BUGS

There is no way to simulate system calls such as *pipe(2)*, which return values in register d1.

NAME

`truncate`, `ftruncate` – truncate a file to a specified length

SYNOPSIS

```
truncate(path, length)
char *path;
unsigned long length;

ftruncate(fd, length)
int fd;
unsigned long length;
```

DESCRIPTION

truncate causes the file named by *path* or referenced by *fd* to be truncated to at most *length* bytes in size. If the file previously was larger than this size, the extra data is lost. With *ftruncate*, the file must be open for writing.

RETURN VALUES

A value of 0 is returned if the call succeeds. If the call fails a -1 is returned, and the global variable *errno* specifies the error.

ERRORS

Truncate succeeds unless:

ENOTDIR	A component of the path prefix of <i>path</i> is not a directory.
EINVAL	<i>path</i> contains a character with the high-order bit set.
ENAMETOOLONG	The length of a component of <i>path</i> exceeds 255 characters, or the length of <i>path</i> exceeds 1023 characters.
ENOENT	The file referred to by <i>path</i> does not exist.
EACCES	Search permission is denied for a component of the path prefix of <i>path</i> .
EACCES	Write permission is denied for the file referred to by <i>path</i> .
ELOOP	Too many symbolic links were encountered in translating <i>path</i> .
EISDIR	The file referred to by <i>path</i> is a directory.
EROFS	The file referred to by <i>path</i> resides on a read-only file system.
ETXTBSY	The file referred to by <i>path</i> is a pure procedure (shared text) file that is being executed.
EIO	An I/O error occurred while reading from or writing to the file system.
EFAULT	<i>path</i> points outside the process's allocated address space.

ftruncate succeeds unless:

EINVAL	<i>fd</i> is not a valid descriptor of a file open for writing.
EINVAL	<i>fd</i> references a socket, not a file.
EIO	An I/O error occurred while reading from or writing to the file system.

SEE ALSO

`open(2V)`

BUGS

Partial blocks discarded as the result of truncation are not zero filled; this can result in holes in files which do not read as zero.

These calls should be generalized to allow ranges of bytes in a file to be discarded.

NAME

umask – set file creation mode mask

SYNOPSIS

```
oumask = umask(numask)  
int oumask, numask;
```

DESCRIPTION

Umask sets the process's file mode creation mask to *numask* and returns the previous value of the mask. The low-order 9 bits of *numask* are used whenever a file is created, clearing corresponding bits in the file mode (see *chmod(2)*). This clearing allows each user to restrict the default access to his files.

The value is initially 022 (write access for owner only). The mask is inherited by child processes.

RETURN VALUE

The previous value of the file mode mask is returned by the call.

SEE ALSO

chmod(2), *mknod(2)*, *open(2V)*

NAME

`uname` – get name of current UNIX system

SYNOPSIS

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

```
int uname (name)
```

```
struct utsname *name;
```

DESCRIPTION

Note: This system call is only available for use with the System V compatibility libraries. These are located in the directory */usr/5lib*, and are compiled using the System V version of the C compiler, */usr/5lib/cc*.

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

uname uses the structure defined in `<sys/utsname.h>` whose members are:

```
char    sysname[9];
char    nodename[9];
char    release[9];
char    version[9];
char    machine[9];
```

uname returns a null-terminated character string naming the current UNIX system in *sysname* and *nodename*. This name will be the name returned by the *gethostname(2)* system call, truncated to 8 characters. *release* and *version* further identify the operating system. *machine* contains a name that identifies the hardware that the UNIX system is running on.

SEE ALSO

`uname(1V)`

NAME

unlink – remove directory entry

SYNOPSIS

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

DESCRIPTION

unlink removes the directory entry named by the path name pointed to by *path*. If this entry was the last link to the file, and no process has the file open, then all resources associated with the file are reclaimed. If, however, the file was open in any process, the actual resource reclamation is delayed until it is closed, even though the directory entry has disappeared.

RETURN VALUE

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

ERRORS

The *unlink* succeeds unless:

ENOTDIR A component of the path prefix of *path* is not a directory.

EINVAL *path* contains a character with the high-order bit set.

ENAMETOOLONG

The length of a component of *path* exceeds 255 characters, or the length of *path* exceeds 1023 characters.

ENOENT The file referred to by *path* does not exist.

EACCES Search permission is denied for a component of the path prefix of *path*.

EACCES Write permission is denied for the directory containing the link to be removed.

ELOOP Too many symbolic links were encountered in translating *path*.

EPERM The file referred to by *path* is a directory and the effective user ID of the process is not the super-user.

EBUSY The entry to be unlinked is the mount point for a mounted file system.

EIO An I/O error occurred while reading from or writing to the file system.

EROFS The file referred to by *path* resides on a read-only file system.

EFAULT *path* points outside the process's allocated address space.

SEE ALSO

close(2), link(2), rmdir(2)

NAME

unmount – remove a file system

SYNOPSIS

```
unmount(name)  
char *name;
```

DESCRIPTION

unmount announces to the system that the directory *name* is no longer to refer to the root of a mounted file system. The directory *name* reverts to its ordinary interpretation.

RETURN VALUE

unmount returns 0 if the action occurred; -1 if the directory is inaccessible or does not have a mounted file system, or if there are active files in the mounted file system.

ERRORS

unmount may fail with one of the following errors:

EPERM	The caller is not the super-user.
ENOTDIR	A component of the path prefix of <i>name</i> is not a directory.
EINVAL	<i>name</i> is not the root of a mounted file system.
EBUSY	A process is holding a reference to a file located on the file system.
EINVAL	The path name contains a character with the high-order bit set.
ENAMETOOLONG	The length of a component of the path name exceeds 255 characters, or the length of the entire path name exceeds 1023 characters.
ENOENT	<i>name</i> does not exist.
EACCES	Search permission is denied for a component of the path prefix.
EFAULT	<i>name</i> points outside the process's allocated address space.
ELOOP	Too many symbolic links were encountered in translating the path name.
EIO	An I/O error occurred while reading from or writing to the file system.

SEE ALSO

mount(2), *mount(8)*, *umount(8)*

BUGS

The error codes are in a state of disarray; too many errors appear to the caller as one value.

NAME

`utimes` – set file times

SYNOPSIS

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

utimes(file, tvp)
char *file;
struct timeval tvp[2];
```

DESCRIPTION

The `utimes` call uses the “accessed” and “updated” times in that order from the `tvp` vector to set the corresponding recorded times for `file`.

The caller must be the owner of the file or the super-user. The “inode-changed” time of the file is set to the current time.

RETURN VALUE

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

ERRORS

`utime` will fail if one or more of the following are true:

ENOTDIR A component of the path prefix of `file` is not a directory.

EINVAL `file` contained a character with the high-order bit set.

ENAMETOOLONG

The length of a component of `file` exceeds 255 characters, or the length of `file` exceeds 1023 characters.

ENOENT The file referred to by `file` does not exist.

EACCES Search permission is denied for a component of the path prefix of `file`.

ELOOP Too many symbolic links were encountered in translating `file`.

EPERM The process is not super-user and not the owner of the file.

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 `file` or `tvp` points outside the process’s allocated address space.

SEE ALSO

`stat(2)`

NAME

`vadvise` – give advice to paging system

SYNOPSIS

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

```
vadvise(param)
```

```
int param;
```

DESCRIPTION

Vadvise is used to inform the system that process paging behavior merits special consideration. Parameters to *vadvise* are defined in the file `<vadvise.h>`. Currently, two calls to *vadvise* are implemented.

The call

```
vadvise(VA_ANOM);
```

advises that the paging behavior is not likely to be well handled by the system's default algorithm, since reference information is collected over macroscopic intervals (e.g. 10-20 seconds) will not serve to indicate future page references. The system in this case will choose to replace pages with little emphasis placed on recent usage, and more emphasis on referenceless circular behavior. It is *essential* that processes which have very random paging behavior (such as LISP during garbage collection of very large address spaces) call *vadvise*, as otherwise the system has great difficulty dealing with their page-consumptive demands.

The call

```
vadvise(VA_NORM);
```

restores default paging replacement behavior after a call to

```
vadvise(VA_ANOM);
```

BUGS

Will go away soon, being replaced by a per-page *madvise* facility.

NAME

vfork – spawn new process in a virtual memory efficient way

SYNOPSIS

```
pid = vfork()
int pid;
```

DESCRIPTION

vfork can be used to create new processes without fully copying the address space of the old process, which is horrendously inefficient in a paged environment. It is useful when the purpose of *fork(2)* would have been to create a new system context for an *execve*. *Vfork* differs from *fork* in that the child borrows the parent's memory and thread of control until a call to *execve(2)* or an exit (either by a call to *exit(2)* or abnormally.) The parent process is suspended while the child is using its resources.

vfork returns 0 in the child's context and (later) the pid of the child in the parent's context.

vfork can normally be used just like *fork*. It does not work, however, to return while running in the child's context from the procedure which called *vfork* since the eventual return from *vfork* would then return to a no longer existent stack frame. Be careful, also, to call *_exit* rather than *exit* if you can't *execve*, since *exit* will flush and close standard I/O channels, and thereby mess up the parent processes standard I/O data structures. (Even with *fork* it is wrong to call *exit* since buffered data would then be flushed twice.)

SEE ALSO

fork(2), *execve(2)*, *sigvec(2)*, *wait(2)*,

DIAGNOSTICS

Same as for *fork*.

BUGS

This system call will be eliminated when proper system sharing mechanisms are implemented. Users should not depend on the memory sharing semantics of *vfork* as it will, in that case, be made synonymous to *fork*.

To avoid a possible deadlock situation, processes that are children in the middle of a *vfork* are never sent SIGTTOU or SIGTTIN signals; rather, output or *ioctl*s are allowed and input attempts result in an end-of-file indication.

NAME

vhangup – virtually “hangup” the current control terminal

SYNOPSIS

vhangup()

DESCRIPTION

Vhangup is used by the initialization process *init*(8) (among others) to arrange that users are given “clean” terminals at login, by revoking access of the previous users’ processes to the terminal. To effect this, *vhangup* searches the system tables for references to the control terminal of the invoking process, revoking access permissions on each instance of the terminal that it finds. Further attempts to access the terminal by the affected processes will yield i/o errors (EBADF). Finally, a hangup signal (SIGHUP) is sent to the process group of the control terminal.

SEE ALSO

init (8)

BUGS

Access to the control terminal via */dev/tty* is still possible.

This call should be replaced by an automatic mechanism that takes place on process exit.

NAME

wait, *wait3* – wait for process to terminate or stop

SYNOPSIS

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

pid = wait(status)
int pid;
union wait *status;

pid = wait(0)
int pid;

#include <sys/time.h>
#include <sys/resource.h>

pid = wait3(status, options, rusage)
int pid;
union wait *status;
int options;
struct rusage *rusage;
```

DESCRIPTION

wait causes its caller to delay until a signal is received or one of its child processes terminates or stops due to tracing. If any child has died or stopped due to tracing and this has not been reported via *wait*, return is immediate, returning the process ID and exit status of one of those children. If that child had died, it is discarded. If there are no children, return is immediate with the value -1 returned. If there are only running or stopped but reported children, the calling processes is blocked.

On return from a successful *wait* call, *status* is nonzero, and the high byte of *status* contains the low byte of the argument to *exit* supplied by the child process; the low byte of *status* contains the termination status of the process. A more precise definition of the *status* word is given in `<sys/wait.h>`.

wait3 is an alternate interface that allows both non-blocking status collection and the collection of the status of children stopped by any means. The *status* parameter is defined as above. The *options* parameter is used to indicate the call should not block if there are no processes that have status to report (WNOHANG), and/or that children of the current process that are stopped due to a SIGTTIN, SIGTTOU, SIGTSTP, or SIGSTOP signal are eligible to have their status reported as well (WUNTRACED). A terminated child is discarded after it reports status, and a stopped process will not report its status more than once. If *rusage* is non-zero, a summary of the resources used by the terminated process and all its children is returned. (This information is currently not available for stopped processes.)

When the WNOHANG option is specified and no processes have status to report, *wait3* returns a *pid* of 0. The WNOHANG and WUNTRACED options may be combined by *or*'ing the two values.

NOTES

See *sigvec(2)* for a list of termination statuses (signals); 0 status indicates normal termination. A special status (0177) is returned for a stopped process that has not terminated and can be restarted; see *ptrace(2)* and *sigvec(2)*. If the 0200 bit of the termination status is set, a core image of the process was produced by the system.

If the parent process terminates without waiting on its children, the initialization process (process ID = 1) inherits the children.

wait and *wait3* are automatically restarted when a process receives a signal while awaiting termination of a child process.

RETURN VALUE

If *wait* returns due to a stopped or terminated child process, the process ID of the child is returned to the calling process. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

wait3 returns -1 if there are no children not previously waited for; 0 is returned if WNOHANG is specified and there are no stopped or exited children.

ERRORS

wait will fail and return immediately if one or more of the following are true:

ECHILD The calling process has no existing unwaited-for child processes.

EFAULT The *status* or *rusage* arguments point to an illegal address.

The call is forced to terminate prematurely due to the arrival of a signal whose SM SV_INTERRUPT bit in sv_flags is set (see *sigvec(2)*). *signal(3V)*, in the System V compatibility library, sets this bit for any signal it catches.

SEE ALSO

exit(2), *getrusage(2)*

NAME

`write`, `writew` – write output

SYNOPSIS

```
cc = write(d, buf, nbytes)
int cc, d;
char *buf;
int nbytes;

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

cc = writew(d, iov, iovcnt)
int cc, d;
struct iovec *iov;
int iovcnt;
```

DESCRIPTION

`write` attempts to write *nbytes* of data to the object referenced by the descriptor *d* from the buffer pointed to by *buf*. `writew` performs the same action, but gathers the output data from the *iovcnt* buffers specified by the members of the *iov* array: *iov*[0], *iov*[1], ..., *iov*[*iovcnt* - 1].

For `writew`, the *iovec* structure is defined as

```
struct iovec {
    caddr_t iov_base;
    int     iov_len;
};
```

Each *iovec* entry specifies the base address and length of an area in memory from which data should be written. `writew` will always write a complete area before proceeding to the next.

On objects capable of seeking, the `write` starts at a position given by the pointer associated with *d*, see *lseek*(2). Upon return from `write`, the pointer is incremented by the number of bytes actually written.

Objects that are not capable of seeking always write from the current position. The value of the pointer associated with such an object is undefined.

If the `O_APPEND` flag of the file status flags is set, the file pointer will be set to the end of the file prior to each write.

If the real user is not the super-user, then `write` clears the set-user-id bit on a file. This prevents penetration of system security by a user who “captures” a writable set-user-id file owned by the super-user.

When using non-blocking I/O on objects that are subject to flow control, such as sockets, pipes (or FIFOs), or terminals, `write` and `writew` may write fewer bytes than requested; the return value must be noted, and the remainder of the operation should be retried when possible. If such an object’s buffers are full, so that it cannot accept any data, then `write` and `writew` will return -1 and set *errno* to `EWOULDBLOCK`. Otherwise, they will block until space becomes available.

SYSTEM V DESCRIPTION

A `write` (but not a `writew`) on an object that cannot accept any data will return a count of 0, rather than returning -1 and setting *errno* to `EWOULDBLOCK`.

RETURN VALUE

Upon successful completion the number of bytes actually written is returned. Otherwise a -1 is returned and the global variable *errno* is set to indicate the error.

ERRORS

`write` and `writew` will fail and the file pointer will remain unchanged if one or more of the following are true:

`EBADF` *d* is not a valid descriptor open for writing.

EPIPE An attempt is made to write to a pipe that is not open for reading by any process (or to a socket of type `SOCK_STREAM` that is connected to a peer socket.) Note: an attempted write of this kind will also cause you to receive a `SIGPIPE` signal from the kernel. If you've not made a special provision to catch or ignore this signal, your process will die.

EFBIG An attempt was made to write a file that exceeds the process's file size limit or the maximum file size.

EFAULT Part of *iov* or data to be written to the file points outside the process's allocated address space.

The call is forced to terminate prematurely due to the arrival of a signal whose `SM SV_INTERRUPT` bit in `sv_flags` is set (see `sigvec(2)`). `signal(3V)`, in the System V compatibility library, sets this bit for any signal it catches.

EINVAL The pointer associated with *d* was negative.

ENOSPC There is no free space remaining on the file system containing the file.

EDQUOT The user's quota of disk blocks on the file system containing the file has been exhausted.

EIO An I/O error occurred while reading from or writing to the file system.

EWOULDBLOCK

The file was marked for non-blocking I/O, and no data could be written immediately.

In addition, *writev* may return one of the following errors:

EINVAL *iovcnt* was less than or equal to 0, or greater than 16.

EINVAL One of the *iov_len* values in the *iov* array was negative.

EINVAL The sum of the *iov_len* values in the *iov* array overflowed a 32-bit integer.

SEE ALSO

`fcntl(2)`, `lseek(2)`, `open(2V)`, `pipe(2)`, `select(2)`

NAME

intro – introduction to library functions

DESCRIPTION

Section 3 describes library routines. The main C library is */lib/libc.a*, which contains all system call entry points described in section 2, as well as functions described in several subsections here. The primary functions are described in the main section 3. Functions associated with the “standard I/O library” used by many C programs are found in section 3S. The main C library also includes Internet network functions, described in section 3N, and routines providing compatibility with other UNIX systems, described in section 3C.

Other sections are:

- (3F) This section, for FORTRAN library routines and functions, is contained in the *FORTRAN Programmer's Guide*.
- (3M) The Math Library. C declarations for the types of functions are be obtained from the include file *<math.h>*. To use these functions with C programs compile them with the *-lm* option with *cc(1)*. They are automatically loaded as needed by the FORTRAN and Pascal compilers *f77(1)* and *pc(1)*.
- (3V) The System V Compatibility Library. System V versions of functions that are not yet merged into the standard Sun libraries. To use these functions, compile programs with */usr/5bin/cc*, instead of */bin/cc*.
- (3X) Various specialized libraries have not been given distinctive captions. Files in which such libraries are found are named on appropriate pages if they don't appear in the *libc* library.

FILES

<i>/lib/libc.a</i>	C Library ((2), (3), (3N) and (3C) routines)
<i>/usr/lib/libc_p.a</i>	Profiling C library (for <i>gprof(1)</i>)
<i>/usr/lib/libm.a</i>	Math Library <i>-lm</i> (see section 3M)
<i>/usr/lib/libm_p.a</i>	Profiling version of <i>-lm</i>
<i>/usr/lib/libcurses.a</i>	screen management routines (see <i>curses(3X)</i>)
<i>/usr/lib/libdbm.a</i>	data base management routines (see <i>dbm(3X)</i>)
<i>/usr/lib/libmp.a</i>	multiple precision math library (see <i>mp(3X)</i>)
<i>/usr/lib/libtermcap.a</i>	terminal handling routines (see <i>termcap(3X)</i>)
<i>/usr/lib/libtermcap_p.a</i>	"
<i>/usr/lib/libtermplib</i>	(link to <i>/usr/lib/libtermcap.a</i>)
<i>/usr/lib/libtermplib_p.a</i>	(link to <i>/usr/lib/libtermcap_p.a</i>)
<i>/usr/lib/libplot*.a</i>	plot routines (see <i>plot(3X)</i>)
<i>/usr/lib/libresolv.a</i>	Internet name server routines (see <i>resolver(3X)</i>)

SEE ALSO

intro(3C), intro(3S), intro(3F), intro(3M), intro(3N), nm(1), ld(1), cc(1), f77(1), intro(2)

DIAGNOSTICS

Functions in the math library (section 3M) may return conventional values when the function is undefined for the given arguments or when the value is not representable. In these cases the external variable *errno* (see *intro(2)*) is set to the value EDOM (domain error) or ERANGE (range error). The values of EDOM and ERANGE are defined in the include file *<errno.h>*.

LIST OF FUNCTIONS

<i>Name</i>	<i>Appears on Page</i>	<i>Description</i>
a64l	a64l(3)	convert base-64 ASCII to long
abort	abort(3)	generate a fault
abs	abs(3)	integer absolute value
acos	sin(3M)	trigonometric functions
acosh	asinh(3M)	inverse hyperbolic function
addmntent	getmntent(3)	get file system descriptor file entry

alarm	alarm(3C)	schedule signal after specified time
alloca	malloc(3)	memory allocator
alphasort	scandir(3)	scan a directory
asctime	ctime(3)	convert date and time to ASCII
asin	sin(3M)	trigonometric functions
asinh	asinh(3M)	inverse hyperbolic function
assert	assert(3)	program verification
atan	sin(3M)	trigonometric functions
atanh	asinh(3M)	inverse hyperbolic function
atof	atof(3)	convert ASCII to numbers
atoi	atof(3)	convert ASCII to numbers
atol	atof(3)	convert ASCII to numbers
bcmp	bstring(3)	bit and byte string operations
bcopy	bstring(3)	bit and byte string operations
bsearch	bsearch(3)	binary search a sorted table
bzero	bstring(3)	bit and byte string operations
cabs	hypot(3M)	Euclidean distance
calloc	malloc(3)	memory allocator
cbc_crypt	des_crypt(3)	fast DES encryption
cbrt	sqrt(3M)	cube root
ceil	floor(3M)	ceiling
cfree	malloc(3)	memory allocator
clearerr	ferror(3S)	stream status inquiries
clock	clock(3C)	report CPU time used
closedir	directory(3)	directory operations
closelog	syslog(3)	control system log
copysign	ieee(3M)	copysign remainder exponent manipulations
cos	sin(3M)	trigonometric functions
cosh	sinh(3M)	hyperbolic functions
crypt	crypt(3)	DES encryption
ctermid	ctermid(3S)	generate filename for terminal
ctime	ctime(3)	convert date and time to ASCII
cuserid	cuserid(3S)	get character login name of user
des_crypt	des_crypt(3)	fast DES encryption
des_setparity	des_crypt(3)	fast DES encryption
dn_comp	resolver(3X)	Internet name server routines
dn_expand	resolver(3X)	Internet name server routines
drand48	drand48(3)	generate uniformly distributed pseudo-random numbers
drem	ieee(3M)	copysign remainder exponent manipulations
dysize	ctime(3)	convert date and time to ASCII
ecb_crypt	des_crypt(3)	fast DES encryption
ecvt	ecvt(3)	output conversion
edata	end(3)	last locations in program,
encrypt	crypt(3)	DES encryption
end	end(3)	last locations in program
endfsent	getfsent(3)	get file system descriptor 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
endpwent	getpwent(3)	get password file entry

endservent	getservent(3N)	get service entry
environ	execl(3)	execute a file
erand48	drand48(3)	generate uniformly distributed pseudo-random numbers
erf	erf(3M)	error functions
errno	perror(3)	system error messages
etext	end(3)	last locations in program
ether	ether(3R)	monitor traffic on the Ethernet
ether_aton	ethers(3N)	Ethernet address mapping
ether_hostton	ethers(3N)	Ethernet address mapping
ether_line(3N)	ethers	Ethernet address mapping
ether_ntoa	ethers(3N)	Ethernet address mapping
ether_ntohost	ethers(3N)	Ethernet address mapping
execl	execl(3)	execute a file
execle	execl(3)	execute a file
execlp	execl(3)	execute a file
execv	execl(3)	execute a file
execvp	execl(3)	execute a file
exit	exit(3)	terminate a process after performing cleanup
exp	exp(3M)	exponential function
fabs	floor(3M)	absolute value
fclose	fclose(3S)	close or flush a stream
fcvt	ecvt(3)	output conversion
fdopen	fopen(3S)	open a stream
feof	ferror(3S)	stream status inquiries
ferror	ferror(3S)	stream status inquiries
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
fgets	gets(3S)	get a string from a stream
fileno	ferror(3S)	stream status inquiries
finite	ieee(3M)	copysign remainder exponent manipulations
floor	floor(3M)	floor function
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
frexp	frexp(3)	split into mantissa and exponent
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
fwrite	fread(3S)	buffered binary input/output
gcvt	ecvt(3)	output conversion
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)	value for environment name
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
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
getopt	getopt(3)	get option letter from argv
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
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
getrpcport	getrpcport(3R)	get RPC port number
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
getw	getc(3S)	get character or integer from stream
getwd	getwd(3)	get current working directory pathname
gmtime	ctime(3)	convert date and time to ASCII
gsignal	signal(3)	software signals
gtty	stty(3C)	set and get terminal state
hasmntopt	getmntent(3)	get file system descriptor file entry
havedisk	rstat(3R)	get remote host performance data
hcreate	hsearch(3)	manage hash search tables
hdestroy	hsearch(3)	manage hash search tables
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
hypot	hypot(3M)	Euclidean distance
ieee	ieee(3M)	copysign remainder exponent manipulations
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

initgroups	initgroups(3)	initialize group access list
initstate	random(3)	better random number generator
innetgr	getnetgrent(3N)	get network group entry
insque	insque(3)	insert/remove element from a queue
isalnum	ctype(3)	character classification and conversion macros
isalpha	ctype(3)	character classification and conversion macros
isascii	ctype(3)	character classification and conversion macros
isatty	ttyname(3)	find name of a terminal
iscntrl	ctype(3)	character classification and conversion macros
isdigit	ctype(3)	character classification and conversion macros
isgraph	ctype(3)	character classification and conversion macros
isinf	isinf(3)	test for indeterminate floating point values
islower	ctype(3)	character classification and conversion macros
isnan	isinf(3)	test for indeterminate floating point values
isprint	ctype(3)	character classification and conversion macros
ispunct	ctype(3)	character classification and conversion macros
isspace	ctype(3)	character classification and conversion macros
isupper	ctype(3)	character classification and conversion macros
isxdigit	ctype(3)	character classification and conversion macros
j0	j0(3M)	Bessel functions
j1	j0(3M)	Bessel functions
jn	j0(3M)	Bessel functions
rand48	drand48(3)	generate uniformly distributed pseudo-random numbers
l64a	a64l(3)	convert long to base-64 ASCII
lcong48	drand48(3)	generate uniformly distributed pseudo-random numbers
ldexp	frexp(3)	split into mantissa and exponent
lfind	lsearch(3)	linear search and update
lgamma	lgamma(3M)	log gamma function
localtime	ctime(3)	convert date and time to ASCII
lockf	lockf(3)	advisory record locking on files
log	exp(3M)	exponential functions
log10	exp(3M)	exponential functions
logb	ieee(3M)	copysign remainder exponent manipulations
longjmp	setjmp(3)	non-local goto
lrand48	drand48(3)	generate uniformly distributed pseudo-random numbers
lsearch	lsearch(3)	linear search and update
malloc	malloc(3)	memory allocator
malloc_debug	malloc(3)	memory allocator
malloc_verify	malloc(3)	memory allocator
matherr	matherr(3M)	math library error-handling function
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
memset	memory(3)	memory operations
mkstemp	mktemp(3)	make a unique file name
mktemp	mktemp(3)	make a unique file name
modf	frexp(3)	split into mantissa and exponent
moncontrol	monitor(3)	prepare execution profile
monitor	monitor(3)	prepare execution profile
monstartup	monitor(3)	prepare execution profile
rand48	drand48(3)	generate uniformly distributed pseudo-random numbers

nice	nice(3C)	set program priority
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	onexit(3)	name termination handler
opendir	directory(3)	directory operations
openlog	syslog(3)	control system log
optarg	getopt(3)	get option letter from argv
optind	getopt(3)	get option letter from argv
pause	pause(3C)	stop until signal
pclose	popen(3S)	initiate I/O to/from a process
perror	perror(3)	system error messages
popen	popen(3S)	initiate I/O to/from a process
pow	exp(3M)	exponential functions
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	utenv(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
qsort	qsort(3)	quicker sort
rand	rand(3C)	random number generator
random	random(3)	better random number generator
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
regexp	regexp(3)	regular expression compile and match routines
remque	insque(3)	insert/remove element from a queue
res_init	resolver(3X)	Internet name server routines
res_mkquery	resolver(3)	Internet name server routines
res_send	resolver(3)	Internet name server routines
rewind	fseek(3S)	reposition a stream
rewinddir	directory(3)	directory operations
rex	rex(3R)	remote execution protocol
rexec	rexec(3N)	return stream to a remote command
rindex	string(3)	string operations
rint	floor(3M)	round to nearest integer
musers	musers(3R)	return info about users on remote hosts
rquota	rquota(3R)	implement quotas on remote hosts
rresvport	rcmd(3N)	routines for returning a stream to a remote command
rstat	rstat(3R)	get remote host performance data
ruserok	rcmd(3N)	routines for returning a stream to a remote command
users	musers(3R)	return info about users on remote hosts
rwall	rwall(3R)	write to remote host
scalb	ieee(3M)	copysign remainder exponent manipulations
scandir	scandir(3)	scan a directory
scanf	scanf(3S)	formatted input conversion

seed48	drand48(3)	generate uniformly distributed pseudo-random numbers
seekdir	directory(3)	directory operations
setbuf	setbuf(3S)	assign buffering to a stream
setbuffer	setbuf(3S)	assign buffering to a stream
setegid	setuid(3)	set user and group ID
seteuid	setuid(3)	set user and group ID
setfsent	getfsent(3)	get file system descriptor file entry
setgid	setuid(3)	set user and group ID
setgrent	getgrent(3)	get group file entry
sethostent	gethostent(3N)	get network host entry
setjmp	setjmp(3)	non-local goto
setkey	crypt(3)	DES encryption
setlinebuf	setbuf(3S)	assign buffering to a stream
setlinebuf	setbuf(3S)	assign buffering to a stream
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
setpwent	getpwent(3)	get password file entry
setrgid	setuid(3)	set user and group ID
setruid	setuid(3)	set user and group ID
setservent	getservent(3N)	get service entry
setstate	random(3)	better random number generator
setuid	setuid(3)	set user and group ID
setvbuf	setbuf(3S)	assign buffering to a stream
siginterrupt	siginterrupt(3)	allow signals to interrupt system calls
signal	signal(3)	simplified software signal facilities
sin	sin(3M)	trigonometric functions
sinh	sinh(3M)	hyperbolic functions
sleep	sleep(3)	suspend execution for interval
spray	spray(3R)	scatter packets to check network
sprintf	printf(3S)	formatted output conversion
sqrt	sqrt(3M)	square root
srand	rand(3C)	random number generator
srand48	drand48(3)	generate uniformly distributed pseudo-random numbers
srandom	random(3)	better random number generator
sscanf	scanf(3S)	formatted input conversion
ssignal	ssignal(3)	software signals
stdio	intro(3S)	standard buffered input/output package
strcat	string(3)	string operations
strcmp	string(3)	string operations
strcpy	string(3)	string operations
strlen	string(3)	string operations
strncat	string(3)	string operations
strncmp	string(3)	string operations
strncpy	string(3)	string operations
strtod	strtod(3)	convert string to double-precision number
strtol	strtol(3)	convert string to integer
stty	stty(3C)	set and get terminal state
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
tan	sin(3M)	trigonometric functions
tanh	sinh(3M)	hyperbolic functions
tdelete	tsearch(3)	manage binary search trees
telldir	directory(3)	directory operations
tfind	tsearch(3)	manage binary search trees
time	time(3C)	get date and time
times	times(3C)	get process times
timezone	ctime(3)	convert date and time to ASCII
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
tolower	ctype(3)	character classification and conversion macros
toupper	ctype(3)	character classification and conversion macros
tsearch	tsearch(3)	manage binary search trees
ttyname	ttyname(3)	find name of a terminal
ttyslot	ttyname(3)	find name of a terminal
twalk	tsearch(3)	manage binary search trees
ualarm	ualarm(3)	schedule signal after microsecond interval
ulimit	ulimit(3C)	get and set user limits
ungetc	ungetc(3S)	push character back into input stream
usleep	usleep(3S)	suspend execution for micorsecond interval
utime	utime(3C)	set file times
valloc	valloc(3)	aligned memory allocator
values	values(3)	machine-dependent values
varargs	varargs(3)	variable argument list
vfprintf	vfprintf(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
y0	j0(3M)	Bessel functions
y1	j0(3M)	Bessel functions
yn	j0(3M)	Bessel functions
yp_all	ypclnt(3N)	YP client interface routines
yp_bind	ypclnt(3N)	YP client interface routines
yp_first	ypclnt(3N)	YP client interface routines
yp_get_default_domain	ypclnt(3N)	YP client interface routines
yp_master	ypclnt(3N)	YP client interface routines
yp_match	ypclnt(3N)	YP client interface routines
yp_next	ypclnt(3N)	YP client interface routines
yp_order	ypclnt(3N)	YP client interface routines
yp_unbind	ypclnt(3N)	YP client interface routines
ypclnt	ypclnt(3N)	YP client interface routines
yperr_string	ypclnt(3N)	YP client interface routines
yppasswd	yppasswd(3R)	update user YP password
ypprot_err	ypclnt(3N)	YP client interface routines

NAME

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

SYNOPSIS

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

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

DESCRIPTION

to long integer" 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 causes an IOT signal to be sent 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(2)* system call.

SEE ALSO

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

DIAGNOSTICS

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

NAME

abs – integer absolute value

SYNOPSIS

abs(i)
int i;

DESCRIPTION

Abs returns the absolute value of its integer operand.

SEE ALSO

floor(3M) for *fabs*

BUGS

Applying the *abs* function to the most negative integer generates a result which is the most negative integer. That is, *abs(0x80000000)* returns *0x80000000* as a result.

NAME

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 causes an *exit(2)* with a diagnostic comment on the standard output when *expression* is false (0). Compiling with the *cc(1)* option *-DNDEBUG* effectively deletes *assert* from the program.

DIAGNOSTICS

'Assertion failed: file *f* line *n*.' *F* is the source file and *n* the source line number of the *assert* statement.

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*;

***bcmp*(*b1*, *b2*, *length*)**

char **b1*, **b2*;

int *length*;

***bzero*(*b*, *length*)**

char **b*;

int *length*;

***ffs*(*i*)**

int *i*;

DESCRIPTION

The functions *bcopy*, *bcmp*, and *bzero* operate on variable length strings of bytes. They do not check for null bytes as the routines in *string*(3) do.

Bcopy copies *length* bytes from string *b1* to the string *b2*. 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.

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 -1 indicates the value passed is zero.

CAVEAT

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

NAME

crypt, *setkey*, *encrypt* – password and data encryption

SYNOPSIS

```
char *crypt(key, salt)
char *key, *salt;

setkey(key)
char *key;

encrypt(block, edflag)
char *block;
```

DESCRIPTION

crypt is the password encryption routine. It is *based* on the NBS Data Encryption Standard, with variations intended (among other things) to frustrate use of hardware implementations of the DES for key search.

The first argument to *crypt* is normally a user's typed password. The second is a 2-character string chosen from the set [a-zA-Z0-9./]. The *salt* string is used to perturb the DES algorithm in one of 4096 different ways, after which the password is used as the key to encrypt repeatedly a constant string. The returned value points to the encrypted password, in the same alphabet as the salt. The first two characters are the salt itself.

The *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

passwd(1), *passwd*(5), *login*(1), *getpass*(3)

BUGS

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

NAME

ctime, *localtime*, *gmtime*, *asctime*, *timezone*, *dysize* – convert date and time to ASCII

SYNOPSIS

```
char *ctime(clock)
long *clock;

#include <time.h>

struct tm *localtime(clock)
long *clock;

struct tm *gmtime(clock)
long *clock;

char *asctime(tm)
struct tm *tm;

char *timezone(zone, dst)

int dysize(y)
int y;
```

DESCRIPTION

ctime converts to ASCII a long integer, pointed to by *clock*, that represents the time in seconds since Jan. 1, 1970, 00:00, Greenwich Mean Time. It returns a pointer to a 26-character string of the form:

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

Each field has a constant width. *localtime* and *gmtime* return pointers to structures containing the broken-down time. *localtime* corrects for the time zone and possible daylight savings time; *gmtime* converts directly to GMT, which is the time UNIX uses. *asctime* converts a broken-down time to ASCII and returns a pointer to a 26-character string.

Declarations of all the functions and externals, and the “tm” structure, are in the `<time.h>` header file. The structure declaration is:

```
struct tm {
    int tm_sec;           /* seconds (0 - 59) */
    int tm_min;          /* minutes (0 - 59) */
    int tm_hour;         /* hours (0 - 23) */
    int tm_mday;         /* day of month (1 - 31) */
    int tm_mon;          /* month of year (0 - 11) */
    int tm_year;         /* year - 1900 */
    int tm_wday;         /* day of week (Sunday = 0) */
    int tm_yday;         /* day of year (0 - 365) */
    int tm_isdst;
};
```

tm_isdst is non-zero if Daylight Savings Time is in effect.

When local time is called for, the program consults the system to determine the time zone and whether the U.S.A., Canadian, Australian, Eastern European, Middle European, or Western European daylight saving time adjustment is appropriate. The program knows about various peculiarities in time conversion over the past 10-20 years.

timezone returns the name of the time zone associated with its first argument, which is measured in minutes westward from Greenwich. If the second argument is 0, the standard name is used, otherwise the Daylight Savings Time version. If the required name does not appear in a table built into the routine, the difference from GMT is produced; e.g., in Afghanistan *timezone(-(60*4+30), 0)* is appropriate because it is 4:30 ahead of GMT and the string **GMT+4:30** is produced.

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

SEE ALSO

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

BUGS

The return values point to static data, whose contents are overwritten by each call.

NAME

`ctype`, `isalpha`, `isupper`, `islower`, `isdigit`, `isxdigit`, `isalnum`, `isspace`, `ispunct`, `isprint`, `isctrl`, `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)`).

isalpha(c) *c* is a letter

isupper(c) *c* is an upper case letter

islower(c) *c* is a lower case letter

isdigit(c) *c* is a digit [0-9].

isxdigit(c) *c* is a hexadecimal digit [0-9], [A-F], or [a-f].

isalnum(c) *c* is an alphanumeric character, that is, *c* is a letter or a digit

isspace(c) *c* is a space, tab, carriage return, newline, vertical tab, or formfeed

ispunct(c) *c* is a punctuation character (neither control nor alphanumeric)

isprint(c) *c* is a printing character, code 040(8) (space) through 0176 (tilde)

isctrl(c) *c* is a delete character (0177) or ordinary control character (less than 040).

isascii(c) *c* is an ASCII character, code less than 0200

isgraph(c) *c* is a visible graphic character, code 041 (exclamation mark) through 0176 (tilde).

CHARACTER CONVERSION MACROS

These macros perform simple conversions on single characters.

toupper(c) converts *c* to its upper-case equivalent. Note that this *only* works where *c* is known to be a lower-case character to start with (presumably checked via `islower`).

tolower(c) converts *c* to its lower-case equivalent. Note that this *only* works where *c* is known to be an upper-case character to start with (presumably checked via `isupper`).

toascii(c) masks *c* with the correct value so that *c* is guaranteed to be an ASCII character in the range 0 thru 0x7f.

DIAGNOSTICS

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

SEE ALSO

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

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 Data Encryption Standard (DES). 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 Electronic Code Book (ECB) mode, which encrypts blocks of data independently. `cbc_crypt` encrypts in Cipher Block Chaining (CBC) 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.

DIAGNOSTICS

```
DESERR_NONE
    no error.
DESERR_NOHWDEVICE
    encryption succeeded, but done in software instead of the requested hardware.
DESERR_HWERR
    an error occurred in the hardware or driver.
DESERR_BADPARAM
    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.

SEE ALSO

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

NAME

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

SYNOPSIS

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

DIR *opendir(filename)
char *filename;

struct direct *readdir(dirp)
DIR *dirp;

long telldir(dirp)
DIR *dirp;

seekdir(dirp, loc)
DIR *dirp;
long loc;

rewinddir(dirp)
DIR *dirp;

closedir(dirp)
DIR *dirp;
```

DESCRIPTION

opendir opens the directory named by *filename* and associates a *directory stream* with it. *opendir* returns a pointer to be used to identify the *directory stream* in subsequent operations. The pointer NULL is returned if *filename* cannot be accessed or 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.

Sample code which searches a directory for entry ‘name’ is:

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

SEE ALSO

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

NOTES

All UNIX programs that examine directories must be converted to use this package in Sun release 3.0 and beyond. Direct reading of directories is no longer allowed. SH BUGS The new directory format is not obvious.

delim \$\$

NAME

drand48, *erand48*, *lrand48*, *nrand48*, *mrand48*, *jrand48*, *srand48*, *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 mrand48 ( )
long jrand48 (xsubi)
unsigned short xsubi[3];
void srand48 (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 *erand48* return non-negative double-precision floating-point values uniformly distributed over the interval $[0.0, 1.0)$.

Functions *lrand48* and *nrand48* return non-negative long integers uniformly distributed over the interval $[0, 2^{31})$.

Functions *mrand48* and *jrand48* return signed long integers uniformly distributed over the interval $[-2^{31}, 2^{31})$.

Functions *srand48*, *seed48*, and *lcong48* are initialization entry points, one of which should be invoked before either *drand48*, *lrand48*, or *mrand48* is called. (Although it is not recommended practice, constant default initializer values will be supplied automatically if *drand48*, *lrand48*, or *mrand48* is called without a prior call to an initialization entry point.) Functions *erand48*, *nrand48*, and *jrand48* do not require an initialization entry point to be called first.

All the routines work by generating a sequence of 48-bit integer values, $X_{sub\ i}$, according to the linear congruential formula

$$X_{sub\{n+1\}} = (aX_{sub\ n} + c)_{sub\{roman\ mod\ m\}} \quad n \geq 0.$$

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 = 5DEECE66Dsub\ 16 = 273673163155sub\ 8
c = Bsub\ 16 = 13sub\ 8.
```

The value returned by any of the functions *drand48*, *erand48*, *lrand48*, *nrand48*, *mrand48*, or *jrand48* is computed by first generating the next 48-bit $X_{sub\ i}$ 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 $X_{sub\ i}$ and transformed into the returned value.

The functions *drand48*, *lrand48*, and *mrnd48* store the last 48-bit $\$X$ sub i generated in an internal buffer; that is why they must be initialized prior to being invoked. The functions *erand48*, *nrnd48*, and *jrnd48* require the calling program to provide storage for the successive $\$X$ sub 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 $\$X$ sub i into the array and pass it as an argument. By using different arguments, functions *erand48*, *nrnd48*, and *jrnd48* allow separate modules of a large program to generate several *independent* streams of pseudo-random numbers, i.e., the sequence of numbers in each stream will *not* depend upon how many times the routines have been called to generate numbers for the other streams.

The initializer function *srand48* sets the high-order 32 bits of $\$X$ sub i to the 32 bits contained in its argument. The low-order 16 bits of $\$X$ sub i are set to the arbitrary value $\$roman\ 330E$ sub 16 . $\$$

The initializer function *seed48* sets the value of $\$X$ sub i to the 48-bit value specified in the argument array. In addition, the previous value of $\$X$ sub 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 $\$X$ sub 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 $\$X$ sub i , $\$$ the multiplier value $\$a$, $\$$ and the addend value $\$c$. $\$$ Argument array elements *param*[0-2] specify $\$X$ sub 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 *srand48* or *seed48* will restore the “standard” multiplier and addend values, $\$a$ and $\$c$, $\$$ specified on the previous page.

SEE ALSO

rand(3C)

NAME

ecvt, *fcvt*, *gcvt* – output conversion

SYNOPSIS

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

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

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

DESCRIPTION

Ecvt converts the *value* to a null-terminated string of *ndigit* ASCII digits and returns a pointer thereto. The position of the decimal point relative to the beginning of the string is stored indirectly through *decpt* (negative means to the left of the returned digits). If the sign of the result is negative, the word pointed to by *sign* is non-zero, otherwise it is zero. The low-order digit is rounded.

Fcvt is identical to *ecvt*, except that the correct digit has been rounded for Fortran F-format output of the number of digits specified by *ndigits*.

Gcvt converts the *value* to a null-terminated ASCII string in *buf* and returns a pointer to *buf*. It attempts to produce *ndigit* significant digits in Fortran F format if possible, otherwise E format, ready for printing. Trailing zeros may be suppressed.

SEE ALSO

isinf(3), *printf*(3S)

BUGS

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

NAME

end, *etext*, *edata* – last locations in program

SYNOPSIS

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

DESCRIPTION

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

When execution begins, the program break (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)*), the profile (**-p**) option of *cc(1)*, and so on. Thus, the current value of the program break should be determined by *sbrk(0)* (see *brk(2)*).

SEE ALSO

brk(2), *malloc(3)*

NAME

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

SYNOPSIS

```

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

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

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

execlp(name, arg0, arg1, ..., argn, 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 *name* argument is a pointer to the name of the file to be executed. The pointers *arg*[0], *arg*[1] ... address null-terminated strings. Conventionally *arg*[0] is the name of the file.

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

The *execv* version is useful when the number of arguments is unknown in advance; the arguments to *execv* are the name of the file to be executed and a vector of strings containing the arguments. The last argument string must be followed by a 0 pointer.

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

/bin/sh shell, invoked if command file found by *execlp* or *execvp*

SEE ALSO

execve(2), *fork(2)*, *environ(5V)*, *csh(1)*, *sh(1)*

UNIX Programming in Programming Utilities for the Sun Workstation,

UNIX Interface Overview

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

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()
```

FILES

/usr/lib/libU77.a

SEE ALSO

ctime(3), *time*(3F), *idate*(3F)

NAME

frexp, ldexp, modf – floating point analysis and synthesis

SYNOPSIS

double frexp(value, eptr)

double value;

int *eptr;

double ldexp(value, exp)

double value;

int exp;

double modf(value, iptr)

double value, *iptr;

DESCRIPTION

Frexp returns the significand of a double *value* as a double quantity, *x*, of magnitude less than 1 and stores an integer *n*, indirectly through *eptr*, such that $value = x * 2^n$.

The results are not defined when *value* is an IEEE infinity or NaN.

ldexp returns the quantity:

$$value * 2^{exp}.$$

modf returns the positive fractional part of *value* and stores the integer part indirectly through *iptr*. Thus the argument *value* and the returned values *modf* and **iptr* would satisfy, in the absence of rounding error,

$$(*iptr + modf) == value$$

and

$$0 \leq modf < abs(value).$$

The results are not defined when *value* is an IEEE infinity or NaN.

Note that the definition of *modf* varies among Unix implementations; avoid *modf* in portable code.

SEE ALSO

isinf(3)

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 that *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

`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 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

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.

setfsent opens and rewinds the file.

endfsent closes the file.

getfsspec and *getfsfile* sequentially search from the beginning of the file until a matching special file name or file system file name is found, or until EOF is encountered. *getfstype* does likewise, matching on the file system type field.

FILES

/etc/fstab

SEE ALSO

fstab(5)

DIAGNOSTICS

Null pointer (0) returned on EOF or error.

BUGS

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

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 *setgrent* has the effect of rewinding the group file to allow repeated searches. *Endgrent* may be called to close the group file when processing is complete.

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

FILES

```
/etc/group
/etc/yp/domainname/group.byname
/etc/yp/domainname/group.bygid
```

SEE ALSO

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

DIAGNOSTICS

A NULL pointer is returned on EOF or error.

WARNING

The above routines use `<stdio.h>`, 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.

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

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_fname; /* 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 that *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

fstab(5), *getfsent(3)*

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

`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* is a string of recognized option letters; if a letter is followed by a colon, the option is expected to have an argument that may or may not be separated from it by white space. *optarg* is set to point to the start of the option argument on return from `getopt`.

`getopt` places in *optind* the *argv* index of the next argument to be processed. Because *optind* is external, it is normally initialized to zero automatically before the first call to `getopt`.

When all options have been processed (i.e., up to the first non-option argument), `getopt` returns EOF. The special option `—` may be used to delimit the end of the options; EOF will be returned, and `—` will be skipped.

DIAGNOSTICS

`getopt` prints an error message on *stderr* and returns a question mark (?) when it encounters an option letter not included in *optstring*. This error message may be disabled by setting *opterr* to zero.

EXAMPLE

The following code fragment shows how one might process the arguments for a command that can take the mutually exclusive options `a` and `b`, and the options `f` and `o`, both of which require arguments:

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

```
        infile = optarg;
        break;
    case 'o':
        ofile = optarg;
        bufsiza = 512;
        break;
    case '?':
        errflg++;
    }
    if (errflg) {
        fprintf(stderr, "usage: . . . ");
        exit(2);
    }
    for (; optind < argc; optind++) {
        if (access(argv[optind], 4)) {
            .
            .
            .
        }
    }
```

SEE ALSO
getopt(1)

NAME

getpass – read a password

SYNOPSIS

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

DESCRIPTION

getpass reads up to a newline or EOF from the file */dev/tty*, 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 causes it to increase 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

`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

`getpwent`, `getpwuid`, `getpwnam`, `setpwent`, `endpwent`, `fgetpwent` – get password file entry

SYNOPSIS

```
#include <pwd.h>

struct passwd *getpwent()

struct passwd *getpwuid(uid)
int uid;

struct passwd *getpwnam(name)
char *name;

int setpwent()

int endpwent()

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

DESCRIPTION

`getpwent`, `getpwuid` and `getpwnam` each return a pointer to an object with the following structure containing the broken-out fields of a line in the 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 `setpwent` has the effect of rewinding the password file to allow repeated searches. `endpwent` may be called to close the password file when processing is complete.

`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
/etc/yp/domainname/passwd.byname
/etc/yp/domainname/passwd.byuid
```


SEE ALSO

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

DIAGNOSTICS

A NULL pointer is returned on EOF or error.

WARNING

The above routines use `<stdio.h>`, 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

`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.

BUGS

Getwd may fail to return to the current directory if an error occurs.

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

initgroups – initialize group access list

SYNOPSIS

```
initgroups(name, basegid)  
char *name;  
int basegid;
```

DESCRIPTION

Initgroups reads through the group file and sets up, using the *setgroups(2)* call, the group access list for the user specified in *name*. The *basegid* is automatically included in the groups list. Typically this value is given as the group number from the password file.

FILES

/etc/group

SEE ALSO

setgroups(2)

DIAGNOSTICS

Initgroups returns *-1* if it was not invoked by the super-user.

BUGS

Initgroups uses the routines based on *getgrent(3)*. If the invoking program uses any of these routines, the group structure will be overwritten in the call to *initgroups*.

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

isinf, *isnan* – test for indeterminate floating-point values

SYNOPSIS

int *isinf*(value)
double value;

int *isnan*(value)
double value;

DESCRIPTION

Isinf returns a value of 1 if its *value* is an IEEE format infinity (two words 0x7ff00000 0x00000000) or an IEEE negative infinity, and returns a zero otherwise.

*Isn*an returns a value of 1 if its *value* is an IEEE format ‘not-a-number’ (two words 0x7ff nnnnn 0x nnnnnnnn) where *n* is not zero) or its negative, and returns a zero otherwise.

Some library routines such as *ecvt*(3) do not handle indeterminate floating-point values gracefully. Prospective arguments to such routines should be checked with *isinf* or *isnan* before calling these routines.

The *Floating-Point Programmer’s Guide for the Sun Workstation* gives details for the format of IEEE standard floating-point.

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 */

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(2)* 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* will cause the function to return a -1 and set *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 end-of-file. 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(2), *lockd(8C)*

BUGS

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

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;

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. They return a null pointer if the request cannot be completed (see DIAGNOSTICS).

Malloc returns a pointer to a block of at least *size* bytes beginning on a word boundary. A null (0) pointer is returned 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*, *valloc*, or *memalign*.

malloc, *calloc*, *realloc*, *valloc*, 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 that 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 that 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.

SEE ALSO

"Fast Fits" by C. J. Stephenson, 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.

DIAGNOSTICS

Malloc, *calloc*, *realloc*, *valloc*, and *memalign* return a null pointer (0) and set *errno* if arguments are invalid, or if there is insufficient available memory, or if the heap has been detectably corrupted, e.g. by storing outside the bounds of a block.

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 that 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 *Malloc*, *calloc*, *realloc*, *valloc*, *memalign*, *cfree*, and *free* abort with a message to *stderr* 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 *malloc*, *calloc*, *realloc*, *valloc*, *memalign*, *cfree*, and *free*.

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.

ERRORS

Malloc, *calloc*, *realloc*, *valloc*, *memalign*, *cfree*, and *free* will set *errno* if:

- EINVAL An invalid argument was given. The value of *ptr* given to *free*, *cfree*, or *realloc* must be a pointer to a block previously allocated by *malloc*, *calloc*, *realloc*, *valloc*, or *memalign*. The EINVAL condition also occurs if the heap is found to have been corrupted. More detailed information may be obtained by enabling range checks using *malloc_debug*.
- ENOMEM *size* bytes of memory could not be allocated.

FILES

/usr/lib/debug/malloc.o diagnostic versions of *malloc*, *free*, etc.

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.

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

memset 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 optional *<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 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

monitor(lowpc, highpc, buffer, bufsize, nfunc)

int (*lowpc)(), (*highpc)();

short buffer[];

monstartup(lowpc, highpc)

int (*lowpc)(), (*highpc)();

moncontrol(mode)

DESCRIPTION

There are two different forms of monitoring available: An executable program created by:

```
cc -p . . .
```

automatically includes calls for the *prof(1)* monitor and includes an initial call to its start-up routine *monstartup* with default parameters; *monitor* need not be called explicitly except to gain fine control over profil buffer allocation. An executable program created by:

```
cc -pg . . .
```

automatically includes calls for the *gprof(1)* monitor.

Monstartup is a high level interface to *profil(2)*. *Lowpc* and *highpc* specify the address range that is to be sampled; the lowest address sampled is that of *lowpc* and the highest is just below *highpc*. *Monstartup* allocates space using *sbrk(2)* and passes it to *monitor* (see below) to record a histogram of periodically sampled values of the program counter, and of counts of calls of certain functions, in the buffer. Only calls of functions compiled with the profiling option *-p* of *cc(1)* are recorded.

To profile the entire program, it is sufficient to use

```
extern etext();
. . .
monstartup(0x8000, etext);
```

Etext lies just above all the program text, see *end(3)*.

To stop execution monitoring and write the results on the file *mon.out*, use

```
monitor(0);
```

then *prof(1)* can be used to examine the results.

Moncontrol is used to selectively control profiling within a program. This works with either *prof(1)* or *gprof(1)* type profiling. When the program starts, profiling begins. To stop the collection of histogram ticks and call counts use *moncontrol(0)*; to resume the collection of histogram ticks and call counts use *moncontrol(1)*. This allows the cost of particular operations to be measured. Note that an output file will be produced upon program exit irregardless of the state of *moncontrol*.

Monitor is a low level interface to *profil(2)*. *Lowpc* and *highpc* are the addresses of two functions; *buffer* is the address of a (user supplied) array of *bufsize* short integers. At most *nfunc* call counts can be kept. For the results to be significant, especially where there are small, heavily used routines, it is suggested that the buffer be no more than a few times smaller than the range of locations sampled. *Monitor* divides the buffer into space to record the histogram of program counter samples over the range *lowpc* to *highpc*, and space to record call counts of functions compiled with the *-p* option to *cc(1)*.

To profile the entire program, it is sufficient to use

```
extern etext();
. . .
monitor(0x8000, etext, buf, bufsize, nfunc);
```

FILES

mon.out

SEE ALSO

cc(1), prof(1), gprof(1), profil(2), sbrk(2)

NAME

nlist – get entries from name list

SYNOPSIS

```
#include <nlist.h>

nlist(filename, nl)
char *filename;
struct nlist nl[];
```

DESCRIPTION

nlist examines the name list in the executable file whose name is pointed to by *filename*, list of values and puts them in the array of *nlist* structures pointed to by *nl*. The name list *nl* consists of an array of structures containing names, types and values. The list is terminated with a null name; that is, a null string is in the name position of the structure. Each name is looked up in the name list of the file. If the name is found, the type and value of the name are inserted in the next two fields. If the name is not found, both entries are set to 0. See *a.out(5)* for the structure declaration.

This subroutine is useful for examining the system name list kept in the file */vmunix*. In this way programs can obtain system addresses that are up to date.

SEE ALSO

a.out(5)

DIAGNOSTICS

All type entries are set to 0 if the file cannot be read or if does not contain a valid name list.

nlist returns -1 upon error.

NAME

`on_exit` – name termination handler

SYNOPSIS

```
int on_exit(procp, arg)
void (*procp)();
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

```
(*procp)(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 *(*procp)*, 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

`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 Sun Unix and should not be used if portability is a concern.

Standard I/O exit processing is always done last.

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. The argument string *s* is printed first, then a colon and a blank, then the message and a new-line. To be of most use, the argument string should include the name of the program that incurred the error. The error number is taken from the external variable *errno* (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

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 new-line. 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

putenv – change or add value to environment

SYNOPSIS

```
int putenv (string)
char *string;
```

DESCRIPTION

string points to a string of the form “*name=value*.” *putenv* makes the value of the environment variable *name* equal to *value* by altering an existing variable or creating a new one. In either case, the string pointed to by *string* becomes part of the environment, so altering the string will change the environment. The space used by *string* is no longer used once a new string-defining *name* is passed to *putenv*.

DIAGNOSTICS

putenv returns non-zero if it was unable to obtain enough space via *malloc* for an expanded environment, otherwise zero.

SEE ALSO

exec(2), *getenv*(3), *malloc*(3), *environ*(7).

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`.

DIAGNOSTICS

`putpwent` returns non-zero if an error was detected during its operation, otherwise zero.

SEE ALSO

`getpwent(3)`.

WARNING

The above routine uses `<stdio.h>`, which causes it to increase 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

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. *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

bsearch(3), lsearch(3), string(3), sort(1)

NAME

random, *srandom*, *initstate*, *setstate* – better random number generator; routines for changing generators

SYNOPSIS

```
long random()
srandom(seed)
int seed;

char *initstate(seed, state, n)
unsigned seed;
char *state;
int n;

char *setstate(state)
char *state;
```

DESCRIPTION

random uses a non-linear additive feedback random number generator employing a default table of size 31 long integers to return successive pseudo-random numbers in the range from 0 to $2^{31}-1$. The period of this random number generator is very large, approximately $16 \times (2^{31}-1)$.

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.

DIAGNOSTICS

If *initstate* is called with less than 8 bytes of state information, or if *setstate* detects that the state information has been garbled, error messages are printed on the standard error output.

SEE ALSO

rand(3C)

BUGS

About 2/3 the speed of *rand(3C)*.

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 0 if the string *s* was compiled successfully; otherwise a string containing an error message is returned. If *re_comp* is passed 0 or a null string, it returns without changing the currently compiled regular expression.

Re_exec returns 1 if the string *s* matches the last compiled regular expression, 0 if the string *s* failed to match the last compiled regular expression, and -1 if the compiled regular expression was invalid (indicating an internal error).

The strings passed to both *re_comp* and *re_exec* may have trailing or embedded newline characters; they are terminated by nulls. The regular expressions recognized are described in the manual entry for *ed(1)*, given the above difference.

SEE ALSO

ed(1), *ex(1)*, *egrep(1)*, *fgrep(1)*, *grep(1)*

DIAGNOSTICS

Re_exec returns -1 for an internal error.

Re_comp returns one of the following strings if an error occurs:

No previous regular expression

Regular expression too long

*unmatched *

missing]

too many \(\) pairs

unmatched \)

NAME

regex – regular expression compile and match routines

SYNOPSIS

```
#define INIT <declarations>
#define GETC() <getc code>
#define PEEKC() <peekc code>
#define UNGETC(c) <ungetc code>
#define RETURN(pointer) <return code>
#define ERROR(val) <error code>

#include <regex.h>

char *compile (instr, expbuf, endbuf, eof)
char *instr, *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(<i>c</i>)	Cause the argument <i>c</i> to be returned by the next call to GETC() (and PEEKC()). No more than one character of pushback is ever needed and this character is guaranteed to be the last character read by GETC(). The value of the macro UNGETC(<i>c</i>) is always ignored.
RETURN(<i>pointer</i>)	This macro is used on normal exit of the <i>compile</i> routine. The value of the argument <i>pointer</i> is a pointer to the character after the last character of the compiled regular expression. This is useful to programs 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 <code>{ }</code> .
49	<code>[]</code> imbalance.
50	Regular expression overflow.

The syntax of the *compile* routine is as follows:

```
compile(instring, expbuf, endbuf, eof)
```

The first parameter *instring* is never used explicitly by the *compile* routine but is useful for programs that pass down different pointers to input characters. It is sometimes used in the INIT declaration (see below). Programs 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 the opening curly brace `{}`. It is used for dependent declarations and initializations. Most often it is used to set a register variable to point the beginning of the regular expression so that this register variable can be used in the declarations for `GETC()`, `PEEKC()` and `UNGETC()`. Otherwise it can be used to declare external variables that might be used by `GETC()`, `PEEKC()` and `UNGETC()`. See the example below of the declarations taken from *grep(1)*.

There are other functions in this file 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 *locs* is equal to the point in the string at sometime during the backing up process, *advance* will break out of

the loop that backs up and will return zero. This 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*(1):

```
#define INIT      register char *sp = instring;
#define GETC()   (*sp++)
#define PEEKC() (*sp)
#define UNGETC(c)  (—sp)
#define RETURN(c)  return;
#define ERROR(c)   regerr()

#include <regexp.h>
...
                (void) compile(*argv, expbuf, &expbuf[ESIZE], ^0');
...
                if (step(linebuf, expbuf)
                    succeed());
```

FILES

/usr/include/regexp.h

BUGS

The handling of *circf* is kludgy.

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

setjmp, longjmp – non-local goto

SYNOPSIS

```
#include <setjmp.h>
```

```
val = setjmp(env)
```

```
jmp_buf env;
```

```
longjmp(env, val)
```

```
jmp_buf env;
```

```
val = _setjmp(env)
```

```
jmp_buf env;
```

```
_longjmp(env, val)
```

```
jmp_buf env;
```

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*. *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*. 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 machine 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.

The following code fragment indicates the flow of control of the *setjmp* and *longjmp* combination:

```
... function declaration
    jmp_buf my_environment;

    ... code ...
    if (setjmp(my_environment)) {
        this is the code after the return from longjmp
        ... more code ...
        register variables have unpredictable values
        ... more code ...
    } else {
        this is the return from setjmp
        ... more code ...
        Do not modify register variables
        in this leg of the code
        ... more code ...
    }
}
```

Setjmp and *longjmp* save and restore the signal mask *sigsetmask(2)*, while *_setjmp* and *_longjmp* manipulate only the C stack and registers.

SEE ALSO

sigsetmask(2), sigvec(2), signal(3)

BUGS

Setjmp does not save current notion of whether the process is executing on the signal stack. The result is that a longjmp to some place on the signal stack leaves the signal stack state incorrect.

longjmp never returns zero in the Sun implementation.

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

setreuid(2), setregid(2), getuid(2), getgid(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

`siginterrupt` – allow signals to interrupt system calls

SYNOPSIS

```
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 4.2 BSD, and on Sun UNIX 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 UNIX systems, such as 4.1 BSD and System V UNIX. It is the default behavior on Sun UNIX 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 that the new 4.2 BSD 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-4.2 BSD 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*(3) 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

sigvec(2), *sigblock*(2), *sigpause*(2), *sigsetmask*(2).

NAME

signal – simplified software signal facilities

SYNOPSIS

```
#include <signal.h>

(*signal(sig, func))()
int (*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 UNIX 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 *tty*(4)). Signals are optionally generated when a process resumes after being stopped, when the status of child processes changes, or when input is ready at the control terminal. Most signals cause termination of the receiving process if no action is taken; some signals instead cause the process receiving them to be stopped, or are simply discarded if the process has not requested otherwise. Except for the SIGKILL and SIGSTOP signals, the *signal* call allows signals either to be ignored or to cause an interrupt to a specified location. The following is a list of all signals with names as in the include file *<signal.h>*:

SIGHUP	1	hangup
SIGINT	2	interrupt
SIGQUIT	3*	quit
SIGILL	4*	illegal instruction (other than A-line or F-line op code)
SIGTRAP	5*	trace trap
SIGIOT	6*	IOT trap (not generated on Suns)
SIGEMT	7*	EMT trap (A-line or F-line op code)
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 with no one to read it
SIGALRM	14	alarm clock
SIGTERM	15	software termination signal
SIGURG	16•	urgent condition present on socket
SIGSTOP	17†	stop (cannot be caught, blocked, or ignored)
SIGTSTP	18†	stop signal generated from keyboard
SIGCONT	19•	continue after stop (cannot be blocked)
SIGCHLD	20•	child status has changed
SIGTTIN	21†	background read attempted from control terminal
SIGTTOU	22†	background write attempted to control terminal
SIGIO	23•	I/O is possible on a descriptor (see <i>fcntl</i> (2))
SIGXCPU	24	cpu time limit exceeded (see <i>setrlimit</i> (2))
SIGXFSZ	25	file size limit exceeded (see <i>setrlimit</i> (2))
SIGVTALRM	26	virtual time alarm (see <i>setitimer</i> (2))
SIGPROF	27	profiling timer alarm (see <i>setitimer</i> (2))
SIGWINCH	28•	window changed (see <i>win</i> (4S))
SIGLOST	29*	resource lost (see <i>lockd</i> (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, causing the call to terminate prematurely, the call is automatically restarted. In particular this can occur during a *read* 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:

```
handler(sig, code, scp)
int sig, code;
struct sigcontext *scp;
```

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.

CODES

The following defines the codes for signals which produce them. All of these symbols are defined in <signal.h>:

Hardware condition	Signal	Code
Illegal instruction	SIGILL	ILL_INSTR_FAULT
Privilege violation	SIGILL	ILL_PRIVVIO_FAULT
Coprocessor protocol error	SIGILL	ILL_INSTR_FAULT
Trap # <i>n</i> (1 ≤ <i>n</i> ≤ 14)	SIGILL	ILL_TRAP_FAULT
A-line op code	SIGEMT	EMT_EMU1010
F-line op code	SIGEMT	EMT_EMU1111
Integer division by zero	SIGFPE	FPE_INTDIV_TRAP
CHK or CHK2 instruction	SIGFPE	FPE_CHKINST_TRAP
TRAPV or TRAPcc or cpTRAPcc	SIGFPE	FPE_TRAPV_TRAP
IEEE floating point compare unordered	SIGFPE	FPE_FLTBSUN_TRAP
IEEE floating point inexact	SIGFPE	FPE_FLTINEX_TRAP
IEEE floating point division by zero	SIGFPE	FPE_FLTDIV_TRAP
IEEE floating point underflow	SIGFPE	FPE_FLTUND_TRAP
IEEE floating point operand error	SIGFPE	FPE_FLTOPERR_TRAP
IEEE floating point overflow	SIGFPE	FPE_FLTOVF_FAULT
IEEE floating point signaling NaN	SIGFPE	FPE_FLTNAN_TRAP

RETURN VALUE

The previous action is returned on a successful call. Otherwise, -1 is returned and *errno* is set to indicate the error.

ERRORS

signal will fail and no action will take place if one of the following occur:

EINVAL *sig* is not a valid signal number.

EINVAL An attempt is made to ignore or supply a handler for SIGKILL or SIGSTOP.

EINVAL An attempt is made to ignore SIGCONT (by default SIGCONT is ignored).

SEE ALSO

kill(1), ptrace(2), kill(2), sigvec(2), sigblock(2), sigsetmask(2), sigpause(2), sigstack(2), setjmp(3),
tty(4)

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

setitimer(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 *gsignal*. Raising a software signal causes the action established for that signal to be *taken*.

The first argument to *ssignal* is a number identifying the type of signal for which an action is to be established. The second argument defines the action; it is either the name of a (user-defined) *action function* or one of the manifest constants SIG_DFL (default) or SIG_IGN (ignore). *ssignal* returns the action previously established for that signal type; if no action has been established or the signal number is illegal, *ssignal* returns SIG_DFL.

gsignal raises the signal identified by its argument, *sig*:

If an action function has been established for *sig*, then that action is reset to SIG_DFL and the action function is entered with argument *sig*. *gsignal* returns the value returned to it by the action function.

If the action for *sig* is SIG_IGN, *gsignal* returns the value 1 and takes no other action.

If the action for *sig* is SIG_DFL, *gsignal* returns the value 0 and takes no other action.

If *sig* has an illegal value or no action was ever specified for *sig*, *gsignal* returns the value 0 and takes no other action.

SEE ALSO

signal(3)

NAME

string, strcat, strncat, 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;

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 <string.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.

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 terminating null 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* (*strrchr*) 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; the reason for this is also historical.

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

`strtod`, `atof` – convert string to double-precision number

SYNOPSIS

double strtod (str, ptr)

char *str, **ptr;

double atof (str)

char *str;

DESCRIPTION

strtod returns as a double-precision floating-point number the value represented by the character string pointed to by *str*. The string is scanned up to the first unrecognized character.

strtod recognizes an optional string of spaces, then an optional sign, then a string of digits optionally containing a decimal point, then an optional e or E followed by an optional sign or space, followed by an integer.

If the value of *ptr* is not (char **)NULL, a pointer to the character terminating the scan is returned in the location pointed to by *ptr*. If no number can be formed, **ptr* is set to *str*, and zero is returned.

atof(str) is equivalent to *strtod(str, (char **)NULL)*.

SEE ALSO

`ctype(3)`, `scanf(3S)`, `strtol(3)`.

DIAGNOSTICS

If the correct value would cause overflow, \pm HUGE is returned (according to the sign of the value), and *errno* is set to ERANGE.

If the correct value would cause underflow, zero is returned and *errno* is set to ERANGE.

NAME

strtol, *atol*, *atoi* – convert string to integer

SYNOPSIS

long *strtol* (*str*, *ptr*, *base*)

char **str*, ****ptr**;

int *base*;

long *atol* (*str*)

char **str*;

int *atoi* (*str*)

char **str*;

DESCRIPTION

strtol returns as a long integer the value represented by the character string pointed to by *str*. The string is scanned up to the first character inconsistent with the base. Leading “white-space” characters (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

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-ender machines (PDP-11's and VAX'es).

Nbytes should be even.

The *from* and *to* addresses should not overlap in portable programs.

NAME

syslog, openlog, closelog – control system log

SYNOPSIS

```
#include <syslog.h>

openlog(ident, logstat)
char *ident;

syslog(priority, message, parameters ... )
char *message;

closelog()
```

DESCRIPTION

Syslog arranges to write the *message* onto the system log maintained by *syslog(8)*. The message is tagged with *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. This message will be read by *syslog(8)* and output to the system console or files as appropriate.

If special processing is needed, *openlog* can be called to initialize the log file. Parameters are *ident* which is prepended to every message, and *logstat* which is a bit field indicating special status; current values are:

LOG_PID log the process id with each message: useful for identifying instantiations of daemons.

Openlog returns zero on success. If *syslog* cannot send datagrams to *syslog(8)*, then it writes on */dev/console* instead. If */dev/console* cannot be written, standard error is used. In either case, it returns -1.

Closelog can be used to close the log file. It is automatically closed on a successful exec system call (see *execve(2)*).

EXAMPLES

```
syslog(LOG_SALERT, "who: internal error 23");

openlog("serverftp", LOG_PID);
syslog(LOG_INFO, "Connection from host %d", CallingHost);
```

SEE ALSO

syslog(8)

NAME

`system` – issue a shell command

SYNOPSIS

```
system(string)  
char *string;
```

DESCRIPTION

System causes the *string* to be given to *sh*(1) as input as if the string had been typed as a command at a terminal. The current process waits until the shell has completed, then returns the exit status of the shell.

SEE ALSO

`popen`(3S), `execve`(2), `wait`(2)

DIAGNOSTICS

Exit status 127 (may be displayed as "32512") indicates the shell couldn't be executed.

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), *ttys*(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

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(2)`.

Ualarm causes signal SIGALRM see *signal(3)*, to be sent 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 (e.g. 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)`, `setitimer(2)`, `sigpause(2)`, `sigvec(2)`, `signal(3)`, `sleep(3)`, `alarm(3)`, `usleep(3)`

NAME

usleep – suspend execution for interval in microseconds

SYNOPSIS

```
usleep(useconds)
unsigned useconds;
```

DESCRIPTION

for interval in microseconds" The current process is suspended from execution 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(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

setitimer(2), *getitimer(2)*, *sigpause(2)*, *ualarm(3)*, *sleep(3)*, *alarm(3)*

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.

BITS (<i>type</i>)	The number of bits in a specified type (e.g., int).
HIBITS	The value of a short integer with only the high-order bit set (in most implementations, 0x8000).
HIBITL	The value of a long integer with only the high-order bit set (in most implementations, 0x80000000).
HIBITI	The value of a regular integer with only the high-order bit set (usually the same as HIBITS or HIBITL).
MAXSHORT	The maximum value of a signed short integer (in most implementations, 0x7FFF \equiv 32767).
MAXLONG	The maximum value of a signed long integer (in most implementations, 0x7FFFFFFF \equiv 2147483647).
MAXINT	The maximum value of a signed regular integer (usually the same as MAXSHORT or MAXLONG).
MAXFLOAT, LN_MAXFLOAT	The maximum value of a single-precision floating-point number, and its natural logarithm.
MAXDOUBLE, LN_MAXDOUBLE	The maximum value of a double-precision floating-point number, and its natural logarithm.
MINFLOAT, LN_MINFLOAT	The minimum positive value of a single-precision floating-point number, and its natural logarithm.
MINDOUBLE, LN_MINDOUBLE	The minimum positive value of a double-precision floating-point number, and its natural logarithm.
FSIGNIF	The number of significant bits in the mantissa of a single-precision floating-point number.
DSIGNIF	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(pvar);
f = va_arg(pvar, type);
va_end(pvar);
```

DESCRIPTION

This set of macros provides a means of writing portable procedures that accept variable argument lists. Routines having variable argument lists (such as *printf(3S)*) but do not use *varargs* are inherently non-portable, since different machines use different argument passing conventions.

`va_alist` is used in a function header to declare a variable argument list.

`va_dcl` is a declaration for `va_alist`. 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*. *type* is the type to which the expected argument will be converted when passed as an argument. 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);
}
```

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

intro – introduction to compatibility library functions

DESCRIPTION

These functions constitute the compatibility library portion of *libc*. They are automatically loaded as needed by the C compiler *cc*(1). The link editor searches this library under the “-lc” option. Use of these routines (instead of newer equivalent routines) is encouraged for the sake of program portability. Manual entries for the functions in this library describe the proper routine to use.

LIST OF FUNCTIONS

<i>Name</i>	<i>Appears on Page</i>	<i>Description</i>
alarm	alarm(3C)	schedule signal after specified time
clock	clock(3C)	report CPU time used
ftime	time(3C)	get date and time
gtty	stty(3C)	set and get terminal state
nice	nice(3C)	set program priority
pause	pause(3C)	stop until signal
rand	rand(3C)	random number generator
srand	rand(3C)	random number generator
stty	stty(3C)	set and get terminal state
time	time(3C)	get date and time
times	times(3C)	get process times
ulimit	ulimit(3C)	get and set user limits
utime	utime(3C)	set file times
vlimit	vlimit(3C)	control maximum system resource consumption
vtimes	vtimes(3C)	get information about resource utilization

NAME

alarm – schedule signal after specified time

SYNOPSIS

alarm(seconds)
unsigned seconds;

DESCRIPTION

Alarm causes signal SIGALRM, see *sigvec(2)*, to be sent to the invoking process in a number of seconds given by the argument. Unless caught or ignored, the signal terminates the process.

Alarm requests are not stacked; successive calls reset the alarm clock. If the argument is 0, any alarm request is canceled. Because of scheduling delays, resumption of execution of when the signal is caught may be delayed an arbitrary amount. The longest specifiable delay time is 2147483647 seconds.

The return value is the amount of time previously remaining in the alarm clock.

SEE ALSO

sigpause(2), *sigvec(2)*, *signal(3)*, *sleep(3)*, *ualarm(3)*, *usleep(3)*

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

nice – change priority of a process

SYNOPSIS

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), *getpriority(2)*, *setpriority(2)*, *fork(2)*, *renice(8)*

NAME

pause – stop until signal

SYNOPSIS

pause()

DESCRIPTION

Pause never returns normally. It is used to give up control while waiting for a signal from *kill(2)* or an interval timer, see *setitimer(2)*. Upon termination of a signal handler started during a *pause*, the *pause* call will return.

RETURN VALUE

Always returns -1.

ERRORS

Pause always returns:

EINTR The call was interrupted.

SEE ALSO

kill(2), *select(2)*, *sigpause(2)*

NAME

rand, *srand* – 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

stty, *gty* – set and get terminal state

SYNOPSIS

```
#include <sgtty.h>
```

```
stty(fd, buf)  
int fd;  
struct sgttyb *buf;
```

```
gty(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*. *Gty* 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 *gty* call is “*ioctl*(fd, TIOCGETP, buf)”. See *ioctl*(2) and *tty*(4) for an explanation.

DIAGNOSTICS

If the call is successful 0 is returned, otherwise -1 is returned and the global variable *errno* contains the reason for the failure.

SEE ALSO

ioctl(2), *tty*(4)

NAME

time, ftime – get date and time

SYNOPSIS

```
timeofday = time(0)
timeofday = time(tloc)
long *tloc;
#include <sys/types.h>
#include <sys/timeb.h>
ftime(tp)
struct timeb *tp;
```

DESCRIPTION

Time returns the time since 00:00:00 GMT, Jan. 1, 1970, measured in seconds.

If *tloc* is nonnull, the return value is also stored in the place to which *tloc* points.

The *ftime* entry fills in a structure pointed to by its argument, as defined by *<sys/timeb.h>*:

```
struct timeb
{
    time_t   time;
    unsigned short millitm;
    short    timezone;
    short    dstflag;
};
```

The structure contains the time since the epoch in seconds, up to 1000 milliseconds of more-precise interval, the local time zone (measured in minutes of time westward from Greenwich), and a flag that, if nonzero, indicates that Daylight Saving time applies locally during the appropriate part of the year.

SEE ALSO

date(1), gettimeofday(2), settimeofday(2), ctime(3)

NAME

`times` – get process times

SYNOPSIS

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

times(buffer)
struct tms *buffer;
```

DESCRIPTION

This interface is obsoleted by `getrusage(2)`.

`Times` returns time-accounting information for the current process and for the terminated child processes of the current process. All times are in 1/HZ seconds, where HZ is 60.

This is the structure returned by `times`:

```
struct tms {
    time_t  tms_utime;           /* user time */
    time_t  tms_stime;           /* system time */
    time_t  tms_cutime;          /* user time, children */
    time_t  tms_cstime;          /* system time, children */
};
```

The children times are the sum of the children's process times and their children's times.

SEE ALSO

`time(1V)`, `getrusage(2)`, `wait3(2)`, `time(3C)`

NAME

ulimit – get and set user limits

SYNOPSIS

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

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)*.

RETURN VALUE

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

SEE ALSO

brk(2), *setrlimit(2)*, *write(2V)*

NAME

utime – set file times

SYNOPSIS

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

```
utime(file, timep)
```

```
char *file;
```

```
time_t timep[2];
```

DESCRIPTION

The *utime* call uses the ‘accessed’ and ‘updated’ times in that order from the *timep* vector to set the corresponding recorded times for *file*.

The caller must be the owner of the file or the super-user. The ‘inode-changed’ time of the file is set to the current time.

SEE ALSO

utimes(2), *stat(2)*

NAME

`vlimit` – control maximum system resource consumption

SYNOPSIS

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

```
vlimit(resource, value)
```

DESCRIPTION

This facility is superseded by `getrlimit(2)`.

Limits the consumption by the current process and each process it creates to not individually exceed *value* on the specified *resource*. If *value* is specified as `-1`, then the current limit is returned and the limit is unchanged. The resources which are currently controllable are:

LIM_NORAISE A pseudo-limit; if set non-zero then the limits may not be raised. Only the super-user may remove the *noraise* restriction.

LIM_CPU the maximum number of cpu-seconds to be used by each process

LIM_FSIZE the largest single file which can be created

LIM_DATA the maximum growth of the data+stack region via *sbrk(2)* beyond the end of the program text

LIM_STACK the maximum size of the automatically-extended stack region

LIM_CORE the size of the largest core dump that will be created.

LIM_MAXRSS a soft limit for the amount of physical memory (in bytes) to be given to the program. If memory is tight, the system will prefer to take memory from processes which are exceeding their declared **LIM_MAXRSS**.

Because this information is stored in the per-process information this system call must be executed directly by the shell if it is to affect all future processes created by the shell; *limit* is thus a built-in command to *csh(1)*.

The system refuses to extend the data or stack space when the limits would be exceeded in the normal way; a *break* call fails if the data space limit is reached, or the process is killed when the stack limit is reached (since the stack cannot be extended, there is no way to send a signal!).

A file i/o operation which would create a file which is too large will cause a signal **SIGXFSZ** to be generated, this normally terminates the process, but may be caught. When the cpu time limit is exceeded, a signal **SIGXCPU** is sent to the offending process; to allow it time to process the signal it is given 5 seconds grace by raising the cpu time limit.

SEE ALSO

csh(1)

BUGS

If **LIM_NORAISE** is set, then no grace should be given when the cpu time limit is exceeded.

There should be *limit* and *unlimit* commands in *sh(1)* as well as in *csh*.

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)`, `wait3(2)`



NAME

intro – introduction to mathematical library functions and constants

SYNOPSIS

```
#include <math.h>
```

DESCRIPTION

The include file `<math.h>` contains declarations of all the functions in the Math Library *libm* (described in Section 3M), as well as various functions in the C Library (Section 3C) that return floating-point values. Functions in this library are automatically loaded as needed by the Fortran compiler *f77*(1). The link editor searches this library under the “-lm” option.

`<math.h>` also defines the structure and constants used by the *matherr*(3M) error-handling mechanisms, including the following constant used as an error-return value:

HUGE The maximum value of a double-precision floating-point number.

The following mathematical constants are defined for user convenience:

M_E The base of natural logarithms (*e*).

M_LOG2E The base-2 logarithm of *e*.

M_LOG10E The base-10 logarithm of *e*.

M_LN2 The natural logarithm of 2.

M_LN10 The natural logarithm of 10.

M_PI The ratio of the circumference of a circle to its diameter. (There are also several fractions of its reciprocal and its square root.)

M_SQRT2 The positive square root of 2.

M_SQRT1_2 The positive square root of 1/2.

For the definitions of various machine-dependent “constants,” see the description of the `<values.h>` header file.

LIST OF FUNCTIONS

<i>Name</i>	<i>Appears on Page</i>	<i>Description</i>
acos	sin(3M)	inverse trigonometric functions
acosh	asinh(3M)	inverse hyperbolic function
asin	sin(3M)	inverse trigonometric function
asinh	asinh(3M)	inverse hyperbolic function
atan	sin(3M)	inverse trigonometric function
atan2	sin(3M)	inverse trigonometric function
atanh	asinh(3M)	inverse hyperbolic function
cabs	hypot(3M)	complex magnitude
cbrt	sqrt(3M)	cube root
ceil	floor(3M)	ceiling function
copysign	ieee(3M)	copy sign bit
cos	sin(3M)	trigonometric function
cosh	sinh(3M)	hyperbolic function
drem	ieee(3M)	remainder
erf	erf(3M)	error function
erfc	erf(3M)	complementary error function
exp	exp(3M)	exponential function
expm1	exp(3M)	exp(X)-1
fabs	floor(3M)	absolute value function
finite	ieee(3M)	test for finite number
floor	floor(3M)	floor function

hypot	hypot(3M)	Euclidean distance
j0	j0(3M)	Bessel function
j1	j0(3M)	Bessel function
jn	j0(3M)	Bessel function
lgamma	lgamma(3M)	log gamma function
log	exp(3M)	natural logarithm
log10	exp(3M)	common logarithm
log1p	exp(3M)	log(1+X)
logb	ieee(3M)	exponent extraction
matherr	matherr(3M)	math library error-handling routines
pow	exp(3M)	power x**y
rint	floor(3M)	round to nearest integral value
scalb	ieee(eM)	exponent adjustment
sin	sin(3M)	trigonometric function
sinh	sinh(3M)	hyperbolic function
sqrt	sqrt(3M)	square root
tan	sin(3M)	trigonometric function
tanh	sinh(3M)	hyperbolic function
y0	j0(3M)	Bessel function
y1	j0(3M)	Bessel function
yn	j0(3M)	Bessel function

NAME

asinh, *acosh*, *atanh* – inverse hyperbolic functions

SYNOPSIS

```
#include <math.h>
```

```
double asinh(x)
```

```
double x;
```

```
double acosh(x)
```

```
double x;
```

```
double atanh(x)
```

```
double x;
```

DESCRIPTION

These functions compute the designated inverse hyperbolic functions for real arguments. They inherit much of their (roundoff, etc.) error from *log1p*, as described in *exp(3M)*.

SEE ALSO

intro(3M), *exp(3M)*

DIAGNOSTICS

Acosh returns a NaN if the argument is less than 1.

Atanh returns a NaN if the argument has absolute value greater than 1.

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 $1.0 - \text{erf}(x)$.

The entry for erfc is provided because of the extreme loss of relative accuracy if erf(x) is called for large x and the result subtracted from 1. (e.g. for x = 10, 12 places are lost).

SEE ALSO

intro(3M)

NAME

exp, log, log10, pow – exponential, logarithm, power

SYNOPSIS

```
#include <math.h>
```

```
double exp(x)
```

```
double x;
```

```
double expm1(x)
```

```
double x;
```

```
double log(x)
```

```
double x;
```

```
double log10(x)
```

```
double x;
```

```
double log1p(x)
```

```
double x;
```

```
double pow(x, y)
```

```
double x, y;
```

DESCRIPTION

Exp returns the exponential function of x .

Expml returns $\exp(x)-1$ accurately even for tiny x .

Log returns the natural logarithm of x .

Log10 returns the base 10 logarithm.

Log1p returns $\log(1+x)$ accurately even for tiny x ;

Pow returns x^y .

SEE ALSO

hypot(3M), sinh(3M), intro(2)

DIAGNOSTICS

These functions handle exceptional arguments in the spirit of IEEE standard P754 for binary floating point arithmetic. $\text{Log}(x)$ for $x < 0$, $\text{log10}(x)$ for $x < 0$, $\text{pow}(0.0,0.0)$, $\text{pow}(\text{infinity},0.0)$, and $\text{pow}(1.0,\text{infinity})$ are invalid, as is $\text{pow}(x,y)$ if $x < 0$ and y is not an integer value or infinite value; in all these cases NaN function values are returned and `errno` is set to `EDOM`.

NAME

floor, ceil, fabs, rint – absolute value, floor, ceiling and round-to-nearest functions

SYNOPSIS

```
#include <math.h>
```

```
double floor(x)
```

```
double x;
```

```
double ceil(x)
```

```
double x;
```

```
double fabs(x)
```

```
double x;
```

```
double rint(x)
```

```
double x;
```

DESCRIPTION

Fabs returns the absolute value $|x|$.

Floor returns the value of the greatest integer less than or equal to x .

Ceil returns the value of the least integer greater than or equal to x .

Rint returns the value of the integer nearest x in the direction of the prevailing rounding mode.

SEE ALSO

abs(3)

NAME

hypot, cabs – Euclidean distance

SYNOPSIS

```
#include <math.h>
double hypot(x, y)
double x, y;
double cabs(z)
struct { double x, y;} z;
```

DESCRIPTION

Hypot and *cabs* return
 $\text{sqrt}(x*x + y*y)$,
taking precautions against unwarranted overflows.

SEE ALSO

exp(3M) for *sqrt*

NAME

ieee, copysign, drem, finite, logb, scalb – copysign, remainder, exponent manipulations

SYNOPSIS

```
#include <math.h>

double copysign(x,y)
double x,y;

double drem(x,y)
double x,y;

int finite(x)
double x;

double logb(x)
double x;

double scalb(x,n)
double x;
int n;
```

DESCRIPTION

These functions are required for, or recommended by the IEEE standard 754 for floating-point arithmetic.

Copysign(x,y) returns x with its sign changed to y's.

Drem(x,y) returns the remainder $r := x - n*y$ where n is the integer nearest the exact value of x/y; moreover if $|n - x/y| = 1/2$ then n is even. Consequently the remainder is computed exactly and $|r| \leq |y|/2$. But drem(x,0) is exceptional; see below under DIAGNOSTICS.

Finite(x) = 1 just when $-\infty < x < +\infty$,
= 0 otherwise (when $|x| = \infty$ or x is NaN.)

Logb(x) returns x's exponent n, a signed integer converted to double-precision floating-point and so chosen that $1 \leq |x|/2^{**n} < 2$ unless $x = 0$ or (only on machines that conform to IEEE 754) $|x| = \infty$ or x lies between 0 and the Underflow Threshold; see below under "BUGS".

Scalb(x,n) = $x*(2^{**n})$ computed, for integer n, without first computing 2^{**n} .

SEE ALSO

floor(3M), intro(3M)

DIAGNOSTICS

IEEE 754 defines drem(x,0) and drem(∞ ,y) to be invalid operations that produce a NaN.

IEEE 754 defines $\text{logb}(\pm\infty) = \pm\infty$ and $\text{logb}(0) = -\infty$, and requires the latter to signal Division-by-Zero.

IEEE 754 currently specifies that $\text{logb}(\text{denormalized no.}) = \text{logb}(\text{tiniest normalized no. } > 0)$ but the consensus has changed to the specification in the new proposed IEEE standard p854, namely that $\text{logb}(x)$ satisfy

$$1 \leq \text{scalb}(|x|, -\text{logb}(x)) < \text{Radix} \quad \dots = 2 \text{ for IEEE 754}$$

for every x except 0, ∞ and NaN. Almost every program that assumes 754's specification will work correctly if logb follows 854's specification instead.

IEEE 754 requires $\text{copysign}(x, \text{NaN}) = \pm x$ but says nothing else about the sign of a NaN - (Not a Number.)

NAME

j0, j1, jn, y0, y1, yn – Bessel functions

SYNOPSIS

```
#include <math.h>  
double j0(x)  
double x;  
double j1(x)  
double x;  
double jn(n, x)  
double x;  
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.

DIAGNOSTICS

Negative arguments cause *y0*, *y1*, and *yn* to return a huge negative value and set *errno* to EDOM.

NAME

lgamma, gamma – log gamma function

SYNOPSIS

```
#include <math.h>
```

```
double lgamma(x)
```

```
double x;
```

```
double gamma(x)
```

```
double x;
```

DESCRIPTION

Lgamma

lgamma

returns $\ln |\Gamma(x)|$ where

$$\Gamma(x) = \int_0^{\infty} t^{x-1} e^{-t} dt \quad \text{for } x > 0 \text{ and}$$

$$\Gamma(x) = \pi / (\Gamma(1-x) \sin(\pi x)) \quad \text{for } x < 1.$$

The external integer *signgam* returns the sign of $\Gamma(x)$.

Gamma

Gamma returns $\ln |\Gamma(|x|)|$. The sign of $\Gamma(|x|)$ is returned in the external integer *signgam*. The following C program might be used to calculate Γ :

```
y = gamma(x);
#ifdef vax
if (y > 88.0)
#endif
#ifdef sun
if (y > 706.0)
#endif
error();
y = exp(y);
if(signgam)
y = -y;
```

IDIOSYNCRASIES

Do **not** use the expression *signgam*exp(lgamma(x))* to compute $g := \Gamma(x)$. Instead use a program like this (in C):

```
lg = lgamma(x); g = signgam*exp(lg);
```

Only after *lgamma* has returned can *signgam* be correct. Note too that $\Gamma(x)$ must overflow when x is large enough, underflow when $-x$ is large enough, and spawn a division by zero when x is a nonpositive integer.

DIAGNOSTICS

For very large arguments over/underflows will occur inside the *lgamma* routine.

gamma returns a huge value for negative integer arguments.

SEE ALSO

intro(3M)

BUGS

gamma should return a positive indication of error.

Only in the UNIX math library for C was the name gamma ever attached to $\ln\Gamma$. Elsewhere, for instance in IBM's FORTRAN library, the name GAMMA belongs to Γ and the name ALGAMA to $\ln\Gamma$ in single precision; in double the names are DGAMMA and DLGAMA. Why should C be different?

NAME

matherr – math library error-handling function

SYNOPSIS

```
#include <math.h>

int matherr (x)
struct exception *x;
```

DESCRIPTION

matherr is invoked by functions in the Math Library when errors are detected. Users may define their own procedures for handling errors, by including a function named *matherr* in their programs. *matherr* must be of the form described above. When an error occurs, a pointer to the exception structure *x* will be passed to the user-supplied *matherr* function. This structure, which is defined in the `<math.h>` header file, is as follows:

```
struct exception {
    int type;
    char *name;
    double arg1, arg2, retval;
};
```

The element *type* is an integer describing the type of error that has occurred, from the following list of constants (defined in the header file):

DOMAIN	argument domain error
SING	argument singularity
OVERFLOW	overflow range error
UNDERFLOW	underflow range error
TLOSS	total loss of significance
PLOSS	partial loss of significance

The element *name* points to a string containing the name of the function that incurred the error. The variables *arg1* and *arg2* are the arguments with which the function was invoked. *retval* is set to the default value that will be returned by the function unless the user's *matherr* sets it to a different value.

If the user's *matherr* function returns non-zero, no error message will be printed, and *errno* will not be set.

If *matherr* is not supplied by the user, the default error-handling procedures, described with the math functions involved, will be invoked upon error. These procedures are also summarized in the table below. In every case, *errno* is set to EDOM or ERANGE and the program continues.

NOTE

In the Sun environment, the facilities provided by *matherr* are only available when a program is built with the software floating point library, as there would be a substantial performance penalty imposed by providing these facilities with the libraries that support various Sun floating point hardware options.

EXAMPLE

```
#include <math.h>

int
matherr(x)
register struct exception *x;
{
    switch (x->type) {
    case DOMAIN:
        /* change sqrt to return sqrt(-arg1), not 0 */
        if (!strcmp(x->name, "sqrt")) {
            x->retval = sqrt(-x->arg1);
            return (0); /* print message and set errno */
        }
    }
}
```

```

    }
case SING:
    /* all other domain or sing errors, print message and abort */
    fprintf(stderr, "domain error in %s\n", x->name);
    abort( );
case PLOSS:
    /* print detailed error message */
    fprintf(stderr, "loss of significance in %s(%g) = %g\n",
            x->name, x->arg1, x->retval);
    return (1); /* take no other action */
}
return (0); /* all other errors, execute default procedure */
}

```

ERROR HANDLING

DEFAULT ERROR HANDLING PROCEDURES						
	<i>Types of Errors</i>					
type	DOMAIN	SING	OVERFLOW	UNDERFLOW	TLOSS	PLOSS
<i>errno</i>	EDOM	EDOM	ERANGE	ERANGE	ERANGE	ERANGE
BESSEL: y0, y1, yn (arg ≤ 0)	- M, -H	- -	- -	- -	M, 0 -	* -
EXP:	-	-	H	0	-	-
LOG, LOG10: (arg < 0) (arg = 0)	M, -H -	- M, -H	- -	- -	- -	- -
POW: neg ** non-int 0 ** non-pos	- M, 0	- -	±H -	0 -	- -	- -
SQRT:	M, 0	-	-	-	-	-
GAMMA:	-	M, H	H	-	-	-
HYPOT:	-	-	H	-	-	-
SINH:	-	-	±H	-	-	-
COSH:	-	-	H	-	-	-
SIN, COS, TAN: -	-	-	-	M, 0	*	-
ASIN, ACOS, ATAN2: M, 0	-	-	-	-	-	-

ABBREVIATIONS	
*	As much as possible of the value is returned.
M	Message is printed (EDOM error).
H	HUGE is returned.
-H	-HUGE is returned.
±H	HUGE or -HUGE is returned.
0	0 is returned.

NAME

sin, *cos*, *tan*, *asin*, *acos*, *atan*, *atan2* – trigonometric functions

SYNOPSIS

```
#include <math.h>
```

```
double sin(x)
```

```
double x;
```

```
double cos(x)
```

```
double x;
```

```
double asin(x)
```

```
double x;
```

```
double acos(x)
```

```
double x;
```

```
double atan(x)
```

```
double x;
```

```
double atan2(y, x)
```

```
double x, y;
```

DESCRIPTION

Sin, *cos* and *tan* return trigonometric functions of radian arguments.

Asin returns the arc sin in the range $-\pi/2$ to $\pi/2$.

Acos returns the arc cosine in the range 0 to π .

Atan returns the arc tangent of x in the range $-\pi/2$ to $\pi/2$.

Atan2 returns the arc tangent of y/x in the range $-\pi$ to π .

DIAGNOSTICS

These functions handle exceptional arguments in the spirit of IEEE standard P754 for binary floating point arithmetic. When x is infinity in *sin*(x), *cos*(x), or *tan*(x), or when $|x| > 1$ in *asin*(x) or *acos*(x), the functions return NaN values and *errno* is set to EDOM.

NAME

sinh, cosh, tanh – hyperbolic functions

SYNOPSIS

```
#include <math.h>
```

```
double sinh(x)
```

```
double x;
```

```
double cosh(x)
```

```
double x;
```

```
double tanh(x)
```

```
double x;
```

DESCRIPTION

These functions compute the designated hyperbolic functions for real arguments.

DIAGNOSTICS

These functions handle exceptional arguments in the spirit of IEEE standard P754 for binary floating point arithmetic. Thus *sinh* and *cosh* return infinity on overflow.

NAME

sqrt, cbrt – cube root, square root

SYNOPSIS

```
#include <math.h>
```

```
double cbrt(x)
```

```
double x;
```

```
double sqrt(x)
```

```
double x;
```

DESCRIPTION

Cbrt(x) returns the cube root of x.

Sqrt(x) returns the square root of x.

SEE ALSO

intro(3M)

DIAGNOSTICS**ERROR** (due to Roundoff etc.)

Cbrt is accurate to within 0.7 *ulps*.

Sqrt on a machine that conforms to IEEE 754 is correctly rounded in accordance with the rounding mode in force; the error is less than half an *ulp* in the default mode (round-to-nearest). An *ulp* is one *Unit* in the *Last Place* carried.



NAME

intro – introduction to network library functions

DESCRIPTION

This section describes functions that are applicable to the DARPA Internet network, which are part of the standard C library.

LIST OF FUNCTIONS

<i>Name</i>	<i>Appears on Page</i>	<i>Description</i>
endhostent	gethostent(3N)	get network host entry
endnetent	getnetent(3N)	get network entry
endprotoent	getprotoent(3N)	get protocol entry
endservent	getservent(3N)	get service entry
gethostbyaddr	gethostent(3N)	get network host entry
gethostbyname	gethostent(3N)	get network host entry
gethostent	gethostent(3N)	get network host entry
getnetbyaddr	getnetent(3N)	get network entry
getnetbyname	getnetent(3N)	get network entry
getnetent	getnetent(3N)	get network entry
getprotobyname	getprotoent(3N)	get protocol entry
getprotobynumber	getprotoent(3N)	get protocol entry
getprotoent	getprotoent(3N)	get protocol entry
getrpcbyname	getrpcent(3N)	get rpc entry
getrpcbynumber	getrpcent(3N)	get rpc entry
getrpcent	getrpcent(3N)	get rpc entry
getservbyname	getservent(3N)	get service entry
getservbyport	getservent(3N)	get service entry
getservent	getservent(3N)	get service entry
htonl	byteorder(3N)	convert values between host and network byte order
htons	byteorder(3N)	convert values between host and network byte order
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
ntohl	byteorder(3N)	convert values between host and network byte order
ntohs	byteorder(3N)	convert values between host and network byte order
rcmd	rcmd(3N)	routines for returning a stream to a remote command
rexec	rexec(3N)	return stream to a remote command
rresvport	rcmd(3N)	routines for returning a stream to a remote command
ruserok	rcmd(3N)	routines for returning a stream to a remote command
sethostent	gethostent(3N)	get network host entry
setnetent	getnetent(3N)	get network entry
setprotoent	getprotoent(3N)	get protocol entry
setservent	getservent(3N)	get service entry
yp_all	ypclnt(3N)	YP client interface routines
yp_bind	ypclnt(3N)	YP client interface routines
yp_first	ypclnt(3N)	YP client interface routines
yp_get_default_domain	ypclnt(3N)	ypclnt(3N)YP client interface routines
yp_master	ypclnt(3N)	YP client interface routines
yp_match	ypclnt(3N)	YP client interface routines
yp_next	ypclnt(3N)	YP client interface routines
yp_order	ypclnt(3N)	YP client interface routines

yp_unbind	ypclnt(3N)	YP client interface routines
ypclnt	ypclnt(3N)	YP client interface routines
yperr_string	ypclnt(3N)	YP client interface routines
ypprot_err	ypclnt(3N)	YP client interface routines

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 machines such as the Sun these routines are defined as null macros in the include file *<netinet/in.h>*.

These routines are most often used in conjunction with Internet addresses and ports as returned by *gethostent(3N)* and *getservent(3N)*.

SEE ALSO

gethostent(3N), *getservent(3N)*

BUGS

The VAX handles bytes backwards from most everyone else in the world. This is not expected to be fixed in the near future.

NAME

`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

`ether_ntoa`, `ether_aton`, `ether_ntohost`, `ether_hostton`, `ether_line`

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: “`x:x:x:x:x:x`” 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 yellowpages’ maps `ethers.byaddr` and `ethers.byname`)

SEE ALSO

`ethers(5)`

NAME

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

SYNOPSIS

```
#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 EOF is encountered. Host addresses are supplied in network order.

FILES

```
/etc/hosts
/etc/yp/domainname/hosts.byname
/etc/yp/domainname/hosts.byaddr
```

SEE ALSO

hosts(5), ypserv(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. Only the Internet address format is currently understood.

NAME

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

SYNOPSIS

```
#include <netdb.h>

struct netent *getnetent()

struct netent *getnetbyname(name)
char *name;

struct netent *getnetbyaddr(net, 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.

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

endnetent closes the file.

Getnetbyname and *getnetbyaddr* sequentially search from the beginning of the file until a matching net name or net address and type is found, or until EOF is encountered. Network numbers are supplied in host order.

FILES

```
/etc/networks
/etc/yp/domainname/networks.byname
/etc/yp/domainname/networks.byaddr
```

SEE ALSO

networks(5), *ypserv(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.

Only Internet network numbers are currently understood.

NAME

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

SYNOPSIS

```
innnetgr(netgroup, machine, user, domain)  
char *netgroup, *machine, *user, *domain;  
  
setnetgrent(netgroup)  
char *netgroup  
  
endnetgrent()  
  
getnetgrent(machinep, userp, domainp)  
char **machinep, **userp, **domainp;
```

DESCRIPTION

Innnetgr 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 malloc space for the name. This space is released when a endnetgrent call is made. Getnetgrent returns 1 if it succeeding in obtaining another member of the network group, 0 if it has reached the end of the group.

Setnetgrent 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  
/etc/yp/domain/netgroup  
/etc/yp/domain/netgroup.byuser  
/etc/yp/domain/netgroup.byhost
```

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.

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

endprotoent closes the file.

getprotobyname and *getprotobynumber* sequentially search from the beginning of the file until a matching protocol name or protocol number is found, or until EOF is encountered.

FILES

```
/etc/protocols
/etc/yp/domainname/protocols.byname
/etc/yp/domainname/protocols.bynumber
```

SEE ALSO

protocols(5), yp(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. Only the Internet protocols are currently understood.

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.

Setrpcent 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 EOF is encountered.

FILES

```
/etc/rpc
/etc/yp/domainname/rpc.bynumber
```

SEE ALSO

`rpc(5)`, `rpcinfo(8)`, `ypservices(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

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.

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

endservent closes the file.

getservbyname and *getservbyport* sequentially search from the beginning of the file until a matching protocol name or port number is found, or until EOF is encountered. If a protocol name is also supplied (non-NULL), searches must also match the protocol.

FILES

```
/etc/services
/etc/yp/domainname/services.byname
```

SEE ALSO

getprotoent(3N), *services*(5), *ypserv*(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. Expecting port numbers to fit in a 32 bit quantity is probably naive.

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 that when an Internet address is viewed as a 32-bit integer quantity on the VAX the bytes referred to above appear as “d.c.b.a”. That is, VAX bytes are ordered from right to left.

When a three part address is specified, the last part is interpreted as a 16-bit quantity and placed in the right most two bytes of the network address. This makes the three part address format convenient for specifying Class B network addresses as “128.net.host”.

When a two part address is supplied, the last part is interpreted as a 24-bit quantity and placed in the right most three bytes of the network address. This makes the two part address format convenient for specifying Class A network addresses as ‘net.host’.

When only one part is given, the value is stored directly in the network address without any byte rearrangement.

All numbers supplied as ‘parts’ in a ‘.’ notation may be decimal, octal, or hexadecimal, as specified in the C language (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

`rcmd`, `rresvport`, `ruserok` – routines for returning a stream to a remote command

SYNOPSIS

```
rem = rcmd(ahost, inport, locuser, remuser, cmd, fd2p);
char **ahost;
u_short inport;
char *locuser, *remuser, *cmd;
int *fd2p;

s = rresvport(port);
int *port;

ruserok(rhost, superuser, ruser, luser);
char *rhost;
int superuser;
char *ruser, *luser;
```

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(3N)`, returning `-1` if the host does not exist. Otherwise `*ahost` is set to the standard name of the host and a connection is established to a server residing at the well-known Internet port `inport`.

If the call succeeds, a socket of type `SOCK_STREAM` is returned to the caller, and given to the remote command as `stdin` and `stdout`. If `fd2p` is non-zero, then an auxiliary channel to a control process will be set up, and a descriptor for it will be placed in `*fd2p`. The control process will return diagnostic output from the command (unit 2) on this channel, and will also accept bytes on this channel as being UNIX signal numbers, to be forwarded to the process group of the command. If `fd2p` is 0, then the `stderr` (unit 2 of the remote command) will be made the same as the `stdout` and no provision is made for sending arbitrary signals to the remote process, although you may be able to get its attention by using out-of-band data.

The protocol is described in detail in `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 addresses consist of a port in the range 0 to 1023. Only the super-user is allowed to bind an address of this sort to a socket.

`Ruserok` takes a remote host's name, as returned by a `gethostent(3N)` routine, two user names and a flag indicating if the local user's name is the super-user. It then checks the files `/etc/hosts.equiv` and, possibly, `.rhosts` in the current working directory (normally the local user's home directory) to see if the request for service is allowed. A 0 is returned if the machine name is listed in the "hosts.equiv" file, or the host and remote user name are found in the ".rhosts" file; otherwise `ruserok` returns -1. If the `superuser` flag is 1, the checking of the "host.equiv" file is bypassed.

SEE ALSO

`rlogin(1C)`, `rsh(1C)`, `rexec(3N)`, `rexecd(8C)`, `rlogind(8C)`, `rshd(8C)`

BUGS

There is no way to specify options to the `socket` call which `rcmd` makes.

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*(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 *stdin* and *stdout*. If *fd2p* is non-zero, then an auxiliary channel to a control process will be setup, and a descriptor for it will be placed in **fd2p*. The control process will return diagnostic output from the command (unit 2) on this channel, and will also accept bytes on this channel as being UNIX signal numbers, to be forwarded to the process group of the command. If *fd2p* is 0, then the *stderr* (unit 2 of the remote command) will be made the same as the *stdout* and no provision is made for sending arbitrary signals to the remote process, although you may be able to get its attention by using out-of-band data.

SEE ALSO

rcmd(3N), *rexecd*(8C)

BUGS

There is no way to specify options to the *socket* call which *rexec* makes.

NAME

rpc – library routines for remote procedure calls

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.

FUNCTIONS

auth_destroy()	destroy authentication information handle
authnone_create()	return RPC authentication handle with no checking
authunix_create()	return RPC authentication handle with UNIX permissions
authunix_create_default()	return default UNIX authentication handle
callrpc()	call remote procedure, given [prognum,versnum,procnum]
clnt_broadcast()	broadcast remote procedure call everywhere
clnt_call()	call remote procedure associated with client handle
clnt_destroy()	destroy client's RPC handle
clnt_freeres()	free data allocated by RPC/XDR system when decoding results
clnt_geterr()	copy error information from client handle to error structure
clnt_pcreateerror()	print message to stderr about why client handle creation failed
clnt_permo()	print message to stderr corresponding to condition given
clnt_perror()	print message to stderr about why RPC call failed
clnt_spermo()	print message to a string corresponding to condition given
clnt_sperro()	print message to a string
clntraw_create()	create toy RPC client for simulation
clnttcp_create()	create RPC client using TCP transport
clntudp_create()	create RPC client using UDP transport
get_myaddress()	get the machine's IP address
pmap_getmaps()	return list of RPC program-to-port mappings
pmap_getport()	return port number on which waits supporting service
pmap_rmtcall()	instructs portmapper to make an RPC call
pmap_set()	establish mapping between [prognum,versnum,procnum] and port
pmap_unset()	destroy mapping between [prognum,versnum,procnum] and port
registerrpc()	register procedure with RPC service package
rpc_createerr	global variable indicating reason why client creation failed
svc_destroy()	destroy RPC service transport handle
svc_fds	global variable with RPC service file descriptor mask
svc_freeargs()	free data allocated by RPC/XDR system when decoding arguments
svc_getargs()	decodes the arguments of an RPC request
svc_getcaller()	get the network address of the caller of a procedure
svc_getreq()	returns when all associated sockets have been serviced
svc_register()	associates prognum and versnum with service dispatch procedure
svc_run()	wait for RPC requests to arrive and call appropriate service
svc_sendreply()	send back results of a remote procedure call
svc_unregister()	remove mapping of [prognum,versnum] to dispatch routines
svcerr_auth()	called when refusing service because of authentication error
svcerr_decode()	called when service cannot decode its parameters
svcerr_noproc()	called when service hasn't implemented the desired procedure
svcerr_noprogram()	called when program is not registered with RPC package
svcerr_progvers()	called when version is not registered with RPC package
svcerr_systemerr()	called when service detects system error
svcerr_weakauth()	called when refusing service because of insufficient authentication
svccraw_create()	creates a toy RPC service transport for testing
svctcp_create()	creates an RPC service based on TCP transport

<code>svcudp_create()</code>	creates an RPC service based on UDP transport
<code>xdr_accepted_reply()</code>	generates RPC-style replies without using RPC package
<code>xdr_authunix_parms()</code>	generates UNIX credentials without using RPC package
<code>xdr_callhdr()</code>	generates RPC-style headers without using RPC package
<code>xdr_callmsg()</code>	generates RPC-style messages without using RPC package
<code>xdr_opaque_auth()</code>	describes RPC messages, externally
<code>xdr_pmap()</code>	describes parameters for portmap procedures, externally
<code>xdr_pmaplist()</code>	describes a list of port mappings, externally
<code>xdr_rejected_reply()</code>	generates RPC-style rejections without using RPC package
<code>xdr_replymsg()</code>	generates RPC-style replies without using RPC package
<code>xprt_register()</code>	registers RPC service transport with RPC package
<code>xprt_unregister()</code>	unregisters RPC service transport from RPC package

SEE ALSO

Remote Procedure Call Programming Guide, in *Networking on the Sun Workstation*.

NAME

xdr – library routines for external data representation

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.

FUNCTIONS

xdr_array()	translate arrays to/from external representation
xdr_bool()	translate Booleans to/from external representation
xdr_bytes()	translate counted byte strings to/from external representation
xdr_destroy()	destroy XDR stream and free associated memory
xdr_double()	translate double precision to/from external representation
xdr_enum()	translate enumerations to/from external representation
xdr_float()	translate floating point to/from external representation
xdr_getpos()	return current position in XDR stream
xdr_inline()	invoke the in-line routines associated with XDR stream
xdr_int()	translate integers to/from external representation
xdr_long()	translate long integers to/from external representation
xdr_opaque()	translate fixed-size opaque data to/from external representation
xdr_reference()	chase pointers within structures
xdr_setpos()	change current position in XDR stream
xdr_short()	translate short integers to/from external representation
xdr_string()	translate null-terminated strings to/from external representation
xdr_u_int()	translate unsigned integers to/from external representation
xdr_u_long()	translate unsigned long integers to/from external representation
xdr_u_short()	translate unsigned short integers to/from external representation
xdr_union()	translate discriminated unions to/from external representation
xdr_void()	always return one (1)
xdr_wrapstring()	package RPC routine for XDR routine, or vice-versa
xdrmem_create()	initialize an XDR stream
xdrrec_create()	initialize an XDR stream with record boundaries
xdrrec_endofrecord()	mark XDR record stream with an end-of-record
xdrrec_eof()	mark XDR record stream with an end-of-file
xdrrec_skiprecord()	skip remaining record in XDR record stream
xdrstdio_create()	initialize an XDR stream as standard I/O FILE stream

SEE ALSO

External Data Representation Protocol Specification, in *Networking on the Sun Workstation*.

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

```
#include <rpcsvc/ypclnt.h>

yp_bind(indomain);
char *indomain;

void yp_unbind(indomain)
char *indomain;

yp_get_default_domain(outdomain);
char **outdomain;

yp_match(indomain, inmap, inkey, inkeylen, outval, outvallen)
char *indomain;
char *inmap;
char *inkey;
int inkeylen;
char **outval;
int *outvallen;

yp_first(indomain, inmap, outkey, outkeylen, outval, outvallen)
char *indomain;
char *inmap;
char **outkey;
int *outkeylen;
char **outval;
int *outvallen;

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_all(indomain, inmap, incallback);
char *indomain;
char *inmap;
struct ypall_callback incallback;

yp_order(indomain, inmap, outorder);
char *indomain;
char *inmap;
int *outorder;

yp_master(indomain, inmap, outname);
char *indomain;
char *inmap;
char **outname;

char *yperr_string(icode)
int icode;
```

```

ypprot_err(incode)
unsigned int incode;

```

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, *lib/libc.a*. Refer to *yfiles(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.

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.

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 (e.g., a local file) in cases when YP services are not available.

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 can't 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_match returns the value associated with a passed key. This key must be exact; no pattern matching is available.

yp_first returns the first key-value pair from the named map in the named domain.

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 processed; 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 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 take 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 doesn't want to see any more key-value pairs.

The third parameter to *yp_all* is

```
struct ypoll_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 `<rpsvc/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 won't 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 main-line 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 returns the order number for a map.

yp_master returns the machine name of the master YP server for a map.

yperr_string returns a pointer to an error message string that is null-terminated but contains no period or newline.

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

SEE ALSO

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 been unbound */
#define YPERR_DOMAIN  3 /* can't bind to server on this domain */
#define YPERR_MAP      4 /* no such map in server's domain */
#define YPERR_KEY      5 /* no such key in map */
#define YPERR_YPERR    6 /* internal yp server or client error */
#define YPERR_RESRC    7 /* resource allocation failure */
#define YPERR_NOMORE   8 /* no more records in map database */
#define YPERR_PMAP     9 /* can't communicate with portmapper */
#define YPERR_YPBIND  10 /* can't communicate with ypbind */
#define YPERR_YPSESV  11 /* can't communicate with ypserv */
#define YPERR_NODOM   12 /* local domain name not set */
```


NAME

intro – introduction to RPC service library functions

DESCRIPTION

These functions constitute the RPC service library, *librpcsvc*. In order to get the link editor to load this library, use the `-lrpcsvc` option of `cc`. Declarations for these functions may be obtained from various include files `<rpcsvc/*.h>`.

LIST OF FUNCTIONS

<i>routine</i>	<i>on page</i>	<i>description</i>
ether	ether(3R)	monitor traffic on the Ethernet
getrpcport	getrpcport(3R)	get RPC port number
havedisk	rstat(3R)	determine if remote machine has disk
rex	rex(3r)	remote execution protocol
musers	musers(3R)	return number of users on remote machine
rquota	rquota(3R)	implement quotas on remote machines
rstat	rstat(3R)	get performance data from remote kernel
rusers	musers(3R)	return information about users on remote machine
rwall	rwall(3R)	write to specified remote machines
spray	spray(3R)	scatter data in order to check the network
yppasswd	yppasswd(3R)	update user password in yellow pages

NAME

ether – monitor traffic on the Ethernet

SYNOPSIS

```
#include <rpcsvc/ether.h>
```

RPC INFO

program number:

ETHERPROG

xdr routines:

```
xdr_etherstat(xdrs, es)
    XDR *xdrs;
    struct etherstat *es;
xdr_etheraddrs(xdrs, ea)
    XDR *xdrs;
    struct etheraddrs *ea;
xdr_etherhtable(xdrs, hm)
    XDR *xdrs;
    struct etherhmem **hm;
xdr_etherhmem(xdrs, hm)
    XDR *xdrs;
    struct etherhmem **hm;
xdr_etherhbody(xdrs, hm)
    XDR *xdrs;
    struct etherhmem *hm;
xdr_addrmask(xdrs, am)
    XDR *xdrs;
    struct addrmask *am;
```

Xdr_etherhmem processes a single *etherhmem* structure. *Xdr_etherhtable* processes an array of `HASHSIZE *struct etherhmems`. The `**etherhmem` field of *etheraddrs* is actually a hashtable, that is, it is a pointer to an array of `HASHSIZE hmem` pointers.

procs:

```
ETHERPROC_GETDATA
    no args, returns struct etherstat
ETHERPROC_ON
    no args or results, puts server in promiscuous mode
ETHERPROC_OFF
    no args or results, puts server in promiscuous mode
ETHERPROC_GETSRCDATA
    no args, returns struct etheraddrs with information
    about source of packets
ETHERPROC_GETDSTDATA
    no args, returns struct etheraddrs with information
    about destination of packets
ETHERPROC_SELECTSRC
    takes struct mask as argument, no results
    sets a mask for source
ETHERPROC_SELECTDST
    takes struct mask as argument, no results
    sets a mask for dst
ETHERPROC_SELECTPROTO
    takes struct mask as argument, no results
    sets a mask for proto
```

ETHERPROC_SELECTLNTH

takes struct mask as argument, no results
sets a mask for lnth

versions:

ETHERVERS_ORIG

structures:

```

/*
 * all ether stat's except src, dst addresses
 */
struct etherstat {
    struct timeval    e_time;
    unsigned long    e_bytes;
    unsigned long    e_packets;
    unsigned long    e_bcast;
    unsigned long    e_size[NBUCKETS];
    unsigned long    e_proto[NPROTOS];
};
/*
 * member of address hash table
 */
struct etherhmem {
    int h_addr;
    unsigned h_cnt;
    struct etherhmem *h_nxt;
};
/*
 * src, dst address info
 */
struct etheraddrs {
    struct timeval    e_time;
    unsigned long    e_bytes;
    unsigned long    e_packets;
    unsigned long    e_bcast;
    struct etherhmem **e_addrs;
};
/*
 * for size, a_addr is lowvalue, a_mask is high value
 */
struct addrmask {
    int a_addr;
    int a_mask;      /* 0 means wild card */
};

```

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 return that port number.

NAME

rex – remote execution protocol

SYNOPSIS

```
#include <sys/ioctl.h>
#include <rpcsvc/rex.h>
```

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 that this service is only provided with the TCP transport.

RPC INFO

program number:

REXPROG

xdr routines:

```
int xdr_rex_start(xdrs, start);
    XDR *xdrs;
    struct rex_start *start;
int xdr_rex_result(xdrs, result);
    XDR *xdrs;
    struct rex_result *result;
int xdr_rex_ttymode(xdrs, mode);
    XDR *xdrs;
    struct rex_ttymode *mode;
int xdr_rex_tysize(xdrs, size);
    XDR *xdrs;
    struct tty_size *size;
```

procs:

REXPROC_START

Takes rex_start structure, starts a command executing, and returns a rex_result structure.

REXPROC_WAIT

Takes no arguments, waits for a command to finish executing, and returns a rex_result structure.

REXPROC_MODES

Takes a rex_ttymode structure, and sends the tty modes.

REXPROC_WINCH

Takes a tty_size structure, and sends window size information.

versions:

REXVERS_ORIG

Original version

structures:

```
#define REX_INTERACTIVE          1      /* Interactive mode */
struct rex_start {
    char **rst_cmd;                /* list of command and args */
    char *rst_host;                /* working directory host name */
    char *rst_fname;               /* working directory file system name */
    char *rst_dirwithin;           /* working directory within file system */
    char **rst_env;                /* list of environment */
    u_short rst_port0;             /* port for stdin */
    u_short rst_port1;             /* port for stdin */
    u_short rst_port2;             /* port for stdin */
```

```
        u_long rst_flags;                /* options - see #defines above */
};

struct rex_result {
    int rlt_stat;                        /* integer status code */
    char *rlt_message;                  /* string message for human consumption */
};

struct rex_ttymode {
    struct sgtyb basic;                 /* standard unix tty flags */
    struct tchars more;                 /* interrupt, kill characters, etc. */
    struct ltchars yetmore;             /* special Berkeley characters */
    u_long andmore;                     /* and Berkeley modes */
};
```

SEE ALSO

on(1C), rexd(8C)

NAME

rnusers, rusers – return information about users on remote machines

SYNOPSIS

```
#include <rpcsvc/rusers.h>

rnusers(host)
    char *host

rusers(host, up)
    char *host
    struct utmpidlearr *up;
```

DESCRIPTION

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 relevant structures are:

```
struct utmparr {
    /* RUSERSVERS_ORIG */
    struct utmp **uta_arr;
    int uta_cnt
};

struct utmpidle {
    struct utmp ui_utmp;
    unsigned ui_idle;
};

struct utmpidlearr {
    /* RUSERSVERS_IDLE */
    struct utmpidle **uia_arr;
    int uia_cnt
};
```

RPC INFO

program number:

RUSERSPROG

xdr routines:

```
int xdr_utmp(xdrs, up)
    XDR *xdrs;
    struct utmp *up;
int xdr_utmpidle(xdrs, ui);
    XDR *xdrs;
    struct utmpidle *ui;
int xdr_utmpptr(xdrs, up);
    XDR *xdrs;
    struct utmp **up;
int xdr_utmpidleptr(xdrs, up);
    XDR *xdrs;
    struct utmpidle **up;
int xdr_utmparr(xdrs, up);
    XDR *xdrs;
    struct utmparr *up;
int xdr_utmpidlearr(xdrs, up);
    XDR *xdrs;
    struct utmpidlearr *up;
```

procs:

RUSERSPROC_NUM

No arguments, returns number of users as an *unsigned long*.

RUSERSPROC_NAMES

No arguments, returns *utmparr* or *utmpidlearr*, depending on version number.

RUSERSPROC_ALLNAMES

No arguments, returns *utmparr* or *utmpidlearr*, depending on version number.

Returns listing even for *utmp* entries satisfying *nonuser()* in *utmp.h*.

versions:

RUSERSVERS_ORIG

RUSERSVERS_IDLE

structures:

SEE ALSO

rusers(1C)

NAME

rquota – implement quotas on remote machines

SYNOPSIS

```
#include <rpcsvc/rquota.h>
```

RPC INFO

program number:

RQUOTAPROG

xdr routines:

```
xdr_getquota_args(xdrs, gqa);
    XDR *xdrs;
    struct getquota_args *gqa;
xdr_getquota_rslt(xdrs, gqr);
    XDR *xdrs;
    struct getquota_rslt *gqr;
xdr_rquota(xdrs, rq);
    XDR *xdrs;
    struct rquota *rq;
```

procs:

```
RQUOTAPROC_GETQUOTA
RQUOTAPROC_GETACTIVEQUOTA
    Arguments of struct getquota_args.
    Returns struct getquota_rslt.
    Uses UNIX authentication.
    Returns quota only on filesystems with quota active.
```

versions:

RQUOTAVERS_ORIG

structures:

```
struct getquota_args {
    char *gqa_pathp;        /* path to filesystem of interest */
    int gqa_uid;           /* inquire about quota for uid */
};
/*
 * remote quota structure
 */
struct rquota {
    int rq_bsize;          /* block size for block counts */
    bool_t rq_active;      /* indicates whether quota is active */
    u_long rq_bhardlimit;  /* absolute limit on disk blks alloc */
    u_long rq_bsoftlimit;  /* preferred limit on disk blks */
    u_long rq_curblocks;   /* current block count */
    u_long rq_fhardlimit;  /* absolute limit on allocated files */
    u_long rq_fsoftlimit;  /* preferred file limit */
    u_long rq_curfiles;    /* current # allocated files */
    u_long rq_btimeleft;   /* time left for excessive disk use */
    u_long rq_ftimeleft;   /* time left for excessive files */
};
enum gqr_status {
    Q_OK = 1,              /* quota returned */
    Q_NOQUOTA = 2,        /* noquota for uid */
    Q_EPERM = 3           /* no permission to access quota */
};
```

```
struct getquota_rslt {
    enum gqr_status gqr_status; /* discriminant */
    struct rquota gqr_rquota; /* valid if status == Q_OK */
};
```

SEE ALSO

quota(1), quotactl(2)

NAME

rstat, havedisk – get performance data from remote kernel

SYNOPSIS

```
#include <rpcsvc/rstat.h>
```

```
havedisk(host)
```

```
    char *host;
```

```
rstat(host, statp)
```

```
    char *host;
```

```
    struct statstime *statp;
```

DESCRIPTION

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 relevant structures are:

```
struct stats {
    int cp_time[CPUSTATES];
    int dk_xfer[DK_NDRIVE];
    unsigned v_pgpgin;      /* these are cumulative sum */
    unsigned v_pgpgout;
    unsigned v_pswpin;
    unsigned v_pswpout;
    unsigned v_intr;
    int if_ipackets;
    int if_ierrors;
    int if_opackets;
    int if_oerrors;
    int if_collisions;
};

struct statsswch {
    int cp_time[CPUSTATES];
    int dk_xfer[DK_NDRIVE];
    unsigned v_pgpgin;      /* these are cumulative sum */
    unsigned v_pgpgout;
    unsigned v_pswpin;
    unsigned v_pswpout;
    unsigned v_intr;
    int if_ipackets;
    int if_ierrors;
    int if_opackets;
    int if_oerrors;
    int if_collisions;
    unsigned v_swch;
    long avenrun[3];
    struct timeval boottime
};

struct statstime {
    int cp_time[CPUSTATES];
    int dk_xfer[DK_NDRIVE];
    unsigned v_pgpgin;      /* these are cumulative sum */
    unsigned v_pgpgout;
    unsigned v_pswpin;
    unsigned v_pswpout;
    unsigned v_intr;
```

```

    int if_ipackets;
    int if_ierrors;
    int if_opackets;
    int if_oerrors;
    int if_collisions;
    unsigned v_swch;
    long avenrun[3];
    struct timeval boottime;
    struct timeval curtime;
};
RPC INFO
program number:
    RSTATPROG

xdr routines:
    int xdr_stats(xdrs, stat)
        XDR *xdrs;
        struct stats *stat;
    int xdr_statsswch(xdrs, stat)
        XDR *xdrs;
        struct statsswch *stat;
    int xdr_statstime(xdrs, stat)
        XDR *xdrs;
        struct statstime *stat;
    int xdr_timeval(xdrs, tv)
        XDR *xdrs;
        struct timeval *tv;

procs:
    RSTATPROC_HAVEDISK
        Takes no arguments, returns long which is true if remote host has a disk.
    RSTATPROC_STATS
        Takes no arguments, return struct statsxxx, depending on version.

versions:
    RSTATVERS_ORIG
    RSTATVERS_SWTCH
    RSTATVERS_TIME

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

Rwall causes *host* to print 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(1), *shutdown*(8), *rwalld*(8C)

NAME

spray – scatter data in order to check the network

SYNOPSIS

```
#include <rpcsvc/spray.h>
```

RPC INFO

program number:

SPRAYPROG

xdr routines:

```
xdr_sprayarr(xdrs, arr);
    XDR *xdrs;
    struct sprayarr *arr;
xdr_spraycumul(xdrs, cumul);
    XDR *xdrs;
    struct spraycumul *cumul;
```

procs:

SPRAYPROC_SPRAY

Takes no arguments, returns no value.

Increments a counter in server daemon.

The server does not return this call, so the caller should have a timeout of 0.

SPRAYPROC_GET

Takes no arguments, returns *struct spraycumul* with value of counter and clock.

SPRAYPROC_CLEAR

Takes no arguments and returns no value.

Zeros out counter and clock.

versions:

SPRAYVERS_ORIG

structures:

```
struct spraycumul {
    unsigned counter;
    struct timeval clock;
};
struct sprayarr {
    int *data,
    int lnth
};
```

SEE ALSO

spray(8), sprayd(8)

NAME

yppasswd – update user password in yellow pages

SYNOPSIS

```
#include <rpcsvc/yppasswd.h>

yppasswd(oldpass, newpw)
char *oldpass
struct passwd *newpw;
```

DESCRIPTION

If *oldpass* is indeed the old user password, this routine replaces the password entry with *newpw*. It returns 0 if successful.

RPC INFO

program number:

YPPASSWDPROG

xdr routines:

```
xdr_yppasswd(xdrs, yp)
XDR *xdrs;
struct yppasswd *yp;
xdr_yppasswd(xdrs, pw)
XDR *xdrs;
struct passwd *pw;
```

procs:

YPPASSWDPROC_UPDATE

Takes *struct yppasswd* as argument, returns integer.

Same behavior as *yppasswd()* wrapper.

Uses UNIX authentication.

versions:

YPPASSWDVERS_ORIG

structures:

```
struct yppasswd {
char *oldpass; /* old (unencrypted) password */
struct passwd newpw; /* new pw structure */
};
```

SEE ALSO

yppasswd(1), yppasswdd(8C)

1

NAME

intro, 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* 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*.

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.

For purposes of efficiency, this implementation of the standard library has been changed to line buffer output to a terminal by default and attempts to do this transparently by flushing the output whenever a *read*(2V) from the standard input is necessary. This is almost always transparent, but may cause confusion or malfunctioning of programs which use standard I/O routines but use *read*(2V) themselves to read from the standard input.

In cases where a large amount of computation is done after printing part of a line on an output terminal, it is necessary to *fflush* (see *fclose*(3S)) the standard output before going off and computing so that the output will appear.

BUGS

The standard buffered functions do not interact well with certain other library and system functions, especially *vfork*.

LIST OF FUNCTIONS

<i>Name</i>	<i>Appears on Page</i>	<i>Description</i>
clearerr	ferror(3S)	stream status inquiries
ctermid	ctermid(3S)	generate filename for terminal
cuserid	cuserid(3S)	get character login name of user
fclose	fclose(3S)	close or flush a stream
fdopen	fopen(3S)	open a stream
feof	ferror(3S)	stream status inquiries
ferror	ferror(3S)	stream status inquiries
fflush	fclose(3S)	close or flush a stream
fgetc	getc(3S)	get character or integer from stream
fgets	gets(3S)	get a string from a stream
fileno	ferror(3S)	stream status inquiries
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
freopen	fopen(3S)	open a stream
fscanf	scanf(3S)	formatted input conversion
fseek	fseek(3S)	reposition a stream
ftell	fseek(3S)	reposition a stream
fwrite	fread(3S)	buffered binary input/output
getc	getc(3S)	get character or integer from stream
getchar	getc(3S)	get character or integer from stream
gets	gets(3S)	get a string from a stream
getw	getc(3S)	get character or integer from stream
pclose	popen(3S)	initiate I/O to/from a process
popen	popen(3S)	initiate I/O to/from a process
printf	printf(3S)	formatted output conversion
putc	putc(3S)	put character or word on a stream
putchar	putc(3S)	put character or word on a stream
puts	puts(3S)	put a string on a stream
putw	putc(3S)	put character or word on a stream
rewind	fseek(3S)	reposition a stream
scanf	scanf(3S)	formatted input conversion
setbuf	setbuf(3S)	assign buffering to a stream
setbuffer	setbuf(3S)	assign buffering to a stream
setlinebuf	setbuf(3S)	assign buffering to a stream
sprintf	printf(3S)	formatted output conversion
sscanf	scanf(3S)	formatted input conversion
ungetc	ungetc(3S)	push character back into input stream
vfprintf	vprintf(3S)	print formatted varargs output
vprintf	vprintf(3S)	print formatted varargs output
vsprintf	vprintf(3S)	print formatted varargs output

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

cuserid – get character login name of the user

SYNOPSIS

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

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

```
char *s;
```

DESCRIPTION

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

DIAGNOSTICS

If the login name cannot be found, *cuserid* returns a NULL pointer; if *s* is not a NULL pointer, a null character (`*e0`) will be placed at *s[0]*.

SEE ALSO

`getlogin(3)`, `getpwent(3)`

NAME

fclose, fflush – close or flush a stream

SYNOPSIS

#include <stdio.h>

fclose(stream)

FILE *stream;

fflush(stream)

FILE *stream;

DESCRIPTION

fclose causes any buffered data for the named *stream* to be written out, and the named *stream* to be closed. Buffers allocated by the standard input/output system are freed.

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

fflush causes any buffered data for the named output *stream* to be written out. 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

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 end-of-file 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

fopen(3S), *open(2V)*

NAME

fopen, *freopen*, *fdopen* – open a stream

SYNOPSIS

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

```
FILE *fopen(filename, type)
```

```
char *filename, *type;
```

```
FILE *freopen(filename, type, stream)
```

```
char *filename, *type;
```

```
FILE *stream;
```

```
FILE *fdopen(fildes, type)
```

```
char *type;
```

DESCRIPTION

fopen opens the file named by *filename* and associates a stream with it. *fopen* returns a pointer to be used to identify the stream in subsequent operations.

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

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

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

freopen substitutes the named file in place of the open *stream*. It returns the original value of *stream*. The original stream is closed, regardless of whether the open ultimately succeeds.

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

fdopen associates a stream with a file descriptor. 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), *fclose(3S)*, *fseek(3S)*, *fopen(3V)*

DIAGNOSTICS

fopen and *freopen* return a NULL pointer on failure.

BUGS

In order to support the same number of open files as the system does, *fopen* must allocate additional memory for data structures using *calloc* after 20 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 end-of-file or error condition is encountered while reading *stream*, or if *nitems* items have been read. *fread* leaves the file pointer in *stream*, if defined, pointing to the byte following the last byte read if there is one. *fread* does not change the contents of *stream*.

If the standard output is line-buffered, *fread* flushes its output before reading from the standard input. *This is also true for the standard error.*

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), *putc*(3S), *gets*(3S), *puts*(3S), *printf*(3S), *scanf*(3S), *fread*(3V)

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)*, *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 via *fopen*; in particular, *fseek* may not be used on a terminal, or on a file opened via *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 non-UNIX systems 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

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

fgetc behaves like *getc*, but is a function rather than a macro. *fgetc* runs more slowly than *getc*, but 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

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

DIAGNOSTICS

These functions return the integer constant `EOF` at end-of-file or upon an error. The end-of-file condition is remembered, even on a terminal, and all subsequent attempts to read will return `EOF` until the condition is cleared with *clearerr(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++)* doesn't 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

`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 new-line character is read or an end-of-file condition is encountered. The new-line character is discarded and the string is terminated with a null character. `gets` returns its argument.

`fgets` reads characters from the `stream` into the array pointed to by `s`, until $n-1$ characters are read, a new-line character is read and transferred to `s`, or an end-of-file condition is encountered. The string is then terminated with a null character. `fgets` returns its first argument.

SEE ALSO

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

DIAGNOSTICS

If end-of-file is encountered and no characters have been read, no characters are transferred to `s` and a NULL pointer is returned. If a read error 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

popen, *pclose* – initiate I/O to/from 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*, never *csh*.

SEE ALSO

pipe(2), *fopen*(3S), *fclose*(3S), *system*(3), *wait*(2), *sh*(1)

DIAGNOSTICS

popen returns a NULL pointer if files or processes cannot be created, or the shell cannot be accessed.

pclose returns *-1* if *stream* is not associated with a “*popen ed*” command.

BUGS

If the original and “*popen ed*” 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 <stdarg.h>
int _doprnt(format, args, stream)
char *format;
va_list *args;
FILE *stream;
```

DESCRIPTION

IX string "number conversion" string "number conversion — printf" *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, while *sprintf* returns a pointer to the string. *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. If the field width for an s conversion is preceded by a 0, the string is right adjusted with zero-padding on the left.

A *precision* that gives the minimum number of digits to appear for the d, o, u, x, or X conversions, the number of digits to appear after the decimal point for the e, 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.

An optional l (ell) specifying that a following d, o, u, x, or X conversion character applies to a long integer *arg*. A l before any other conversion character is ignored.

A character that indicates the type of conversion to be applied.

A field width or precision may be indicated by an asterisk (*) instead of a digit string. In this case, an integer *arg* supplies the field width or precision. The *arg* that is actually converted is not fetched until the conversion letter is seen, so the *args* specifying field width or precision must appear *before* the *arg* (if any) to be converted.

The flag characters and their meanings are:

- The result of the conversion will be left-justified within the field.
- + The result of a signed conversion will always begin with a sign (+ or –).
- blank If the first character of a signed conversion is not a sign, a blank will be prefixed to the result. This implies that if the blank and + flags both appear, the blank flag will be ignored.
- # This flag specifies that the value is to be converted to an “alternate form.” For *c*, *d*, *s*, and *u* conversions, the flag has no effect. For *o* conversion, it increases the precision to force the first digit of the result to be a zero. For *x* 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,o,u,x,X* The integer *arg* is converted to signed decimal, unsigned octal, unsigned decimal, 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 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 *d*, in style *f*, or in style *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 *sprintf* are printed as if *putc*(3S) had been called.

EXAMPLES

To print a date and time in the form “Sunday, July 3, 10:02,” where *weekday* and *month* are pointers to null-terminated strings:


```
printf("%s, %s %d, %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 *prinif*, not all implementations have such a routine. We strongly recommend that you use the routines described in *vprintf(3S)* instead.

SEE ALSO

putc(3S), *scanf(3S)*, *ecvt(3)*, *printf(3V)*

BUGS

Very wide fields (>128 characters) fail.

The values "Infinity" and "Nan" cannot be read by *scanf(3S)*.

NAME

`putc`, `putchar`, `fputc`, `putw` – put character or word on a stream

SYNOPSIS

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

```
int putc(c, stream)
```

```
char c;
```

```
FILE *stream;
```

```
putchar(c)
```

```
fputc(c, stream)
```

```
FILE *stream;
```

```
putw(w, stream)
```

```
FILE *stream;
```

DESCRIPTION

`putc` writes the character *c* onto the named output *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 output *stream* (at the position at which the file pointer, if defined, is pointing). The size of a word is the size of an integer and varies from machine to machine. It returns the integer written. `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 new-line character is written or terminal input is requested). `setbuf(3S)`, `setbuffer(3S)`, or `setvbuf(3S)` may be used to change the stream's buffering strategy.

SEE ALSO

`fopen(3S)`, `fclose(3S)`, `getc(3S)`, `puts(3S)`, `printf(3S)`, `fread(3S)`

DIAGNOSTICS

On success, these functions each return the value they have written. On error, these functions return the constant EOF. Because EOF is a valid integer, `error(3S)` should be used to detect `putw` errors.

BUGS

Because it is implemented as a macro, `putc` treats a *stream* argument with side effects improperly. In particular, `putc(c, *f++)`; doesn't work sensibly. `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

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 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.

SEE ALSO

fopen(3S), *putc(3S)*, *printf(3S)*, *ferror(3S)*, *fread(3S)*

NOTES

puts appends a newline while *fputs* does not.

NAME

scanf, fscanf, sscanf – formatted input conversion

SYNOPSIS

```
#include <stdio.h>

scanf(format [, pointer ] ... )
char *format;

fscanf(stream, format [, pointer ] ... )
FILE *stream;
char *format;

sscanf(s, format [, pointer ] ... )
char *s, *format;
```

DESCRIPTION

scanf reads from the standard input stream *stdin*. *fscanf* reads from the named input *stream*. *sscanf* reads from the character string *s*. Each function reads characters, interprets them according to a format, and stores the results in its arguments. Each expects, as arguments, a control string *format*, described below, and a set of *pointer* arguments indicating where the converted input should be stored.

The control string usually contains conversion specifications, which are used to direct interpretation of input sequences. The control string may contain:

1. White-space characters (blanks, tabs, or new-lines) 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 a integer pointer.
- x a hexadecimal integer is expected; the corresponding argument should be an integer pointer.
- e,f,g a floating point number is expected; the next field is converted accordingly and stored through the corresponding argument, which should be a pointer to a *float*. The input format for floating point numbers is an optionally signed string of digits, possibly containing a decimal point, followed by an optional exponent field consisting of an E or e followed by an optional +, -, or space, followed by an integer.
- s a character string is expected; the corresponding argument should be a character pointer 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**, and **x** may be capitalized or 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.

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 that this is different from 0, which means that no conversion was done; if conversion was intended, it was frustrated by an inappropriate character in the input.

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; float x; char name[50];
(void) scanf ("%2d%f%*d %[0-9]", &i, &x, name);
```

with input:

```
56789 0123 56a72
```

will assign 56 to *i*, 789.0 to *x*, skip 0123, and place the string 56\0 in *name*. The next call to *getchar* (see *getc*(3S)) will return a.

SEE ALSO

getc(3S), *printf*(3S), *strtod*(3), *strtol*(3), *scanf*(3V)

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.

scanf cannot read the strings which *printf*(3S) generates for IEEE indeterminate floating point values.

scanf provides no way to convert a number in any arbitrary base (decimal, hex or octal) based on the traditional C conventions (leading 0 or 0x).

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. If the standard stream *stderr* refers to a terminal it is line 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 causes the character array *buf* whose size is determined by the *size* argument to be used 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:

<code>_IOFBF</code>	causes input/output to be fully buffered.
<code>_IOLBF</code>	causes output to be line buffered; the buffer will be flushed when a newline is written, the buffer is full, or input is requested.
<code>_IONBF</code>	causes input/output to be completely unbuffered. If <i>buf</i> is not the NULL pointer, the array it points to will be used for buffering, instead of an automatically allocated buffer. <i>Size</i> 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.

SEE ALSO

`fopen(3S)`, `getc(3S)`, `putc(3S)`, `malloc(3)`, `fclose(3S)`, `puts(3S)`, `printf(3S)`, `fread(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.

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.

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, i.e. *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* 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

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

getc(3S), *setbuf*(3S), *fseek*(3S)

DIAGNOSTICS

Ungetc returns EOF if it can't push a character back.

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`, `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...);
 */
/*VARARGS0*/
void
error(va_alist)
/* Note that the function_name and format arguments cannot be
 * separately declared because of the definition of varargs.
 */
va_dcl
{
    va_list args;
    char *fmt;

    va_start(args);
    /* print out name of function causing error */
    (void)fprintf(stderr, "ERROR in %s: ", va_arg(args, char *));
    fmt = va_arg(args, char *);
    /* print out remainder of message */
    (void)vfprintf(fmt, args);
    va_end(args);
    (void)abort( );
}
```

SEE ALSO
varargs(3)

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 FUNCTIONS

<i>Name</i>	<i>Appears on Page</i>	<i>Description</i>
<code>_tolower</code>	<code>ctype(3V)</code>	character classification and conversion
<code>_toupper</code>	<code>ctype(3V)</code>	character classification and conversion
<code>asctime</code>	<code>ctime(3V)</code>	convert date and time to ASCII
<code>assert</code>	<code>assert(3V)</code>	verify program assertion
<code>ctime</code>	<code>ctime(3V)</code>	convert date and time to ASCII
<code>curses</code>	<code>curses(3V)</code>	CRT screen handling and optimization package
<code>endpwent</code>	<code>getpwent(3V)</code>	get password file entry
<code>fdopen</code>	<code>fopen(3V)</code>	open a stream
<code>feof</code>	<code>ferror(3V)</code>	stream status inquiry
<code>ferror</code>	<code>ferror(3V)</code>	stream status inquiry
<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 inquiry
<code>fopen</code>	<code>fopen(3V)</code>	open a stream
<code>fprintf</code>	<code>printf(3V)</code>	formatted output conversion
<code>fread</code>	<code>fread(3V)</code>	buffered binary input/output
<code>freopen</code>	<code>fopen(3V)</code>	open a stream
<code>fscanf</code>	<code>scanf(3V)</code>	formatted input conversion
<code>fwrite</code>	<code>fread(3V)</code>	buffered binary input/output
<code>getc</code>	<code>getc(3V)</code>	get character or integer from stream
<code>getchar</code>	<code>getc(3V)</code>	get character or integer from stream
<code>getpass</code>	<code>getpass(3V)</code>	read a password
<code>getpwent</code>	<code>getpwent(3V)</code>	get password file entry
<code>getpwnam</code>	<code>getpwent(3V)</code>	get password file entry
<code>getpwuid</code>	<code>getpwent(3V)</code>	get password file entry
<code>getw</code>	<code>getc(3V)</code>	get character or integer from stream
<code>gmtime</code>	<code>ctime(3V)</code>	convert date and time to ASCII
<code>isalnum</code>	<code>ctype(3V)</code>	character classification and conversion
<code>isalpha</code>	<code>ctype(3V)</code>	character classification and conversion
<code>isascii</code>	<code>ctype(3V)</code>	character classification and conversion
<code>iscntrl</code>	<code>ctype(3V)</code>	character classification and conversion
<code>isdigit</code>	<code>ctype(3V)</code>	character classification and conversion
<code>isgraph</code>	<code>ctype(3V)</code>	character classification and conversion
<code>islower</code>	<code>ctype(3V)</code>	character classification and conversion
<code>isprint</code>	<code>ctype(3V)</code>	character classification and conversion
<code>ispunct</code>	<code>ctype(3V)</code>	character classification and conversion
<code>isspace</code>	<code>ctype(3V)</code>	character classification and conversion
<code>isupper</code>	<code>ctype(3V)</code>	character classification and conversion
<code>isxdigit</code>	<code>ctype(3V)</code>	character classification and conversion
<code>localtime</code>	<code>ctime(3V)</code>	convert date and time to ASCII
<code>nice</code>	<code>nice(3V)</code>	change priority of a process
<code>printf</code>	<code>printf(3V)</code>	formatted output conversion

rand	rand(3V)	simple random number generator
scanf	scanf(3V)	formatted input conversion
setbuf	setbuf(3V)	assign buffering to a stream
setbuffer	setbuf(3V)	assign buffering to a stream
setlinebuf	setbuf(3V)	assign buffering to a stream
setpwent	getpwent(3V)	get password file entry
setuid	setuid(3V)	set user ID
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
times	times(3V)	get process and child process times
toascii	ctype(3V)	character classification and conversion
tolower	ctype(3V)	character classification and conversion
toupper	ctype(3V)	character classification and conversion
ttyslot	ttyslot(3V)	find the slot in the utmp file of the current process
tzset	ctime(3V)	convert date and time to ASCII

NAME

assert – verify program assertion

SYNOPSIS

```
#include <assert.h>  
assert (expression)  
int expression; System V"
```

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*(1) 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(1), *abort*(3)

NAME

ctime, *localtime*, *gmtime*, *asctime*, *tzset* – convert date and time to ASCII

SYNOPSIS

```
char *ctime(clock)
long *clock;

#include <time.h>

struct tm *localtime(clock)
long *clock;

struct tm *gmtime(clock)
long *clock;

char *asctime(tm)
struct tm *tm;

extern long timezone;
extern int daylight;
extern char *tzname[2];
void tzset ( )
```

DESCRIPTION

ctime converts to ASCII a long integer, pointed to by *clock*, that represents the time in seconds since Jan. 1, 1970, 00:00, Greenwich Mean Time. It returns a pointer to a 26-character string of the form:

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

Each field has a constant width. *localtime* and *gmtime* return pointers to structures containing the time broken down. *localtime* corrects for the time zone and possible daylight savings time; *gmtime* converts directly to GMT, which is the time UNIX uses. *asctime* converts the broken-down time to ASCII and returns a pointer to a 26-character string.

Declarations of all the functions and externals, and the “tm” structure, are in the *<time.h>* header file. The structure declaration is:

```
struct tm {
    int tm_sec;           /* seconds (0 - 59) */
    int tm_min;          /* minutes (0 - 59) */
    int tm_hour;         /* hours (0 - 23) */
    int tm_mday;         /* day of month (1 - 31) */
    int tm_mon;          /* month of year (0 - 11) */
    int tm_year;         /* year - 1900 */
    int tm_wday;         /* day of week (Sunday = 0) */
    int tm_yday;         /* day of year (0 - 365) */
    int tm_isdst;
};
```

tm_isdst is non-zero if Daylight Savings Time is in effect.

The external *long* variable *timezone* contains the difference, in seconds, between GMT and local standard time (in PST, *timezone* is 8*60*60); the external variable *daylight* is non-zero if and only if Daylight Savings Time conversion should be applied. Its value indicates the type of conversion to apply; it is normally the value returned by *gettimeofday (2)* in the *tz_dsttime* field of the *timezone* structure. The program knows about various peculiarities in time conversion over the past 10-20 years.

The external variable *tzname* is an array of two pointers which contains the names of the current time zone. The first pointer points to a character string which 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 which is the name of the current time zone when Daylight Savings Time is in effect.

If an environment variable named `TZ` is present, *asctime* uses the contents of the variable to override the time zone and conversion rule type supplied by the system. The value of `TZ` must be a three-letter time zone name, followed by a signed number representing the difference between local time and Greenwich Mean Time in hours, followed by an optional three-letter name for a daylight time zone. For example, the setting for California would be `PST8PDT`. The effects of setting `TZ` are thus to change the values of the external variables *timezone*, *daylight*, and *tzname*. The function *tzset* sets these external variables from `TZ` or, if `TZ` is not present in the environment, the values supplied by the system. *tzset* is called by *asctime* and may also be called explicitly by the user.

SEE ALSO

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

BUGS

The return values point to static data, whose contents are overwritten by each call.

NAME

`ctype`, `isalpha`, `isupper`, `islower`, `isdigit`, `isxdigit`, `isalnum`, `isspace`, `ispunct`, `isprint`, `iscntrl`, `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(3S)*).

`isalpha(c)` *c* is a letter

`isupper(c)` *c* is an upper case letter

`islower(c)` *c* is a lower case letter

`isdigit(c)` *c* is a digit [0-9].

`isxdigit(c)` *c* is a hexadecimal digit [0-9], [A-F], or [a-f].

`isalnum(c)` *c* is an alphanumeric character, that is, *c* is a letter or a digit

`isspace(c)` *c* is a space, tab, carriage return, newline, vertical tab, or formfeed

`ispunct(c)` *c* is a punctuation character (neither control nor alphanumeric)

`isprint(c)` *c* is a printing character, code 040(8) (space) through 0176 (tilde)

`iscntrl(c)` *c* is a delete character (0177) or ordinary control character (less than 040).

`isascii(c)` *c* is an ASCII character, code less than 0200

`isgraph(c)` *c* 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.

`toupper(c)` converts *c* to its upper-case equivalent. If *c* is not a lower-case letter, it is returned unchanged.

`tolower(c)` converts *c* to its lower-case equivalent. If *c* is not an upper-case letter, it is returned unchanged.

`toascii(c)` masks *c* with the correct value so that *c* is guaranteed to be an ASCII character in the range 0 thru 0x7f.

These macros perform simple conversions on single characters.

`_toupper(c)` converts *c* to its upper-case equivalent. Note that this *only* works where *c* is known to be a lower-case character to start with (presumably checked via `islower`).

`_tolower(c)` converts *c* to its lower-case equivalent. Note that this *only* works where *c* is known to be an upper-case character to start with (presumably checked via `isupper`).

DIAGNOSTICS

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

SEE ALSO

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

NAME

curses – CRT screen handling and optimization package

SYNOPSIS

```
#include <curses.h>
/usr/5bin/cc [ flags ] files -lcurses [ libraries ]
```

DESCRIPTION

These routines give the user a method of updating screens with reasonable optimization. In order to initialize the routines, the routine *initscr()* must be called before any of the other routines that deal with windows and screens are used. The routine *endwin()* should be called before exiting. To get character-at-a-time input without echoing, (most interactive, screen oriented-programs want this) after calling *initscr()* you should call “*nonl(); cbreak(); noecho();*”

The full curses interface permits manipulation of data structures called *windows* which can be thought of as two dimensional arrays of characters representing all or part of a CRT screen. A default window called *stdscr* is supplied, and others can be created with *newwin*. Windows are referred to by variables declared “WINDOW *”, the type WINDOW is defined in *curses.h* to be a C structure. These data structures are manipulated with functions described below, among which the most basic are *move* and *addch*. (More general versions of these functions are included with names beginning with ‘w’, allowing you to specify a window. The routines not beginning with ‘w’ affect *stdscr*.) Then *refresh* is called, telling the routines to make the user’s CRT screen look like *stdscr*.

Mini-Curses is a subset of curses which does not allow manipulation of more than one window. To invoke this subset, use *-DMINICURSES* as a cc option. This level is smaller and faster than full curses.

If the environment variable *TERMINFO* is defined, any program using curses will check for a local terminal definition before checking in the standard place. For example, if the standard place is */usr/5lib/terminfo*, and *TERM* is set to “vt100”, then normally the compiled file is found in */usr/5lib/terminfo/v/vt100*. (The “v” is copied from the first letter of “vt100” to avoid creation of huge directories.) However, if *TERMINFO* is set to */usr/mark/myterms*, curses will first check */opusr/mark/myterms/v/vt100*, and if that fails, will then check */usr/5lib/terminfo/v/vt100*. This is useful for developing experimental definitions or when write permission in */usr/5lib/terminfo* is not available.

SEE ALSO

ioctl(2), *getenv(3)*, *tty(4)*, *terminfo(5V)*

FUNCTIONS

Routines listed here may be called when using the full curses. Those marked with an asterisk may be called when using Mini-Curses.

<i>addch(ch)*</i>	add a character to <i>stdscr</i> (like <i>putchar</i>) (wraps to next line at end of line)
<i>addstr(str)*</i>	calls <i>addch</i> with each character in <i>str</i>
<i>attroff(attrs)*</i>	turn off attributes named
<i>attron(attrs)*</i>	turn on attributes named
<i>attrset(attrs)*</i>	set current attributes to <i>attrs</i>
<i>baudrate()*</i>	current terminal speed
<i>beep()*</i>	sound beep on terminal
<i>box(win, vert, hor)</i>	draw a box around edges of <i>win</i> <i>vert</i> and <i>hor</i> are chars to use for <i>vert</i> . and <i>hor</i> . edges of box
<i>clear()</i>	clear <i>stdscr</i>
<i>clearok(win, bf)</i>	clear screen before next redraw of <i>win</i>
<i>clrtoobot()</i>	clear to bottom of <i>stdscr</i>
<i>clrtoeol()</i>	clear to end of line on <i>stdscr</i>
<i>cbreak()*</i>	set <i>cbreak</i> mode

delay_output(ms)*	insert ms millisecond pause in output
delch()	delete a character
deleteln()	delete a line
delwin(win)	delete <i>win</i>
doupdate()	update screen from all wnooutrefresh
echo()*	set echo mode
endwin()*	end window modes
erase()	erase <i>stdscr</i>
erasechar()	return user's erase character
fixterm()	restore tty to "in curses" state
flash()	flash screen or beep
flushinp()*	throw away any typeahead
getch()*	get a char from tty
getstr(str)	get a string through <i>stdscr</i>
getmode()	establish current tty modes
getyx(win, y, x)	get (y, x) co-ordinates
has_ic()	true if terminal can do insert character
has_il()	true if terminal can do insert line
idlok(win, bf)*	use terminal's insert/delete line if bf != 0
inch()	get char at current (y, x) co-ordinates
initscr()*	initialize screens
insch(c)	insert a char
insertln()	insert a line
intrflush(win, bf)	interrupts flush output if bf is TRUE
keypad(win, bf)	enable keypad input
killchar()	return current user's kill character
leaveok(win, flag)	OK to leave cursor anywhere after refresh if flag!=0 for <i>win</i> , otherwise cursor must be left at current position.
longname()	return verbose name of terminal
meta(win, flag)*	allow meta characters on input if flag != 0
move(y, x)*	move to (y, x) on <i>stdscr</i>
mvaddch(y, x, ch)	move(y, x) then addch(ch)
mvaddstr(y, x, str)	similar...
mvcur(oldrow, oldcol, newrow, newcol)	low level cursor motion
mvdelch(y, x)	like delch, but move(y, x) first
mvgetch(y, x)	etc.
mvgetstr(y, x)	
mvinch(y, x)	
mvinsch(y, x, c)	
mvprintw(y, x, fmt, args)	
mvscanw(y, x, fmt, args)	
mvwaddch(win, y, x, ch)	
mvwaddstr(win, y, x, str)	
mvwdelch(win, y, x)	
mvwgetch(win, y, x)	
mvwgetstr(win, y, x)	
mvwin(win, by, bx)	
mvwinch(win, y, x)	
mvwinsch(win, y, x, c)	
mvwprintw(win, y, x, fmt, args)	
mvwscanw(win, y, x, fmt, args)	

<code>newpad(nlines, ncols)</code>	create a new pad with given dimensions
<code>newterm(type, fd)</code>	set up new terminal of given type to output on fd
<code>newwin(lines, cols, begin_y, begin_x)</code>	
<code>nl()*</code>	create a new window
<code>nocbreak()*</code>	set newline mapping
<code>nodelay(win, bf)</code>	unset cbreak mode
<code>noecho()*</code>	enable nodelay input mode through getch
<code>nonl()*</code>	unset echo mode
<code>noraw()*</code>	unset newline mapping
<code>overlay(win1, win2)</code>	unset raw mode
<code>overwrite(win1, win2)</code>	overlay win1 on win2
<code>pnoutrefresh(pad, pminrow, pmincol, sminrow, smincol, smaxrow, smaxcol)</code>	overwrite win1 on top of win2
	like prefresh but with no output until doupdate called
<code>prefresh(pad, pminrow, pmincol, sminrow, smincol, smaxrow, smaxcol)</code>	
	refresh from pad starting with given upper left corner of pad with output to given portion of screen
<code>printw(fmt, arg1, arg2, ...)</code>	
	printf on <i>stdscr</i>
<code>raw()*</code>	set raw mode
<code>refresh()*</code>	make current screen look like <i>stdscr</i>
<code>resetterm()*</code>	set tty modes to "out of curses" state
<code>resetty()*</code>	reset tty flags to stored value
<code>saveterm()*</code>	save current modes as "in curses" state
<code>savetty()*</code>	store current tty flags
<code>scanw(fmt, arg1, arg2, ...)</code>	
	scanf through <i>stdscr</i>
<code>scroll(win)</code>	scroll <i>win</i> one line
<code>scrollok(win, flag)</code>	allow terminal to scroll if flag != 0
<code>set_term(new)</code>	now talk to terminal new
<code>setscrreg(t, b)</code>	set user scrolling region to lines t through b
<code>setterm(type)</code>	establish terminal with given type
<code>setupterm(term, filenum, errret)</code>	
<code>standend()*</code>	clear standout mode attribute
<code>standout()*</code>	set standout mode attribute
<code>subwin(win, lines, cols, begin_y, begin_x)</code>	
	create a subwindow
<code>touchwin(win)</code>	"change" all of <i>win</i>
<code>traceoff()</code>	turn off debugging trace output
<code>traceon()</code>	turn on debugging trace output
<code>typeahead(fd)</code>	use file descriptor fd to check typeahead
<code>unctrl(ch)*</code>	printable version of <i>ch</i>
<code>waddch(win, ch)</code>	add char to <i>win</i>
<code>waddstr(win, str)</code>	add string to <i>win</i>
<code>wattroff(win, attrs)</code>	turn off <i>attrs</i> in <i>win</i>
<code>wattron(win, attrs)</code>	turn on <i>attrs</i> in <i>win</i>
<code>wattrset(win, attrs)</code>	set <i>attrs</i> in <i>win</i> to <i>attrs</i>
<code>wclear(win)</code>	clear <i>win</i>
<code>wclrtoBot(win)</code>	clear to bottom of <i>win</i>
<code>wclrtoeol(win)</code>	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 (y, x) in <i>win</i>
winsch(win, c)	insert char into <i>win</i>
winsertln(win)	insert line into <i>win</i>
wmove(win, y, x)	set current (y, x) co-ordinates on <i>win</i>
wnoutrefresh(win)	refresh but no screen output
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>
wsetscreg(win, t, b)	set scrolling region of <i>win</i>
wstandend(win)	clear standout attribute in <i>win</i>
wstandout(win)	set standout attribute in <i>win</i>

TERMINFO LEVEL ROUTINES

These routines should be called by programs wishing to deal directly with the terminfo database. Due to the low level of this interface, it is discouraged. Initially, *setupterm* should be called. This will define the set of terminal dependent variables defined in terminfo(4). The include files < curses.h> and < term.h> should be included to get the definitions for these strings, numbers, and flags. Parameterized strings should be passed through *tparm* to instantiate them. All terminfo strings (including the output of *tparm*) should be printed with *tputs* or *putp*. Before exiting, *resetterm* should be called to restore the tty modes. (Programs desiring shell escapes or suspending with control Z can call *resetterm* before the shell is called and *fixterm* after returning from the shell.)

fixterm()	restore tty modes for terminfo use (called by <i>setupterm</i>)
resetterm()	reset tty modes to state before program entry
setupterm(term, fd, rc)	read in database. Terminal type is the character string <i>term</i> , all output is to UNIX System file descriptor <i>fd</i> . A status value is returned in the integer pointed to by <i>rc</i> : 1 is normal. The simplest call would be <i>setupterm(0, 1, 0)</i> which uses all defaults.
tparm(str, p1, p2, ..., p9)	instantiate string <i>str</i> with parms p_i .
tputs(str, affcnt, putc)	apply padding info to string <i>str</i> . <i>affcnt</i> is the number of lines affected, or 1 if not applicable. <i>putc</i> is a putchar-like function to which the characters are passed, one at a time.
putp(str)	handy function that calls <i>tputs</i> (<i>str</i> , 1, <i>putc</i>)
vidputs(attrs, putc)	output the string to put terminal in video attribute mode <i>attrs</i> , which is any combination of the attributes listed below. Chars are passed to putchar-like function <i>putc</i> .
vidattr(attrs)	Like <i>vidputs</i> but outputs through <i>putc</i>

TERMCAP COMPATIBILITY ROUTINES

These routines were included as a conversion aid for programs that use *termcap*. Their parameters are the same as for *termcap*. They are emulated using the *terminfo* database. They may go away at a later date.

tgetent(bp, name)	look up termcap entry for name
tgetflag(id)	get boolean entry for id
tgetnum(id)	get numeric entry for id
tgetstr(id, area)	get string entry for id
tgoto(cap, col, row)	apply parms to given cap

ATTRIBUTES

The following video attributes can be passed to the functions *attron*, *attroff*, *attrset*.

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_BLANK	Blanking (invisible)
A_PROTECT	Protected
A_ALTCHARSET	Alternate character set

FUNCTION KEYS

The following function keys might be returned by *getch* if *keypad* has been enabled. Note that not all of these are currently supported, due to lack of definitions in *terminfo* or the terminal not transmitting a unique code when the key is pressed.

<i>Name</i>	<i>Value</i>	<i>Key name</i>
KEY_BREAK	0401	break key (unreliable)
KEY_DOWN	0402	The four arrow keys ...
KEY_UP	0403	
KEY_LEFT	0404	
KEY_RIGHT	0405	...
KEY_HOME	0406	Home key (upward+left arrow)
KEY_BACKSPACE	0407	backspace (unreliable)
KEY_F0	0410	Function keys. Space for 64 is reserved.
KEY_F(n)	(KEY_F0+(n))	Formula for fn.
KEY_DL	0510	Delete line
KEY_IL	0511	Insert line
KEY_DC	0512	Delete character
KEY_IC	0513	Insert char or enter insert mode
KEY_EIC	0514	Exit insert char mode
KEY_CLEAR	0515	Clear screen
KEY_EOS	0516	Clear to end of screen
KEY_EOL	0517	Clear to end of line
KEY_SF	0520	Scroll 1 line forward
KEY_SR	0521	Scroll 1 line backwards (reverse)
KEY_NPAGE	0522	Next page
KEY_PPAGE	0523	Previous page
KEY_STAB	0524	Set tab
KEY_CTAB	0525	Clear tab
KEY_CATAB	0526	Clear all tabs
KEY_ENTER	0527	Enter or send (unreliable)
KEY_SRESET	0530	soft (partial) reset (unreliable)
KEY_RESET	0531	reset or hard reset (unreliable)
KEY_PRINT	0532	print or copy
KEY_LL	0533	home down or bottom (lower left)

WARNING

The plotting library *plot(3X)* and the curses library *curses(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 `erase()` in the *plot(3X)* code.

NAME

error, 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 end-of-file indication lasts until the stream is closed; however, operations which attempt to read from the stream will ignore the current state of the end-of-file 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

fopen(3S), *open(2V)*

NAME

fopen, *freopen*, *fdopen* – open a stream

SYNOPSIS

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

```
FILE *fopen(filename, type)
```

```
char *filename, *type;
```

```
FILE *freopen(filename, type, stream)
```

```
char *filename, *type;
```

```
FILE *stream;
```

```
FILE *fdopen(fildes, type)
```

```
char *type;
```

DESCRIPTION

fopen opens the file named by *filename* and associates a stream with it. *fopen* returns a pointer to be used to identify the stream in subsequent operations.

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

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

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

freopen substitutes the named file in place of the open *stream*. It returns the original value of *stream*. The original stream is closed, regardless of whether the open ultimately succeeds.

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

fdopen associates a stream with a file descriptor. 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 (i.e., when *type* is "a" or "a+"), it is impossible to overwrite information already in the file. *fseek* may be used to reposition the file pointer to any position in the file, but when output is written to the file, the current file pointer is disregarded. All output is written at the end of the file and causes the file pointer to be repositioned at the end of the output. If two separate processes open the same file for append, each process may write freely to the file without fear of destroying output being written by the other. The output from the two processes will be intermixed in the file in the order in which it is written.

SEE ALSO

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

DIAGNOSTICS

fopen and *freopen* return a NULL pointer on failure.

BUGS

In order to support the same number of open files as does the system, *fopen* must allocate additional memory for data structures using *calloc* after 20 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 end-of-file or error condition is encountered while reading *stream*, or if *nitems* items have been read. *fread* leaves the file pointer in *stream*, if defined, pointing to the byte following the last byte read if there is one. *fread* does not change the contents of *stream*.

When input is read from *any* line-buffered stream, output to *all* line-buffered streams is flushed (including the standard error). Input read from a stream that is not line-buffered does not cause flushing of these streams.

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), *putc*(3S), *gets*(3S), *puts*(3S), *printf*(3S), *scanf*(3S), *fread*(3S)

DIAGNOSTICS

fread and *fwrite* return 0 upon end of file or error.

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 (i.e., byte) from the named input *stream*, as an integer. It also moves the file pointer, if defined, ahead one character in *stream*. *getchar* is defined as *getc(stdin)*. *Getc* and *getchar* are macros.

fgetc behaves like *getc*, but is a function rather than a macro. *fgetc* runs more slowly than *getc*, but 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

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

DIAGNOSTICS

These functions return the integer constant EOF at end-of-file or upon an error. Because EOF is a valid integer, *ferror(3S)* should be used to detect *getw* errors.

WARNING

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

BUGS

Because it is implemented as a macro, *getc* treats a *stream* argument with side effects incorrectly. In particular, *getc(*f++)* doesn't 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 causes it to increase 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*, *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()

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 *setpwent* has the effect of rewinding the password file to allow repeated searches. *endpwent* may be called to close the password file when processing is complete.

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
/etc/yp/domainname/passwd.byname
/etc/yp/domainname/passwd.byuid
```


SEE ALSO

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

DIAGNOSTICS

A NULL pointer is returned on EOF or error.

WARNING

The above routines use `<stdio.h>`, 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

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 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), *getpriority*(2), *setpriority*(2), *fork*(2), *renice*(8)

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. If the field width for an *s* conversion is preceded by a 0, the string is right adjusted with zero-padding on the left.

A *precision* that gives the minimum number of digits to appear for the *d*, *o*, *u*, *x*, or *X* conversions, the number of digits to appear after the decimal point for the *e*, *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.

An optional *l* (ell) specifying that a following *d*, *o*, *u*, *x*, or *X* conversion character applies to a long integer *arg*. A *l* before any other conversion character is ignored.

A character that indicates the type of conversion to be applied.

A field width or precision may be indicated by an asterisk (*) instead of a digit string. In this case, an integer *arg* supplies the field width or precision. The *arg* that is actually converted is not fetched until the conversion letter is seen, so the *args* specifying field width or precision must appear *before* the *arg* (if any) to be converted.

The flag characters and their meanings are:

- The result of the conversion will be left-justified within the field.
- + The result of a signed conversion will always begin with a sign (+ or –).
- blank If the first character of a signed conversion is not a sign, a blank will be prefixed to the result. This implies that if the blank and + flags both appear, the blank flag will be ignored.
- # This flag specifies that the value is to be converted to an “alternate form.” For *c*, *d*, *s*, and *u* conversions, the flag has no effect. For *o* conversion, it increases the precision to force the first digit of the result to be a zero. For *x* 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,o,u,x,X* The integer *arg* is converted to signed decimal, unsigned octal, unsigned decimal, 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 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 *d*, in style *f*, or in style *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 *sprintf* are printed as if *putc(3S)* had been called.

EXAMPLES

To print a date and time in the form “Sunday, July 3, 10:02,” where *weekday* and *month* are pointers to null-terminated strings:

```
printf("%s, %s %d, %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

putc(3S), *scanf(3V)*, *ecvt(3)*, *printf(3V)*

BUGS

Very wide fields (>128 characters) fail.

The values "Infinity" and "Nan" cannot be read by *scanf(3V)*.

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 control string usually contains conversion specifications, which are used to direct interpretation of input sequences. The control string may contain:

1. White-space characters (blanks, tabs, or new-lines) 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 a integer pointer.
- x a hexadecimal integer is expected; the corresponding argument should be an integer pointer.
- e,f,g a floating point number is expected; the next field is converted accordingly and stored through the corresponding argument, which should be a pointer to a *float*. The input format for floating point numbers is an optionally signed string of digits, possibly containing a decimal point, followed by an optional exponent field consisting of an E or e followed by an optional +, -, or space, followed by an integer.
- s a character string is expected; the corresponding argument should be a character pointer 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**, and **x** may be capitalized or 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.

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 that this is different from 0, which means that no conversion was done; if conversion was intended, it was frustrated by an inappropriate character in the input.

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; float x; char name[50];
(void) scanf("%2d%f%*d %[0-9]", &i, &x, name);
```

with input:

```
56789 0123 56a72
```

will assign 56 to *i*, 789.0 to *x*, skip 0123, and place the string 56\0 in *name*. The next call to *getchar* (see *getc*(3S)) will return a.

SEE ALSO

getc(3S), *printf*(3V) *strtod*(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.

scanf cannot read the strings which *printf*(3V) generates for IEEE indeterminate floating point values.

scanf provides no way to convert a number in any arbitrary base (decimal, hex or octal) based on the traditional C conventions (leading 0 or 0x).

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. *flush* (see *fclose(3S)*) may be used to force the block out early. Normally all files are block buffered. A buffer is obtained from *malloc(3)* upon the first *getc* or *putc(3S)* on the file.

By default, output to a terminal is line buffered 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 causes the character array *buf* whose size is determined by the *size* argument to be used 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      causes input/output to be fully buffered.
_IOLBF      causes output to be line buffered; the buffer will be flushed when a newline is written, the
            buffer is full, or input is requested.
_IONBF      causes input/output to be completely unbuffered. If buf is not the NULL pointer, the array it
            points to will be used for buffering, instead of an automatically allocated buffer. Size
            specifies the size of the buffer to be used.
```

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.

SEE ALSO

`fopen(3V)`, `getc(3S)`, `putc(3S)`, `malloc(3)`, `fclose(3S)`, `puts(3S)`, `printf(3V)`, `fread(3V)`, `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.

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.

NAME

setuid – set user ID

SYNOPSIS

setuid(uid)

DESCRIPTION

setuid is used to set the real user ID and effective user ID of the calling process.

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

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

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

SEE ALSO

setreuid(2), getuid(2)

DIAGNOSTICS

Zero is returned if the user ID is set; -1 is returned otherwise, with the global variable *errno* set as for *setreuid*.

NAME

signal – simplified software signal facilities

SYNOPSIS

```
#include <signal.h>

(*signal(sig, func))()
int (*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 UNIX 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 *tty*(4)). Signals are optionally generated when a process resumes after being stopped, when the status of child processes changes, or when input is ready at the control terminal. Most signals cause termination of the receiving process if no action is taken; some signals instead cause the process receiving them to be stopped, or are simply discarded if the process has not requested otherwise. Except for the SIGKILL and SIGSTOP signals, the *signal* call allows signals either to be ignored or to cause an interrupt to a specified location. The following is a list of all signals with names as in the include file *<signal.h>*:

SIGHUP	1	hangup
SIGINT	2	interrupt
SIGQUIT	3*	quit
SIGILL	4*	illegal instruction (other than A-line or F-line op code)
SIGTRAP	5*	trace trap
SIGIOT	6*	IOT trap (not generated on Suns)
SIGEMT	7*	EMT trap (A-line or F-line op code)
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 with no one to read it
SIGALRM	14	alarm clock
SIGTERM	15	software termination signal
SIGURG	16*	urgent condition present on socket
SIGSTOP	17†	stop (cannot be caught, blocked, or ignored)
SIGTSTP	18†	stop signal generated from keyboard
SIGCONT	19*	continue after stop (cannot be blocked)
SIGCHLD	20*	child status has changed
SIGTTIN	21†	background read attempted from control terminal
SIGTTOU	22†	background write attempted to control terminal
SIGIO	23*	I/O is possible on a descriptor (see <i>fcntl</i> (2))
SIGXCPU	24	cpu time limit exceeded (see <i>setrlimit</i> (2))
SIGXFSZ	25	file size limit exceeded (see <i>setrlimit</i> (2))
SIGVTALRM	26	virtual time alarm (see <i>setitimer</i> (2))
SIGPROF	27	profiling timer alarm (see <i>setitimer</i> (2))
SIGWINCH	28*	window changed (see <i>win</i> (4S))
SIGLOST	29*	resource lost (see <i>lockd</i> (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* 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:

```
handler(sig, code, scp)
int sig, code;
struct sigcontext *scp;
```

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.

CODES

The following defines the codes for signals which produce them. All of these symbols are defined in <signal.h>:

Hardware condition	Signal	Code
Illegal instruction	SIGILL	ILL_INSTR_FAULT
Privilege violation	SIGILL	ILL_PRIVVIO_FAULT
Coprocessor protocol error	SIGILL	ILL_INSTR_FAULT
Trap # <i>n</i> (1 ≤ <i>n</i> ≤ 14)	SIGILL	ILL_TRAP_FAULT
A-line op code	SIGEMT	EMT_EMU1010
F-line op code	SIGEMT	EMT_EMU1111
Integer division by zero	SIGFPE	FPE_INTDIV_TRAP
CHK or CHK2 instruction	SIGFPE	FPE_CHKINST_TRAP
TRAPV or TRAPcc or cpTRAPcc	SIGFPE	FPE_TRAPV_TRAP
IEEE floating point compare unordered	SIGFPE	FPE_FLTBSUN_TRAP
IEEE floating point inexact	SIGFPE	FPE_FLTINEX_TRAP
IEEE floating point division by zero	SIGFPE	FPE_FLTDIV_TRAP
IEEE floating point underflow	SIGFPE	FPE_FLTUND_TRAP
IEEE floating point operand error	SIGFPE	FPE_FLTOPERR_TRAP
IEEE floating point overflow	SIGFPE	FPE_FLTOVF_FAULT
IEEE floating point signaling NaN	SIGFPE	FPE_FLTNAN_TRAP

RETURN VALUE

The previous action is returned on a successful call. Otherwise, -1 is returned and *errno* is set to indicate the error.

ERRORS

signal will fail and no action will take place if one of the following occur:

EINVAL *sig* is not a valid signal number.

EINVAL An attempt is made to ignore or supply a handler for SIGKILL or SIGSTOP.

EINVAL An attempt is made to ignore SIGCONT (by default SIGCONT is ignored).

SEE ALSO

kill(1), ptrace(2), kill(2), sigvec(2), sigblock(2), sigsetmask(2), sigpause(2), sigstack(2), setjmp(3), tty(4)

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

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)*, *wait3(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

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

intro – introduction to other libraries

DESCRIPTION

This section contains manual pages describing other libraries, which are available only from C. The list below includes libraries which provide device independent plotting functions, terminal independent screen management routines for two dimensional non-bitmap display terminals, and functions for managing data bases with inverted indexes. All functions are located in separate libraries indicated in each manual entry.

FILES

/usr/lib/libcurses.a	screen management routines (see <i>curses</i> (3X))
/usr/lib/libdbm.a	data base management routines (see <i>dbm</i> (3X))
/usr/lib/libmp.a	multiple precision math library (see <i>mp</i> (3X))
/usr/lib/libplot.a	plot routines (see <i>plot</i> (3X))
/usr/lib/lib300.a	"
/usr/lib/lib300s.a	"
/usr/lib/lib450.a	"
/usr/lib/lib4014.a	"
/usr/lib/libtermcap.a	terminal handling routines (see <i>termcap</i> (3X))
/usr/lib/libtermcap_p.a	
/usr/lib/libtermlib.a	(link to /usr/lib/libtermcap.a)
/usr/lib/libtermlib_p.a	(link to /usr/lib/libtermcap_p.a)
/usr/lib/libresolv.a	Internet server routines (see <i>resolver</i> (3X))

NAME

curses – screen functions with “optimal” cursor motion

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

ioctl(2), getenv(3), tty(4), termcap(5)

Programmer's Reference Manual for Curses

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

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

SYNOPSIS

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

dbminit(file)
char *file;

datum fetch(key)
datum key;

store(key, content)
datum key, content;

delete(key)
datum key;

datum firstkey()

datum nextkey(key)
datum key;

dbmclose()
```

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. 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.)

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))
```

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

DIAGNOSTICS

All functions that return an *int* indicate errors with negative values. A zero return indicates ok. 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 UNIX systems may create real file blocks for these holes when touched. These files cannot be copied by normal means (`cp`, `cat`, `tp`, `tar`, `ar`) 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

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. *Rpow* 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;

dbm_close(db)
    DBM *db;

datum dbm_fetch(db, key)
    DBM *db;
    struct key;
    datum key;

dbm_store(db, key, content, flags)
    DBM *db;
    datum key, content;
    int flags;

dbm_delete(db, key)
    DBM *db;
    datum key;

datum dbm_firstkey(db)
    DBM *db;

datum dbm_nextkey(db)
    DBM *db;

datum dbm_error(db)
    DBM *db;

datum 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.

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)*).

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 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.

DIAGNOSTICS

All functions that return an *int* indicate errors with negative values. A zero return indicates ok. 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 UNIX systems may create real file blocks for these holes when touched. These files cannot be copied by normal means (cp, cat, tp, tar, ar) 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.

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 newlines.

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).

```

SEE ALSO

plot(5), *plot(1G)*, *graph(1G)*

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

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 can't write the file */etc/termcap*. Note that if the window size changes, the "lines" and "columns" entries in *bp* are no longer correct. See the *Sunwindows Reference Manual* 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 that programs using *termcap* should in general turn off XTABS anyway since some terminals use control I for other functions, such as nondestructive space.) If a % sequence is given which is not understood, then *tgoto* returns "OOPS".

Tputs decodes the leading padding information of the string *cp*; *affcnt* gives the number of lines affected by the operation, or 1 if this is not applicable, *outc* 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 -ltermcap library
/etc/termcap data base

SEE ALSO

ex(1), curses(3X), tty(4), termcap(5)

NAME

intro – introduction to special files and hardware support

DESCRIPTION

This section describes device interfaces (drivers) in the operating system for disks, tapes, serial communications, high-speed network communications, and other devices such as mice, frame buffers and windows. The section is divided into a few subsections:

- Sun-specific drivers are grouped in '4S'.
- Protocol families are grouped in '4F'.
- Protocols and raw interfaces are treated in '4P'.
- Network interfaces are grouped in '4N'.

The operating system can be built with or without many of the drivers listed here. For most of them, the SYNOPSIS section of the manual page gives the syntax of the line to include in a kernel configuration file if you wish to include the driver in a system. See *config(8)* for a description of this process.

Several manual pages will contain SYNOPSIS sections specific to the Sun-2 and Sun-3 architectures. Where a SYNOPSIS section appears without any specific architecture against it, it applies to both the Sun-2 and Sun-3 architectures. Where a SYNOPSIS section appears with only one specific architecture against it, it applies only to that specific architecture.

The pages for most drivers also include a DIAGNOSTICS section listing error messages the driver may produce. These messages appear on the system console, and also in the system error log file */usr/adm/messages*.

DEVICES ALWAYS PRESENT

Drivers which are present in every kernel include a driver for the paging device, *drum(4)*; drivers for accessing physical, virtual, and I/O space, *mem(4S)*; and drivers for the data sink, *null(4)*.

COMMUNICATIONS DEVICES

Communications lines are most often used with the terminal driver described in *tty(4)*. The terminal driver runs on communications lines provided either by a communications driver such as *mti(4S)* or *zs(4S)* or by a virtual terminal. The virtual terminal may be provided either by the Sun console monitor, *cons(4S)*, or by a true pseudo-terminal, *pty(4)*, used in applications such as windowing or remote networking.

MAGNETIC TAPE DEVICES

Magnetic tapes all provide the interface described in *mtio(4)*. Tape devices for the Sun include *ar(4S)*, *tm(4S)*, *st(4S)*, and *xt(4S)*.

DISK DEVICES

Disk controllers provide standard block and raw interfaces, as well as a set of ioctl's defined in *dkio(4S)*, which support getting and setting disk geometry and partition information. Drivers available for the Sun include *xy(4S)*, *ip(4S)*, and *sd(4S)*.

PROTOCOL FAMILIES

The operating system supports one or more protocol families for local network communications. The only complete protocol family in this version of the system is the Internet protocol family; see *inet(4F)*. Each protocol family provides basic services — packet fragmentation and reassembly, routing, addressing, and basic transport — to each protocol implementation. A protocol family is normally composed of a number of protocols, one per *socket(2)* type. A protocol family is not required to support all socket types.

The primary network support is for the Internet protocol family described in *inet(4F)*. Major protocols in this family include the Internet Protocol, *ip(4P)*, describing the universal datagram format, the stream Transmission Control Protocol *tcp(4P)*, the User Datagram Protocol *udp(4P)*, the Address Resolution Protocol *arp(4P)*, the Internet Control Message Protocol *icmp(4P)*, and the Network Interface Tap *nit(4P)*. The primary network interface is for the 10 Megabit Ethernet; see *ec(4S)*, *ie(4S)*, and *le(4S)*. A software loopback interface, *lo(4)* also exists. General properties of these (and all) network interfaces are described in *if(4N)*.

The general support in the system for local network routing is described in *routing(4N)*; these facilities apply to all protocol families.

MISCELLANEOUS DEVICES

Miscellaneous devices include color frame buffers *cg*(4S)*, monochrome frame buffers *bw*(4S)*, the console frame buffer *fb(4S)*, the graphics processor interface *gpone(4S)*, the console mouse *mouse(4S)*, and the window devices *win(4S)*.

GENERAL IOCTL CALLS

In general, *ioctl* calls relating to a specific device are mentioned with the description for that device. There are however a bunch of *ioctl* calls that apply to files in general. These are described here. The form of the *ioctl* call for file control is:

```
#include <sys/ioctl.h>
ioctl(fd, request, argp)
int fd, request;
int *argp;
```

- | | |
|-----------|---|
| FIOCLEX | Set set close-on-exec flag for the file descriptor specified by <i>fd</i> . This flag is also manipulated by the <i>F_SETFD</i> command of <i>fcntl(2)</i> . The <i>argp</i> argument is not used in this call. |
| FIONCLEX | Remove close-on-exec flag for the file descriptor specified by <i>fd</i> . The <i>argp</i> argument is not used in this call. |
| FIONREAD | Returns in the long integer whose address is <i>argp</i> the number of immediately readable characters from whatever the descriptor specified by <i>fd</i> refers to. This works for files, pipes, and terminals. |
| FIONBIO | Set or clear non-blocking I/O. If the value pointed to by <i>argp</i> is a 1 (one) the descriptor is set for non-blocking I/O. If the value pointed to by <i>argp</i> is a 0 (zero) the descriptor is cleared for non-blocking I/O. |
| FIOASYNC | Set or clear asynchronous I/O. If the value pointed to by <i>argp</i> is a 1 (one) the descriptor is set for asynchronous I/O. If the value pointed to by <i>argp</i> is a 0 (zero) the descriptor is cleared for asynchronous I/O. |
| FIOSETOWN | Set the process-group ID that will subsequently receive <i>SIGIO</i> or <i>SIGURG</i> signals for this descriptor. |
| FIOGETOWN | Get the process-group ID that is receiving <i>SIGIO</i> or <i>SIGURG</i> signals for this descriptor. |

SEE ALSO

fcntl(2)

NAME

ar – Archive 1/4 inch Streaming Tape Drive

SYNOPSIS — SUN-2

device ar0 at mbio ? csr 0x200 priority 3

device ar1 at mbio ? csr 0x208 priority 3

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 tape.

ar*: %b error at block # %d punted.

ar*: %b error at block # %d.

ar: giving up on Rdy, try again.

BUGS

The tape cannot reverse direction so the BSF and BSR ioctls 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 end of file mark on close if the last operation was a write, without regard for the mode used when opening the file. This will cause empty files to be deleted on a raw tape copy operation.

NAME

arp – Address Resolution Protocol

SYNOPSIS**pseudo-device ether****DESCRIPTION**

ARP is a protocol used to dynamically map between DARPA Internet and 10Mb/s Ethernet addresses. It is used by all the 10Mb/s Ethernet interface drivers.

ARP caches Internet-Ethernet address mappings. When an interface requests a mapping for an address not in the cache, ARP queues the message which requires the mapping and broadcasts a message on the associated network requesting the address mapping. If a response is provided, the new mapping is cached and any pending messages are transmitted. ARP 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 enable communications with systems which do not use ARP, ioctls are provided to enter and delete entries in the Internet-to-Ethernet tables. Usage:

```
#include <sys/ioctl.h>
#include <sys/socket.h>
#include <net/if.h>
struct arpreq arpreq;

ioctl(s, SIOCSARP, (caddr_t)&arpreq);
ioctl(s, SIOCGARP, (caddr_t)&arpreq);
ioctl(s, SIOCDAARP, (caddr_t)&arpreq);
```

Each ioctl takes the same structure as an argument. SIOCSARP sets an ARP entry, SIOCGARP gets an ARP entry, and SIOCDAARP deletes an ARP entry. These ioctls 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      2          /* completed entry (arp_ha valid) */
#define ATF_PERM    4          /* permanent entry */
#define ATF_PUBL    8          /* publish (respond for other 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 and ATF_PUBL. ATF_PERM causes the entry to be permanent if the ioctl call succeeds. The peculiar nature of the ARP tables may cause the ioctl to fail if more than 4 (permanent) Internet host 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 Sun to act as an "ARP server" which may be useful in convincing an ARP-only machine to talk to a non-ARP machine.

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).

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.

SEE ALSO

ec(4S), ie(4S), inet(4F), arp(8C), ifconfig(8C)

An Ethernet Address Resolution Protocol, RFC826, Dave Plummer, MIT (Sun 800-1059-01)

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

tty(4)

NAME

bwone – Sun-1 black and white frame buffer

SYNOPSIS — SUN-2

device bwone0 at mbmem ? csr 0xc0000 priority 3

DESCRIPTION

The *bwone* interface provides access to Sun-1 black and white graphics controller boards. It supports the **FBIOGTYPE** ioctl which programs can use to determine the characteristics of the display device; see *fbio*(4S).

bwone also supports the **FBIOGPIXRECT** ioctl which allows SunWindows to be run on it; see *fbio*(4S).

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/bwone[0-9]

SEE ALSO

mmap(2), *fb*(4S), *fbio*(4S)

BUGS

Use of vertical-retrace interrupts is not supported.

The **FBVIDEO_ON** value returned by the **FBIOGVIDEO** ioctl may be incorrect. See *fbio*(4S).

NAME

bwtwo – Sun-3/Sun-2 black and white frame buffer

SYNOPSIS — SUN-3

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

The first synopsis line given above should be used to generate a kernel for a Sun-3/160; the second, for a Sun-3/75M; the third, for a Sun-3/260; and the fourth, for a Sun-3/110.

SYNOPSIS — SUN-2

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; the second, for a Sun-2/50 or Sun-2/160.

DESCRIPTION

The *bwtwo* interface provides access to Sun Monochrome Video Controller boards.

bwtwo supports the `FBIOGTYPE` ioctl, which may be used to determine the characteristics of the display device, and the `FBIOGPIXRECT` ioctl, which allows SunWindows to be run on it (see *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 the write operations, but consumes an extra 128K bytes of memory.

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), *fb(4S)*, *fbio(4S)*, *cgfour(4S)*

BUGS

Use of vertical-retrace interrupts is not supported.

The `FBVIDEO_ON` value returned by the `FBIOGVIDEO` ioctl may be incorrect. See *fbio(4S)*.

NAME

cgfour – Sun-3 color graphics interface

SYNOPSIS — SUN-3

***cgfour0* at obmem 4 csr 0xff000000**

DESCRIPTION

The *cgfour* is the Sun-3/110 color frame buffer, normally supplied with a 19'' color, 19'' grayscale, or 15'' color 66 Hz non-interlaced color monitor. 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 via 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.

Currently the *ioctl*s FBIOSATTR and FBIOGATTR are only supported by the *cgfour* frame buffer. See *fbio(4S)*.

FILES

/dev/cgfour0

SEE ALSO

mmap(2), *fbio(4S)*

Sun-3/1xx CPU Board Hardware Engineering Manual

BUGS

The FBVIDEO_ON value returned by the FBIOGVIDEO *ioctl* may be incorrect. See *fbio(4S)*.

NAME

cgone – Sun-1 color graphics interface

SYNOPSIS — SUN-2

device *cgone0* at *mbmem* ? *csr 0xec000* priority 3

DESCRIPTION

The *cgone* interface provides access to the Sun-1 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 *FBIOPIXRECT* ioctl which allows SunWindows 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

cgtwo – Sun-3/Sun-2 color graphics interface

SYNOPSIS — SUN-3

cgtwo0 at *vme24d16* ? *csr 0x400000*

SYNOPSIS — SUN-2

cgtwo0 at *vme24* ? *csr 0x400000*

DESCRIPTION

The *cgtwo* interface provides access to the Sun-3/Sun-2 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

console – console driver and terminal emulator for the Sun workstation

SYNOPSIS

None; included in standard system.

DESCRIPTION

Cons is an indirect driver for the Sun workstation console, which implements a standard UNIX system terminal. *Cons* is implemented by calling the PROM resident monitor or other kernel UART drivers (*zs(4S)*) to perform I/O to and from the current system console, which is either a Sun frame buffer or an RS232 port.

When the Sun window system *win(4S)* is active, console input is directed through the window system rather than being read from the hardware console.

An ioctl TIOCCONS can be applied to serial devices other than the console to route output which would normally appear on the console to the other devices instead. Thus, the window system does a TIOCCONS on a pseudoterminal to route console output to the pseudoterminal 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 however, that when you use TIOCCONS in this way, the console *input* is routed from the pseudoterminal as well.

ANSI STANDARD TERMINAL EMULATION

The Sun Workstation's PROM monitor provides routines that emulate a standard ANSI X3.64 terminal.

Note that 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 that 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 linefeed 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 linefeed 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 awaiting printing. 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 linefeeds 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* at least 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 that, 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 (control-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 unchanged.

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 Sun-2 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

/dev/ttya

alternate console (serial port)

SEE ALSO

kb(4S), tty(4), zs(4S), tektool(1)

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

SYNOPSIS — SUN-3

des0 at obio ? csr 0x1c0000

#include <sys/des.h>

SYNOPSIS — SUN-2

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. The *ioctls* 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_crypt(3), des(1)

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 labelling 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_ctrl;      /* 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

/* 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 DKIOCGINFO ioctl returns a `dk_info` structure which tells the kind 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 DKIOCGGEO and DKIOCSGEO ioctls do similar things for the per-drive geometry information.

SEE ALSO

`ip(4S)`, `sd(4S)`, `xy(4S)`

NAME

drum – paging device

SYNOPSIS

None; included with standard system.

DESCRIPTION

This file refers to the paging device in use by the system. This may actually be a subdevice of one of the disk drivers, but in a system with paging interleaved across multiple disk drives it provides an indirect driver for the multiple drives.

FILES

/dev/drum

BUGS

Reads from the drum are not allowed across the interleaving boundaries. Since these only occur every .5Mbytes or so, and since the system never allocates blocks across the boundary, this is usually not a problem.

NAME

ec – 3Com 10 Mb/s Ethernet interface

SYNOPSIS — SUN-2

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

SYNOPSIS

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 Monitor Proms 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

fbio(4S), *bwone(4S)*, *bwtwo(4S)*, *cgone(4S)*, *cgtwo(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

```

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 a 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 SunWindows) 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 this 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), cgone(4S), cgtwo(4S), cgfour(4S), gpone(4S), fb(4S), win(4S)

BUGS

FBIOSATTR and FBIOGATTR are only supported by the *cgfour*(4S) frame buffer.

The FBVIDEO_ON flag may be incorrect for the *bwone*(4S), *bwtwo*(4S), and *cgfour*(4S) frame buffers since the drivers for these devices do not test the hardware, but simply report the last state stored by FBIOSVIDEO. The stored state and actual hardware state can get out of sync for several reasons: (1) processes having the same *bwone* frame buffer mapped can directly enable or disable the video, unknown to the driver; (2) */dev/bwtwo0* and */dev/cgfour0* on a *cgfour* CPU refer to the same frame buffer hardware; the video state of */dev/bwtwo0* may change, unknown to the *cgfour* driver; (3) if the hardware is not the default frame buffer */dev/fb*, the hardware's initial state is "video off", unknown to the driver.

NAME

`gpone` – Sun-3/Sun-2 graphics processor

SYNOPSIS — SUN-3

`device gpone0 at vme24d16 ? csr`

SYNOPSIS — SUN-2

`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.

GP IOCTL

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 *devfb* 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.

GP1IO_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/{gplcmds.h,gplreg.h,gplvar.h}

SEE ALSO

fbio(4S), *mmap*(2), *gpconfig*(8)

Hardware Reference Manual for the Sun Graphics Processor (Sun 800-1190-01)

Software Interface Manual for the Sun Graphics Processor (Sun 800-1571-01)

DIAGNOSTICS

The Graphics Processor has been restarted. You may see display garbage as a result.

NAME

icmp – Internet Control Message Protocol

SYNOPSIS

None; included automatically with *inet*(4F).

DESCRIPTION

The Internet Control Message Protocol, ICMP, is used by gateways and destination hosts which process datagrams to communicate errors in datagram-processing to source hosts. The datagram level of Internet is discussed in *ip*(4P). ICMP uses the basic support of IP as if it were a higher level protocol; however, ICMP is actually an integral part of IP. ICMP messages are sent in several situations; for example: when a datagram cannot reach its destination, when the gateway does not have the buffering capacity to forward a datagram, and when the gateway can direct the host to send traffic on a shorter route.

The Internet protocol is not designed to be absolutely reliable. The purpose of these control messages is to provide feedback about problems in the communication environment, not to make IP reliable. There are still no guarantees that a datagram will be delivered or that a control message will be returned. Some datagrams may still be undelivered without any report of their loss. The higher level protocols which use IP must implement their own reliability mechanisms if reliable communication is required.

The ICMP messages typically report errors in the processing of datagrams; for fragmented datagrams, ICMP messages are sent only about errors in handling fragment 0 of the datagram. To avoid the infinite regress of messages about messages etc., no ICMP messages are sent about ICMP messages. ICMP may however be sent in response to ICMP messages (for example, ECHOREPLY). There are eleven types of ICMP packets which can be received by the system. They are defined in this excerpt from `<netinet/ip_icmp.h>`, which also defines the values of some additional codes specifying the cause of certain errors.

```

/*
 * Definition of type and code field values
 */
#define ICMP_ECHOREPLY          0    /* echo reply */
#define ICMP_UNREACH           3    /* dest unreachable, codes: */
#define ICMP_UNREACH_NET       0    /* bad net */
#define ICMP_UNREACH_HOST      1    /* bad host */
#define ICMP_UNREACH_PROTOCOL  2    /* bad protocol */
#define ICMP_UNREACH_PORT      3    /* bad port */
#define ICMP_UNREACH_NEEDFRAG  4    /* IP_DF caused drop */
#define ICMP_UNREACH_SRCFAIL   5    /* src route failed */
#define ICMP_SOURCEQUENCH      4    /* packet lost, slow down */
#define ICMP_REDIRECT          5    /* shorter route, codes: */
#define ICMP_REDIRECT_NET      0    /* for network */
#define ICMP_REDIRECT_HOST     1    /* for host */
#define ICMP_REDIRECT_TOSNET   2    /* for tos and net */
#define ICMP_REDIRECT_TOSHOST  3    /* for tos and host */
#define ICMP_ECHO              8    /* echo service */
#define ICMP_TIMXCEED          11   /* time exceeded, code: */
#define ICMP_TIMXCEED_INTRANS  0    /* ttl==0 in transit */
#define ICMP_TIMXCEED_REASS    1    /* ttl==0 in reass */
#define ICMP_PARAMPROB        12   /* ip header bad */
#define ICMP_TSTAMP            13   /* timestamp request */
#define ICMP_TSTAMPREPLY      14   /* timestamp reply */
#define ICMP_IREQ              15   /* information request */
#define ICMP_IREQREPLY        16   /* information reply */

```

Arriving ECHO and TSTAMP packets cause the system to generate ECHOREPLY and TSTAMPREPLY packets. IREQ packets are not yet processed by the system, and are discarded. UNREACH, SOURCE-QUENCH, TIMXCEED and PARAMPROB packets are processed internally by the protocols implemented in the system, or reflected to the user if a raw socket is being used; see *ip(4P)*. REDIRECT, ECHOREPLY, TSTAMPREPLY and IREQREPLY are also reflected to users of raw sockets. In addition, REDIRECT messages cause the kernel routing tables to be updated; see *routing(4N)*.

SEE ALSO

inet(4F), *ip(4P)*

Internet Control Message Protocol, RFC792, J. Postel, USC-ISI (Sun 800-1064-01)

BUGS

IREQ messages are not processed properly: the address fields are not set.

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

SYNOPSIS — SUN-3

device ie0 at obio ? csr 0xc0000 priority 3
device ie1 at vme24d16 ? csr 0x88000 priority 3 vector ieintr 0x75

SYNOPSIS — SUN-2

device ie0 at obio 2 csr 0x7f0800 priority 3
device ie0 at mbmem ? csr 0x88000 priority 3
device ie1 at mbmem ? csr 0x8c000 flags 2 priority 3
device ie1 at vme24 ? csr 0x88000 priority 3 vector ieintr 0x75

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 synopsis—Sun-3 lines above, the first line specifies the first, CPU board-resident, Intel Ethernet interface; the second line specifies a second Intel interface present on a Sun Ethernet board.

In the synopsis—Sun-2 lines above, the first line specifies the first Intel Ethernet controller on a Sun-2/50 or Sun-2/160; the second line specifies the first Intel Ethernet controller on a Sun-2/120 or Sun-2/170.

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, causing the 82586 to give 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 which has underlying hardware support makes itself known to the system during the autoconfiguration process. Once the interface has acquired its address it is expected to install a routing table entry so that messages may be routed through it. Most interfaces require some part of their address specified with an *SIOCSIFADDR* *ioctl* before they will allow traffic to flow through them. On interfaces where the network-link layer address mapping is static, only the network number is taken from the *ioctl*; the remainder is found in a hardware specific manner. On interfaces which provide dynamic network-link layer address mapping facilities (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 */
};
```

SEE ALSO

arp(4P), ec(4S), lo(4)

NAME

inet – Internet protocol family

SYNOPSIS

options INET

DESCRIPTION

The Internet protocol family is a collection of protocols layered atop the *Internet Protocol* (IP) transport layer, and using the Internet address format. The Internet family provides protocol support for the SOCK_STREAM, SOCK_DGRAM, and SOCK_RAW socket types; the SOCK_RAW interface provides access to the IP protocol.

ADDRESSING

Internet addresses are four byte quantities, stored in network standard format (on the VAX these are word and byte reversed; on the Sun they are not reversed). The include file `<netinet/in.h>` defines the Internet address as a discriminated union.

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 to return and manipulate structures of this form are in section 3N of the manual; see *intro*(3N) and the other section 3 entries mentioned under SEE ALSO below). Each socket has a local address which may be specified in this form, which can be established with *bind*(2); the *getsockname*(2) call returns this address. Each socket also may be bound to a peer socket with an address specified in this form; this peer address can be specified in a *connect*(2) call, or transiently with a single message in a *sendto* or *sendmsg* call; see *send*(2). The peer address of a socket is returned by the *getpeername*(2) call.

The *sin_addr* field of the socket address specifies the Internet address of the machine on which the socket is located. A special value may be specified or returned for this field, *sin_addr.s_addr*==INADDR_ANY. This address is a “wildcard” and matches any of the legal internet addresses on the local machine. This address is useful when a process neither knows (nor cares) what the local Internet address is, and even more useful for server processes which wish to service all requests of the current machine. Since a machine can have several addresses (one per hardware network interface), specifying a single address would restrict access to the service to those clients which specified the address of that interface. By specifying INADDR_ANY, the server can arrange to service clients from all interfaces.

When a socket address is bound, the networking system checks for an interface with the address specified on the current machine (unless, of course, a wildcard address is specified), and returns an error EADDRNOTAVAIL if no such interface is found.

The local port address specified in a *bind*(2) call is restricted to be greater than IPPORT_RESERVED (=1024, in `<netinet/in.h>`) unless the creating process is running as the super-user, providing a space of protected port numbers. The local port address is also required to not be in use in order for it to be assigned. This is checked by looking for another socket of the same type which has the same local address and local port number. If such a socket already exists, you will not be able to create another socket at the same address, and will instead get the error EADDRINUSE. If the local port address is specified as 0, then the system picks a unique port address not less than IPPORT_RESERVED and assigns it to the port. A unique local port address is also picked for a socket which is not bound but which is used with *connect*(2) or *sendto*(2); this allows *tcp*(4P) connections to be made by simply doing *socket*(2) and then *connect*(2) in the case where the local port address is not significant; it is defaulted by the system. Similarly if you are sending datagrams with *udp*(4P) and do not care which port they come from, you can just do *socket*(2) and *sendto*(2) and let the system pick a port number.

Let us say that two sockets are incompatible if they have the same port number, are not connected to other sockets, and do not have different local host addresses. (It is possible to have two sockets with the same port number and different local host addresses because a machine may have several local addresses from its different network interfaces.) The Internet system does not allow such incompatible sockets to exist on a single machine. Consider a socket which has a specific local host and local port number on the current machine. If another process tries to create a socket with a wildcard local host address and the same port number then that request will be denied. For connection based sockets this prevents these two sockets from attempting to connect to the same foreign host/socket, and thereby causing great havoc. For connectionless sockets this prevents the dilemma which would result from trying to determine who to deliver an incoming datagram to (since more than one socket could match an address given on a datagram). The same restriction applies if the wildcard socket exists first. (If both sockets are wildcard, then the normal restrictions on duplicate addresses apply.)

A socket option `SO_REUSEADDR` exists to allow incompatible sockets to be created. This option is needed to implement the File Transfer Protocol (FTP) which requires that a connection be made from an existing port number (the port number of its primary connection) to a different port number on the same remote host. The danger here is that the user would attempt to connect this second port to the same remote host/port that the primary connection was using. In using `SO_REUSEADDR` the user is pledging not to do this, since this will cause the first connection to abort.

When a `connect(2)` is done, the Internet system first checks that the socket is not already connected. It does not allow connections to port number 0 on another host, nor does it allow connections to a wildcard host (`sin_addr.s_addr==INADDR_ANY`); attempts to do this yield `EADDRINUSE`. If the socket from which the connection is being made currently has a wildcard local address (either because it was bound to a specific port with a wildcard address, or was never subjected to `bind(2)`), then the system picks a local Internet address for the socket from the set of addresses of interfaces on the local machine. If there is an interface on the local machine on the same network as the machine being connected to, then that address is used. Otherwise, the "first" local network interface is used (this is the one that prints out first in `"netstat -i"`; see `netstat(8)`). Although it is not supposed to matter which interface address is used, in practice it would probably be better to select the address of the interface through which the packets are to be routed. This is not currently done (as it would involve a fair amount of additional overhead for datagram transmission).

PROTOCOLS

The Internet protocol family supported by the operating system is comprised of the Internet Datagram Protocol (IP) `ip(4P)`, Address Resolution Protocol (ARP) `arp(4P)`, Internet Control Message Protocol (ICMP) `icmp(4P)`, Transmission Control Protocol (TCP) `tcp(4P)`, and User Datagram Protocol (UDP) `udp(4P)`.

TCP is used to support the `SOCK_STREAM` abstraction while UDP is used to support the `SOCK_DGRAM` abstraction. A raw interface to IP is available by creating an Internet socket of type `SOCK_RAW`; see `ip(4P)`. The ICMP message protocol is most often used by the kernel to handle and report errors in protocol processing; it is, however, accessible to user programs. The ARP protocol is used to translate 32-bit Internet host numbers into the 48-bit addresses needed for an Ethernet.

SEE ALSO

`intro(3N)`, `byteorder(3N)`, `gethostent(3N)`, `getnetent(3N)`, `getprotoent(3N)`, `getservent(3N)`, `inet(3N)`, `arp(4P)`, `tcp(4P)`, `udp(4P)`, `ip(4P)`

Internet Protocol Transition Workbook, Network Information Center, SRI (Sun 800-1056-01)

Internet Protocol Implementation Guide, Network Information Center, SRI (Sun 800-1055-01)

A 4.2BSD Interprocess Communication Primer

NAME

ip – Internet Protocol

SYNOPSIS

None; included by default with *inet*(4F).

DESCRIPTION

The Internet Protocol is designed for use in interconnected systems of packet-switched computer communication networks. It provides for transmitting blocks of data called “datagrams” from sources to destinations, where sources and destinations are hosts identified by fixed-length addresses. It also provides for fragmentation and reassembly of long datagrams, if necessary, for transmission through “small packet” networks.

IP is specifically limited in scope. There are no mechanisms to augment end-to-end data reliability, flow control, sequencing, or other services commonly found in host-to-host protocols. IP can capitalize on the services of its supporting networks to provide various types and qualities of service.

IP is called on by host-to-host protocols, including *tcp*(4P) a reliable stream protocol, *udp*(4P) a socket-socket datagram protocol, and *nd*(4P) the network disk protocol. Other protocols may be layered on top of IP using the *raw* protocol facilities described here to receive and send datagrams with a specific IP protocol number. The IP protocol calls on local network drivers to carry the internet datagram to the next gateway or destination host.

When a datagram arrives at a UNIX system host, the system performs a checksum on the header of the datagram. If this fails, or if the datagram is unreasonably short or the header length specified in the datagram is not within range, then the datagram is dropped. Checksumming of Internet datagrams may be disabled for debugging purposes by patching the kernel variable *ipcksum* to have the value 0.

Next the system scans the IP options of the datagram. Options allowing for source routing (see *routing*(4N)) and also the collection of time stamps as a packet follows a particular route (for network monitoring and statistics gathering purposes) are handled; other options are ignored. Processing of source routing options may result in an UNREACH *icmp*(4P) message because the source routed host is not accessible.

After processing the options, IP checks to see if the current machine is the destination for the datagram. If not, then IP attempts to forward the datagram to the proper host. Before forwarding the datagram, IP decrements the time to live field of the datagram by IPTTLDEC seconds (currently 5 from <netinet/ip.h>), and discards the datagram if its lifetime has expired, sending an ICMP TIMXCEED error packet back to the source host. Similarly if the attempt to forward the datagram fails, then ICMP messages indicating an unreachable network, datagram too large, unreachable port (datagram would have required broadcasting on the target interface, and IP does not allow directed broadcasts), lack of buffer space (reflected as a source quench), or unreachable host. Note however, in accordance with the ICMP protocol specification, ICMP messages are returned only for the first fragment of fragmented datagrams.

It is possible to disable the forwarding of datagrams by a host by patching the kernel variable *ipforwarding* to have value 0.

If a packet arrives and is destined for this machine, then IP must check to see if other fragments of the same datagram are being held. If this datagram is complete, then any previous fragments of it are discarded. If this is only a fragment of a datagram, it may yield a complete set of pieces for the datagram, in which case IP constructs the complete datagram and continues processing with that. If there is yet no complete set of pieces for this datagram, then all data thus far received is held (but only one copy of each data byte from the datagram) in hopes that the rest of the pieces of the fragmented datagram will arrive and we will be able to proceed. We allow IPFRAGTTL (currently 15 in <netinet/ip.h>) seconds for all the fragments of a datagram to arrive, and discard partial fragments then if the datagram has not yet been completely assembled.

When we have a complete input datagram it is passed out to the appropriate protocol's input routine: either *tcp*(4P), *udp*(4P), *nd*(4P), *icmp*(4P) or a user process through a raw IP socket as described below.

Datagrams are output by the system-implemented protocols *tcp*(4P), *udp*(4P), *nd*(4P), and *icmp*(4P); as well as by packet forwarding operations and user processes through raw IP sockets. Output packets are normally subjected to routing as described in *routing*(4N). However, special processes such as the routing daemon *routed*(8C) occasionally use the `SO_DONTROUTE` socket option to make packets avoid the routing tables and go directly to the network interface with the network number which the packet is addressed to. This may be used to test the ability of the hardware to transmit and receive packets even when we believe that the hardware is broken and have therefore deleted it from the routing tables.

If there is no route to a destination address or if the `SO_DONTROUTE` option is given and there is no interface on the network specified by the destination address, then the IP output routine returns a `ENETUNREACH` error. (This and the other IP output errors are reflected back to user processes through the various protocols, which individually describe how errors are reported.)

In the (hopefully normal) case where there is a suitable route or network interface, the destination address is checked to see if it specifies a broadcast (address `INADDR_ANY`; see *inet*(4F)); if it does, and the hardware interface does not support broadcasts, then an `EADDRNOTAVAIL` is returned; if the caller is not the super-user then a `EACCESS` error will be returned. IP also does not allow broadcast messages to be fragmented, returning a `EMSGSIZE` error in this case.

If the datagram passes all these tests, and is small enough to be sent in one chunk, then the system calls the output routine for the particular hardware interface to transmit the packet. The interface may give an error indication, which is reflected to IP output's caller; see the documentation for the specific interface for a description of errors it may encounter. If a datagram is to be fragmented, it may have the `IP_DF` (don't fragment) flag set (although currently this can happen only for forwarded datagrams). If it does, then the datagram will be rejected (and result in an ICMP error datagram). If the system runs out of buffer space in fragmenting a datagram then a `ENOBUFS` error will be returned.

IP provides a space of 255 protocols. The known protocols are defined in `<netinet/in.h>`. The ICMP, TCP, UDP and ND protocols are processed internally by the system; others may be accessed through a raw socket by doing:

```
s = socket(AF_INET, SOCK_RAW, IPPROTO_XXX);
```

Datagrams sent from this socket will have the current host's address and the specified protocol number; the raw IP driver will construct an appropriate header. When IP datagrams are received for this protocol they are queued on the raw socket where they may be read with *recvfrom*; the source IP address is reflected in the received address.

SEE ALSO

send(2), *recv*(2), *inet*(4F)

Internet Protocol, RFC791, USC-ISI (Sun 800-1063-01)

BUGS

One should be able to send and receive IP options.

Raw sockets should receive ICMP error packets relating to the protocol; currently such packets are simply discarded.

NAME

ip – Disk driver for Interphase 2180 SMD Disk Controller

SYNOPSIS — SUN-2

controller ipc0 at mbio ? csr 0x040 priority 2
 controller ipc1 at mbio ? csr 0x044 priority 2
 disk ip0 at ipc0 drive0
 disk ip1 at ipc0 drive1
 disk ip2 at ipc1 drive0
 disk ip3 at ipc1 drive1

DESCRIPTION

Special files *ip** refer to disk devices controlled by an Interphase SMD 2180 disk controller.

The standard *ip* device names begin with the letters “ip”, followed by the drive unit number, followed by a letter from the series a – h to name one of the eight partitions on the drive. For example, */dev/ip1c* refers to partition c on the second drive controlled by the Interphase controller.

The device names provide the binding into the minor device numbers for the driver software. Files with minor device numbers 0 through 7 refer to the eight partitions (a – h) of unit 0; files with device numbers 8 through 15 refer to the eight partitions of drive 1, and so on.

The block files access the disk via 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 results in exactly one I/O operation and therefore raw I/O is considerably more efficient when many words are transmitted. Raw files conventionally have a leading “r” — */dev/rip0c*, for instance.

In raw I/O, counts should be a multiple of 512 bytes (a disk sector). Likewise *seek* calls should specify a multiple of 512 bytes.

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 *ip?a* partition is normally used for the root file system on a disk, the *ip?b* partition as a paging area, and the *ip?c* partition for pack-pack copying (it normally maps the entire disk). The rest of the disk is normally the *ip?g* partition.

FILES

/dev/ip[0-7][a-h] block files
/dev/rip[0-7][a-h] raw files

SEE ALSO

dkio(4S), xy(4S)

DIAGNOSTICS

ipn: SMD-2180

When booting tells the controller type.

ipn: initialization failed

Because the controller didn’t respond; perhaps another device is at the address the system expected an Interphase controller at.

ipn: error *n* reading label on head *n*

Error reading drive geometry/partition table information.

ipn: Corrupt label on head *n*

The geometry/partition label checksum was incorrect.

ipn: Misplaced label on head *n*

A disk label was copied to the wrong head on the disk; shouldn’t happen.

ipn : Unsupported phys partition # *n*
This indicates a bad label.

ipn : unit not online

ipn c : *cmd how (msg) blk n*

A command such as *read*, *write*, or *format* 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", "sector not found" or "disk write protected".

BUGS

In raw I/O *read* and *write(2V)* truncate file offsets to 512-byte block boundaries, and *write* scribbles on the tail of incomplete blocks. Thus, in programs that are likely to access raw devices, *read*, *write* and *lseek(2)* should always deal in 512-byte multiples.

NAME

kb, kbd – Sun keyboard

SYNOPSIS

pseudo-device *kbnumber*

DESCRIPTION

kb provides access to the Sun workstation keyboard and its translation tables. 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. In addition, the kernel recognizes the special keyboard device `/dev/kbd`; it is a synonym for the system keyboard, which is physically attached to the system console. This keyboard's keystrokes are normally treated as input to the `/dev/console` device.

The UNIX kernel 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 call `KIOCTRANS` controls the presence of keyboard translation, for which the following values are defined:

```
#define TR_NONE          0
#define TR_ASCII        1
#define TR_EVENT        2
#define TR_UNTRANS_EVENT 3
```

```
int x;
err = ioctl(fd, KIOCTRANS, &x);
```

When *x* is `TR_NONE`, keyboard translation is turned off and up/down key codes are reported. Specifying *x* as `TR_ASCII` causes ASCII to be reported. Specifying *x* as `TR_EVENT` causes *Firm_events* to be reported (see below). Specifying *x* as `TR_UNTRANS_EVENT` gives unencoded keyboard values for all input events within the window system.

Keyboard Translation-Table Entries

The call `KIOCSETKEY` changes a keyboard translation table entry, using the *kiockey* struct:

```
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) */
};

struct kiockey key;
err = ioctl(fd, KIOCSETKEY, &key);
```

To alter a keyboard translation-table entry, set *kio_tablemask* table's *kio_station* to *kio_entry*. Copy *kio_string* to the string table if *kio_entry* is between `STRING` and `STRING+15`. 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.

The call `KIOCGETKEY` determines the current value of a keyboard translation table entry:

```
struct kiockey key;
err = ioctl(fd, KIOCGETKEY, &key);
```

Get `kio_tablemask` table's `kio_station` to `kio_entry`. Get `kio_string` from string table if `kio_entry` is between `STRING` and `STRING+15`. This call may return `EINVAL` if there are invalid arguments.

Keyboard Type

The call `KIOCTYPE` returns the type of the keyboard:

```
#define KB_KLUNK      0x00      /* Micro Switch 103SD32-2 */
#define KB_VT100     0x01      /* Keytronics VT100 compatible */
#define KB_SUN2      0x02      /* Sun-2 custom keyboard */
#define KB_SUN3      0x03      /* Sun-3 custom keyboard */
#define KB_ASCII     0x0F      /* ASCII terminal masquerading as kbd */
```

```
int x;
err = ioctl(fd, KIOCTYPE, &x);
```

When `x` is `-1`, the keyboard type is unknown.

Reading From The Keyboard

Normally, keystrokes are discarded, except for those typed at the system keyboard; those are translated and treated as input on the system console device `/dev/console`. In order to read keystrokes directly, the call `KIOCSDIRECT` must be used to set the keyboard to "direct input" mode. In this mode, keystrokes are translated and queued to be read by a process that has opened a keyboard device:

```
int on = 1;
err = ioctl(fd, KIOCSDIRECT, &on);
```

The `KIOCSDIRECT` call turns "direct input" mode on or off, depending on whether the variable pointed to by its argument has the value 1 or 0. The call `KIOCGDIRECT` sets the variable pointed to by its argument to the current state of this mode:

```
int direct_state;
err = ioctl(fd, KIOCGDIRECT, &direct_state);
```

When the keyboard device is closed, "direct mode" is turned off.

Keyboard Commands

The call `KIOCCMD` sends a command to the keyboard:

```
/*
 * Commands to the Sun-2 keyboard.
 */
#define KBD_CMD_RESET      0x01      /* Reset keyboard as if power-up */
#define KBD_CMD_BELL      0x02      /* Turn on the bell */
#define KBD_CMD_NOBELL    0x03      /* Turn off the bell */

/*
 * Commands to the Sun-3 keyboard. KBD_CMD_BELL & KBD_CMD_NOBELL work
 * as well.
 */
#define KBD_CMD_CLICK      0x0A      /* Turn on the click annunciator */
#define KBD_CMD_NOCLICK   0x0B      /* Turn off the click annunciator */

int x;
err = ioctl(fd, KIOCCMD, &x);
```

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 can't query the keyboard, and also because a process could do writes to the appropriate serial driver — thus going around this *ioctl*) we don't provide an equivalent *ioctl* to query its state.

Shift Masks

When shift keys are pressed or locked, a different translation table is used to translate keyboard actions. The shift mask indicates which translation table to use. Since there may be more than one bit in the shift mask at any given time, they are prioritized as follows:

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.

That is: if the event corresponds to a key-up, use the “Key Up” table. If the CTRL is down, use the “Controlled” table, and so on.

Special-Entry Values

Special-entry values are classified according to the value of the high-order bits (with the exception of ALT, which is defined as 0x6). 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:
CAPSLock 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 — 0xBF

The low-order bits index a table of strings. Each null-terminated string is returned character by character. The maximum length is defined as:

KTAB_STRLEN 10

Individual string numbers are defined as:

HOMEARROW 0x00

UPARROW 0x01

DOWNARROW 0x02

LEFTARROW 0x03

RIGHTARROW 0x04

String numbers 5 — F are available for custom entries.

Function Key Groups

In the following groups, the low-order bits indicate the function key number within the group:

LEFTFUNC	0xC0
RIGHTFUNC	0xD0
TOPFUNC	0xE0
BOTTOMFUNC	0xF0
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.

Normally, 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.

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 Sun-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

Sun-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 '\n'
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

**Sun-3 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 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 ''	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 ']'	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 SHIFTKEY CAPSLOCK
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 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	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 '\n'
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 SHIFTKEY CAPSLOCK
78 BUCKYBITS+ METABIT	79 ''	7A BUCKYBITS+ METABIT	7B HOLE	7C HOLE	7D HOLE	7E ERROR	7F IDLE

**Sun-3 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 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 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 '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('l')	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 '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 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

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

**Sun-2 Keyboard
Unshifted**

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(
08 TF(3)	09 TF(12)	0A TF(4)	0B TF(13)	0C TF(5)	0D TF(14)	0E TF(6)	0F TF(
10 TF(7)	11 TF(8)	12 TF(9)	13 TF(10)	14 HOLE	15 RF(1)	16 RF(2)	17 RF(
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(
30 HOLE	31 LF(5)	32 LF(13)	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 'j'	42 0x7F	43 HOLE	44 RF(7)	45 STRING+ UPARROW	46 RF(9)	47 HO(
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 '\'	59 '\r'	5A HOLE	5B STRING+ LEFTARROW	5C RF(11)	5D STRING+ RIGHTARROW	5E HOLE	5F LF(
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 HO(
70 BUCKYBITS+ METABIT	71 ''	72 BUCKYBITS+ METABIT	73 HOLE	74 HOLE	75 HOLE	76 ERROR	77 IDL

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(
08 TF(3)	00 TF(12)	0A TF(4)	0B TF(13)	0C TF(5)	0D TF(14)	0E TF(6)	0F TF(
10 TF(7)	11 TF(8)	12 TF(9)	13 TF(10)	14 HOLE	15 RF(1)	16 RF(2)	17 RF(
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(
30 HOLE	31 LF(5)	32 LF(13)	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 'J'	42 0x7F	43 HOLE	44 RF(7)	45 STRING+ UPARROW	46 RF(9)	47 HO
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 '\r'	5A HOLE	5B STRING+ LEFTARROW	5C RF(11)	5D STRING+ RIGHTARROW	5E HOLE	5F LF(
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 HO
78 BUCKYBITS+ METABIT	79 ''	7A BUCKYBITS+ METABIT	7B HOLE	7C HOLE	7D HOLE	7E ERROR	7F IDI

**Sun-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(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 c('')	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 'u'	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 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 '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 HOLE
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+	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 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 LF(13)	33 LF(6)	34	HOLE	35 'u'	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('f')	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 LF(14)	4B HOLE	4C	SHIFTKEYS+ 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('v')	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	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 HOLE
78 BUCKYBITS+ METABIT	79 c('')	7A BUCKYBITS+ METABIT	7B HOLE	7C	HOLE	7D HOLE	7E	ERROR	7F IDLE

**Sun-2 Keyboard
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 OOPS	03 OOPS	04 HOLE	05 OOPS	06 OOPS	07 OO		
08 OOPS	09 OOPS	0A OOPS	0B OOPS	0C OOPS	0D OOPS	0E OOPS	0F OO		
10 OOPS	11 OOPS	12 OOPS	13 OOPS	14 HOLE	15 OOPS	16 OOPS	17 NO		
18 HOLE	19 OOPS	1A OOPS	1B OOPS	1C HOLE	1D NOP	1E NOP	1F NO		
20 NOP	21 NOP	22 NOP	23 NOP	24 NOP	25 NOP	26 NOP	27 NO		
28 NOP	29 NOP	2A NOP	2B NOP	2C HOLE	2D OOPS	2E OOPS	2F NO		
30 HOLE	31 OOPS	32 OOPS	33 OOPS	34 HOLE	35 NOP	36 NOP	37 NO		
38 NOP	39 NOP	3A NOP	3B NOP	3C NOP	3D NOP	3E NOP	3F NO		
40 NOP	41 NOP	42 NOP	43 HOLE	44 OOPS	45 OOPS	46 NOP	47 HO		
48 OOPS	49 OOPS	4A OOPS	4B HOLE	4C SHIFTKEYS+ LEFTCTRL	4D NOP	4e NOP	4F NO		
50 NOP	51 NOP	52 NOP	53 NOP	54 NOP	55 NOP	56 NOP	57 NO		
58 NOP	59 NOP	5A HOLE	5B OOPS	5C OOPS	5D NOP	5E HOLE	5F OO		
60 OOPS	61 OOPS	62 HOLE	63 SHIFTKEYS+ LEFTSHIFT	64 NOP	65 NOP	66 NOP	67 NO		
68 NOP	69 NOP	6A NOP	6B NOP	6C NOP	6D NOP	6E SHIFTKEYS+ RIGHTSHIFT	6F NO		
70 OOPS	71 OOPS	72 NOP	73 HOLE	74 HOLE	75 HOLE	76 HOLE	77 HO		
78 BUCKYBITS+ METABIT	79 NOP	7A BUCKYBITS+ METABIT	7B HOLE	7C HOLE	7D HOLE	7E HOLE	7F RES		

**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('l')	14 HOLE	15 RF(1)	16 '+'	17 '-'	
18 HOLE	19 LF(4)	1A 'f'	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 't'	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 'v'	5A HOLE	5B '1'	5C '2'	5D '3'	5E HOLE	5F NOSCROL	
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 SHIFTKEY RIGHTSHI	
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	

**Micro Switch 103SD32-2 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(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 '!	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 '\t'	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 '\r'	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

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('')	14 HOLE	15 RF(1)	16 '+'	17 '-'
18 HOLE	19 LF(4)	1A '\'	1B LF(6)	1C HOLE	1D SHIFKEYS+ 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 '\t'	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 '\r'	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 SHIFKEYS+ LEFTCTRL	7B ''	7C SHIFKEYS RIGHTCTRL	7D HOLE	7E HOLE	7F IDLE

**Micro Switch 103SD32-2 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 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 'v'	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 'w'	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('\')	59 'v'	5A HOLE	5B OOPS	5C OOPS	5D OOPS	5E HOLE	5F NOSCROI	
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 SHIFTKE RIGHTSH	
70 NOP	71 0x7F	72 OOPS	73 NOP	74 OOPS	75 HOLE	76 HOLE	77 HOLE	
78 HOLE	79 HOLE	7A SHIFTKEYS+ LEFTCTRL	7B '0'	7C SHIFTKEYS+ RIGHTCTRL	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
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 SHIFTKE RIGHTSH	
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	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	'\t'				
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	'\r'	47	'\n'				
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	'\r'	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				

Shifted

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	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	'!'	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	'\t'				
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	'\r'	47	'\n'				
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	'\r'	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
Caps Locked**

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
08	HOLE	09	HOLE	0A	STRING+ UPARROW	0B	STRING+ DOWNARROW	0C	STRING+ LEFTARROW	0D	STRING+ RIGHTARROW	0E	HOLE	0F	TF	10	TF(2)	11	TF(3)
10	TF(2)	11	TF(3)	12	TF(4)	13	c('I')	14	'1'	15	'2'	16	'3'	17	'4'	18	'5'	19	'6'
18	'5'	19	'6'	1A	'7'	1B	'8'	1C	'9'	1D	'0'	1E	'-'	1F	'='	20	'"'	21	c('H')
20	'"'	21	c('H')	22	BUCKYBITS+ METABIT	23	'7'	24	'8'	25	'9'	26	'-'	27	'\u0304'	28	'Q'	29	'W'
28	'Q'	29	'W'	2A	'E'	2B	'R'	2C	'T'	2D	'Y'	2E	'U'	2F	'I'	30	'O'	31	'P'
30	'O'	31	'P'	32	'I'	33	'J'	34	0x7F	35	'4'	36	'5'	37	'6'	38	'.'	39	SHIFTKEYS+ LEFTCTRL
38	'.'	39	SHIFTKEYS+ LEFTCTRL	3A	SHIFTKEYS+ CAPSLOCK	3B	'A'	3C	'S'	3D	'D'	3E	'F'	3F	'G'	40	'H'	41	'J'
40	'H'	41	'J'	42	'K'	43	'L'	44	';'	45	'\"'	46	'r'	47	'\u0304'	48	'1'	49	'2'
48	'1'	49	'2'	4A	'3'	4B	NOP	4C	NOSCROLL	4D	SHIFTKEYS+ LEFTSHIFT	4E	'Z'	4F	'X'	50	'C'	51	'V'
50	'C'	51	'V'	52	'B'	53	'N'	54	'M'	55	'.'	56	'.'	57	'/'	58	SHIFTKEYS+ RIGHTSHIFT	59	'n'
58	SHIFTKEYS+ RIGHTSHIFT	59	'n'	5A	'0'	5B	HOLE	5C	'.'	5D	'r'	5E	HOLE	5F	HOLE	60	HOLE	61	HOLE
60	HOLE	61	HOLE	62	' '	63	HOLE	64	HOLE	65	HOLE	66	HOLE	67	HOLE	68	HOLE	69	HOLE
68	HOLE	69	HOLE	6A	HOLE	6B	HOLE	6C	HOLE	6D	HOLE	6E	HOLE	6F	HOLE	70	HOLE	71	HOLE
70	HOLE	71	HOLE	72	HOLE	73	HOLE	74	HOLE	75	HOLE	76	HOLE	77	HOLE	78	HOLE	79	HOLE
78	HOLE	79	HOLE	7A	HOLE	7B	HOLE	7C	HOLE	7D	HOLE	7E	HOLE	7F	ID				

Controlled

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
08	HOLE	09	HOLE	0A	STRING+ UPARROW	0B	STRING+ DOWNARROW	0C	STRING+ LEFTARROW	0D	STRING+ RIGHTARROW	0E	HOLE	0F	TF	10	TF(2)	11	TF(3)
10	TF(2)	11	TF(3)	12	TF(4)	13	c('I')	14	'1'	15	c('@')	16	'3'	17	'4'	18	'5'	19	c('"'
18	'5'	19	c('"'	1A	'7'	1B	'8'	1C	'9'	1D	'0'	1E	c('_')	1F	'='	20	c('"'	21	c('H')
20	c('"'	21	c('H')	22	BUCKYBITS+ METABIT	23	'7'	24	'8'	25	'9'	26	'-'	27	'\u0304'	28	CTRLQ	29	c('W')
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')
30	c('O')	31	c('P')	32	c('I')	33	c('J')	34	0x7F	35	'4'	36	'5'	37	'6'	38	'.'	39	SHIFTKEYS+ LEFTCTRL
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')
40	c('H')	41	c('J')	42	c('K')	43	c('L')	44	';'	45	'\"'	46	'r'	47	c('\'	48	'1'	49	'2'
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')
50	c('C')	51	c('V')	52	c('B')	53	c('N')	54	c('M')	55	'.'	56	'.'	57	c('/'	58	SHIFTKEYS+ RIGHTSHIFT	59	'n'
58	SHIFTKEYS+ RIGHTSHIFT	59	'n'	5A	'0'	5B	HOLE	5C	'.'	5D	HOLE	5E	HOLE	5F	HOLE	60	HOLE	61	HOLE
60	HOLE	61	HOLE	62	c(' ')	63	HOLE	64	HOLE	65	HOLE	66	HOLE	67	HOLE	68	HOLE	69	HOLE
68	HOLE	69	HOLE	6A	HOLE	6B	HOLE	6C	HOLE	6D	HOLE	6E	HOLE	6F	HOLE	70	HOLE	71	HOLE
70	HOLE	71	HOLE	72	HOLE	73	HOLE	74	HOLE	75	HOLE	76	HOLE	77	HOLE	78	HOLE	79	HOLE
78	HOLE	79	HOLE	7A	HOLE	7B	HOLE	7C	HOLE	7D	HOLE	7E	HOLE	7F	ID				

**VT100-Style 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	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

FILES

/dev/kbd

SEE ALSO

setkeys(1), click(1)

The SunView System Programmer's Guide - Appendix: Writing a Virtual User Input Device Driver
(describes *Firm_event* format)

NAME

le – Sun-3/50 10 Mb/s Ethernet interface

SYNOPSIS

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.

DIAGNOSTICS

le%d: transmitter frozen -- resetting A bug in the LANCE chip has caused the chip's transmitter section to stop. 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 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 ENP bit in rmd cleared This error "should never happen," as it occurs only in conjunction with a LANCE feature that the driver does not use.

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, causing the LANCE chip to give 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.

DIAGNOSTICS

lo%d: can't handle af%d. The interface was handed a message with addresses formatted in an unsuitable address family; the packet was dropped.

SEE ALSO

if(4N), *inet(4F)*

BUGS

It should handle all address and protocol families. An approved network address should be reserved for this interface.

NAME

mem, *kmem*, *vme16d16*, *vme24d16*, *vme32d16*, *vme16d32*, *vme24d32*, *vme32d32*, *mbmem*, *mbio*, – main memory and bus I/O space

SYNOPSIS

None; included with standard system.

DESCRIPTION

These devices are special files that map memory and bus I/O space. They may be read, written, seek'ed and (except for *kmem*) *mmap(2)*'ed.

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.

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 ONLY

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 ONLY

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 1 Megabyte. *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.

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

NAME

mouse – Sun mouse

SYNOPSIS

pseudo-device ms3

DESCRIPTION

The *mouse* interface provides access to the Sun Workstation mouse.

The mouse incorporates a microprocessor which generates a byte-stream protocol encoding mouse motions.

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* delta's 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 $0x80-0x87$.

The mouse can be used as a device that emits *Firm_events* as specified by the protocol of a *Virtual User Input Device*. It understands `VUIDSFORMAT`, `VUIDGFORMAT`, `VUIDSADDR` and `VUIDGADDR` ioctls (see reference below).

FILES

`/dev/mouse`

SEE ALSO

`win(4S)`

The SunView System Programmer's Guide

NAME

mti – Systech MTI-800/1600 multi-terminal interface

SYNOPSIS — SUN-3

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

SYNOPSIS — SUN-2

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

DESCRIPTION

The Systech MTI card provides 8 (MTI-800) or 16 (MTI-1600) serial communication lines with modem control. Each line behaves as described in *ty(4)*. Input and output for each line may independently be set to run at any of 16 speeds; see *ty(4)* for the encoding.

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 cause line *ty02* to be treated in this way.

To allow a single *ty* 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 0 – 127 correspond directly to the normal *ty* lines and are named *ty**. Minor device numbers in the range 128 – 256 correspond to the same physical lines as those above (i.e. the same line as the minor device number minus 128) and are (conventionally) named *cua**. The *cua* lines are special in that they can be opened even when there is no carrier on the line. Once a *cua* line is opened, the corresponding *ty* line can not be opened until the *cua* line is closed. Also, if the *ty* line has been opened successfully (usually only when carrier is recognized on the modem) the corresponding *cua* line can not be opened. This allows a modem to be attached to */dev/tty00* (usually renamed to */dev/ttyd0*) and used for dialin (by enabling the line for login in */etc/ttys*) and also used for dialout (by *tip(1C)* or *uucp(1C)*) as */dev/cua0* when no one is logged in on the line. Note that the bit in the flags word in the config file (see above) must be zero for this line.

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.

FILES

/dev/tty0[0-9a-f] hardwired *ty* lines
/dev/ttyd[0-9a-f] dialin *ty* lines
/dev/cua[0-9a-f] dialout *ty* lines

SEE ALSO

ty(4), *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 : error *n* .

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 couldn't understand.

NAME

mtio – UNIX system magnetic tape interface

SYNOPSIS

```
#include <sys/ioctl.h>
#include <sys/mtio.h>
```

DESCRIPTION

The files *mt0*, ..., *mt15* refer to the UNIX system magnetic tape drives, which read and write magnetic tape in 2048 byte blocks (the 2048 is actually `BLKDEV_IOSIZE` in `<sys/param.h>`). The following description applies to any of the transport/controller pairs. The files *mt0*, ..., *mt3* and *mt8*, ..., *mt11* are rewound when closed; the others are not. When a nine track tape file, open for writing or just written, is closed, two end-of-files are written; if the tape is not to be rewound it is positioned with the head between the two tape-marks. 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.

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 (see below) must be a multiple of the underlying block-size, 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.

The *mt* files discussed above are useful when it you want to access the tape in a way compatible with ordinary files. When using foreign tapes, and especially when reading or writing long records, the 'raw' interface is appropriate. The associated files are named *rmt0*, ..., *rmt15*, but the same minor-device considerations as for the regular files still apply. Each *read* or *write* call reads or writes the next record on the tape. In the write case the record has the same length as the buffer given. During a read, the record size is passed back as the number of bytes read, provided it is no greater than the buffer size. 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.

A number of additional `ioctl` operations are available on raw magnetic tape. 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 */

/* 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 */

/* 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

```

FILES

```

/dev/mt*
/dev/rmt*
/dev/rar*

```

SEE ALSO

```

mt(1), tar(1), ar(4S), tm(4S), st(4S), xt(4S)

```

NAME

nd – network disk driver

SYNOPSIS

pseudo-device nd

DESCRIPTION

The network disk device, */dev/nd**, allows a client workstation to perform disk I/O operations on a server system over the network. To the client system, this device looks like any normal disk driver: it allows read/write operations at a given block number and byte count. Note that this provides a network *disk block* access service rather than a network *file* access service.

Typically the client system will have no disks at all. In this case */dev/nd0* contains the client's root file system (including */usr* files), and *ndl* is used as a paging area. Client access to these devices is converted to *net disk protocol* requests and sent to the server system over the network. The server receives the request, performs the actual disk I/O, and sends a response back to the client.

The server contains a table which lists the net address of each of his clients and the server disk partition which corresponds to each client unit number (*nd0,1,...*). This table resides in the server kernel in a structure owned by the *nd* device. The table is initialized by running the program */etc/nd* with text file */etc/nd.local* as its input. */etc/nd* then issues *ioctl(2)* functions to load the table into the kernel.

In addition to the read/write units */dev/nd**, there are *public* read-only units which are named */dev/ndp**. The correspondence to server partitions is specified by the */etc/nd.local* text file, in a similar manner to the private partitions. The public units can be used to provide shared access to binaries or libraries (*/bin*, */usr/bin*, */usr/ucb*, */usr/lib*) so that each diskless client does not have to consume space in his private partitions for these files. This is done by providing a public file system at the server (*/dev/ndp0*) which is mounted on */pub* of each diskless client. The clients then use symbolic links to read the public files: */bin -> /pub/bin*, */usr/ucb -> /pub/usr/ucb*. One requirement in this case is that the server (who has read/write access to this file system) should not perform write activity with any public filesystem. This is because each client is locally cacheing blocks, and may get out of sync with the physical disk image. In certain cases, the client will detect an inconsistency and panic.

One last type of unit is provided for use by the server. These are called *local* units and are named */dev/ndl**. The Sun physical disk sector 0 label only provides a limited number of partitions per physical disk (eight). Since this number is small and these partitions have somewhat fixed meanings, the *nd* driver itself has a *subpartitioning* capability built-in. This allows the large server physical disk partition (e.g. */dev/xy0g*) to be broken up into any number of diskless client partitions. Of course on the client side these would be referenced as */dev/nd0,1,...*; but the server needs to reference these client partitions from time to time, to do *mkfs(8)* and *fsck(8)* for example. The */dev/ndl** entries allow the server 'local' access to his subpartitions without causing any net activity. The actual local unit number to client unit number correspondence is again recorded in the */etc/nd.local* text file.

The *nd* device driver is the same on both the client and server sides. There are no user level processes associated with either side, thus the latency and transfer rates are close to maximal.

The minor device and *ioctl* encoding used is given in file *<sun/ndio.h>*. The low six bits of the minor number are the unit number. The 0x40 bit indicates a *public* unit; the 0x80 bit indicates a *local* unit.

INITIALIZATION

No special initialization is required on the client side; he finds the server by broadcasting the initial request. Upon getting a response, he locks onto that server address.

At the server, the *nd(8C)* command initializes the network disk service by issuing *ioctl*'s to the kernel.

ERRORS

Generally physical disk I/O errors detected at the server are returned to the client for action. If the server is down or unaccessable, the client will see the console message:

nd: file server not responding: still trying.

The client continues (forever) making his request until he gets positive acknowledgement from the server.

This means the server can crash or power down and come back up without any special action required of the user at the client machine. It also means the process performing the I/O to *nd* will block, insensitive to signals, since the process is sleeping inside the kernel at PRIBIO.

PROTOCOL AND DRIVER INTERNALS

The protocol packet is defined in `<sun/ndio.h>` and also included below:

```

/*
 * 'nd' protocol packet format.
 */
struct ndpack {
    struct ip np_ip; /* ip header, proto IPPROTO_ND */
    u_char np_op;    /* operation code, see below */
    u_char np_min;   /* minor device */
    char np_error;   /* b_error */
    char np_ver;     /* version number */
    long np_seq;     /* sequence number */
    long np_blkno;   /* b_blkno, disk block number */
    long np_bcount;  /* b_bcount, byte count */
    long np_resid;   /* b_resid, residual byte count */
    long np_caddr;   /* current byte offset of this packet */
    long np_ccount;  /* current byte count of this packet */
}; /* data follows */

/*
 * np_oe operation codes.
 */
#define NDOPREAD 1 /* read */
#define NDOPWRITE 2 /* write */
#define NDOPERROR 3 /* error */
#define NDOPCODE 7 /* op code mask */
#define NDOPWAIT 010 /* waiting for DONE or next request */
#define NDOPDONE 020 /* operation done */

/*
 * misc protocol defines.
 */
#define NDMAXDATA 1024 /* max data per packet */
#define NDMAXIO 63*1024 /* max np_bcount */

```

IP datagrams were chosen instead of UDP datagrams because only the IP header is checksummed, not the entire packet as in UDP. Also the kernel level interface to the IP layer is simpler. The *min*, *blkno*, and *bcount* fields are copied directly from the client's strategy request. The sequence number field *seq* is incremented on each new client request and is matched with incoming server responses. The server essentially echos the request header in his responses, altering certain fields. The *caddr* and *ccount* fields show the current byte address and count of the data in this packet, or the data expected to be sent by the other side.

The protocol is very simple and driven entirely from the client side. As soon as the client `ndstrategy` routine is called, the request is sent to the server; this allows disk sorting to occur at the server as soon as possible. Transactions which send data (client writes on the client side, client reads on the server side) can only send a set number of packets of NDMAXDATA bytes each, before waiting for an acknowledgement. The defaults are currently set at 6 packets of 1K bytes each; the `NDIOCETHER` ioctl allows setting this value on the server side. This allows the normal 4K byte case to occur with just one 'transaction'. The NDOPWAIT bit is set in the *op* field by the sender to indicate he will send no more until acknowledged (or requested) by the other side. The NDOPDONE bit is set by the server side to indicate the request operation has completed; for both the read and write cases this means the requested disk I/O has actually occurred.

Requests received by the server are entered on an active list which is timed out and discarded if not completed within NDXTIMER seconds. Requests received by the server allocate a *bcount* size buffer to minimize buffer copying. Contiguous DMA disk I/O thus occurs in the same size chunks it would if requested from a local physical disk.

BOOTSTRAP

The Sun workstation has PROM code to perform a net boot using this driver. Usually, the boot files are obtained from public device 0 (*/dev/ndp0*) on the server with which the client is registered; this allows multiple servers to exist on the same net (even running different releases of kernel and boot software). If the station you are booting is not registered on any of the servers, you will have to specify the hex Internet host number of the server in a boot command string like: 'bec(0,5,0)vmunix'.

This booting performs exactly the same steps involved in a real disk boot:

- 1) User types 'b' to PROM monitor.
- 2) PROM loads blocks 1 thru 15 of */dev/ndp0* (*bootnd*).
- 3) *bootnd* loads */boot*.
- 4) */boot* loads */vmunix*.

SEE ALSO

ioctl(2), *nd*(8C)

BUGS

The operations described in *dkio*(4) are not supported.

The local host's disk buffer cache is not used by network disk access. This means that if either a local host or a remote host is writing, the changes will be visible at random based on the cache hit frequency on the local host. Use *sync* on the server to force the data out to disk. If both the local and remote hosts are writing to the same filesystem, one machine's changes can be randomly lost, based again on cache hit and deferred write timings.

If an R/O remote file system is mounted R/W by mistake, it is impossible to umount it.

NAME

nfs, NFS – network file system

SYNOPSIS

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 *Networking on the Sun Workstation*. 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 the *NFS Protocol Spec.* in *Networking on the Sun Workstation*.

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.

SEE ALSO

exports(5), *fstab(5)*, *mntent(5)*, *mount(2)*, *mount(8)*, *nfsd(8)*

NAME

nit – Network Interface Tap Protocol

SYNOPSIS

options NIT

DESCRIPTION

nit is a provisional protocol family which runs on top of the kernel raw socket code and provides the superuser with a tee connection into a specified network interface. For example, it provides the unprocessed packet read and write capability on the Ethernet interface *ie(4S)*.

nit uses two structures to communicate information, the *nit_ioc* structure, which contains the ioctl information used to set parameter values; and the *nit_hdr* structure, which contains per packet statistics and is prepended to every delivered packet. When setting parameters, values that are otherwise impossible mean "don't change".

nit collects incoming packets into chunks to reduce the per packet overhead. The chunks are returned by *read(2V)* and *recv(2)* system calls. Outgoing packets are not buffered. The ioctl value *nioc_chunksize* sets the size of the incoming chunk. *Nioc_bufalign* and *nioc_bufoffset* control packet placement within buffers. The (nit) header for each packet in a buffer starts *nioc_bufoffset* bytes past some multiple of *nioc_bufalign* bytes from the beginning. The packet itself appears immediately beyond the header. *nit* also limits the amount of buffer space consumed. To change the default, set *nioc_bufspace*.

nit performs packet filtering and data selection on incoming packets. The data selection criterion is the length of the initial portion of the data packet to return to the user. The filtering criteria are packet destination and packet type. The filtering and data selection criteria are set via *nioc_snaplen*, *nioc_flags*, and *nioc_typetomatch*. The choices for destination are either normal or promiscuous. Normal destination filtering considers only those packets that are normally received by the machine running *nit* (both host specific and broadcast packets). Promiscuous destination filtering considers every packet visible on the network; this can place a large demand on the processor if there are many packets to receive. The packets are further filtered on type, an interface specific quantity. For the Ethernet interfaces, the type field is the packet type from the Ethernet header. See *<netinet/if_ether.h>*.

Outgoing packets are not (yet) handled in a general way, since there is no one address family which says "send the packet as is", where the data portion of the packet contains a complete packet to be transmitted without further processing. Therefore, in general, you can't send arbitrary packets. For the Ethernet, however, the address family AF_UNSPEC is defined so that the remaining 14 bytes of the *sockaddr* correspond to the first 14 bytes of the outgoing packet, which are the (6 byte) destination address, the (6 byte) source address (possibly overridden), and the (2 byte) type. See struct *ether_header* in *<netinet/if_ether.c>*. Therefore, for Ethernet in particular, it is possible to transmit an arbitrary packet. In the example which follows, *rarp_write* accepts an arbitrary packet and performs the interface specific manipulations required to transmit that packet.

The following definitions are taken from *<net/nit.h>*.

```
#define NITIFSIZ 10          /* size of ifname in sockaddr */
#define NITBUFSIZ 1024      /* buffers are rounded up to a
                             * multiple of this size (MCLBYTES) */

struct sockaddr_nit {
    u_short    snit_family;
    caddr_t    snit_cookie;          /* link to filtering */
    char    snit_ifname[NITIFSIZ];  /* interface name (eg, ie0) */
};

/* Header preceding each packet returned to user */
struct nit_hdr {
    int    nh_state;          /* state of tap -- see below */
    struct timeval nh_timestamp; /* time of arriving packet */
    int    nh_wirelen;       /* length (with header) off wire */
};
```

```

        union {
            int    info;          /* generic information */
            int    datalen;       /* length of saved packet portion */
            int    dropped;       /* number of dropped matched packets */
            int    seqno;         /* sequence number */
        } nh_un;
    };
#define nh_info          nh_un.info
#define nh_datalen       nh_un.datalen
#define nh_dropped       nh_un.dropped
#define nh_seqno         nh_un.seqno

/* Ioctl parameter block */
struct nit_ioc {
    int    nioc_bufspace;        /* total buffer space to use */
    int    nioc_chunksize;       /* size of chunks to send */
    u_int  nioc_typedtomatch;     /* magic type with which to match */
    int    nioc_snaplen;         /* length of packet portion to sna */
    int    nioc_bufalign;        /* packet header alignment multipl */
    int    nioc_bufoffset;       /* packet header alignment offset */
    struct timeval nioc_timeout; /* delay after packet before drain */
    int    nioc_flags;           /* see below */
};
#define NT_NOTYPES      ((u_int)0) /* match no packet types */
#define NT_ALLTYPES     ((u_int)-1) /* match all packet types */

#define NF_PROMISC      0x01       /* enter promiscuous mode */
#define NF_TIMEOUT      0x02       /* timeout value valid */
#define NF_BUSY         0x04       /* buffer is busy (has data) */

/*
 * States for the packet capture portion of nit,
 * some of which are passed to the user.
 */
#define NIT_QUIET       0          /* inactive */
#define NIT_CATCH       1          /* capturing packets */
#define NIT_NOMBUF      2          /* discarding -- out of mbufs */
#define NIT_NOCLUSTER   3          /* discarding -- out of mclusters */
#define NIT_NOSPACE     4          /* discarding -- would exceed buf space */
/* Pseudo-states returned in information packets */
#define NIT_SEQNO       5          /* sequence number of chunk */

```

To use *nit*:

- o Include definitions and declare needed variables, for example

```

#include <sys/types.h>
#include <sys/time.h>
#include <sys/socket.h>
#include <sys/ioctl.h>
#include <net/nit.h>
#include <net/if.h>

struct sockaddr_nit snit;
struct nit_ioc nioc;

```

- o Create a socket with the call


```
s = socket(AF_NIT, SOCK_RAW, NITPROTO_RAW);
```
- o Bind it to an interface with a code fragment like


```
snit.snit_family = AF_NIT;
strncpy(snit.snit_ifname, "ie0", sizeof (snit.snit_ifname));
bind(s, (struct sockaddr *)&snit, sizeof (snit));
```
- o To establish the operating modes, issue an `ioctl`; for example


```
bzero(&nioc, sizeof(nioc));
nioc.nioc_bufspace = NITBUFSIZ;
nioc.nioc_chunksize = NITBUFSIZ;
nioc.nioc_typedmatch = NT_ALLTYPES;
nioc.nioc_snaplen = 32767;
nioc.nioc_flags = NF_TIMEOUT;
nioc.nioc_timeout.tv_usec = 200;
if (ioctl(if_fd, SIOCSNIT, &nioc) != 0) {
    perror("nit ioctl");
    exit(2);
}
```
- o To receive packets, issue `reads` (or `recvs`). To transmit packets, issue `writes` (or `sends`). For example, the following routine will transmit an arbitrary packet (including address information) on the Ethernet. Note that the Ethernet addresses and type are provided in the incoming buffer `buf`, and must be moved into the `sockaddr` destination address to satisfy the kernel.


```
rarp_write(fd, buf, len)
    int fd, len;
    char *buf;

    {
        struct sockaddr sa;
        int offset = sizeof(sa.sa_data);
        int result;

        sa.sa_family = AF_UNSPEC;
        bcopy(buf, sa.sa_data, offset);
        result = sendto(fd, buf+offset, len-offset,
            0, &sa, sizeof(sa));
        return (result+offset);
    }
```

SEE ALSO

`bind(2)`, `config(8)`, `ec(4S)`, `ie(4S)`, `if(4N)`, `ioctl(2)`, `read(2V)`, `recv(2)`, `send(2)`, `socket(2)`, `write(2V)`.

Network Implementation in Networking on the Sun Workstation.

BUGS

This protocol is provisional, and is subject to change.

Buffering is limited to 32767 bytes.

Interface `ioctl`'s may have different semantics on a nit socket.

`nit` is unable to see outgoing transmissions on some interfaces.

The selection criteria is very simplistic. Therefore, many packets may be passed to the user program, especially in promiscuous mode.

NAME

null – data sink

SYNOPSIS

None; included with standard system.

DESCRIPTION

Data written on a null special file is discarded.

Reads from a null special file always return an end-of-file indication.

FILES

/dev/null

NAME

pty – pseudo terminal driver

SYNOPSIS

pseudo-device pty

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 *master* and a *slave*. The slave device provides an interface identical to that described in *tty(4)*, but instead of having a hardware interface such as the Zilog chip and associated hardware used by *zs(4S)* supporting the terminal functions, the functions of the terminal are implemented by another process manipulating the master side of the pseudo-terminal.

The master and the slave sides of the pseudo-terminal are tightly connected. Any data written on the master 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 master 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.

A few special *ioctl*'s are provided on the control-side devices of pseudo-terminals to provide the functionality needed by applications programs to emulate real hardware interfaces:

TIOCSTOP

Stops output to a terminal (that is, like typing ^S). Takes no parameter.

TIOCPKT

Restarts output (stopped by TIOCSTOP or by typing ^Q). Takes no parameter.

There are also two independent modes which can be used by applications programs:

TIOCPKT

Enable/disable *packet* mode. Packet mode is enabled by specifying (by reference) a nonzero parameter and disabled by specifying (by reference) a zero parameter. When applied to the master side of a pseudo terminal, each subsequent *read* from the terminal will return data written on the slave part of the pseudo terminal preceded by a zero byte (symbolically defined as TIOCPKT_DATA), or a single byte reflecting control status information. In the latter case, the byte is an inclusive-or of zero or more of the bits:

TIOCPKT_FLUSHREAD

whenever the read queue for the terminal is flushed.

TIOCPKT_FLUSHWRITE

whenever the write queue for the terminal is flushed.

TIOCPKT_STOP

whenever output to the terminal is stopped a la ^S.

TIOCPKT_START

whenever output to the terminal is restarted.

TIOCPKT_DOSTOP

whenever *t_stopc* is ^S and *t_startc* is ^Q.

TIOCPKT_NOSTOP

whenever the start and stop characters are not ^S/^Q.

This mode is used by *rlogin(1C)* and *rlogind(8C)* to implement a remote-echoed, locally ^S/^Q flow-controlled remote login with proper back-flushing of output when interrupts occur; it can be used by other similar programs.

TIOCREMOTE

A mode for the master half of a pseudo terminal, independent of TIOCPKT. This mode causes

input to the pseudo terminal to be flow controlled and not input edited (regardless of the terminal mode). Each write to the control terminal produces a record boundary for the process reading the terminal. In normal usage, a write of data is like the data typed as a line on the terminal; a write of 0 bytes is like typing an end-of-file character. TIOCREMOTE can be used when doing remote line editing in a window manager, or whenever flow controlled input is required.

FILES

/dev/pty[p-r][0-9a-f]	master pseudo terminals
/dev/tty[p-r][0-9a-f]	slave pseudo terminals

BUGS

It is apparently not possible to send an EOT by writing zero bytes in TIOCREMOTE mode.

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 (i.e. the packet is forwarded).

Routing table entries installed by a user process may not specify the hash, reference count, use, or interface fields; these are filled in by the routing routines. If a route is in use when it is deleted (*rt_refcnt* is non-zero), the resources associated with it will not be reclaimed until 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.

SEE ALSO

route(8C), routed(8C)

NAME

sd – Disk driver for SCSI Disk Controllers

SYNOPSIS — SUN-3

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. The third controller line above specifies the first and only SCSI host adapter on a Sun-3/50. 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.

SYNOPSIS — SUN-2

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. The third controller line above specifies the first host adapter on a Sun-2/160. 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).

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

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.

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

/dev/sd[0-7][a-h] block files
/dev/rsd[0-7][a-h] raw files

SEE ALSO

dkio(4S)
Adaptec ACB 4000 and 5000 Series Disk Controllers OEM Manual
Emulex MD21 SCSI Disk Controller Programmer Reference Manual

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

st – Driver for Sysgen SC 4000 (Archive) and the Emulex MT-02 Tape Controller

SYNOPSIS — SUN-3

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. The third controller line above specifies the first and only SCSI controller on a Sun-3/50. 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.

SYNOPSIS — SUN-2

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. The third controller line specifies the first controller on a Sun-2/160. 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.

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).

FILES

/dev/rst[0-3] QIC 11 Format
 /dev/rst[8-11] QIC 24 Format
 /dev/nrst[0-3] non-rewinding QIC 11 Format
 /dev/nrst[8-11] non-rewinding QIC 24 Format

SEE ALSO

mtio(4)

Sysgen SC4000 Intelligent Tape Controller Product Specification

DIAGNOSTICS

st*: tape not online.
 st*: no cartridge in drive.
 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. This will cause empty files to 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

tcp – Internet Transmission Control Protocol

SYNOPSIS

None; included automatically with *inet*(4F).

DESCRIPTION

TCP is a connection-oriented, end-to-end reliable protocol designed to fit into a layered hierarchy of protocols which support multi-network applications. TCP provides for reliable inter-process communication between pairs of processes in host computers attached to distinct but interconnected computer communication networks. Very few assumptions are made as to the reliability of the communication protocols below TCP layer. TCP assumes it can obtain a simple, potentially unreliable datagram service from the lower level protocols. In principle, TCP should be able to operate above a wide spectrum of communication systems ranging from hard-wired connections to packet-switched or circuit switched networks.

TCP fits into a layered protocol architecture just above the basic Internet Protocol (IP) described in *ip*(4P) which provides a way for TCP to send and receive variable-length segments of information enclosed in Internet datagram “envelopes.” The Internet datagram provides a means for addressing source and destination TCPs in different networks, deals with any fragmentation or reassembly of the TCP segments required to achieve transport and delivery through multiple networks and interconnecting gateways, and has the ability to carry information on the precedence, security classification and compartmentalization of the TCP segments (although this is not currently implemented under the UNIX system.)

An application process interfaces to TCP through the *socket*(2) abstraction and the related calls *bind*(2), *listen*(2), *accept*(2), *connect*(2), *send*(2) and *recv*(2). The primary purpose of TCP is to provide a reliable bidirectional virtual circuit service between pairs of processes. In general, the TCP's decide when to block and forward data at their own convenience. In the UNIX system implementation, it is assumed that any buffering of data is done at the user level, and the TCP's transmit available data as soon as possible to their remote peer. They do this and always set the PUSH bit indicating that the transferred data should be made available to the user process at the remote end as soon as practicable.

To provide reliable data TCP must recover from data that is damaged, lost, duplicated, or delivered out of order by the underlying internet communications system. This is achieved by assigning a sequence number to each byte of data transmitted and requiring a positive acknowledgement from the receiving TCP. If the ACK is not received within an (adaptively determined) timeout interval, the data is retransmitted. At the receiver, the sequence numbers are used to correctly order segments that may be received out of order and to eliminate duplicates. Damage is handled by adding a checksum to each segment transmitted, checking it at the receiver, and discarding damaged segments. As long as the TCP's continue to function properly and the internet system does not become disjoint, no transmission errors will affect the correct delivery of data, as TCP recovers from communications errors.

TCP provides flow control over the transmitted data. The receiving TCP is allowed to specify the amount of data which may be sent by the sender, by returning a *window* with every acknowledgement indicating a range of acceptable sequence numbers beyond the last segment successfully received. The window indicates an allowed number of bytes that the sender may transmit before receiving further permission.

TCP extends the standard 32-bit Internet host addresses with a 16-bit port number space; the combined addresses are available at the UNIX system process level in the standard *sockaddr_in* format described in *inet*(4F).

Sockets utilizing the tcp protocol are either “active” or “passive”. Active sockets initiate connections to passive sockets. By default TCP sockets are created active; to create a passive socket the *listen*(2) system call must be used after binding the socket to an address with the *bind*(2) system call. Only passive sockets may use the *accept*(2) call to accept incoming connections. Only active sockets may use the *connect*(2) call to initiate connections.

Passive sockets may “underspecify” their location to match incoming connection requests from multiple networks. This technique, termed “wildcard addressing”, allows a single server to provide service to clients on multiple networks. To create a socket which listens on all networks, the Internet address

INADDR_ANY must be bound. The TCP port may still be specified at this time; if the port is not specified the system will assign one. Once a connection has been established the socket's address is fixed by the peer entity's location. The address assigned the socket is the address associated with the network interface through which packets are being transmitted and received. Normally this address corresponds to the peer entity's network. See *inet(4F)* for a complete description of addressing in the Internet family.

A TCP connection is created at the server end by doing a *socket(2)*, a *bind(2)* to establish the address of the socket, a *listen(2)* to cause connection queueing, and then an *accept(2)* which returns the descriptor for the socket. A client connects to the server by doing a *socket(2)* and then a *connect(2)*. Data may then be sent from server to client and back using *read(2V)* and *write(2V)*.

TCP implements a very weak out-of-band mechanism, which may be invoked using the out-of-band provisions of *send(2)*. This mechanism allows setting an urgent pointer in the data stream; it is reflected to the TCP user by making the byte after the urgent pointer available as out-of-band data and providing a *SIOCATMARK* ioctl which returns an integer indicating whether the stream is at the urgent mark. The system never returns data across the urgent mark in a single read. Thus, when a SIGURG signal is received indicating the presence of out-of-band data, and the out-of-band data indicates that the data to the mark should be flushed (as in remote terminal processing), it suffices to loop, checking whether you are at the out-of-band mark, and reading data while you are not at the mark.

SEE ALSO

inet(4F), *ip(4P)*

BUGS

It should be possible to send and receive TCP options.

The system always tries to negotiate the maximum TCP segment size to be 1024 bytes. This can result in poor performance if an intervening network performs excessive fragmentation.

SIOCSETH and *SIOCGHWM* ioctls 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

None; included by default.

DESCRIPTION

This section describes the special file `/dev/tty` and the terminal drivers used for interactive I/O by devices such as `zs(4S)`, `cons(4S)`, and `pty(4)`.

Opening a Terminal File

When a terminal file is opened, it normally causes the process to wait until a connection is established. In practice, users' programs seldom open these files; they are opened by `init(8)` and become a user's standard input, output, and error files.

The Controlling Terminal

A terminal may belong to a process, in which case it is known as its *controlling terminal*. This controlling terminal may have a distinguished process group associated with it which plays a special role in handling QUIT and INT (interrupt) signals, as discussed below. The controlling terminal is inherited by a child process during a `fork(2)`.

If a process that has no controlling terminal opens a terminal file, then the device or pseudo-device associated with that terminal file becomes the controlling terminal for the process; the terminal's distinguished process group is set to that of the process.

The file `/dev/tty` is, for each process, a synonym for its controlling terminal. This is useful for programs that wish to be sure of writing messages on the terminal directly, no matter how output has been redirected. It can also be used for programs that demand a filename for output when typed output is desired and it is tiresome to find out which terminal is currently in use.

Implementation Restrictions

Due to restrictions imposed by the current terminal driver, some features are not fully supported:

1. Certain terminal driver features are always enabled, except when the driver is in "RAW mode". If the character size is 8 bits, no parity is specified, no output processing is selected (i.e., either OPOST is on or none of OLCUC, ONLCR, or any of the delays are selected), and no input process is selected (as with BRKINT, IGNPAR, INPCK, ISTRIP, ICRNL, IUCLC, and IXON are all off in the `c_iflag` word and ISIG, ICANON, and XCASE are all off in the `c_lflag` word), the driver is in "RAW mode". If a terminal port is being used for transferring binary data (such as when `uucp(1)` or some microcomputer data transfer program like KERMIT is using the port), it is usually in "RAW mode". If it is being used to give a user interactive access to the computer, it is usually not in "RAW mode".

BRKINT and IGNPAR are disabled only in "RAW mode"; if they are to be disabled, the other modes listed must also be disabled. The WERASE, REPRINT, DISCARD, and LNEXT characters are also disabled only in raw mode.

2. IUCLC, OLCUC, and XCASE must either all be on or all be off, and ICRNL and ONLCR must either both be on or both be off.
3. The MIN and TIME values supported by other implementations can be set, but this has no effect on the terminal driver. The driver behaves as if MIN were 1 and TIME were 0.
4. Character sizes CS5 and CS6 may not be selected; size CS7 may only be selected when PARENB is set, and size CS8 may only be selected when PARENB is not set. Furthermore, if size CS8 is selected, unless OPOST is not set, it only applies to input, not output.
5. IGNBRK, PARMRK, INLCR, IGNCR, ONOCR, OFILL, OFDEL and ECHONL are treated as if they were always not set. CREAD and ECHOK are treated as if they were always set.
6. The TCSETAW call will flush any pending input, just as TCSETAF does. TCSBRK will also flush any pending input.

7. The EOF character may not be escaped with a backslash (\) (unless XCASE is set; see below).
8. If XCASE is set, a backslash (\) followed by any character other than a letter or one of the special characters listed in the description of XCASE will be read as the character; the backslash will not be read. This includes the special characters ERASE, WERASE, KILL, REPRINT, EOF, NEWLINE (the ASCII NL character), EOL, and DISCARD.

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. When the input limit is reached, if the terminal port is in "RAW mode", all the saved characters are thrown away without notice. Otherwise, any further input is discarded and an ASCII BEL character is echoed.

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).

Canonical Mode Input Processing

In canonical mode, terminal input is processed in units of lines. A line is delimited by a NL (ASCII LF) character, an end-of-file (ASCII EOT) character, or an end-of-line character. This means that a program attempting to read will be suspended until an entire line has been typed. 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 can be requested in a read, even one, without losing information.

Erase and kill processing is normally done during input. The ERASE character (by default, the character DEL) erases the last character typed. The WERASE character (the character ^W) erases the last "word" typed (but not any preceding spaces or tabs). A "word" is defined as a sequence of nonblank characters, with tabs counted as blanks. Neither ERASE nor WERASE will erase beyond the beginning of the line. The KILL character (by default, the character ^U) kills (deletes) the entire input line, and optionally produces a NL character. These special characters operate on a keystroke basis, independently of any backspacing or tabbing that may have been done.

The REPRINT character (the character ^R) prints a NL followed by all characters which have not been read. Reprinting also occurs automatically if characters which 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 by such commands as *stty(1)*, and *tset(1)*.

Noncanonical Mode Input Processing

In non-canonical mode, input characters are not assembled into lines, and erase and kill processing does not occur. Characters are read as soon as they are typed.

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 (Control-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** (Control-| 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, or EOL character.
- WERASE** (^W or ASCII ETB) erases the preceding "word". It will not erase beyond the start of a line, as delimited by a NL, EOF, or EOL character.
- KILL** (^U or ASCII NAK) deletes the entire line, as delimited by a NL, EOF, or EOL character.
- REPRINT** (^R or ASCII DC2) reprints all characters which have not been read, preceded by a NL.
- EOF** (^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 NL, 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 or escaped.
- EOL** (Off by default) is an additional line delimiter, like NL. It is not normally used.
- SUSP** (^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.
- DSUSP** (^Y or ASCII SUB) is used by the job control facility to change the current job to return to the controlling job. It generates a SIGTSTP signal as SUSP does, but the signal is sent when a program attempts to read the DSUSP character, rather than when it is typed.
- STOP** (^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** (^Q or ASCII DC1) is used to resume output which has been suspended by a STOP character. While output is not suspended, START characters are ignored and not read. The start/stop characters can not be changed or escaped.
- DISCARD** (^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** (^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 (such as KILL, or QUIT).

The character values for INTR, QUIT, ERASE, KILL, EOF, and EOL may be changed to suit individual tastes. The ERASE and KILL 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.

When in "RAW mode", none of the special characters perform any special function.

Modem Disconnect

When the carrier signal from the data-set drops, 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 causes the processes to terminate. If SIGHUP is ignored or caught, any subsequent read returns with an end-of-file indication. Thus, programs that read a terminal and test for end-of-file can terminate

appropriately when hung up on.

ioctl Calls

Several *ioctl*(2) system calls apply to terminal files. The primary calls use the following structure, defined in `<termio.h>`:

```
#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 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   (disabled)
6  reserved
7  reserved
```

Input Modes

The `c_iflag` field describes the basic terminal input control:

```
BRKINT    0000002  Signal interrupt on break.
IGNPAR    0000004  Ignore characters with parity errors.
INPCK     0000020  Enable input parity check.
ISTRIP    0000040  Strip character.
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.
```

If BRKINT is set, the break condition will generate an interrupt signal and flush both the input and output queues. If IGNPAR is set, characters with other framing and parity errors are ignored. A framing or parity error which is not ignored is read as the 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 eight bits are processed.

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. Note: if this bit is set, the OLCUC bit in the `c_oflag` word and the XCASE bit in the `c_lflag` word must also be set.

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. All start/stop characters are not read, but merely perform flow control functions. If IXANY is set, any input character, will restart output which 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.

The initial input control value is undefined.

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.
ONLRET	0000040	NL performs CR function.
NLDLY	0000400	Select
NL		
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	
TAB3	0014000	Expand tabs to spaces.
BSDLY	0020000	Select backspace delays:
BS0	0	
BS1	0020000	
VDLY	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. Note: if this bit is set, the IUCLC bit in the *c_iflag* word and the XCASE bit in the *c_lflag* word must also be set.

If ONLCR is set, the NL character is transmitted as the CR-NL character pair. 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 a form-feed or vertical-tab delay is specified, it lasts for about 2 seconds.

NEWLINE

delay lasts about 0.10 seconds. If ONLRET is set, the carriage-return delays are used instead of the NL delays.

Carriage-return delay type 1 is dependent on the current column position, type 2 is about 0.08 seconds, and type 3 is about 0.16 seconds.

Horizontal-tab delay type 1 is dependent on the current column position. Type 2 is not supported. Type 3 specifies that tabs are to be expanded into spaces.

Backspace delay is not supported.

The actual delays depend on line speed and system load.

The initial output control value is undefined.

Control Modes

The *c_flag* 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
EXTA	0000016	19200 baud
EXTB	0000017	External B
CSIZE	0000060	Character size:
CS7	0000040	7 bits
CS8	0000060	8 bits
CSTOPB	0000100	Send two stop bits, else one.
PARENB	0000400	Parity enable.
PARODD	0001000	Odd parity, else even.
HUPCL	0002000	Hang up on last close.
CLOCAL	0004000	Local line, else dial-up.

The CBAUD bits specify the baud rate. The zero baud rate, B0, is used to hang up the connection. If B0 is specified, the data-terminal-ready signal will not be asserted. Normally, this will disconnect the line. 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. The only combinations that are supported are CS7 with PARENB and CS8 without PARENB.

If HUPCL is set, the line will be disconnected when the last process with the line open closes it or terminates. That is, the data-terminal-ready signal will not be asserted.

If CLOCAL is set, the line is assumed to be a local, direct connection with no modem control. Otherwise modem control is assumed.

The initial hardware control value after open is undefined.

Local Modes

The *c_iflag* 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.
NOFLSH	0000200	Disable flush after interrupt or quit.

If ISIG is set, each input character is checked against the special control characters INTR, QUIT, SUSP, and DSUSP. 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. These functions may be disabled individually by changing the value of the control character to an unlikely or impossible value (e.g., 0377).

If ICANON is set, canonical processing is enabled. This enables the erase and kill edit functions, and the assembly of input characters into lines delimited by NL, EOF, and EOL. If ICANON is not set, read requests are satisfied directly from the input queue. A read will be satisfied as soon as one character is received; the values of MIN and TIME are ignored.

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>
\	\/
	!\
~	\/
{	\{
}	\}

Any other character, when preceded on input by \, will be read as itself, and the \ will not be read. This means a \ must be entered as \\. For example, A is input as \a, \n as \\n, and \N as \\N. Note: if this bit is set, the IUCLC bit in the *c_iflag* word and the OLCUC bit in the *c_oflag* word must also be set.

If ECHO is set, characters are echoed as received. If ECHO is not set, input characters are not echoed.

When ICANON is set, the following echo functions are possible. If ECHO and ECHOE are set, the erase character is echoed as a sequence of ASCII BS SP BS, which will clear the last character from a CRT screen. If the baud rate is greater than 1200 baud, the kill character is echoed as a sequence of ASCII BS SP BS, which will clear all the characters on the current line from a CRT screen; otherwise, it is echoed as itself (if it is a control character, it will be echoed as described below) followed by an NL character. If ECHOE is not set, the erase character is echoed by printing the character being erased; erased characters are echoed between a backslash (\) and a slash (/). The NL character is always echoed after the kill character to emphasize that the line will be deleted. Note that an escape character preceding the erase or kill character removes any special function.

Non-printing (control) characters are normally 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 ^?. In "RAW mode", control characters and DEL are echoed as themselves.

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.

The initial line-discipline control value is undefined.

The primary *ioctl(2)* system calls have the form:

```
ioctl (fdes, command, arg)
struct termio *arg;
```

The commands using this form are:

TCGETA Get the parameters associated with the terminal and store in the *termio* structure referenced by *arg*.

TCSETA Set the parameters associated with the terminal from the structure referred to by *arg*. The change is immediate.

TCSETAW

TCSETAF Wait for the output to drain, then flush the input queue and set the new parameters. This form should be used when changing parameters that will affect output.

Additional *ioctl(2)* calls have the form:

```
ioctl (fdes, command, arg)
int arg;
```

The commands using this form are:

TCSBRK Wait for the output to drain. If *arg* is 0, then send a break (zero bits for 0.25 seconds).
Note: this call will flush any pending input.

TCXONC Start/stop control. If *arg* is 0, suspend output; if 1, restart suspended output.

TCFLSH If *arg* is 0, flush the input queue; if 1, flush the output queue; if 2, flush both the input and output queues.

FILES

*/dev/tty**

SEE ALSO

stty(1V), *tset(1)*, *fork(2)*, *ioctl(2)*, *setpgrp(2V)*, *signal(2)*.

NAME

tm – tapemaster 1/2 inch tape drive

SYNOPSIS — SUN-3

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

SYNOPSIS — SUN-2

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)

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

tty – general terminal interface

SYNOPSIS

None; included by default.

DESCRIPTION

This section describes the special file `/dev/tty` and the terminal drivers used for conversational computing by devices such as `zs(4S)`, `cons(4S)`, and `pty(4)`.

Line disciplines.

The system provides different *line disciplines* for controlling communications lines. In this version of the system there are three disciplines available:

- old The old (standard) terminal driver. This is used when using the standard shell `sh(1)` and for compatibility with Version 7 UNIX systems.
- new A newer terminal driver, with features for job control; this must be used when using `csh(1)`.
- net A line discipline used for networking and loading data into the system over communications lines. It allows high speed input at very low overhead, and is described in `bk(4)`.

Line discipline switching is accomplished with the TIOCSETD *ioctl*:

```
int ldisc = LDISC;
ioctl(f, TIOCSETD, &ldisc);
```

where LDISC is OTTYDISC for the standard tty driver, NTTYDISC for the new driver and NETLDISC for the networking discipline. The standard (currently old) tty line discipline is 0 by convention. The current line discipline can be obtained with the TIOCGETD *ioctl*. Pending input is discarded when the line discipline is changed.

All of the low-speed asynchronous communications ports can use any of the available line disciplines, no matter what hardware is involved. The remainder of this section discusses the “old” and “new” disciplines.

Opening a terminal device file.

When a terminal file is opened, it causes the process to wait until a connection is established. In practice, user programs seldom open these files; they are opened by `init(8)` and become a user’s standard input, output, and error files.

The controlling terminal.

A terminal may belong to a process as its *controlling terminal*, and may have a distinguished process group associated with it. This distinguished process group plays a special role in handling quit and interrupt signals, as discussed below.

If a process which has no controlling terminal opens a terminal file, then 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. The control terminal is thereafter inherited by a child process during a `fork(2)`, even if the control terminal is closed.

The file `/dev/tty` is, in each process, a synonym for that process’ controlling terminal. It is useful for programs that wish to be sure of writing messages on the terminal no matter how output has been redirected. It can also be used for programs that demand a file name for output, when typed output is desired and it is tiresome to find out which terminal is currently in use.

A process can remove the association it has with its controlling terminal by opening the file `/dev/tty` and issuing a

```
ioctl(f, TIOCNOTTY, 0);
```

This is often desirable in server processes.

Process groups.

Command processors such as *cs(1)* can arbitrate the terminal between different *jobs* by placing related jobs in a single process group and associating this process group with the terminal. A terminal's associated process group may be set using the *TIOCSPGRP ioctl(2)*:

```
ioctl(fildes, TIOCSPGRP, &pgrp);
```

or examined using *TIOCGPGRP*, which returns the current process group in *pgrp*. The new terminal driver aids in this arbitration by restricting access to the terminal by processes which are not in the current process group; see **Job access control** below.

Modes.

The terminal line disciplines have three major modes, characterized by the amount of processing on the input and output characters:

- cooked** The normal mode. In this mode lines of input are collected and input editing is done. The edited line is made available when it is completed by a newline or when the *t_brkc* character, normally an EOT (control-D, hereafter *^D*), is entered. A carriage return is usually made synonymous with newline in this mode, and replaced with a newline whenever it is typed. All line discipline functions (input editing, interrupt generation, output processing such as delay generation and tab expansion, etc.) are available in this mode.
- CBREAK** This mode eliminates the character, word, and line editing input facilities, making the input character available to the user program as it is typed. Flow control, literal-next and interrupt processing are still done in this mode. Output processing is done.
- RAW** This mode eliminates all input processing and makes all input characters available as they are typed; no output processing is done either.

The style of input processing can also be very different when the terminal is put in non-blocking I/O mode; see the *FNDELAY* flag as described in *fcntl(2)*. In this case a *read(2V)* from the control terminal will never block, but rather return an error indication (*EWOULDBLOCK*) if there is no input available.

A process may also request a *SIGIO* signal be sent it whenever input is present. To enable this mode the *FASYNC* flag should be set using *fcntl(2)*.

Input editing.

A UNIX system terminal ordinarily operates in full-duplex mode. Characters may be typed at any time, even while output is occurring, and are only lost when the system's character input buffers become completely choked, which is rare, or when the user has accumulated the maximum allowed number of input characters that have not yet been read by some program. Currently this limit is 256 characters. In *RAW* mode, the terminal driver throws away all input and output without notice when the limit is reached. In *CBREAK* or *cooked* mode it refuses to accept any further input and, if in the new line discipline, rings the terminal bell.

Input characters are normally accepted in either even or odd parity with the parity bit being stripped off before the character is given to the program. By clearing either the *EVEN* or *ODD* bit in the flags word it is possible to have input characters with that parity discarded (see the **Summary** below.)

In all of the line disciplines, it is possible to simulate terminal input using the *TIOCSTI ioctl*, which takes, as its third argument, the address of a character. The system pretends that this character was typed on the argument terminal, which must be the control terminal except for the super-user (this call is not in standard Version 7 UNIX systems).

Input characters are normally echoed by putting them in an output queue as they arrive. This may be disabled by clearing the *ECHO* bit in the flags word using the *stty(3C)* call or the *TIOCSETN* or *TIOCSETP ioctls* (see the **Summary** below).

In cooked mode, terminal input is processed in units of lines. A program attempting to read will normally be suspended until an entire line has been received (but see the description of SIGTTIN in **Job access control** and of FIONREAD in **Summary**, both below.) No matter how many characters are requested in the read call, at most one line will be returned. It is not, however, necessary to read a whole line at once; any number of characters may be requested in a read, even one, without losing information.

During input, line editing is normally done, with the DELETE character logically erasing the last character typed and a ^U (control-U) logically erasing the entire current input line. These characters never erase beyond the beginning of the current input line or an ^D. These characters may be entered literally by preceding them with '\'; the '\' will normally be erased when the character is typed.

The line disciplines normally treat either a carriage return or a newline character as terminating an input line, replacing the return with a newline and echoing a return and a line feed. If the CRMOD bit is cleared in the local mode word then the processing for carriage return is disabled, and it is simply echoed as a return, and does not terminate cooked mode input.

In the new line discipline there is a literal-next character ^V which can be typed in both cooked and CBREAK mode preceding *any* character to prevent its special meaning. This is to be preferred to the use of '\' escaping erase and kill characters, but '\' is retained with its old function in the new line discipline.

The new terminal line discipline also provides two other editing characters in normal mode. The word-erase character, normally ^W, erases the preceding word, but not any spaces before it. For the purposes of ^W, a word is defined as a sequence of non-blank characters, with tabs counted as blanks. Finally, the reprint character, normally ^R, retypes the pending input beginning on a new line. Retyping occurs automatically in cooked mode if characters which would normally be erased from the screen are fouled by program output.

Input echoing and redisplay

The terminal driver has several modes (not present in standard UNIX Version 7 systems) for handling the echoing of terminal input, controlled by bits in a local mode word.

Hardcopy terminals. When a hardcopy terminal is in use, the LPRTERA bit is normally set in the local mode word. Characters which are logically erased are then printed out backwards preceded by '\' and followed by '/' in this mode.

CRT terminals. When a CRT terminal is in use, the LCRTBS bit is normally set in the local mode word. The terminal line discipline then echoes the proper number of backspace characters when input is erased to reposition the cursor. If the input has become fouled due to interspersed asynchronous output, the input is automatically retyped.

Erasing characters from a CRT. When a CRT terminal is in use, the LCRTERA bit may be set to cause input to be erased from the screen with a "backspace-space-backspace" sequence when character or word deleting sequences are used. A LCRTKIL bit may be set as well, causing the input to be erased in this manner on line kill sequences as well.

Echoing of control characters. If the LCTLECH bit is set in the local state word, then non-printing (control) characters are normally echoed as ^X (for some X) rather than being echoed unmodified; delete is echoed as ^?.

The normal modes for use on CRT terminals are speed dependent. At speeds less than 1200 baud, the LCRTERA and LCRTKIL processing is painfully slow, so *stty(1)* normally just sets LCRTBS and LCTLECH; at speeds of 1200 baud or greater all of these bits are normally set. The *stty(1)* command summarizes these option settings and the use of the new terminal line discipline as "newcrt."

Output processing.

When one or more characters are written, they are actually transmitted to the terminal as soon as previously-written characters have finished typing. (As noted above, input characters are normally echoed by putting them in the output queue as they arrive.) When a process produces characters more rapidly than they can be typed, it will be suspended when its output queue exceeds some limit. When the queue has

drained down to some threshold the program is resumed. Even parity is normally generated on output. The EOT character is not transmitted in cooked mode to prevent terminals that respond to it from hanging up; programs using RAW or CBREAK mode should be careful.

The terminal line disciplines provide necessary processing for cooked and CBREAK mode output including delay generation for certain special characters and parity generation. Delays are available after backspaces ^H, form feeds ^L, carriage returns ^M, tabs ^I and newlines ^J. The line disciplines will also optionally expand tabs into spaces, where the tab stops are assumed to be set every eight columns, and optionally convert newlines to carriage returns followed by newline. These functions are controlled by bits in the tty flags word; see **Summary** below.

The terminal line disciplines provide for mapping between upper and lower case on terminals lacking lower case, and for other special processing on deficient terminals.

Finally, in the new terminal line discipline, there is an output flush character, normally ^O, which sets the LFLUSHO bit in the local mode word, causing subsequent output to be flushed until it is cleared by a program or more input is typed. This character has effect in both cooked and CBREAK modes and causes pending input to be retyped if there is any pending input. An *ioctl* to flush the characters in the input or output queues, TIOCFLUSH, is also available.

Upper case terminals and Hazeltines

If the LCASE bit is set in the tty flags, then all upper-case letters are mapped into the corresponding lower-case letter. The upper-case letter may be generated by preceding it by '^'. Upper case letters are preceded by a '^' when output. In addition, the following escape sequences can be generated on output and accepted on input:

for	^		~	{	}
use	\^	\	\~	\{	\}

To deal with Hazeltine terminals, which do not understand that ~ has been made into an ASCII character, the LTILDE bit may be set in the local mode word; in this case the character ~ will be replaced with the character ` on output.

Flow control.

There are two characters (the stop character, normally ^S, and the start character, normally ^Q) which cause output to be suspended and resumed respectively. Extra stop characters typed when output is already stopped have no effect, unless the start and stop characters are made the same, in which case output resumes.

A bit in the flags word may be set to put the terminal into TANDEM mode. In this mode the system produces a stop character (default ^S) when the input queue is in danger of overflowing, and a start character (default ^Q) when the input has drained sufficiently. This mode is useful when the terminal is actually another machine that obeys the conventions.

Line control and breaks.

There are several *ioctl* calls available to control the state of the terminal line. The TIOCSBRK *ioctl* will set the break bit in the hardware interface causing a break condition to exist; this can be cleared (usually after a delay with *sleep(3)*) by TIOCCBRK. Break conditions in the input are reflected as a null character in RAW mode or as the interrupt character in cooked or CBREAK mode. The TIOCCDTR *ioctl* will clear the data terminal ready condition; it can be set again by TIOCSDTR.

When the carrier signal from the dataset drops (usually because the user has hung up his terminal) a SIGHUP hangup signal is sent to the processes in the distinguished process group of the terminal; this usually causes them to terminate (the SIGHUP can be suppressed by setting the LNOHANG bit in the local state word of the driver.) Access to the terminal by other processes is then normally revoked, so any further reads will fail, and programs that read a terminal and test for end-of-file on their input will terminate appropriately.

When using an ACU it is possible to ask that the phone line be hung up on the last close with the `TIOCHPCL` *ioctl*; this is normally done on the outgoing line.

Interrupt characters.

There are several characters that generate interrupts in cooked and CBREAK mode; all are sent to the processes in the control group of the terminal, as if a `TIOCGPGRP` *ioctl* were done to get the process group and then a `killpg(2)` system call were done, except that these characters also flush pending input and output when typed at a terminal (*a la* `TIOCFUSH`). The characters shown here are the defaults; the field names in the structures (given below) are also shown. The characters may be changed.

- `^C` `t_intrc` (ETX) generates a SIGINT signal. This is the normal way to stop a process which is no longer interesting, or to regain control in an interactive program.
- `^\
^Z` `t_quitc` (FS) generates a SIGQUIT signal. This is used to cause a program to terminate and produce a core image, if possible, in the file `core` in the current directory.
- `^Z` `t_suspc` (EM) generates a SIGTSTP signal, which is used to suspend the current process group.
- `^Y` `t_dsuspc` (SUB) generates a SIGTSTP signal as `^Z` does, but the signal is sent when a program attempts to read the `^Y`, rather than when it is typed.

Job access control.

When using the new terminal line discipline, if a process which is not in the distinguished process group of its control terminal attempts to read from that terminal its process group is sent a SIGTTIN signal. This signal normally causes the members of that process group to stop. If, however, the process is ignoring SIGTTIN, has SIGTTIN blocked, or is in the middle of process creation using `vfork(2)`, the read will return `-1` and set `errno` to `EIO`.

When using the new terminal line discipline with the `LTOSTOP` bit set in the local modes, a process is prohibited from writing on its control terminal if it is not in the distinguished process group for that terminal. Processes which are holding or ignoring SIGTTOU signals or which are in the middle of a `vfork(2)` are excepted and allowed to produce output.

Summary of modes.

Unfortunately, due to the evolution of the terminal drivers and line disciplines, there are 4 different structures which contain various portions of the driver and line discipline data. The first of these (`sgttyb`) contains that part of the information largely common between Version 6 and Version 7 UNIX systems. The second contains additional control characters added in Version 7. The third is a word of local state added in 4BSD, and the fourth is another structure of special characters added for the new line discipline. In the future a single structure may be made available to programs which need to access all this information; most programs need not concern themselves with all this state.

Basic modes: `sgtty`.

The basic *ioctls* use the structure defined in `<sgtty.h>`:

```
struct 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 according to the following table, which corresponds to the DEC DH-11 interface. If other hardware is used, impossible speed changes are ignored. Symbolic values in the table are as defined in `<sys/ttydev.h>`.

B0	0	(hang up dataphone)
B50	1	50 baud
B75	2	75 baud
B110	3	110 baud
B134	4	134.5 baud
B150	5	150 baud
B200	6	200 baud
B300	7	300 baud
B600	8	600 baud
B1200	9	1200 baud
B1800	10	1800 baud
B2400	11	2400 baud
B4800	12	4800 baud
B9600	13	9600 baud
EXTA	14	19200 baud
EXTB	15	External B

Code conversion and line control required for IBM 2741's (134.5 baud) must be implemented by the user's program. The half-duplex line discipline required for the 202 dataset (1200 baud) is not supplied; full-duplex 212 datasets work fine.

The *sg_erase* and *sg_kill* fields of the argument structure specify the erase and kill characters respectively. (Defaults are DELETE and ^U.)

The *sg_flags* field of the argument structure contains several bits that determine the system's treatment of the terminal:

ALLDELAY	0177400	Delay algorithm selection
BSDELAY	0100000	Select backspace delays (not implemented):
BS0	0	
BS1	0100000	
VTDELAY	0040000	Select form-feed and vertical-tab delays:
FF0	0	
FF1	0040000	
CRDELAY	0030000	Select carriage-return delays:
CR0	0	
CR1	0010000	
CR2	0020000	
CR3	0030000	
TBDELAY	0006000	Select tab delays:
TAB0	0	
TAB1	0002000	
TAB2	0004000	
XTABS	0006000	
NLDELAY	0001400	Select new-line delays:
NL0	0	
NL1	0000400	
NL2	0001000	
NL3	0001400	
EVENP	0000200	Even parity allowed on input and generated on output
ODDP	0000100	Odd parity allowed on input and generated on output
RAW	0000040	Raw mode: wake up on all characters, 8-bit interface
CRMOD	0000020	Map CR into LF; output LF as CR-LF
ECHO	0000010	Echo (full duplex)
LCASE	0000004	Map upper case to lower on input and lower to upper on output
CBREAK	0000002	Return each character as soon as typed

TANDEM 0000001 Automatic flow control

The delay bits specify how long transmission stops to allow for mechanical or other movement when certain characters are sent to the terminal. In all cases a value of 0 indicates no delay.

Backspace delays are currently ignored but might be used for Terminet 300's.

If a form-feed/vertical tab delay is specified, it lasts for about 2 seconds.

Carriage-return delay type 1 lasts about .08 seconds and is suitable for the Terminet 300. Delay type 2 lasts about .16 seconds and is suitable for the VT05 and the TI 700. Delay type 3 is suitable for the concept-100 and pads lines to be at least 9 characters at 9600 baud.

New-line delay type 1 is dependent on the current column and is tuned for Teletype model 37's. Type 2 is useful for the VT05 and is about .10 seconds. Type 3 is unimplemented and is 0.

Tab delay type 1 is dependent on the amount of movement and is tuned to the Teletype model 37. Type 3, called XTABS, is not a delay at all but causes tabs to be replaced by the appropriate number of spaces on output.

Input characters with the wrong parity, as determined by bits 200 and 100, are ignored in cooked and CBREAK mode.

RAW disables all processing save output flushing with LFLUSHO; full 8 bits of input are given as soon as it is available; all 8 bits are passed on output. A break condition in the input is reported as a null character. If the input queue overflows in raw mode all data in the input and output queues are discarded; this applies to both the new and the old line disciplines.

CRMOD causes input carriage returns to be turned into new-lines, and output and echoed new-lines to be output as a carriage return followed by a line feed.

CBREAK is a sort of half-cooked (rare?) mode. Programs can read each character as soon as typed, instead of waiting for a full line; all processing is done except the input editing: character and word erase and line kill, input reprint, and the special treatment of \ and EOT are disabled.

TANDEM mode causes the system to produce a "stop" character (default ^S) whenever the input queue is in danger of overflowing, and a "start" character (default ^Q) when the input queue has drained sufficiently. It is useful for flow control when the 'terminal' is really another computer which understands the conventions.

Note: The same "stop" and "start" characters are used for both directions of flow control; the *t_stop* character is accepted on input as the character that stops output and is produced on output as the character to stop input, and the *t_start* character is accepted on input as the character that restarts output and is produced on output as the character to restart input.

Basic ioctls

A large number of *ioctl(2)* calls apply to terminals. Some have the general form:

```
#include <sgtty.h>
```

```
ioctl(fildes, code, arg)
```

```
struct sgttyb *arg;
```

The applicable codes are:

TIOCGETP Fetch the basic parameters associated with the terminal, and store in the pointed-to *sgttyb* structure.

TIOCSETP Set the parameters according to the pointed-to *sgttyb* structure. The interface delays until output is quiescent, then throws away any unread characters, before changing the modes.

TIOCSETN Set the parameters like TIOCSETP but do not delay or flush input. Input is not preserved, however, when changing to or from RAW.

With the following codes *arg* is ignored.

TIOCEXCL	Set “exclusive-use” mode: no further opens are permitted until the file has been closed.
TIOCNXCL	Turn off “exclusive-use” mode.
TIOCHPCL	When the file is closed for the last time, hang up the terminal. This is useful when the line is associated with an ACU used to place outgoing calls.

With the following codes *arg* is a pointer to an int.

TIOCGETD	<i>arg</i> is a pointer to an int into which is placed the current line discipline number.
TIOCSETD	<i>arg</i> is a pointer to an int whose value becomes the current line discipline number.
TIOCFLUSH	If the int pointed to by <i>arg</i> has a zero value, 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 defined in <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.

The remaining calls are not available in vanilla Version 7 UNIX systems. In cases where arguments are required, they are described; *arg* should otherwise be given as 0.

TIOCSTI	the argument points to a character which the system pretends had been typed on the terminal.
TIOCSBRK	the break bit is set in the terminal.
TIOCCBRK	the break bit is cleared.
TIOCSDTR	data terminal ready is set.
TIOCCDTR	data terminal ready is cleared.
TIOCSTOP	output is stopped as if the “stop” character had been typed.
TIOCSTART	output is restarted as if the “start” character had been typed.
TIOCGPGRP	<i>arg</i> is a pointer to an int into which is placed the process group ID of the process group for which this terminal is the control terminal.
TIOCSPGRP	<i>arg</i> is a pointer to an int (typically a process ID); the process group whose process group ID is the value of this int becomes the process group for which this terminal is the control terminal.
TIOCOUTQ	returns in the int pointed to by <i>arg</i> the number of characters queued up to be output to the terminal.
FIONREAD	returns in the int pointed to by <i>arg</i> the number of immediately readable characters from the argument unit. This works for files, pipes, and terminals.

Tchars

The second structure associated with each terminal specifies characters that are special in both the old and new terminal interfaces: The following structure is defined in <sys/ioctl.h>, which is automatically included by <sgtty.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 default values for these characters are ^C, ^\, ^Q, ^S, ^D, and ^\377. A character value of ^\377 eliminates the effect of that character. The *t_brkc* character, by default ^\377, acts like a new-line in that it terminates a 'line,' is echoed, and is passed to the program. The 'stop' and 'start' characters may be the same, to produce a toggle effect. It is probably counterproductive to make other special characters (including erase and kill) identical. The applicable *ioctl* calls are:

TIOCGETC Get the special characters and put them in the specified structure.

TIOCSETC Set the special characters to those given in the structure.

Local mode

The third structure associated with each terminal is a local mode word. The bits of the local mode word are:

LCRTBS	000001	Backspace on erase rather than echoing erase
LPRTERA	000002	Printing terminal erase mode
LCRTERA	000004	Erase character echoes as backspace-space-backspace
LTLDE	000010	Convert ~ to ` on output (for Hazeltine terminals)
LLITOUT	000040	Suppress output translations
LTOSTOP	000100	Send SIGTTOU for background output
LFLUSHO	000200	Output is being flushed
LNOHANG	000400	Don't send hangup when carrier drops
	001000	Unimplemented.
LCRTKIL	002000	BS-space-BS erase entire line on line kill
LPASS8	004000	Pass all 8 bits through on input, in any mode
LCTLECH	010000	Echo input control chars as ^X, delete as ^?
LPENDIN	020000	Retype pending input at next read or input character
LDECCTQ	040000	Only ^Q restarts output after ^S, like DEC systems
LNOFLSH	100000	Inhibit flushing of pending I/O when an interrupt character is typed.

The applicable *ioctl* functions are:

TIOCLBIS *arg* is a pointer to an *int* whose value is a mask containing the bits to be set in the local mode word.

TIOCLBIC *arg* is a pointer to an *int* whose value is a mask containing the bits to be cleared in the local mode word.

TIOCLSET *arg* is a pointer to an *int* whose value is stored in the local mode word.

TIOCLGET *arg* is a pointer to an *int* into which the current local mode word is placed.

Local special chars

The final structure associated with each terminal is the *ltchars* structure which defines control characters for the new line discipline. 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_lnextc;       /* literal next character */
};
```

The default values for these characters are ^Z, ^Y, ^R, ^O, ^W, and ^V. A value of ^\377 disables the character.

The applicable *ioctl* functions are:

TIOCSLTC *arg* is a pointer to an *lchars* structure which defines the new local special characters.

TIOCGLTC *arg* is a pointer to an *lchars* structure into which is placed the current set of local special characters.

FILES

/dev/tty
/dev/tty*
/dev/console

SEE ALSO

csh(1), stty(1), ioctl(2), sigvec(2), stty(3C), getty(8), init(8)

BUGS

Half-duplex terminals are not supported.

Processes that are not invoked with a control terminal, but open a *dialout* line can hang indefinitely. Once the *dialout* line is opened, it becomes the control terminal. Should the process then open */dev/tty*, it will hang because */dev/tty* resolves to the corresponding *dialin* line. The process will wait for the *dialin* sequence to complete, even though the line is already connected.

NAME

udp – Internet User Datagram Protocol

SYNOPSIS

None; comes automatically with *inet*(4F).

DESCRIPTION

The User Datagram Protocol (UDP) is defined to make available a datagram mode of packet switched computer communication in the environment of an interconnected set of computer networks. The protocol assumes that the Internet Protocol (IP) as described in *ip*(4P) is used as the underlying protocol.

The protocol provides a procedure for application programs to send messages to other programs with a minimum of protocol mechanism. The protocol is transaction oriented, and delivery and duplicate protection are not guaranteed. Applications requiring ordered reliable delivery of streams of data should use the Transmission Control Protocol (TCP) as described in *tcp*(4P).

The UNIX system implementation of UDP makes it available as a socket of type SOCK_DGRAM. UDP sockets are normally used in a connectionless fashion, with the *sendto* and *recvfrom* calls described in *send*(2) and *recv*(2).

A UDP socket is created with a *socket*(2) call:

```
s = socket(AF_INET, SOCK_DGRAM, 0);
```

The socket initially has no address associated with it, and may be given an address with a *bind*(2) call as described in *inet*(4F). If no *bind* call is done, then the address assignment procedure described in *inet*(4F) is repeated as each datagram is sent.

When datagrams are sent the system encapsulates the user supplied data with UDP and IP headers. Unless the invoker is the super-user datagrams which would become broadcast packets on the network to which they are addressed are not allowed. Unless the socket has had a SO_DONTROUTE option enabled (see *socket*(2)) the outgoing datagram is routed through the routing tables as described in *routing*(4N). If there is insufficient system buffer space to temporarily hold the datagram while it is being transmitted, the *sendto* may result in a ENOBUFS error. Other errors (ENETUNREACH, EADDRNOTAVAIL, EACCES, EMSGSIZE) may be generated by *icmp*(4P) or by the network interfaces themselves, and are reflected back in the *send* call.

As each UDP datagram arrives at a host the system strips out the IP options and checksums the data field, discarding the datagram if the checksum indicates that the datagram has been damaged. If no socket exists for the datagram to be sent to then an ICMP error is returned to the originating socket. If a socket exists for this datagram to be sent to, then we will append the datagram and the address from which it came to a queue associated with the datagram socket. This queue has limited capacity (2048 bytes of datagrams) and arriving datagrams which will not fit within its *high-water* capacity are silently discarded.

UDP processes ICMP errors reflected to it by *icmp*(4P). QUENCH errors are ignored (this is well considered a bug); UNREACH, TIMXCEED and PARAMPROB errors cause the socket to be disconnected from its peer if it was bound to a peer using *bind*(2) so that subsequent attempts to send datagrams via that socket will give an error indication.

The UDP datagram protocol differs from IP datagrams in that it adds a checksum over the data bytes and contains a 16-bit socket address on each machine rather than just the 32-bit machine address; UDP datagrams are addressed to sockets; IP packets are addressed to hosts.

SEE ALSO

recv(2), *send*(2), *inet*(4F)

"User Datagram Protocol," RFC768, John Postel, USC-ISI (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 QUENCH errors if the socket is bound to a peer socket.

NAME

vp – Ikon 10071-5 Versatec parallel printer interface

SYNOPSIS — SUN-2

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

BUGS

If you use the standard i/o library on the Versatec, be sure to explicitly set a buffer using *seibuf*, 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

SYNOPSIS — SUN-2

`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/vp0` 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/vp0`

`/dev/lp0`

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

SYNOPSIS

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 *Programmer's Reference Manual for SunWindows*.

FILES

/dev/win[0-9]

/dev/win[0-9][0-9]

SEE ALSO

Programmer's Reference Manual for SunWindows

NAME

xt - Xylogics 472 1/2 inch tape controller

SYNOPSIS — SUN-3

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

SYNOPSIS — SUN-2

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 (tm(4)).

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), tm(4), mtio(4)

NAME

xy – Disk driver for Xylogics SMD Disk Controllers

SYNOPSIS — SUN-3

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 SMD disk controller in a Sun system.

SYNOPSIS — SUN-2

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 SMD disk controllers in a Sun-2/160 VMEbus based system. The third and fourth **controller** lines specify the first and second Xylogics 450 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 via 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 *seek(2)* calls should specify a multiple of 512 bytes.

If **flags 0x1** is specified, the overlapped seeks feature for that drive is turned off. Note that 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

dkio(4S)

Xylogics Model 450 Peripheral Processor SMD Disk Subsystem Maintenance and Reference Manual (Sun 800-1025-01)

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. See the subsection on the Xylogics controller in the appropriate Sun *Hardware Installation Manual* for your machine(s) for instructions on how to set the jumpers on the 450.

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 does 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*(2) and *write*(2) truncate file offsets to 512-byte block boundaries, and *write*(2) scribbles on the tail of incomplete blocks. Thus, in programs that are likely to access raw devices, *read*(2), *write*(2) 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

zs – zilog 8530 SCC serial communications driver

SYNOPSIS — SUN-3

device zs0 at obio ? csr 0x20000 flags 3 priority 3

device zs1 at obio ? csr 0x00000 flags 0x103 priority 3

SYNOPSIS — SUN-2

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

DESCRIPTION

The Zilog 8530 provides 2 serial communication ports with full modem control in asynchronous mode. Each port behaves as described in *tty(4)*. Input and output for each line may independently be set to run at any of 16 speeds; see *tty(4)* for the encoding.

Of the synopsis lines above, the line for zs0 specifies the serial I/O ports 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 cause line *ttyb* to be treated in this way.

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 0 – 127 correspond directly to the normal *tty* lines and are named *tty**. Minor device numbers in the range 128 – 256 correspond to the same physical lines as those above (i.e. the same line as the minor device number minus 128) and are (conventionally) named *cua**. The *cua* lines are special in that they can be opened even when there is no carrier on the line. Once a *cua* line is opened, the corresponding *tty* line can not be opened until the *cua* line is closed. Also, if the *tty* line has been opened successfully (usually only when carrier is recognized on the modem) the corresponding *cua* line can not be opened. This allows a modem to be attached to */dev/ttya* (usually renamed to */dev/ttyd0*) and used for dialin (by enabling the line for login in */etc/ttys*) and also used for dialout (by *tip(1C)* or *uucp(1C)*) as */dev/cua0* when no one is logged in on the line. Note that the bit in the flags word in the config file (see above) must be zero for this line.

FILES

/dev/tty[a, b, s0-s3]

/dev/ttyd[0-9, a-f]

/dev/cua[0-9, a-f]

SEE ALSO

tty(4)

DIAGNOSTICS

zsn c : silo overflow.

The character input silo overflowed before it could be serviced.



NAME

file formats formats of files used by various programs

DESCRIPTION

This section describes formats of files used by various programs.

a.out	a.out(5)	assembler and link editor output format
acct	acct(5)	execution accounting file
addresses	aliases(5)	addresses and mailing lists for sendmail(8)
aliases	aliases(5)	addresses and mailing lists for sendmail(8)
ar	ar(5)	archive (library) file format
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
directory	dir(5)	format of directories
dump	dump(5)	incremental dump format
dumpdates	dump(5)	incremental dump format
environment	environ(5V)	user (process) environment
ethers	ethers(5)	Ethernet address to hostname database or YP domain
exports	exports(5)	NFS file systems being exported
fcntl	fcntl(5)	file control options
fs	fs(5)	format of file system volume
fspec	fspec(5)	format specification in text files
fstab	mntent(5)	static information about filesystems
ftpusers	ftpusers(5)	list of users prohibited by ftp
gettytab	gettytab(5)	simplified terminal configuration data base
group	group(5)	local host's group file
hosts	hosts(5)	host name data base or YP domain
hosts.equiv	hosts.equiv(5)	list of trusted hosts
inode	fs(5)	format of file system index node
lastlog	usracct(5)	login records
magic	magic(5)	file command's magic number file
mntent	mntent(5)	static information about filesystems
mtab	mtab(5)	currently mounted file system table
netgroup	netgroup(5)	list of network groups
networks	networks(5)	network name data base or YP domain
passwd	passwd(5)	password file or YP domain
phones	phones(5)	remote host phone number data base
plot	plot(5)	graphics interface
printcap	printcap(5)	printer capability data base
protocols	protocols(5)	protocol name data base or YP domain
rasterfile	rasterfile(5)	Sun's file format for raster images
remote	remote(5)	remote host description file
resolver	resolver(5)	configuration file for name server routines
rmtab	rmtab(5)	local file system mount statistics
rpc	rpc(5)	rpc program number data base
sccsfile	sccsfile(5)	format of sccs(1) history file
servers	servers(5)	Internet server data base
services	services(5)	service name data base or YP domain
statmon	statmon(5)	statd directory and file structures
tar	tar(5)	tape archive file format
term	term(5)	terminal driving tables for nroff

termcap	termcap(5)	terminal capability data base
tp	tp(5)	DEC/mag tape formats
ttys	ttys(5)	terminal initialization data
ttytype	ttytype(5)	data base of terminal types by port
types	types(5)	primitive system data types
utmp	usracct(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	usracct(5)	login records
ypfiles	ypfiles(5)	the yellowpages database and directory structure

NAME

a.out – assembler and link editor output format

SYNOPSIS

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

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 if there were no errors and no unresolved external references. Layout information as given in the include file for the Sun system is:

```
/*
 * Header prepended to each a.out file.
 */
struct exec {
    unsigned short a_machtype; /* machine type */
    unsigned short a_magic; /* magic number */
    unsigned a_text; /* size of text segment */
    unsigned a_data; /* size of initialized data */
    unsigned a_bss; /* size of uninitialized data */
    unsigned a_syms; /* size of symbol table */
    unsigned a_entry; /* entry point */
    unsigned a_trsize; /* size of text relocation */
    unsigned a_drsize; /* size of data relocation */
};

#define M_68010 1 /* runs on either MC68010 or MC68020 */
#define M_68020 2 /* runs only on MC68020 */

#define OMAGIC 0407 /* magic number for old impure format */
#define NMAGIC 0410 /* magic number for read-only text */
#define ZMAGIC 0413 /* magic number for demand load format */

#define PAGESIZ 0x2000 /* page size - same for Sun-2 and Sun-3 */
#define SEGSIZ 0x20000 /* segment size - same for Sun-2 and Sun-3 */

/*
 * The following macros take exec structures as arguments. N_BADMAG(x) returns
 * 0 if the file has a reasonable magic number.
 */
#define N_BADMAG(x) \
    (((x).a_magic)!=OMAGIC && ((x).a_magic)!=NMAGIC && ((x).a_magic)!=ZMAGIC)

/*
 * Offsets to text | symbols | strings.
 */
#define N_TXTOFF(x) \
    ((x).a_magic==ZMAGIC ? 0 : sizeof (struct exec))
#define N_SYMOFF(x) \
    (N_TXTOFF(x) + (x).a_text+(x).a_data + (x).a_trsize+(x).a_drsize)
#define N_STROFF(x) \
    (N_SYMOFF(x) + (x).a_syms)

/*
 * Macros which take exec structures as arguments and tell where the
 * various pieces will be loaded.
 */
```

```

*/
#define      N_TXTADDR(x) PAGESIZ
#define      N_DATADDR(x) \
            (((x).a_magic==OMAGIC)? (N_TXTADDR(x)+(x).a_text) \
            : (SEGSIZ+((N_TXTADDR(x)+(x).a_text-1) & ~(SEGSIZ-1))))
#define N_BSSADDR(x) (N_DATADDR(x)+(x).a_data)

```

The *a.out* file has five sections: a header, the program 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 may be executed. Sun-2 code may be executed on Sun-3 systems, but not vice versa. Program files predating release 3.0 are recognized by a machine type of 0.

If the magic number in the header is OMAGIC (0407), it means that this is a non-sharable text which is not to be write-protected, so the data segment is immediately contiguous with the text segment. This is rarely used. If the magic number is NMAGIC (0410) or ZMAGIC (0413), the data segment begins at the first segment boundary following the text segment, and the text segment is not writable by the program; other processes executing the same file will share the text segment. For ZMAGIC format, the text and data sizes must both be 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 suitable for large programs and is the default format produced by *ld(1)*. 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. In all file types other than XMAGIC, this is the same as N_TXTADDR(x). In ZMAGIC files it is N_TXTADDR + sizeof(struct exec).

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(2)*.

After the header in the file follow the text, data, text relocation data relocation, symbol table and string table in that order. The text begins at the beginning 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 the name of an exec structure as argument. The data segment is contiguous with the text and immediately followed by the text relocation and then the data relocation information. The symbol table follows all this; its position is computed by the N_SYMOFF macro. Finally, the string table immediately follows the symbol table at a position which can be gotten easily using N_STROFF. The first 4 bytes of the string table are not used for string storage, but rather contain the size of the string table; this size *includes* the 4 bytes; thus, the minimum string table size is 4.

RELOCATION

The value of a byte in the text or data which is not a portion of a reference to an undefined external symbol is exactly that value which will appear in memory when the file is executed. If a byte in the text or data involves a reference to an undefined external symbol, as indicated by the relocation information, then the value stored in the file is an offset from the associated external symbol. When the file is processed by the link editor and the external symbol becomes defined, the value of the symbol is added to the bytes in the file.

If relocation information is present, it amounts to eight bytes per relocatable datum as in the following structure:

```

/*
 * Format of a relocation datum.
 */
struct relocation_info {
    int            r_address;        /* address which is relocated */
    unsigned      r_symbolnum:24, /* local symbol ordinal */
    r_pcrel:1,     /* was relocated pc relative already */
    r_length:2,    /* 0=byte, 1=word, 2=long */
    r_extern:1,    /* does not include value of sym referenced */
    :4;           /* nothing, yet */
};

```

There is no relocation information if `a_trsize+a_drsize==0`. If `r_extern` is 0, then `r_symbolnum` is actually a `n_type` for the relocation (that is, `N_TEXT` meaning relative to segment text origin.)

SYMBOL TABLE

The layout of a symbol table entry and the principal flag values that distinguish symbol types are given in the include file as follows:

```

/*
 * Format of a symbol table entry.
 */
struct nlist {
    union {
        char            *n_name;        /* for use when in-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 adb 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.

SEE ALSO

`adb(1)`, `as(1)`, `cc(1V)`, `dbx(1)`, `ld(1)`, `nm(1)`, `pc(1)`, `strip(1)`

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 1.1 86/07/07 SMI; from UCB 6.1 83/07/29*/

/*
 * Accounting structures;
 * these use a comp_t type which is a 3 bits base 8
 * exponent, 13 bit fraction "floating point" number.
 */
typedef u_short comp_t;

struct acct
{
    char        ac_comm[10]; /* Accounting command name */
    comp_t      ac_ftime;    /* Accounting user time */
    comp_t      ac_sftime;   /* Accounting system time */
    comp_t      ac_etime;    /* Accounting elapsed time */
    time_t      ac_btime;    /* Beginning time */
    short       ac_uid;      /* Accounting user ID */
    short       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 */

#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
/usr/lib/aliases
/usr/lib/aliases.dir
/usr/lib/aliases.pag
`~/forward`

DESCRIPTION

These files contain mail addresses or aliases recognized by *sendmail*(8):

/etc/passwd Mail addresses (usernames) of local users.

/usr/lib/aliases Local aliases in ASCII format. This file can be edited to add, update, or delete mail aliases for the local host.

/usr/lib/aliases.{dir,pag}

The aliasing information from */usr/lib/aliases*, in binary, *dbm(3X)* format for use by *sendmail*(8). The program *newaliases*(8), which is invoked automatically by *sendmail*(8), maintains these files.

`~/forward` 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 *Sendmail Installation and Operation* in *System Administration for the Sun Workstation* for details. Standard addresses are recommended.

ALIASES

Local Aliases

/usr/lib/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. 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 */usr/lib/aliases** files on the various hosts in a network will one day be obsolete. Domain-wide aliases should ultimately 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 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 a the name of a user on the local host, *sendmail* checks for a *forward* file in that user's home directory. This file can contain one or more addresses or aliases as described above; each recipient is sent a copy of the mail destined for the original user.

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.

SEE ALSO

newaliases(8), dbm(3X), sendmail(8), uucp(1C), vacation(1)

System Administration for the Sun Workstation

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*.

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:

```
/*      @(#)ar.h 1.1 86/07/07 SMI; from UCB 4.1 83/05/03*/
```

```
#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 a even (0 mod 2) boundary; a new-line is inserted between files if necessary. Nevertheless the size given reflects the actual size of the file exclusive of padding.

There is no provision for empty areas in an archive file.

The encoding of the header is portable across machines. If an archive contains printable files, the archive itself is printable.

SEE ALSO

ar(1), *ld*(1), *nm*(1)

BUGS

File names lose trailing blanks. Most software dealing with archives takes even an included blank as a name terminator.

NAME

core – format of memory image file

SYNOPSIS

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

DESCRIPTION

The UNIX System writes out a memory image of a terminated process when any of various errors occur. See *sigvec(2)* for the list of reasons; the most common are memory violations, illegal instructions, bus errors, and user-generated quit signals. The memory image is called 'core' and is written in the process's working directory (provided it can be; normal access controls apply).

The maximum size of a *core* file is limited by *setrlimit(2)*. Files which would be larger than the limit are not created.

Set-user-id programs do not produce core files when they terminate as this would be a security loophole.

The core file consists of a *core* structure defined in the *<sys/core.h>* file. The *core* structure includes the registers, the floating point status, 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.

The remainder of the core file consists first of the data pages and then the stack pages of the process image. The amount of data space image in the core file is given (in bytes) by the *c_dsize* member of the *core* structure. The amount of stack image in the core file is given (in bytes) by the *c_ssize* member of the *core* structure.

SEE ALSO

adb(1), *dbx(1)*, *sigvec(2)*, *setrlimit(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%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)*

NAME

crontab – table of times to run periodic jobs

SYNOPSIS

/usr/lib/crontab

DESCRIPTION

The *letc/cron* utility is a permanent process, started by *letc/rc.local*, that wakes up once every minute. *letc/cron* consults the file *usr/lib/crontab* to find out what tasks are to be done, and at what time.

Each line in *usr/lib/crontab* 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 1 through 7. Monday is day 1 in this scheme of things.
6. (the remainder of the line) is the command to be run. A percent character in this field is translated to a new-line character. Only the first line (up to a % or end of line) of the command field is executed by the Shell. The other lines are made available to the command as standard input.

Any of fields 1 through 5 can be a list of values separated by commas. A field can be 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.

FILES

/usr/lib/crontab

SEE ALSO

cron(8), rc(8)

EXAMPLE

```
0 0 * * * calendar -
15 0 * * * /etc/sa -s >/dev/null
15 4 * * * find /usr/preserve -mtime +7 -a -exec rm -f {} ;
40 4 * * * find / -name '#*' -atime +3 -exec rm -f {} ;
0,15,30,45 * * * * /etc/atrun
0,10,20,30,40,50 * * * * /etc/dmesg - >>/usr/adm/messages
5 4 * * * sh /etc/newsyslog
```

The *calendar* command runs at minute 0 of hour 0 (midnight) of every day. The *letc/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 *atrun* command (which processes shell scripts users have set up with *at*) runs every 15 minutes. The *letc/dmesg* command appends kernel messages to the *usr/adm/messages* file every ten minutes, and finally, the *usr/adm/syslog* script runs at five minutes after four every day.

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 i-node 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 */
};

struct _dirdesc {
    int     dd_fd;
    long    dd_loc;
    long    dd_size;
    char    dd_buf[DIRBLKSIZ];
};
```

By convention, the first two entries in each directory are for '.' and '..'. The first is an entry for the directory itself. The second is for the parent directory. The meaning of '..' is modified for the root directory of the master file system ('/'), for which '..' has the same meaning as '.'.

SEE ALSO

fs(5), readdir(3)

NAME

dump, dumpdates – incremental dump format

SYNOPSIS

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

DESCRIPTION

Tapes used by *dump* and *restore(8)* contain:

- a header record
- two groups of bit map records
- a group of records describing directories
- a group of records describing files

The format of the header record and of the first record of each description as given in the include file *<dumprest.h>* is:

```
#define NTREC          10
#define MLEN          16
#define MSIZ          4096

#define TS_TAPE        1
#define TS_INODE       2
#define TS_BITS        3
#define TS_ADDR        4
#define TS_END         5
#define TS_CLRI        6
#define MAGIC          (int) 60011
#define CHECKSUM       (int) 84446

struct  spcl {
    int          c_type;
    time_t      c_date;
    time_t      c_ddate;
    int          c_volume;
    daddr_t     c_tapea;
    ino_t        c_inumber;
    int          c_magic;
    int          c_checksum;
    struct      dinode      c_dinode;
    int          c_count;
    char         c_addr[BSIZE];
} spcl;

struct  idates {
    char        id_name[16];
    char        id_incno;
    time_t      id_ddate;
};

#define DUMPOUTFMT "%-16s %c %s"      /* for printf */
#define DUMPINFMT  "%16s %c %[\n]\n"  /* name, incno, ctime(date) */
/* 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:

<i>c_type</i>	The type of the header.
<i>c_date</i>	The date the dump was taken.
<i>c_ddate</i>	The date the file system was dumped from.
<i>c_volume</i>	The current volume number of the dump.
<i>c_tapea</i>	The current number of this (1024-byte) record.
<i>c_inumber</i>	The number of the inode being dumped if this is of type TS_INODE.
<i>c_magic</i>	This contains the value MAGIC above, truncated as needed.
<i>c_checksum</i>	This contains whatever value is needed to make the record sum to CHECKSUM.
<i>c_dinode</i>	This is a copy of the inode as it appears on the file system; see <i>fs(5)</i> .
<i>c_count</i>	The count of characters in <i>c_addr</i> .
<i>c_addr</i>	An array of characters describing the blocks of the dumped file. A character is zero if the block associated with that character was not present on the file system, otherwise the character is non-zero. If the block was not present on the file system, no block was dumped; the block will be restored as a hole in the file. If there is not sufficient space in this record to describe all of the blocks in a file, TS_ADDR records will be scattered through the file, each one picking up where the last left off.

Each volume except the last ends with a tapemark (read as an end of file). The last volume ends with a TS_END record and then the tapemark.

The structure *idates* describes an entry in the file */etc/dumpdates* where dump history is kept. The fields of the structure are:

<i>id_name</i>	The dumped filesystem is <i>'/dev/id_nam'</i> .
<i>id_incno</i>	The level number of the dump tape; see <i>dump(8)</i> .
<i>id_ddate</i>	The date of the incremental dump in system format see <i>types(5)</i> .

FILES

/etc/dumpdates

SEE ALSO

dump(8), *restore(8)*, *fs(5)*, *types(5)*

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 <i>sh</i> , <i>time</i> , <i>nice(1)</i> , etc., apply in searching for a file known by an incomplete path name. The prefixes are separated by colons (:). The <i>login(1)</i> process sets PATH=:/usr/ucb:/bin:/usr/bin .
HOME	The name of the user’s login directory, set by <i>login(1)</i> from the password file <i>/etc/passwd</i> (see <i>passwd(5)</i>).
TERM	The kind of terminal for which output is to be prepared. This information is used by commands, such as <i>nroff</i> or <i>plot(1G)</i> , which may exploit special terminal capabilities. See <i>/etc/termcap</i> (<i>termcap(5)</i>) for a list of terminal types.
SHELL	The file name of the user’s login shell.
TERMCAP	The string describing the terminal in TERM , or the name of the termcap file, see <i>termcap(3)</i> , <i>termcap(5)</i> ,
EXINIT	A startup list of commands read by <i>ex(1)</i> , <i>edit(1)</i> , and <i>vi(1)</i> .
USER	
LOGNAME	The login name of the user.

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 *cs(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 the System V environment.

TZ	Time zone information. The format is <i>xxx n zzz</i> where <i>xxx</i> is standard local time zone abbreviation, <i>n</i> is the difference in hours from GMT, and <i>zzz</i> is the abbreviation for the daylight-saving local time zone, if any; for example, EST 5 EDT .
-----------	--

SEE ALSO

csh(1), *ex(1)*, *login(1)*, *sh(1)*, *getenv(3)*, *execve(2)*, *system(3)*, *termcap(3X)*, *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 tabs. 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" 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 – NFS file systems being exported

SYNOPSIS

/etc/exports

DESCRIPTION

The file */etc/exports* describes the file systems which are being exported to nfs clients. It is created by the system administrator using a text editor and processed by the *mount* request daemon *mountd*(8C) each time a mount request is received.

The file consists of a list of file systems and the *netgroup*(5) or machine names allowed to remote mount each file system. The file system names are left justified and followed by a list of names separated by white space. The names will be looked up in */etc/netgroup* and then in */etc/hosts*. A file system name with no name list following means export to everyone. A '#' anywhere in the file indicates a comment extending to the end of the line it appears on. Lines beginning with white space are continuation lines.

EXAMPLE

```
/usr      clients                # export to my clients
/usr/local                # export to the world
/usr2     phoenix sun sundae    # export to only these machines
```

FILES

/etc/exports

SEE ALSO

mountd(8C)

NAME

fcntl – file control options

SYNOPSIS

#include <fcntl.h>

DESCRIPTION

The *fcntl(2)* 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 files flags */
#define F_SETFD      2    /* Set files flags */
#define F_GETFL      3    /* Get file flags */
#define F_SETFL      4    /* Set file flags */
#define F_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(2), open(2V)

NAME

fs, inode – format of file system volume

SYNOPSIS

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

DESCRIPTION

Every file system storage volume (disk, nine-track tape, for instance) has a common format for certain vital information. Every such volume is divided into a certain number of blocks. The block size is a parameter of the file system. Sectors 0 to 15 on a file system are used to contain primary and secondary bootstrapping programs.

The actual file system begins at sector 16 with the *super block*. The layout of the super block as defined by the include file *<sys/fs.h>* is:

```
#define FS_MAGIC    0x011954
struct fs {
    struct fs *fs_link;           /* linked list of file systems */
    struct fs *fs_rlink;         /* used for incore super blocks */
    daddr_t fs_sblkno;           /* addr of super-block in filesys */
    daddr_t fs_cblkno;           /* offset of cyl-block in filesys */
    daddr_t fs_iblkno;           /* offset of inode-blocks in filesys */
    daddr_t fs_dblkno;           /* offset of first data after cg */
    long fs_cgoffset;            /* cylinder group offset in cylinder */
    long fs_cgmask;              /* used to calc mod fs_ntrak */
    time_t fs_time;              /* last time written */
    long fs_size;                 /* number of blocks in fs */
    long fs_dsize;                /* number of data blocks in fs */
    long fs_ncg;                  /* number of cylinder groups */
    long fs_bsize;                /* size of basic blocks in fs */
    long fs_fsize;                /* size of frag blocks in fs */
    long fs_frag;                 /* number of frags in a block in fs */

    /* these are configuration parameters */
    long fs_minfree;              /* minimum percentage of free blocks */
    long fs_rotdelay;            /* num of ms for optimal next block */
    long fs_rps;                  /* disk revolutions per second */

    /* these fields can be computed from the others */
    long fs_bmask;                /* "blkoff" calc of blk offsets */
    long fs_fmask;                /* "fragoff" calc of frag offsets */
    long fs_bshift;               /* "lblkno" calc of logical blkno */
    long fs_fshift;               /* "numfrags" calc number of frags */

    /* these are configuration parameters */
    long fs_maxcontig;            /* max number of contiguous blks */
    long fs_maxbpg;               /* max number of blks per cyl group */

    /* these fields can be computed from the others */
    long fs_fragshift;            /* block to frag shift */
    long fs_fsbtoadb;             /* fsbtoadb and dbtofsb shift constant */
    long fs_sbsize;               /* actual size of super block */
    long fs_csmask;               /* csum block offset */
    long fs_csshift;              /* csum block number */
    long fs_nindir;               /* value of NINDIR */
    long fs_inopb;                /* value of INOPB */
    long fs_nspf;                 /* value of NSPF */
    long fs_optim;                /* optimization preference, see below */
};
```

```

    long    fs_sparecon[3];          /* reserved for future constants */
/* a unique id for this filesystem (currently unused and unmaintained) */
    long    fs_id[2];               /* file system id */
/* sizes determined by number of cylinder groups and their sizes */
    daddr_t fs_csaddr;              /* blk addr of cyl grp summary area */
    long    fs_cssize;              /* size of cyl grp summary area */
    long    fs_cgsize;              /* cylinder group size */
/* these fields should be derived from the hardware */
    long    fs_ntrak;               /* tracks per cylinder */
    long    fs_nsect;               /* sectors per track */
    long    fs_spc;                 /* sectors per cylinder */
/* this comes from the disk driver partitioning */
    long    fs_ncyl;                /* cylinders in file system */
/* these fields can be computed from the others */
    long    fs_cpg;                 /* cylinders per group */
    long    fs_ipg;                 /* inodes per group */
    long    fs_fpg;                 /* blocks per group * fs_frag */
/* this data must be re-computed after crashes */
    struct  csum fs_cstotal; /* cylinder summary information */
/* these fields are cleared at mount time */
    char    fs_fmod;                /* super block modified flag */
    char    fs_clean;               /* file system is clean flag */
    char    fs_ronly;               /* mounted read-only flag */
    char    fs_flags;               /* currently unused flag */
    char    fs_fsmnt[MAXMNTLEN];    /* name mounted on */
/* these fields retain the current block allocation info */
    long    fs_cgrotor;              /* last cg searched */
    struct  csum *fs_csp[MAXCSBUFS]; /* list of fs_cs info buffers */
    long    fs_cpc;                 /* cyl per cycle in postbl */
    short   fs_postbl[MAXCPG][NRPOS]; /* head of blocks for each rotation */
    long    fs_magic;               /* magic number */
    u_char  fs_rotbl[1];            /* list of blocks for each rotation */
/* actually longer */
};

```

Each disk drive contains some number of file systems. A file system consists of a number of cylinder groups. Each cylinder group has inodes and data.

A file system is described by its super-block, which in turn describes the cylinder groups. The super-block is critical data and is replicated in each cylinder group to protect against catastrophic loss. This is done at file system creation time and the critical super-block data does not change, so the copies need not be referenced further unless disaster strikes.

Addresses stored in inodes are capable of addressing fragments of 'blocks'. File system blocks of at most size MAXBSIZE can be optionally broken into 2, 4, or 8 pieces, each of which is addressable; these pieces may be DEV_BSIZE, or some multiple of a DEV_BSIZE unit.

Large files consist of exclusively large data blocks. To avoid undue wasted disk space, the last data block of a small file is allocated as only as many fragments of a large block as are necessary. The file system format retains only a single pointer to such a fragment, which is a piece of a single large block that has been divided. The size of such a fragment is determinable from information in the inode, using the "blksize(fs, ip, lbn)" macro.

The file system records space availability at the fragment level; to determine block availability, aligned fragments are examined.

The root inode is the root of the file system. Inode 0 can't be used for normal purposes and historically bad blocks were linked to inode 1, thus the root inode is 2 (inode 1 is no longer used for this purpose, however numerous dump tapes make this assumption, so we are stuck with it). The *lost+found* directory is given the next available inode when it is initially created by *mkfs*.

fs_minfree gives the minimum acceptable percentage of file system blocks which may be free. If the freelist drops below this level only the super-user may continue to allocate blocks. This may be set to 0 if no reserve of free blocks is deemed necessary, however severe performance degradations will be observed if the file system is run at greater than 90% full; thus the default value of *fs_minfree* is 10%.

Empirically the best trade-off between block fragmentation and overall disk utilization at a loading of 90% comes with a fragmentation of 4, thus the default fragment size is a fourth of the block size.

Cylinder group related limits: Each cylinder keeps track of the availability of blocks at different rotational positions, so that sequential blocks can be laid out with minimum rotational latency. NRPOS is the number of rotational positions which are distinguished. With NRPOS 8 the resolution of the summary information is 2ms for a typical 3600 rpm drive.

fs_rotdelay gives the minimum number of milliseconds to initiate another disk transfer on the same cylinder. It is used in determining the rotationally optimal layout for disk blocks within a file; the default value for *fs_rotdelay* is 2ms.

Each file system has a statically allocated number of inodes. An inode is allocated for each NBPI bytes of disk space. The inode allocation strategy is extremely conservative.

MAXIPG bounds the number of inodes per cylinder group, and is needed only to keep the structure simpler by having the only a single variable size element (the free bit map).

N.B.: MAXIPG must be a multiple of INOPB(fs).

MINBSIZE is the smallest allowable block size. With a MINBSIZE of 4096 it is possible to create files of size 2^{32} with only two levels of indirection. MINBSIZE must be big enough to hold a cylinder group block, thus changes to (struct cg) must keep its size within MINBSIZE. MAXCPG is limited only to dimension an array in (struct cg); it can be made larger as long as that structure's size remains within the bounds dictated by MINBSIZE. Note that super blocks are never more than size SBSIZE.

The path name on which the file system is mounted is maintained in *fs_fsmnt*. MAXMNTLEN defines the amount of space allocated in the super block for this name. The limit on the amount of summary information per file system is defined by MAXCSBUFS. It is currently parameterized for a maximum of two million cylinders.

Per cylinder group information is summarized in blocks allocated from the first cylinder group's data blocks. These blocks are read in from *fs_csaddr* (size *fs_cssize*) in addition to the super block.

N.B.: sizeof (struct csum) must be a power of two in order for the "fs_cs" macro to work.

Super block for a file system: MAXBPC bounds the size of the rotational layout tables and is limited by the fact that the super block is of size SBSIZE. The size of these tables is inversely proportional to the block size of the file system. The size of the tables is increased when sector sizes are not powers of two, as this increases the number of cylinders included before the rotational pattern repeats (*fs_cpc*). The size of the rotational layout tables is derived from the number of bytes remaining in (struct fs).

MAXBPG bounds the number of blocks of data per cylinder group, and is limited by the fact that cylinder groups are at most one block. The size of the free block table is derived from the size of blocks and the number of remaining bytes in the cylinder group structure (struct cg).

Inode: The inode is the focus of all file activity in the UNIX file system. There is a unique inode allocated for each active file, each current directory, each mounted-on file, text file, and the root. An inode is 'named' by its device/i-number pair. For further information, see the include file <sys/inode.h>.

NAME

fspec – format specification in text files

DESCRIPTION

It is sometimes convenient to maintain text files on the UNIX system with non-standard tabs, (i.e., tabs which are not set at every eighth column). Such files must generally be converted to a standard format, frequently by replacing all tabs with the appropriate number of spaces, before they can be processed by UNIX system commands. A format specification occurring in the first line of a text file specifies how tabs are to be expanded in the remainder of the file.

A format specification consists of a sequence of parameters separated by blanks and surrounded by the brackets <: and :>. Each parameter consists of a keyletter, possibly followed immediately by a value. The following parameters are recognized:

- ttabs* The *t* parameter specifies the tab settings for the file. *ttabs* must be one of the following:
1. a list of column numbers separated by commas, indicating tabs set at the specified columns;
 2. *a* – followed immediately by an integer *n*, indicating tabs at intervals of *n* columns;
 3. *a* – followed by the name of a “canned” tab specification.
- Standard tabs are specified by *t*–8, or equivalently, *t*1,9,17,25,etc. The canned tabs which are recognized are as follows:
- a* 1,10,16,36,72
Assembler, IBM S/370, first format
- a*2 1,10,16,40,72
Assembler, IBM S/370, second format
- c* 1,8,12,16,20,55
COBOL, normal format
- c*2 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 setup should include a format specification as follows:
<:t–c2 m6 s66 d:>
- c*3 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 tabs than *c*2. 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
- ssize* The *s* parameter specifies a maximum line size. The value of *size* must be an integer. Size checking is performed after tabs have been expanded, but before the margin is prepended.

mmargin

The **m** parameter specifies a number of spaces to be prepended to each line. The value of *margin* must be an integer.

- d** The **d** parameter takes no value. Its presence indicates that the line containing the format specification is to be deleted from the converted file.
- e** The **e** parameter takes no value. Its presence indicates that the current format is to prevail only until another format specification is encountered in the file.

Default values, which are assumed for parameters not supplied, are **t-8** and **m0**. If the **s** parameter is not specified, no size checking is performed. If the first line of a file does not contain a format specification, the above defaults are assumed for the entire file. The following is an example of a line containing a format specification:

```
* <:t5,10,15 s72:> *
```

If a format specification can be disguised as a comment, it is not necessary to code the **d** parameter.

Several UNIX system commands correctly interpret the format specification for a file. Among them is *newform(1)* which may be used to convert files to a standard format acceptable to other UNIX system commands.

SEE ALSO

newform(1)

NAME

ftpusers – list of users prohibited by ftp

SYNOPSIS

/usr/etc/ftpusers

DESCRIPTION

Ftpusers contains a list of users who cannot access this system using the *ftp*(1) program. *Ftpusers* 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.

Name	Type	Default	Description
ap	bool	false	terminal uses 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 \n after login prompt
ds	str	^Y	delayed suspend character
ec	bool	false	leave echo OFF
ep	bool	false	terminal uses even parity
er	str	^?	erase character
et	str	^D	end of text (EOF) character
ev	str	NULL	initial enviroment
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	/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 odd parity
os	num	unused	output speed
pc	str	\0	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 enviroment)
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.

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 (a **la** termcap). 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** to obtain the hostname. (**%%** 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, causes one character from the real hostname to be copied to the final hostname. A '#' in the **he** string, causes the next character of the real hostname to be skipped. Surplus '@' and '#' characters are ignored.

When *getty* execs the login process, given in the **lo** string (usually **"/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* 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

Group contains for each group the following information:

- group name
- encrypted password
- numerical group ID
- a comma separated list of all users allowed in the group

This is an ASCII file. The fields are separated by colons; each group is separated from the next by a new-line. If the password field is null, no password is demanded.

This file resides in the */etc* directory. Because of the encrypted passwords, it can and does have general read permission and can be used, for example, to map numerical group ID's to names.

A group file can have a line beginning with a plus (+), 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 yellow pages group 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 or group member field, the contents of that field will override what is contained in the yellow pages. The numerical group ID field cannot be overridden.

EXAMPLE

```
+myproject:::bill, steve
+:
```

If these entries appear at the end of a group file, then 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 the groups listed in the yellow pages will be pulled in and placed after the entry for *myproject*.

FILES

/etc/group */etc/yp/group*

SEE ALSO

setgroups(2), *initgroups(3)*, *crypt(3)*, *passwd(1)*, *passwd(5)*

BUGS

The *passwd(1)* command won't change group passwords.

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           # Alison Shanks
```

FILES

/etc/hosts

SEE ALSO

gethostent(3N)

NAME

hosts.equiv – list of trusted hosts

DESCRIPTION

Hosts.equiv resides in directory */etc* and contains a list of trusted hosts. When an *rlogin(1)* or *rsh(1)* request from such a host is made, and the initiator of the request is in */etc/passwd*, then no further validity checking is done. That is, *rlogin* does not prompt for a password, and *rsh* completes successfully. So a remote user is “equivalenced” to a local user with the same user ID when the remote user is in *hosts.equiv*.

The format of *hosts.equiv* is a list of names, as in this example:

```
host1
host2
+@group1
-@group2
```

A line consisting of a simple host name means that anyone logging in from that host is trusted. A line consisting of *+@group* means that all hosts in that network group are trusted. A line consisting of *-@group* means that hosts in that group are not trusted. Programs scan *hosts.equiv* linearly, and stop at the first hit (either positive for hostname and *+@* entries, or negative for *-@* entries). A line consisting of a single *+* means that everyone is trusted.

The *.rhosts* file has the same format as *hosts.equiv*. When user *XXX* executes *rlogin* or *rsh*, the *.rhosts* file from *XXX*'s home directory is conceptually concatenated onto the end of *hosts.equiv* for permission checking. However, *-@* entries are not sticky. If a user is excluded by a minus entry from *hosts.equiv* but included in *.rhosts*, then that user is considered trusted. In the special case when the user is root, then only the *.rhosts* file is checked.

It is also possible to have two entries (separated by a single space) on a line of these files. In this case, if the remote host is equivalenced by the first entry, then the user named by the second entry is allowed to log in as anyone, that is, specify any name to the *-l* flag (provided that name is in the */etc/passwd* file, of course). Thus

```
sundown john
```

allows *john* to log in from *sundown* as anyone. The usual usage would be to put this entry in the *.rhosts* file in the home directory for *bill*. Then *john* may log in as *bill* when coming from *sundown*. The second entry may be a netgroup, thus

```
+@group1 +@group2
```

allows any user in *group2* coming from a host in *group1* to log in as anyone.

FILES

/etc/hosts.equiv

SEE ALSO

rlogin(1), *rsh(1)*, *netgroup(5)*

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; e.g. **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 (e.g. *\n* for new-line).

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, 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.

BUGS

There should be more than one level of subtests, with the level indicated by the number of **>** at the beginning of the line.

SEE ALSO

file(1)

NAME

`mntent`, `fstab` – static information about filesystems

SYNOPSIS

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

DESCRIPTION

The file `/etc/fstab` describes the file systems and swapping partitions used by the local machine. It is created by the system administrator using a text editor and processed by commands which mount, unmount, check consistency of, dump and restore file systems, and by the system in providing swap space.

It consists of a number of lines of the form:

```
fsname dir type opts freq passno
```

an example of which would be:

```
/dev/xy0a / 4.2 rw,noquota 1 2
```

The entries in this file are accessed using the routines in `getmntent(3)`, which returns a structure of the following form:

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

There is one entry per line in the file, and the fields are separated by white space. A '#' as the first non-white character indicates a comment.

The `mnt_opts` field consists of a string of comma separated options. Some of the options are common to all filesystem types, others only make sense for a single filesystem type. See `mount(8)` for a more complete description of the options available.

The `mnt_type` field determines how the `mnt_fsname`, and `mnt_opts` fields will be interpreted. Below is a list of the file system types currently supported and the way each of them interprets these fields.

4.2

`mnt_fsname` Must be a block special device.
`mnt_opts` Valid opts are: ro, rw, suid, nosuid, quota, noquota.

NFS

`mnt_fsname` The path on the server of the directory to be served.
`mnt_opts` Valid opts are: ro, rw, suid, nosuid, hard, soft, bg, fg, retry, rsize, wsize, timeo, retrans, port, intr.

SWAP

`mnt_fsname` Must be a block special device swap partition.
`mnt_opts` Ignored.

If the `mnt_type` is specified as "ignore" the entry is ignored. This is useful to show disk partitions which are currently not used.

The field `mnt_freq` indicates how often each partition should be dumped by the `dump(8)` command (and triggers that commands `w` option which tells which file systems should be dumped). Most systems set the `mnt_freq` field to 1 indicating that the file systems are dumped each day.

The final field *mnt_passno* is used by the disk consistency check program *fsck*(8) to allow overlapped checking of file systems during a reboot. All file systems with *mnt_passno* of 1 are first checked simultaneously, then all file systems with *mnt_passno* of 2, and so on. It is usual to make the *mnt_passno* of the root file system have the value 1 and then check one file system on each available disk drive in each subsequent pass to the exhaustion of file system partitions.

/etc/fstab is only *read* by programs, and not written; it is the duty of the system administrator to properly create and maintain this file. The order of records in */etc/fstab* is important because *fsck*, *mount*, and *umount* process the file sequentially; file systems must appear *after* file systems they are mounted within.

FILES

/etc/fstab

SEE ALSO

fsck(8), *getmntent*(3), *mount*(8), *quotacheck*(8), *quotaon*(8), *umount*(8)

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* 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
```

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 *member_i* is either another group name, or a triple:

```
(hostname, username, domainname)
```

Any of three fields can be empty, in which case it signifies a wild card. Thus

```
universal (,,)
```

defines a group to which everyone belongs. 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.

Network groups are contained in the yellow pages, and are accessed through these files:

```
/etc/yp/domainname/netgroup.dir
/etc/yp/domainname/netgroup.pag
/etc/yp/domainname/netgroup.byuser.dir
/etc/yp/domainname/netgroup.byuser.pag
/etc/yp/domainname/netgroup.byhost.dir
/etc/yp/domainname/netgroup.byhost.pag
```

These files can be created from */etc/netgroup* using *makedbm(8)*.

FILES

```
/etc/netgroup
/etc/yp/domainname/netgroup.dir
/etc/yp/domainname/netgroup.pag
/etc/yp/domainname/netgroup.byuser.dir
/etc/yp/domainname/netgroup.byuser.pag
/etc/yp/domainname/netgroup.byhost.dir
/etc/yp/domainname/netgroup.byhost.pag
```

SEE ALSO

getnetgrent(3), *exportfs(8)*, *makedbm(8)*, *ypserv(8)*

NAME

netrc – .netrc file for ftp(1) remote login data

DESCRIPTION

The *.netrc* file contains data for logging in to a remote host over the network for file transfers by *ftp(1)*. This file resides in the user's home directory on the machine initiating the file transfer. It's permissions should be set to disallow read access by group and others (see *chmod(1V)*).

Each line of the *.netrc* file defines options for a specific remote host. A line in the *.netrc* file can be either a **machine** line or a **default** line. The **default** line indicates a remote host to use as the default destination, and must be the first line in the *.netrc* file if present. The **machine** lines contain login information for each remote host to which files can be transferred.

default *default-machine-name*
machine *machine-name options*

Fields on each line are separated by SPACE or TAB characters.

The *options* for a **machine** line are:

Option	Parameter	Default	Description
login	name	localname	login name for remote machine
password	password	(none)	password for remote login name
command	command	(none)	default command to be executed
write	yes/no	yes	write to user if possible
force	yes/no	no	always prompt for login name and password
quiet	yes/no	no	like the -q option

EXAMPLE

```
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

ftp(1), ftpd(8)

NAME

`networks` – network name data base

DESCRIPTION

The `networks` file contains information regarding the known networks which comprise the DARPA Internet. For each network a single line should be present with the following information:

official network name

network number

aliases

Items are separated by any number of blanks and/or tab characters. A “#” indicates the beginning of a comment; characters up to the end of the line are not interpreted by routines which search the file. This file is normally created from the official network data base maintained at the Network Information Control Center (NIC), though local changes may be required to bring it up to date regarding unofficial aliases and/or unknown networks.

Network number may be specified in the conventional “.” notation using the `inet_network()` routine from the Internet address manipulation library, `inet(3N)`. Network names may contain any printable character other than a field delimiter, newline, or comment character.

FILES

`/etc/networks`

SEE ALSO

`getnetent(3N)`

BUGS

A name server should be used instead of a static file. A binary indexed file format should be available for fast access.

NAME

`passwd` – password file

SYNOPSIS

`/etc/passwd`

DESCRIPTION

The *passwd* file contains for each user the following information:

name User's login name — contains no upper case characters and must not be greater than eight characters long.

password encrypted password

numerical user ID

This is the user's ID in the system and it must be unique.

numerical group ID

This is the number of the group that the user belongs to.

user's real name

In some versions of UNIX, this field also contains the user's office, extension, home phone, and so on. For historical reasons this field is called the GCOS field.

initial working directory

The directory that the user is positioned in when they log in — this is known as the 'home' directory.

shell program to use as Shell when the user logs in.

The user's real name field may contain '&', meaning insert the login name.

The password file is an ASCII file. Each field within each user's entry is separated from the next by a colon. Each user is separated from the next by a new-line. If the password field is null, no password is demanded; if the Shell field is null, */bin/sh* is used.

The *passwd* file can also have line beginning with a plus (+), 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 password 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, directory, gecos, or shell field, they will override what is contained in the yellow pages. The numerical user ID and group ID fields cannot be overridden.

EXAMPLE

Here is a sample */etc/passwd* file:

```
root:q.mJzTnu8icF.:0:10:God:/:/bin/csh
tut:6k/7KCFRPNVXg:508:10:Bill Tuthill:/usr2/tut:/bin/csh
+john:
+@documentation:no-login:
+:::Guest
```

In this example, there are specific entries for users *root* *tut*, in case the yellow pages are out of order. The user 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 gecos field of *Guest*.

The password file resides in the */etc* directory. Because of the encrypted passwords, it has general read permission and can be used, for example, to map numerical user ID's to names.

Appropriate precautions must be taken to lock the */etc/passwd* file against simultaneous changes if it is to be edited with a text editor; *vipw*(8) does the necessary locking.

FILES

/etc/passwd

SEE ALSO

getpwent(3), login(1), crypt(3), passwd(1), group(5), vipw(8), 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 isn’t 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

plot(1G), plot(3X), graph(1G)

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 the *Line Printer Spooler Manual* in the *Sun System Administration Manual* for a discussion of how to set up the database for a given printer.

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 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 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 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'.

CAPABILITIES

Name	Type	Default	Description
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 (sgtty.h)
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 (<i>plot</i> (3X) format)
ic	bool	false	driver supports (non standard) ioctl call for indenting 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
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
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
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	"/usr/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
tf	str	NULL	troff data filter (cat 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 (tty (4))
xs	num	0	like 'xc' but set bits

Error messages sent to the console have a carriage return and a line feed appended to them, rather than just a line feed.

If the local line printer driver supports indentation, the daemon must understand how to invoke it.

Note that 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 lineprinter daemon if the device is a tty. 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). For example, to set exactly the flags 06300 in the 'fs' field, do:

```
:fc#0177777:fs#06300:
```

The same process applies to the 'xc' and 'xs' fields.

SEE ALSO

termcap(5), lpc(8), lpd(8), pac(8), lpr(1), lpq(1), lprm(1)

The *Line Printer Spooler Manual* in the *Sun System Administration Manual*.

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 the Sun UNIX system.

```
#
# Internet (IP) protocols
#
ip           0           IP           # internet protocol, pseudo protocol number
icmp        1           ICMP        # internet control message protocol
ggp         2           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

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 layed 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 that 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

Programmer's Reference Manual for SunWindows

NAME

remote – remote host description file

SYNOPSIS

/etc/remoted

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) UNIX device(s) to open to establish a connection. 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 end-of-file marks. The default is NULL.
- oe** (str) Output end-of-file string. The default is NULL. When *tip* is transferring a file, this string is sent at end-of-file.
- 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

SEE ALSO

tip(1C), phones(5)

NAME

resolver – 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

followed by 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 followed by a domain name, that is 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 followed by 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 (e.g. **nameserver**) must start the line. The value follows the keyword, separated by white space.

FILES

/etc/resolv.conf

SEE ALSO

domainname(1), gethostent(3N), named(8C)

NAME

rmtab – local file system mount statistics

DESCRIPTION

Rmtab resides in directory */etc* and contains a record of all clients that have done remote mounts of file systems from this machine. Whenever a remote *mount* is done, an entry is made in the *rmtab* file of the machine serving up that file system. *Umount* removes entries, if of a remotely mounted file system. *Umount -a* broadcasts to all servers, and informs them that they should remove all entries from *rmtab* created by the sender of the broadcast message. By placing a *umount -a* command in */etc/rc.boot*, *rmtab* tables can be purged of entries made by a crashed host, which upon rebooting did not remount the same file systems it had before. The table is a series of lines of the form

hostname:directory

This table is used only to preserve information between crashes, and is read only by *mountd*(8) when it starts up. *Mountd* keeps an in-core table, which it uses to handle requests from programs like *showmount*(1) and *shutdown*(8).

FILES

/etc/rmtab

SEE ALSO

showmount(1), *mountd*(8), *mount*(8), *umount*(8), *shutdown*(8)

BUGS

Although the *rmtab* table is close to the truth, it is not always 100% accurate.

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 Sun UNIX System.

```
#
#          rpc 1.7 86/04/24
#
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
etherstat      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
```

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 new-lines. The lines containing these login names and/or numerical group IDs are surrounded by the bracketing lines @u and @U. An empty list allows anyone to make a delta.

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>
```

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.

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 don’t 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

servers – inet server data base

DESCRIPTION

The *servers* file contains the list of servers that *inetd*(8) operates. For each server a single line should be present with the following information:

name of server
 protocol
 server location

If the server is rpc based, then the name field should be *rpc*, and following the server location are two additional fields, one with the rpc program number, the second with either a version number or a range of version numbers.

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.

The name of the server should be the official service name as contained in *services*(5). The protocol entry is either *udp* or *tcp*. The server location is the full path name of the server program.

EXAMPLE

The following example is taken from the Sun UNIX system.

```

tcp      tcp  /usr/etc/in.tcpd
telnet   tcp  /usr/etc/in.telnetd
shell    tcp  /etc/in.rshd
login    tcp  /etc/in.rlogind
exec     tcp  /usr/etc/in.rexecd
tcp      udp  /usr/etc/in.ttcpd
syslog   udp  /usr/etc/in.syslog
comsat   udp  /usr/etc/in.comsat
talk     udp  /usr/etc/in.talkd
time     tcp  /usr/etc/in.timed
rpc      udp  /usr/etc/rpc.rstatd    100001  1-2
rpc      udp  /usr/etc/rpc.rusersd  100002  1
rpc      udp  /usr/etc/rpc.rwalld   100008  1
rpc      udp  /usr/etc/rpc.mountd     100005  1

```

FILES

/etc/servers

SEE ALSO

services(5), *inetd*(8)

BUGS

Because of a limitation on the number of open files, this file must contain fewer than 27 lines.

NAME

services – service name data base

SYNOPSIS

/etc/services

DESCRIPTION

The *services* file contains information regarding the known services available in the DARPA Internet. For each service a single line should be present with the following information:

official service name
port number
protocol name
aliases

Items are separated by any number of blanks and/or tab characters. The port number and protocol name are considered a single *item*; a “/” is used to separate the port and protocol (for instance, “512/tcp”). A “#” indicates the beginning of a comment; characters up to the end of the line are not interpreted by routines which search the file.

Service names may contain any printable character other than a field delimiter, newline, or comment character.

EXAMPLE

Here is an example of the */etc/services* file from the Sun UNIX System.

```
#
# @(#)services 1.7 86/02/28 SMI
#
# Network services, Internet style
# This file is never consulted when the yellow pages are running
#
echo                7/udp
discard             9/udp                sink null
systat              11/tcp
daytime             13/tcp
netstat            15/tcp
ftp-data           20/tcp
ftp                 21/tcp
telnet              23/tcp
smtp                25/tcp                mail
time                37/tcp                timserver
time                37/udp               timserver
name                42/udp                nameserver
whois               43/tcp                nickname# usually to sri-nic
domain              53/udp
domain              53/tcp
hostnames           101/tcp                hostname # usually to sri-nic
sunrpc              111/udp
sunrpc              111/tcp
#
# Host specific functions
#
tftp                69/udp
rje                 77/tcp
finger              79/tcp
link                87/tcp                ttylink
supdup              95/tcp
```

csnet-ns	105/tcp		
uucp-path	117/tcp		
untp		119/tcp	usenet
ntp		123/tcp	
ingreslock	1524/tcp		
#			
# UNIX specific services			
#			
exec		512/tcp	
login		513/tcp	
shell		514/tcp	cmd# no passwords used
printer		515/tcp	spooler# experimental
courier		530/tcp	rpc# experimental
biff		512/udp	comsat
who		513/udp	whod
syslog		514/udp	
talk		517/udp	
route		520/udp	router routed
new-rwho	550/udp		# experimental
rmonitor	560/udp		# experimental
monitor		561/udp	# experimental
		new-who	
		rmonitord	

FILES

/etc/services

SEE ALSO

getservent(3N)

BUGS

A name server should be used instead of a static file. A binary indexed file format should be available for fast access.

NAME

statmon, current, backup, state – statd directory and file structures

SYNOPSIS

/etc/statmon/current /etc/statmon/backup, /etc/statmon/state

DESCRIPTION

/etc/statmon/current and */etc/statmon/backup* are directories generated by *statd*. Each entry in */etc/statmon/current* represents the name of the machine to be monitored by the *statd* daemon. Each entry in */etc/statmon/backup* represents the name of the machine to be notified by the *statd* daemon upon its recovery.

/etc/statmon/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.

SEE ALSO

statd(8C), lockd(8C)

NAME

tar – tape archive file format

DESCRIPTION

Tar, (the tape archive command) dumps several files into one, in a medium suitable for transportation.

A “tar tape” or file is a series of blocks. Each block is of size TBLOCK. A file on the tape is represented by a header block which describes the file, followed by zero or more blocks which give the contents of the file. At the end of the tape are two blocks filled with binary zeros, as an end-of-file indicator.

The blocks are grouped for physical I/O operations. Each group of *n* blocks (where *n* is set by the *b* keyletter on the *tar*(1) command line — default is 20 blocks) is written with a single system call; on nine-track tapes, the result of this write is a single tape record. The last group is always written at the full size, so blocks after the two zero blocks contain random data. On reading, the specified or default group size is used for the first read, but if that read returns less than a full tape block, the reduced block size is used for further reads, unless the *B* keyletter is used.

The header block looks like:

```
#define TBLOCK      512
#define NAMSIZ 100

union hblock {
    char dummy[TBLOCK];
    struct header {
        char name[NAMSIZ];
        char mode[8];
        char uid[8];
        char gid[8];
        char size[12];
        char mtime[12];
        char chksum[8];
        char linkflag;
        char linkname[NAMSIZ];
    } dbuf;
};
```

Name is a null-terminated string. The other fields are zero-filled octal numbers in ASCII. Each field (of width *w*) contains *w*-2 digits, a space, and a null, except *size* and *mtime*, which do not contain the trailing null. *Name* is the name of the file, as specified on the *tar* command line. Files dumped because they were in a directory which was named in the command line have the directory name as prefix and *filename* as suffix. *Mode* is the file mode, with the top bit masked off. *Uid* and *gid* are the user and group numbers which own the file. *Size* is the size of the file in bytes. Links and symbolic links are dumped with this field specified as zero. *Mtime* is the modification time of the file at the time it was dumped. *Chksum* is a decimal ASCII value which represents the sum of all the bytes in the header block. When calculating the checksum, the *chksum* field is treated as if it were all blanks. *Linkflag* is ASCII ‘0’ if the file is “normal” or a special file, ASCII ‘1’ if it is a hard link, and ASCII ‘2’ if it is a symbolic link. The name linked-to, if any, is in *linkname*, with a trailing null. Unused fields of the header are binary zeros (and are included in the checksum).

The first time a given i-node number is dumped, it is dumped as a regular file. The second and subsequent times, it is dumped as a link instead. Upon retrieval, if a link entry is retrieved, but not the file it was linked to, an error message is printed and the tape must be manually re-scanned to retrieve the linked-to file.

The encoding of the header is designed to be portable across machines.

SEE ALSO

tar(1)

BUGS

Names or linknames longer than NAMSIZ produce error reports and cannot be dumped.

NAME

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 *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 *tty*(4).
breset bits to reset in the *sg_flags* field of the *sgtty* structure after output; see *tty*(4).
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.

<i>Adj</i>	quantum of white space, in fractions of an inch. (that is, white spaces are a multiple of <i>Adj</i> units wide)										
	Note: if this is less than the size of the space character (in units of <i>Char</i> ; see below for how the sizes of characters are defined), <i>nroff</i> will output fractional spaces using plot mode. Also, if the <i>-e</i> switch to <i>nroff</i> is used, <i>Adj</i> is set equal to <i>Hor</i> by <i>nroff</i> .										
<i>twinit</i>	set of characters used to initialize the terminal in a mode suitable for <i>nroff</i> .										
<i>twrest</i>	set of characters used to restore the terminal to normal mode.										
<i>twnl</i>	set of characters used to move down one line.										
<i>hlf</i>	set of characters used to move up one-half line.										
<i>hlf</i>	set of characters used to move down one-half line.										
<i>flr</i>	set of characters used to move up one line.										
<i>bdon</i>	set of characters used to turn on hardware boldface mode, if any.										
<i>bdoff</i>	set of characters used to turn off hardware boldface mode, if any.										
<i>ploton</i>	set of characters used to turn on hardware plot mode (for Diablo type mechanisms), if any.										
<i>plotoff</i>	set of characters used to turn off hardware plot mode (for Diablo type mechanisms), if any.										
<i>up</i>	set of characters used to move up one resolution unit (<i>Vert</i>) in plot mode, if any.										
<i>down</i>	set of characters used to move down one resolution unit (<i>Vert</i>) in plot mode, if any.										
<i>right</i>	set of characters used to move right one resolution unit (<i>Hor</i>) in plot mode, if any.										
<i>left</i>	set of characters used to move left one resolution unit (<i>Hor</i>) in plot mode, if any.										
<i>codetab</i>	definition of characters needed to print an <i>nroff</i> character on the terminal. The first byte is the number of character units (<i>Char</i>) 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 (<i>.ul</i>). 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.										
<i>zzz</i>	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.

FILES

<i>/usr/lib/term/tabname</i>	driving tables
<i>/usr/lib/term/README</i>	list of terminals supported by <i>nroff</i> (1)

SEE ALSO

nroff(1)

NAME

term – format of compiled term file

SYNOPSIS

term

DESCRIPTION

Compiled *terminfo* descriptions are placed under the directory */usr/5lib/terminfo*. In order to avoid a linear search of a huge UNIX system directory, a two-level scheme is used: */usr/5lib/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/5lib/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* program, and read by the routine *setupterm*. 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 that this format corresponds to the hardware of the VAX and PDP-11. Machines where this does not correspond to the hardware read the integers as two bytes and compute the result.

The terminal names section comes next. It contains the first line of the *terminfo* description, listing the various names for the terminal, separated by the `|` character. The section is terminated with an ASCII NUL character.

The boolean flags have one byte for each flag. This byte is either 0 or 1 as the flag is present or absent. The capabilities are in the same order as the file `<term.h>`.

Between the boolean section and the number section, a null byte will be inserted, if necessary, to ensure that the number section begins on an even byte. All short integers are aligned on a short word boundary.

The numbers section is similar to the flags section. Each capability takes up two bytes, and is stored as a short integer. If the value represented is -1 , the capability is taken to be missing.

The strings section is also similar. Each capability is stored as a short integer, in the format above. A value of -1 means the capability is missing. Otherwise, the value is taken as an offset from the beginning of the string table. Special characters in `^X` or `\c` notation are stored in their interpreted form, not the printing representation. Padding information `$<nn>` and parameter information `%x` are stored intact in uninterpreted form.

The final section is the string table. It contains all the values of string capabilities referenced in the string section. Each string is null terminated.

Note that it is possible for *setupterm* to expect a different set of capabilities than are actually present in the file. Either the database may have been updated since *setupterm* has been recompiled (resulting in extra unrecognized entries in the file) or the program may have been recompiled more recently than the database was updated (resulting in missing entries). The routine *setupterm* must be prepared for both possibilities – this is why the numbers and sizes are included. Also, new capabilities must always be added at the end of the lists of boolean, number, and string capabilities.

As an example, an octal dump of the description for the Microterm ACT 4 is included:

```

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 \0 377 377 377 377 377 377 377 377
120 377 377 377 377 \0 \0 002 \0 377 377 377 377 004 \0 006 \0
140 \b \0 377 377 377 377 \n \0 026 \0 030 \0 377 377 032 \0
160 377 377 377 377 034 \0 377 377 036 \0 377 377 377 377 377 377
200 377 377 377 377 377 377 377 377 377 377 377 377 377 377 377
*
520 377 377 377 377      \0 377 377 377 377 377 377 377 377 377 377
540 377 377 377 377 377 377 007 \0 \r \0 \f \0 036 \0 037 \0
560 024 % p 1 % c % p 2 % c \0 \n \0 035 \0
600 \b \0 030 \0 032 \0 \n \0

```

Some limitations: total compiled entries cannot exceed 4096 bytes. The name field cannot exceed 128 bytes.

FILES

/usr/5lib/terminfo/*/* compiled terminal capability data base

SEE ALSO

curses(3V), terminfo(5V).

NAME

termcap – terminal capability data base

SYNOPSIS

/etc/termcap

DESCRIPTION

Termcap is a data base describing terminals, used, for example, by *vi*(1) and *curses*(3X). Terminals are described in *termcap* by giving a set of capabilities which they have, and by describing how operations are performed. Padding requirements and initialization sequences are included in *termcap*.

Each entry in the *termcap* file describes a terminal, and is a line consisting of a number of fields separated by ':' characters. The first entry for each terminal gives the names which are known for the terminal, separated by '|' characters. The first name is always 2 characters long and is used by older version 6 systems which store the terminal type in a 16 bit word in a systemwide data base. The second name given is the most common abbreviation for the terminal, and the last name given should be a long name fully identifying the terminal. The second name should contain no blanks; the last name may well contain blanks for readability. 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 *termcap* are all introduced by two-character codes, and are of three types:

Boolean capabilities indicate that the terminal 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 supply information such as the size of the terminal or the size of particular delays. 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 and then the numeric value. For example: :co#80: is a numeric entry stating that this terminal has 80 columns.

String capabilities give a sequence which can be used to perform particular terminal 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, :ce=16\E^S: is a sample entry for clear to end-of-line.

CAPABILITIES

(P) indicates padding may be specified

(P*) indicates that padding may be based on the number of lines affected

Name	Type	Pad?	Description
ae	str	(P)	End alternate character set
al	str	(P*)	Add new blank line
am	bool		Terminal has automatic margins
as	str	(P)	Start alternate character set
bc	str		Backspace if not ^H
bl	str		Audible bell character
bs	bool		Terminal can backspace with ^H
bt	str	(P)	Back tab
bw	bool		Backspace wraps from column 0 to last column
CC	str		Command character in prototype if terminal settable
cd	str	(P*)	Clear to end of display
ce	str	(P)	Clear to end of line
ch	str	(P)	Like cm but horizontal motion only, line stays same
cl	str	(P*)	Clear screen
cm	str	(P)	Cursor motion
co	num		Number of columns in a line
cr	str	(P*)	Carriage return, (default ^M)

cs	str	(P)	Change scrolling region (vt100), like cm
ct	str		Clear all tab stops
cv	str	(P)	Like ch but vertical only.
da	bool		Display may be retained above
dB	num		Number of millisecc of bs delay needed
db	bool		Display may be retained below
dC	num		Number of millisecc of cr delay needed
dc	str	(P*)	Delete character
dF	num		Number of millisecc of ff delay needed
dl	str	(P*)	Delete line
dm	str		Delete mode (enter)
dN	num		Number of millisecc of nl delay needed
do	str		Down one line
dT	num		Number of millisecc of tab delay needed
ed	str		End delete mode
ei	str		End insert mode; give “:ei=:” if ic
eo	str		Can erase overstrikes with a blank
ff	str	(P*)	Hardcopy terminal page eject (default ^L)
hc	bool		Hardcopy terminal
hd	str		Half-line down (forward 1/2 linefeed)
ho	str		Home cursor (if no cm)
hu	str		Half-line up (reverse 1/2 linefeed)
hz	str		Hazeltine; can't print ~'s
ic	str	(P)	Insert character
if	str		Name of file containing is
im	bool		Insert mode (enter); give “:im=:” if ic
in	bool		Insert mode distinguishes nulls on display
ip	str	(P*)	Insert pad after character inserted
is	str		Terminal initialization string
k0-k9	str		Sent by “other” function keys 0-9
kb	str		Sent by backspace key
kd	str		Sent by terminal down arrow key
ke	str		Out of “keypad transmit” mode
kh	str		Sent by home key
kl	str		Sent by terminal left arrow key
kn	num		Number of “other” keys
ko	str		Termcap entries for other non-function keys
kr	str		Sent by terminal right arrow key
ks	str		Put terminal in “keypad transmit” mode
ku	str		Sent by terminal up arrow key
l0-l9	str		Labels on “other” function keys
le	str		Move cursor left one place
li	num		Number of lines on screen or page
ll	str		Last line, first column (if no cm)
ma	str		Arrow key map, used by vi version 2 only
mb	str		Turn on blinking
md	str		Enter bold (extra-bright) mode
me	str		Turn off all attributes, normal mode
mh	str		Enter dim (half-bright) mode
mi	bool		Safe to move while in insert mode
ml	str		Memory lock on above cursor.
mr	str		Enter reverse mode
ms	bool		Safe to move while in standout and underline mode

mu	str	Memory unlock (turn off memory lock).
nc	bool	No correctly working carriage return (DM2500,H2000)
nd	str	Non-destructive space (cursor right)
nl	str (P*)	Newline character (default <code>\n</code>)
ns	bool	Terminal is a CRT but doesn't scroll.
os	bool	Terminal overstrikes
pc	str	Pad character (rather than null)
pt	bool	Has hardware tabs (may need to be set with <code>is</code>)
rf	str	Reset file, like <code>if</code> but for <code>reset</code> (1)
rs	str	Reset string, like <code>is</code> but for <code>reset</code> (1)
se	str	End stand out mode
sf	str (P)	Scroll forwards
sg	num	Number of blank chars left by <code>so</code> or <code>se</code>
so	str	Begin stand out mode
sr	str (P)	Scroll reverse (backwards)
st	str	Set a tab in all rows, current column
ta	str (P)	Tab (other than <code>^I</code> or with padding)
tc	str	Entry of similar terminal - must be last
te	str	String to end programs that use <code>cm</code>
ti	str	String to begin programs that use <code>cm</code>
uc	str	Underscore one char and move past it
ue	str	End underscore mode
ug	num	Number of blank chars left by <code>us</code> or <code>ue</code>
ul	bool	Terminal underlines even though it doesn't overstrike
up	str	Upline (cursor up)
us	str	Start underscore mode
vb	str	Visible bell (may not move cursor)
ve	str	Sequence to end open/visual mode
vs	str	Sequence to start open/visual mode
vt	num	Virtual terminal number (CB/UNIX)
xb	bool	Beehive (<code>f1=escape</code> , <code>f2=ctrl C</code>)
xn	bool	A newline is ignored after a wrap (Concept)
xr	bool	Return acts like <code>ce \r \n</code> (Delta Data)
xs	bool	Standout not erased by writing over it (HP 264?)
xt	bool	Tabs are destructive, magic so char (Telaray 1061)

A Sample Entry

The following example describes the wyse terminal entry.

```
wv|wyse-vp|wyse|Wyse 50 in ADDS Viewpoint emulation mode with "enhance" on:
:am:cr=^M:do=^J:nl=^J:bl=^G:if=/usr/lib/tabset/wyse-adds:      :le=^H:bs:li#24:co#80:cm=EY%+  %+
:cd=Ek:ce=EK:nd=^F:                                       :up=^Z:cl=^L:ll=^A:kl=^U:kr=^F:kd=^J:ku=^Z:kh=^A:
:pt:so=^N:se=^O:us=^N:ue=^O:                               :dl=El:al=EM:im=Eq:ei=Er:dc=EW:
:is=E' 72E'9^OEr:rs=E' 72E'9^OEr:
```

Types of Capabilities

Capabilities in *termcap* are of three types: Boolean capabilities which indicate that the terminal has some particular feature, numeric capabilities giving the size of the terminal or the size of particular delays, and string capabilities, which give a sequence which can be used to perform particular terminal operations. All capabilities have two letter codes.

Boolean capabilities are introduced simply by stating the two-character capability code in the field between ':' characters. 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 `co` which indicates the number of columns the terminal has gives the value '80' for the Concept.

String valued capabilities, such as `ce` (clear to end of line sequence) are given by the two character code, an '=', and then a string ending at the next following ':'. A delay in milliseconds may appear after the '=' in such a capability, and padding characters are supplied by the editor after the remainder of the string is sent to provide this delay. The delay can be either a integer, for instance, '20', or an integer followed by an '*', that is, '3*'. A '*' indicates that the padding required is proportional to the number of lines affected by the operation, and the amount given is the per-affected-unit padding required. When a '*' is specified, it is sometimes useful to give a delay of the form '3.5' to specify a delay per unit to tenths of milliseconds.

A number of escape sequences are provided in the string valued capabilities for easy encoding of characters there. A `\E` maps to an ESCAPE character, `^x` maps to a control-x for any appropriate x, and the sequences `\n` `\r` `\t` `\b` `\f` give a newline, return, tab, backspace and formfeed. Finally, characters may be given as three octal digits after a `\`, and the characters `^` and `\` may be given as `\^` and `\\`. If it is necessary to place a `:` in a capability it must be escaped in octal as `\072`. If it is necessary to place a null character in a string capability it must be encoded as `\200`. The routines which deal with *termcap* use C strings, and strip the high bits of the output very late so that a `\200` comes out as a `\000` would.

Preparing Descriptions

We now outline how to prepare descriptions of terminals. The most effective way to prepare a terminal description is by imitating the description of a similar terminal in *termcap* and to build up a description gradually, using partial descriptions with *ex* to check that they are correct. Be aware that a very unusual terminal may expose deficiencies in the ability of the *termcap* file to describe it or bugs in *ex*. To easily test a new terminal description you can set the environment variable `TERMCAP` to a pathname of a file containing the description you are working on and the editor will look there rather than in *etc/termcap*. `TERMCAP` can also be set to the *termcap* entry itself to avoid reading the file when starting up the editor.

Basic capabilities

The number of columns on each line for the terminal is given by the `co` numeric capability. If the terminal is a CRT, then the number of lines on the screen is given by the `li` capability. If the terminal wraps around to the beginning of the next line when it reaches the right margin, then it should have the `am` capability. If the terminal can clear its screen, then this is given by the `cl` string capability. If the terminal can backspace, then it should have the `bs` capability, unless a backspace is accomplished by a character other than `^H` (ugh) in which case you should give this character as the `bc` string capability. If it overstrikes (rather than clearing a position when a character is struck over) then it should have the `os` capability.

A very important point here is that the local cursor motions encoded in *termcap* are undefined at the left and top edges of a CRT terminal. The editor will never attempt to backspace around the left edge, nor will it attempt to go up locally off the top. The editor assumes that feeding off the bottom of the screen will cause the screen to scroll up, and the `am` capability tells whether the cursor sticks at the right edge of the screen. If the terminal has switch selectable automatic margins, the *termcap* file usually assumes that this is on, that is, `am`.

These capabilities suffice to describe hardcopy and "glass-tty" terminals. Thus the model 33 teletype is described as

```
t3|33|tty33:co#72:os
```

while the Lear Siegler ADM-3 is described as

```
cl|adm3|3|lsi adm3:am:bs:cl=~Z:li#24:co#80
```

Cursor addressing

Cursor addressing in the terminal is described by a **cm** string capability, with *printf*(3S) like escapes **%x** in it. These substitute to encodings of the current line or column position, while other characters are passed through unchanged. If the **cm** string is thought of as being a function, then its arguments are the line and then the column to which motion is desired, and the **%** encodings have the following meanings:

%d	as in <i>printf</i> , 0 origin
%2	like %2d
%3	like %3d
%.	like %c
%+x	adds <i>x</i> to value, then % .
%>xy	if value > <i>x</i> adds <i>y</i> , no output.
%r	reverses order of line and column, no output
%i	increments line/column (for 1 origin)
%%	gives a single %
%n	exclusive or row and column with 0140 (DM2500)
%B	BCD (16*(<i>x</i> /10) + (<i>x</i> %10)), no output.
%D	Reverse coding ($x-2*(x\%16)$), no output. (Delta Data).

Consider the HP2645, which, to get to row 3 and column 12, needs to be sent `\E&a12c03Y` padded for 6 milliseconds. Note that the order of the rows and columns is inverted here, and that the row and column are printed as two digits. Thus its **cm** capability is `"cm=6\E&%r%2c%2Y"`. The Microterm ACT-IV needs the current row and column sent preceded by a `^T`, with the row and column simply encoded in binary, `"cm=^T%.%."`. Terminals which use `"%.."` need to be able to backspace the cursor (**bs** or **bc**), and to move the cursor up one line on the screen (**up** introduced below). This is necessary because it is not always safe to transmit `\t`, `\n ^D` and `\r`, as the system may change or discard them.

A final example is the LSI ADM-3a, which uses row and column offset by a blank character, thus `"cm=\E=%+ %+"`.

Cursor motions

If the terminal can move the cursor one position to the right, leaving the character at the current position unchanged, then this sequence should be given as **nd** (non-destructive space). If it can move the cursor up a line on the screen in the same column, this should be given as **up**. If the terminal has no cursor addressing capability, but can home the cursor (to very upper left corner of screen) then this can be given as **ho**; similarly a fast way of getting to the lower left hand corner can be given as **ll**; this may involve going up with **up** from the home position, but the editor will never do this itself (unless **ll** does) because it makes no assumption about the effect of moving up from the home position.

Area clears

If the terminal can clear from the current position to the end of the line, leaving the cursor where it is, this should be given as **ce**. If the terminal can clear from the current position to the end of the display, then this should be given as **cd**. The editor only uses **cd** from the first column of a line.

Insert/delete line

If the terminal can open a new blank line before the line where the cursor is, this should be given as **al**; this is done only from the first position of a line. The cursor must then appear on the newly blank line. If the terminal can delete the line which the cursor is on, then this should be given as **dl**; this is done only from the first position on the line to be deleted. If the terminal can scroll the screen backwards, then this can be given as **sb**, but just **al** suffices. If the terminal can retain display memory above then the **da** capability should be given; if display memory can be retained below then **db** should be given. These let the editor understand that deleting a line on the screen may bring non-blank lines up from below or that scrolling back with **sb** may bring down non-blank lines.

Insert/delete character

There are two basic kinds of intelligent terminals with respect to insert/delete character which can be described using *termcap*. The most common insert/delete character operations affect only the characters on the current line and shift characters off the end of the line rigidly. Other terminals, such as the Concept 100 and the Perkin Elmer Owl, make a distinction between typed and untyped blanks on the screen, shifting upon an insert or delete only to an untyped blank on the screen which is either eliminated, or expanded to two untyped blanks. You can find out which kind of terminal you have by clearing the screen and then typing text separated by cursor motions. Type “abc def” using local cursor motions (not spaces) between the “abc” and the “def”. Then position the cursor before the “abc” and put the terminal in insert mode. If typing characters causes the rest of the line to shift rigidly and characters to fall off the end, then your terminal does not distinguish between blanks and untyped positions. If the “abc” shifts over to the “def” which then move together around the end of the current line and onto the next as you insert, you have the second type of terminal, and should give the capability **in**, which stands for “insert null”. If your terminal does something different and unusual then you may have to modify the editor to get it to use the insert mode your terminal defines. We have seen no terminals which have an insert mode not falling into one of these two classes.

The editor can handle both terminals which have an insert mode, and terminals which send a simple sequence to open a blank position on the current line. Give as **im** the sequence to get into insert mode, or give it an empty value if your terminal uses a sequence to insert a blank position. Give as **ei** the sequence to leave insert mode (give this, with an empty value also if you gave **im** so). Now give as **ic** any sequence needed to be sent just before sending the character to be inserted. Most terminals with a true insert mode will not give **ic**, terminals which send a sequence to open a screen position should give it here. (Insert mode is preferable to the sequence to open a position on the screen if your terminal has both.) If post insert padding is needed, give this as a number of milliseconds in **ip** (a string option). Any other sequence which may need to be sent after an insert of a single character may also be given in **ip**.

It is occasionally necessary to move around while in insert mode to delete characters on the same line (for example, if there is a tab after the insertion position). If your terminal allows motion while in insert mode you can give the capability **mi** to speed up inserting in this case. Omitting **mi** will affect only speed. Some terminals (notably Datamedia's) must not have **mi** because of the way their insert mode works.

Finally, you can specify delete mode by giving **dm** and **ed** to enter and exit delete mode, and **dc** to delete a single character while in delete mode.

Highlighting, underlining, and visible bells

If your terminal has sequences to enter and exit standout mode these can be given as **so** and **se** respectively. If there are several flavors of standout mode (such as inverse video, blinking, or underlining – half bright is not usually an acceptable “standout” mode unless the terminal is in inverse video mode constantly) the preferred mode is inverse video by itself. If the code to change into or out of standout mode leaves one or even two blank spaces on the screen, as the TVI 912 and Teleray 1061 do, then **sg** should be given to tell how many spaces are left.

Codes to begin underlining and end underlining can be given as **us** and **ue** respectively. If they leave blank spaces on the screen, set **ug**. If the terminal has a code to underline the current character and move the cursor one space to the right, such as the Microterm Mime, this can be given as **uc**. (If the underline code does not move the cursor to the right, give the code followed by a nondestructive space.)

Many terminals, such as the HP 2621, automatically leave standout mode when they move to a new line or the cursor is addressed. Programs using standout mode should exit standout mode before moving the cursor or sending a newline.

If the terminal has a way of flashing the screen to indicate an error quietly (a bell replacement) then this can be given as **vb**; it must not move the cursor. If the terminal should be placed in a different mode during open and visual modes of **ex**, this can be given as **vs** and **ve**, sent at the start and end of these modes respectively. These can be used to change, for example, from an underline to a block cursor and back.

If the terminal needs to be in a special mode when running a program that addresses the cursor, the codes to enter and exit this mode can be given as **ti** and **te**. This arises, for example, from terminals like the Concept with more than one page of memory. If the terminal has only memory relative cursor addressing and not screen relative cursor addressing, a one screen-sized window must be fixed into the terminal for cursor addressing to work properly.

If your terminal correctly generates underlined characters (with no special codes needed) even though it does not overstrike, then you should give the capability **ul**. If overstrikes are erasable with a blank, then this should be indicated by giving **eo**.

ANSI terminals have modes for the character highlighting. Dim characters may be generated in dim mode, entered by **mh**; reverse video characters in reverse mode, entered by **mr**; bold characters in bold mode, entered by **md**; and normal mode characters restored by turning off all attributes with **me**.

Keypad

If the terminal has a keypad that transmits codes when the keys are pressed, this information can be given. Note that it is not possible to handle terminals where the keypad only works in local (this applies, for example, to the unshifted HP 2621 keys). If the keypad can be set to transmit or not transmit, give these codes as **ks** and **ke**. Otherwise the keypad is assumed to always transmit. The codes sent by the left arrow, right arrow, up arrow, down arrow, and home keys can be given as **kl**, **kr**, **ku**, **kd**, and **kh** respectively. If there are function keys such as **f0**, **f1**, ..., **f9**, the codes they send can be given as **k0**, **k1**, ..., **k9**. If these keys have labels other than the default **f0** through **f9**, the labels can be given as **l0**, **l1**, ..., **l9**. If there are other keys that transmit the same code as the terminal expects for the corresponding function, such as clear screen, the *termcap* 2 letter codes can be given in the **ko** capability, for example, `“:ko=cl,ll,sf,sb:”`, which says that the terminal has clear, home down, scroll down, and scroll up keys that transmit the same thing as the **cl**, **ll**, **sf**, and **sb** entries.

The **ma** entry is also used to indicate arrow keys on terminals which have single character arrow keys. It is obsolete but still in use in version 2 of **vi**, which must be run on some minicomputers due to memory limitations. This field is redundant with **kl**, **kr**, **ku**, **kd**, and **kh**. It consists of groups of two characters. In each group, the first character is what an arrow key sends, the second character is the corresponding **vi** command. These commands are **h** for **kl**, **j** for **kd**, **k** for **ku**, **l** for **kr**, and **H** for **kh**. For example, the mime would be `:ma=^Kj^Zk^Xl`: indicating arrow keys left (^H), down (^K), up (^Z), and right (^X). (There is no home key on the mime.)

Miscellaneous

If the terminal requires other than a null (zero) character as a pad, then this can be given as **pc**.

If tabs on the terminal require padding, or if the terminal uses a character other than ^I to tab, then this can be given as **ta**.

Hazeltine terminals, which don't allow “” characters to be printed should indicate **hz**. Datamedia terminals, which echo carriage-return linefeed for carriage return and then ignore a following linefeed should indicate **nc**. Early Concept terminals, which ignore a linefeed immediately after an **am** wrap, should indicate **xn**. If an erase-eol is required to get rid of standout (instead of merely writing on top of it), **xs** should be given. Teleray terminals, where tabs turn all characters moved over to blanks, should indicate **xt**. Other specific terminal problems may be corrected by adding more capabilities of the form **xx**.

Other capabilities include **is**, an initialization string for the terminal, and **if**, the name of a file containing long initialization strings. These strings are expected to properly clear and then set the tabs on the terminal, if the terminal has settable tabs. If both are given, **is** will be printed before **if**. This is useful where **if** is `/usr/lib/tabset/std` but **is** clears the tabs first.

Similar Terminals

If there are two very similar terminals, one can be defined as being just like the other with certain exceptions. The string capability **tc** can be given with the name of the similar terminal. This capability must be *last* and the combined length of the two entries must not exceed 1024. Since *termlib* routines search the entry from left to right, and since the **tc** capability is replaced by the corresponding entry, the capabilities

given at the left override the ones in the similar terminal. A capability can be canceled with **xx@** where **xx** is the capability. For example, the entry

```
hn|2621nl:ks@:ke@:tc=2621:
```

defines a `2621nl` that does not have the `ks` or `ke` capabilities, and hence does not turn on the function key labels when in visual mode. This is useful for different modes for a terminal, or for different user preferences.

FILES

`/etc/termcap` file containing terminal descriptions

SEE ALSO

`ex(1)`, `curses(3X)`, `termcap(3X)`, `tset(1)`, `vi(1)`, `ul(1)`, `more(1)`

BUGS

Ex allows only 256 characters for string capabilities, and the routines in `termcap(3X)` do not check for overflow of this buffer. The total length of a single entry (excluding only escaped newlines) may not exceed 1024.

The `ma`, `vs`, and `ve` entries are specific to the `vi` program.

Not all programs support all entries. There are entries that are not supported by any program.

NAME

terminfo – terminal capability data base

SYNOPSIS

/usr/5lib/terminfo/*/*

DESCRIPTION

terminfo is a data base describing terminals, used by *curses*(3V). Terminals are described in *terminfo* by giving a set of capabilities which they have, and by describing how operations are performed. Padding requirements and initialization sequences are included in *terminfo*.

Entries in *terminfo* consist of a number of ‘,’ separated fields. White space after each ‘,’ is ignored. The first entry for each terminal gives the names which are known for the terminal, separated by ‘|’ characters. The first name given is the most common abbreviation for the terminal, the last name given should be a long name fully identifying the terminal, and all others are understood as synonyms for the terminal name. All names but the last should be in lower case and contain no blanks; the last name may well contain upper case and blanks for readability.

Terminal names (except for the last, verbose entry) should be chosen using the following conventions. The particular piece of hardware making up the terminal should have a root name chosen, thus “hp2621”. This name should not contain hyphens, except that synonyms may be chosen that do not conflict with other names. Modes that the hardware can be in, or user preferences, should be indicated by appending a hyphen and an indicator of the mode. Thus, a VT100 in 132 column mode would be vt100-w. The following suffixes should be used where possible:

Suffix	Meaning	Example
-w	Wide mode (more than 80 columns)	vt100-w
-am	With auto. margins (usually default)	vt100-am
-nam	Without automatic margins	vt100-nam
-n	Number of lines on the screen	aaa-60
-na	No arrow keys (leave them in local)	c100-na
-np	Number of pages of memory	c100-4p
-rv	Reverse video	c100-rv

CAPABILITIES

The variable is the name by which the programmer (at the terminfo level) accesses the capability. The cap-name is the short name used in the text of the database, and is used by a person updating the database. The i.code is the two letter internal code used in the compiled database, and always corresponds to the old *termcap* capability name.

Capability names have no hard length limit, but an informal limit of 5 characters has been adopted to keep them short and to allow the tabs in the source file *caps* to line up nicely. Whenever possible, names are chosen to be the same as or similar to the ANSI X3.64-1979 standard. Semantics are also intended to match those of the specification.

- (P) indicates that padding may be specified
- (G) indicates that the string is passed through *tparm* with parms as given (*#i*).
- (*) indicates that padding may be based on the number of lines affected
- (#.i) indicates the *i*th parameter.

Variable Booleans	Cap- name	I. Code	Description
auto_left_margin,	bw	bw	cub1 wraps from column 0 to last column
auto_right_margin,	am	am	Terminal has automatic margins
beehive_glitch,	xs	xb	Beehive (f1=escape, f2=ctrl C)
ceol_standout_glitch,	xhp	xs	Standout not erased by overwriting (hp)

eat_newline_glitch,	xenl	xn	newline ignored after 80 cols (Concept)
erase_overstrike,	eo	eo	Can erase overstrikes with a blank
generic_type,	gn	gn	Generic line type (e.g., dialup, switch).
hard_copy,	hc	hc	Hardcopy terminal
has_meta_key,	km	km	Has a meta key (shift, sets parity bit)
has_status_line,	hs	hs	Has extra "status line"
insert_null_glitch,	in	in	Insert mode distinguishes nulls
memory_above,	da	da	Display may be retained above the screen
memory_below,	db	db	Display may be retained below the screen
move_insert_mode,	mir	mi	Safe to move while in insert mode
move_standout_mode,	msgr	ms	Safe to move in standout modes
over_strike,	os	os	Terminal overstrikes
status_line_esc_ok,	eslok	es	Escape can be used on the status line
teleray_glitch,	xt	xt	Tabs ruin, magic so char (Teleray 1061)
tilde_glitch,	hz	hz	Hazeltine; can not print ~'s
transparent_underline,	ul	ul	underline character overstrikes
xon_xoff,	xon	xo	Terminal uses xon/xoff handshaking
Numbers:			
columns,	cols	co	Number of columns in a line
init_tabs,	it	it	Tabs initially every # spaces
lines,	lines	li	Number of lines on screen or page
lines_of_memory,	lm	lm	Lines of memory if > lines. 0 means varies
magic_cookie_glitch,	xmc	sg	Number of blank chars left by smso or rmso
padding_baud_rate,	pb	pb	Lowest baud where cr/nl padding is needed
virtual_terminal,	vt	vt	Virtual terminal number (UNIX system)
width_status_line,	wsl	ws	No. columns in status line
Strings:			
back_tab,	cbt	bt	Back tab (P)
bell,	bel	bl	Audible signal (bell) (P)
carriage_return,	cr	cr	Carriage return (P*)
change_scroll_region,	csr	cs	change to lines #1 through #2 (VT100) (PG)
clear_all_tabs,	tbc	ct	Clear all tab stops (P)
clear_screen,	clear	cl	Clear screen and home cursor (P*)
clr_eol,	el	ce	Clear to end of line (P)
clr_eos,	ed	cd	Clear to end of display (P*)
column_address,	hpa	ch	Set cursor column (PG)
command_character,	cmdch	CC	Term. settable cmd char in prototype
cursor_address,	cup	cm	Screen rel. cursor motion row #1 col #2 (PG)
cursor_down,	cudl	do	Down one line
cursor_home,	home	ho	Home cursor (if no cup)

cursor_invisible,	civis	vi	Make cursor invisible
cursor_left,	cub1	le	Move cursor left one space
cursor_mem_address,	mrcup	CM	Memory relative cursor addressing
cursor_normal,	cnorm	ve	Make cursor appear normal (undo vs/vi)
cursor_right,	cuf1	nd	Non-destructive space (cursor right)
cursor_to_ll,	ll	ll	Last line, first column (if no cup)
cursor_up,	cuu1	up	Upline (cursor up)
cursor_visible,	cvvis	vs	Make cursor very visible
delete_character,	dch1	dc	Delete character (P*)
delete_line,	dll	dl	Delete line (P*)
dis_status_line,	dsl	ds	Disable status line
down_half_line,	hd	hd	Half-line down (forward 1/2 linefeed)
enter_alt_charset_mode,	smacs	as	Start alternate character set (P)
enter_blink_mode,	blink	mb	Turn on blinking
enter_bold_mode,	bold	md	Turn on bold (extra bright) mode
enter_ca_mode,	smcup	ti	String to begin programs that use cup
enter_delete_mode,	smdc	dm	Delete mode (enter)
enter_dim_mode,	dim	mh	Turn on half-bright mode
enter_insert_mode,	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 stand out mode
enter_underline_mode,	smul	us	Start underscore mode
erase_chars	ech	ec	Erase #1 characters (PG)
exit_alt_charset_mode,	rmacs	ae	End alternate character set (P)
exit_attribute_mode,	sgr0	me	Turn off all attributes
exit_ca_mode,	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 stand out mode
exit_underline_mode,	rmul	ue	End underscore mode
flash_screen,	flash	vb	Visible bell (may not move cursor)
form_feed,	ff	ff	Hardcopy terminal page eject (P*)
from_status_line,	fsl	fs	Return from status line
init_1string,	is1	i1	Terminal initialization string
init_2string,	is2	i2	Terminal initialization string
init_3string,	is3	i3	Terminal initialization string
init_file,	if	if	Name of file containing is
insert_character,	ich1	ic	Insert character (P)
insert_line,	ill	al	Add new blank line (P*)
insert_padding,	ip	ip	Insert pad after character inserted (p*)
key_backspace,	kbs	kb	Sent by backspace key
key_catab,	ktbc	ka	Sent by clear-all-tabs key
key_clear,	kclr	kC	Sent by clear screen or erase key
key_ctab,	kctab	kt	Sent by clear-tab key
key_dc,	kdch1	kD	Sent by delete character key
key_dl,	kdll	kL	Sent by delete line key
key_down,	kcud1	kd	Sent by terminal down arrow key
key_eic,	krmir	kM	Sent by rmir or smir in insert mode
key_eol,	kel	kE	Sent by clear-to-end-of-line key
key_eos,	ked	kS	Sent by clear-to-end-of-screen key

key_f0,	kf0	k0	Sent by function key f0
key_f1,	kf1	k1	Sent by function key f1
key_f10,	kf10	ka	Sent by function key f10
key_f2,	kf2	k2	Sent by function key f2
key_f3,	kf3	k3	Sent by function key f3
key_f4,	kf4	k4	Sent by function key f4
key_f5,	kf5	k5	Sent by function key f5
key_f6,	kf6	k6	Sent by function key f6
key_f7,	kf7	k7	Sent by function key f7
key_f8,	kf8	k8	Sent by function key f8
key_f9,	kf9	k9	Sent by function key f9
key_home,	khome	kh	Sent by home key
key_ic,	kich1	kl	Sent by ins char/enter ins mode key
key_il,	kil1	kA	Sent by insert line
key_left,	kcub1	kl	Sent by terminal left arrow key
key_ll,	kl1	kH	Sent by home-down key
key_npage,	knp	kN	Sent by next-page key
key_ppage,	kpp	kP	Sent by previous-page key
key_right,	kcufl	kr	Sent by terminal right arrow key
key_sf,	kind	kF	Sent by scroll-forward/down key
key_sr,	kri	kR	Sent by scroll-backward/up key
key_stab,	khts	kT	Sent by set-tab key
key_up,	kcuu1	ku	Sent by terminal up arrow key
keypad_local,	rmkx	ke	Out of "keypad transmit" mode
keypad_xmit,	smkx	ks	Put terminal in "keypad transmit" mode
lab_f0,	lf0	l0	Labels on function key f0 if not f0
lab_f1,	lf1	l1	Labels on function key f1 if not f1
lab_f10,	lf10	la	Labels on function key f10 if not f10
lab_f2,	lf2	l2	Labels on function key f2 if not f2
lab_f3,	lf3	l3	Labels on function key f3 if not f3
lab_f4,	lf4	l4	Labels on function key f4 if not f4
lab_f5,	lf5	l5	Labels on function key f5 if not f5
lab_f6,	lf6	l6	Labels on function key f6 if not f6
lab_f7,	lf7	l7	Labels on function key f7 if not f7
lab_f8,	lf8	l8	Labels on function key f8 if not f8
lab_f9,	lf9	l9	Labels on function key f9 if not f9
meta_on,	smm	mm	Turn on "meta mode" (8th bit)
meta_off,	rmm	mo	Turn off "meta mode"
newline,	nel	nw	Newline (behaves like cr followed by lf)
pad_char,	pad	pc	Pad character (rather than null)
parm_dch,	dch	DC	Delete #1 chars (PG*)
parm_delete_line,	dl	DL	Delete #1 lines (PG*)
parm_down_cursor,	cud	DO	Move cursor down #1 lines (PG*)
parm_ich,	ich	IC	Insert #1 blank chars (PG*)
parm_index,	indn	SF	Scroll forward #1 lines (PG)
parm_insert_line,	il	AL	Add #1 new blank lines (PG*)
parm_left_cursor,	cub	LE	Move cursor left #1 spaces (PG)
parm_right_cursor,	cuf	RI	Move cursor right #1 spaces (PG*)
parm_rindex,	rin	SR	Scroll backward #1 lines (PG)
parm_up_cursor,	cuu	UP	Move cursor up #1 lines (PG*)
pkey_key,	pfkey	pk	Prog funct key #1 to type string #2
pkey_local,	pfloc	pl	Prog funct key #1 to execute string #2

pkey_xmit,	pfx	px	Prog funct key #1 to xmit string #2
print_screen,	mc0	ps	Print contents of the screen
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. (PG*)
reset_lstring,	rs1	r1	Reset terminal completely to sane modes.
reset_2string,	rs2	r2	Reset terminal completely to sane modes.
reset_3string,	rs3	r3	Reset terminal completely to sane modes.
reset_file,	rf	rf	Name of file containing reset string
restore_cursor,	rc	rc	Restore cursor to position of last sc
row_address,	vpa	cv	Vertical position absolute (set row) (PG)
save_cursor,	sc	sc	Save cursor position (P)
scroll_forward,	ind	sf	Scroll text up (P)
scroll_reverse,	ri	sr	Scroll text down (P)
set_attributes,	sgr	sa	Define the video attributes (PG9)
set_tab,	hts	st	Set a tab in all rows, current column
set_window,	wind	wi	Current window is lines #1-#2 cols #3-#4
tab,	ht	ta	Tab to next 8 space hardware tab stop
to_status_line,	tsl	ts	Go to status line, column #1
underline_char,	uc	uc	Underscore one char and move past it
up_half_line,	hu	hu	Half-line up (reverse 1/2 linefeed)
init_prog,	ipro	iP	Path name of program for init
key_a1,	ka1	K1	Upper left of keypad
key_a3,	ka3	K3	Upper right of keypad
key_b2,	kb2	K2	Center of keypad
key_c1,	kc1	K4	Lower left of keypad
key_c3,	kc3	K5	Lower right of keypad
prtr_non,	mc5p	pO	Turn on the printer for #1 bytes

A Sample Entry

The following entry, which describes the Concept 100, is among the more complex entries in the *terminfo* file as of this writing.

```
concept100 | c100 | concept | c104 | c100-4p | concept 100,
am, bel=^G, blank=\EH, blink=\EC, clear=^L$<2*>, cnorm=\Ew,
cols#80, cr=^M$<9>, cubl=^H, cudl=^J, cuf1=\E=,
cup=\Ea%p1%' '%+%c%p2%' '%+%c,
cuul=\E;, cvvis=\EW, db, dchl=\E^A$<16*>, dim=\EE, dll=\E^B$<3*>,
ed=\E^C$<16*>, el=\E^U$<16>, eo, flash=\Ek$<20>\EK, ht=\t$<8>,
ill=\E^R$<3*>, in, ind=^J, .ind=^J$<9>, ip=$<16*>,
is2=\EU\Ef\E7\E5\E8\E1\ENH\EK\E\200\Eo&\200\Eo\47\E,
kbs=^h, kcubl=\E>, kcudl=\E<, kcuf1=\E=, kcuul=\E;,
kfl=\E5, kf2=\E6, kf3=\E7, khome=\E?,
lines#24, mir, pb#9600, prot=\EI, rep=\Er%p1%' '%+%c$<.2*>,
rev=\ED, rmcup=\Ev $<6>\Ep\r\n, rmir=\E\200, rmkx=\Ex,
rmso=\Ed\Ee, rmul=\Eg, rmul=\Eg, sgr0=\EN\200,
smcup=\EU\Ev 8p\Ep\r, smir=\E^P, smkx=\EX, smso=\EE\ED,
smul=\EG, tabs, ul, vt#8, xenl,
```

Entries may continue onto multiple lines by placing white space at the beginning of each line except the first. Comments may be included on lines beginning with “#”. Capabilities in *terminfo* are of three types: Boolean capabilities which indicate that the terminal has some particular feature, numeric capabilities giving the size of the terminal or the size of particular delays, and string capabilities, which give a sequence

which can be used to perform particular terminal operations.

Types of Capabilities

All capabilities have names. For instance, the fact that the Concept has *automatic margins* (i.e., an automatic return and linefeed when the end of a line is reached) is indicated by the capability **am**. Hence the description of the Concept includes **am**. Numeric capabilities are followed by the character '#' and then the value. Thus **cols**, which indicates the number of columns the terminal has, gives the value '80' for the Concept.

Finally, string valued capabilities, such as **el** (clear to end of line sequence) are given by the two-character code, an '=', and then a string ending at the next following ','. A delay in milliseconds may appear anywhere in such a capability, enclosed in \$<.> brackets, as in **el=\EK\$<3>**, and padding characters are supplied by *tputs* to provide this delay. The delay can be either a number, e.g., '20', or a number followed by an '*', i.e., '3*'. A '*' indicates that the padding required is proportional to the number of lines affected by the operation, and the amount given is the per-affected-unit padding required. (In the case of insert character, the factor is still the number of *lines* affected. This is always one unless the terminal has **xenl** and the software uses it.) When a '*' is specified, it is sometimes useful to give a delay of the form '3.5' to specify a delay per unit to tenths of milliseconds. (Only one decimal place is allowed.)

A number of escape sequences are provided in the string valued capabilities for easy encoding of characters there. Both **\E** and **\e** map to an ESCAPE character, **\x** maps to a control-x for any appropriate x, and the sequences **\n** **\l** **\r** **\t** **\b** **\f** **\s** give a newline, linefeed, return, tab, backspace, formfeed, and space. Other escapes include **\^** for **^**, **\,** for **,**, **\:** for **:**, and **\0** for null. (**\0** will produce **\200**, which does not terminate a string but behaves as a null character on most terminals.) Finally, characters may be given as three octal digits after a ****.

Sometimes individual capabilities must be commented out. To do this, put a period before the capability name. For example, see the second **ind** in the example above.

Preparing Descriptions

We now outline how to prepare descriptions of terminals. The most effective way to prepare a terminal description is by imitating the description of a similar terminal in *terminfo* and to build up a description gradually, using partial descriptions with 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 easily test a new terminal description you can set the environment variable **TERMINFO** to a pathname of a directory containing the compiled description you are working on and programs will look there rather than in **/usr/5lib/terminfo**. To get the padding for insert line right (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 terminal messes up, more padding is usually needed. A similar test can be used for insert character.

Basic Capabilities

The number of columns on each line for the terminal is given by the **cols** numeric capability. If the terminal is a CRT, then the number of lines on the screen is given by the **lines** capability. If the terminal wraps around to the beginning of the next line when it reaches the right margin, then it should have the **am** capability. If the terminal can clear its screen, leaving the cursor in the home position, then this is given by the **clear** string capability. If the terminal overstrikes (rather than clearing a position when a character is struck over) then it should have the **os** capability. If the terminal is a printing terminal, with no soft copy unit, give it both **hc** and **os**. (**os** applies to storage scope terminals, such as Tektronix 4010 series, as well as hard copy and APL terminals.) If there is a code to move the cursor to the left edge of the current row, give this as **cr**. (Normally this will be carriage return, control M.) If there is a code to produce an audible signal (bell, beep, etc) give this as **bel**.

If there is a code to move the cursor one position to the left (such as backspace) that capability should be given as **cub1**. Similarly, codes to move to the right, up, and down should be given as **cuf1**, **cuu1**, and **cud1**. These local cursor motions should not alter the text they pass over, for example, you would not normally use **'cuf1= '** because the space would erase the character moved over.

A very important point here is that the local cursor motions encoded in *terminfo* are undefined at the left and top edges of a CRT terminal. Programs should never attempt to backspace around the left edge, unless *bw* is given, and never attempt to go up locally off the top. In order to scroll text up, a program will go to the bottom left corner of the screen and send the *ind* (index) string.

To scroll text down, a program goes to the top left corner of the screen and sends the *ri* (reverse index) string. The strings *ind* and *ri* are undefined when not on their respective corners of the screen.

Parameterized versions of the scrolling sequences are *indn* and *rin* which have the same semantics as *ind* and *ri* except that they take one parameter, and scroll that many lines. They are also undefined except at the appropriate edge of the screen.

The *am* capability tells whether the cursor sticks at the right edge of the screen when text is output, but this does not necessarily apply to a *cuf1* from the last column. The only local motion which is defined from the left edge is if *bw* is given, then a *cub1* from the left edge will move to the right edge of the previous row. If *bw* is not given, the effect is undefined. This is useful for drawing a box around the edge of the screen, for example. If the terminal has switch selectable automatic margins, the *terminfo* file usually assumes that this is on; i.e., *am*. If the terminal has a command which moves to the first column of the next line, that command can be given as *nel* (newline). It does not matter if the command clears the remainder of the current line, so if the terminal has no *cr* 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 “glass-tty” terminals. Thus the model 33 teletype is described as

```
33 | tty33 | tty | model 33 teletype,
bel=^G, cols#72, cr=^M, cudl=^J, hc, ind=^J, os,
```

while the Lear Siegler ADM-3 is described as

```
adm3 | 3 | lsi adm3,
am, bel=^G, clear=^Z, cols#80, cr=^M, cub1=^H, cudl=^J,
ind=^J, lines#24,
```

Parameterized Strings

Cursor addressing and other strings requiring parameters in the terminal are described by a parameterized string capability, with *printf*(3S) like escapes *%x* in it. For example, to address the cursor, the *cup* capability is given, using two parameters: the row and column to address to. (Rows and columns are numbered from zero and refer to the physical screen visible to the user, not to any unseen memory.) If the terminal has memory relative cursor addressing, that can be indicated by *mrcup*.

The parameter mechanism uses a stack and special *%* codes to manipulate it. Typically a sequence will push one of the parameters onto the stack and then print it in some format. Often more complex operations are necessary.

The *%* encodings have the following meanings:

<i>%%</i>	outputs ‘%’
<i>%d</i>	print pop() as in printf
<i>%2d</i>	print pop() like %2d
<i>%3d</i>	print pop() like %3d
<i>%02d</i>	
<i>%03d</i>	as in printf
<i>%c</i>	print pop() gives %c
<i>%s</i>	print pop() gives %s
<i>%p[1-9]</i>	push ith parm
<i>%P[a-z]</i>	set variable [a-z] to pop()
<i>%g[a-z]</i>	get variable [a-z] and push it
<i>%'c'</i>	char constant c

<code>%{nn}</code>	integer constant nn
<code>%+ %- %* %/ %m</code>	arithmetic (%m is mod): push(pop() op pop())
<code>%& % %^</code>	bit operations: push(pop() op pop())
<code>%= %> %<</code>	logical operations: push(pop() op pop())
<code>%! %^-</code>	unary operations push(op pop())
<code>%i</code>	add 1 to first two parms (for ANSI terminals)
<code>%? expr %t thenpart %e elsepart %;</code>	if-then-else, %e elsepart is optional. else-if's are possible ala Algol 68: <code>%? c₁ %t b₁ %e c₂ %t b₂ %e c₃ %t b₃ %e c₄ %t b₄ %e %;</code> c _i are conditions, b _i are bodies.

Binary operations are in postfix form with the operands in the usual order. That is, to get x-5 one would use "`%gx%{5}%-`".

Consider the HP 2645, which, to get to row 3 and column 12, needs to be sent `\E&a12c03Y` padded for 6 milliseconds. Note that the order of the rows and columns is inverted here, and that the row and column are printed as two digits. Thus its `cup` capability is "`cup=6\E&%p2%2dc%p1%2dY`".

The Microterm ACT-IV needs the current row and column sent preceded by a `^T`, with the row and column simply encoded in binary, "`cup=^T%p1%c%p2%c`". Terminals which use "`%c`" need to be able to back-space the cursor (`cuB1`), and to move the cursor up one line on the screen (`cuu1`). This is necessary because it is not always safe to transmit `\n ^D` and `\r`, as the system may change or discard them. (The library routines dealing with terminfo set tty modes so that tabs are never expanded, so `\t` is safe to send. This turns out to be essential for the Ann Arbor 4080.)

A final example is the LSI ADM-3a, which uses row and column offset by a blank character, thus "`cup=\E=%p1%' '%+%c%p2%' '%+%c`". After sending `\E=`, this pushes the first parameter, pushes the ASCII value for a space (32), adds them (pushing the sum on the stack in place of the two previous values) and outputs that value as a character. Then the same is done for the second parameter. More complex arithmetic is possible using the stack.

If the terminal has row or column absolute cursor addressing, these can be given as single parameter capabilities `hpa` (horizontal position absolute) and `vpa` (vertical position absolute). Sometimes these are shorter than the more general two parameter sequence (as with the hp2645) and can be used in preference to `cup`. If there are parameterized local motions (e.g., move *n* spaces to the right) these can be given as `cuD`, `cuB`, `cuF`, and `cuu` with a single parameter indicating how many spaces to move. These are primarily useful if the terminal does not have `cup`, such as the Tektronix 4025.

Cursor Motions

If the terminal has a fast way to home the cursor (to very upper left corner of screen) then this can be given as `home`; similarly a fast way of getting to the lower left-hand corner can be given as `ll`; this may involve going up with `cuu1` from the home position, but a program should never do this itself (unless `ll` does) because it can make no assumption about the effect of moving up from the home position. Note that the home position is the same as addressing to (0,0): to the top left corner of the screen, not of memory. (Thus, the `\EH` sequence on HP terminals cannot be used for `home`.)

Area Clears

If the terminal can clear from the current position to the end of the line, leaving the cursor where it is, this should be given as `el`. If the terminal can clear from the current position to the end of the display, then this should be given as `ed`. `Ed` is only defined from the first column of a line. (Thus, it can be simulated by a request to delete a large number of lines, if a true `ed` is not available.)

Insert/delete line

If the terminal can open a new blank line before the line where the cursor is, this should be given as **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 **dll**; this is done only from the first position on the line to be deleted. Versions of **il1** and **dll** which take a single parameter and insert or delete that many lines can be given as **il** and **dl**. If the terminal has a settable scrolling region (like the VT100) the command to set this can be described with the **csr** capability, which takes two parameters: the top and bottom lines of the scrolling region. The cursor position is, alas, undefined after using this command. It is possible to get the effect of insert or delete line using this command – the **sc** and **rc** (save and restore cursor) commands are also useful. Inserting lines at the top or bottom of the screen can also be done using **ri** or **ind** on many terminals without a true insert/delete line, and is often faster even on terminals with those features.

If the terminal has the ability to define a window as part of memory, which all commands affect, it should be given as the parameterized string **wind**. The four parameters are the starting and ending lines in memory and the starting and ending columns in memory, in that order.

If the terminal can retain display memory above, then the **da** capability should be given; if display memory can be retained below, then **db** should be given. These indicate that deleting a line or scrolling may bring non-blank lines up from below or that scrolling back with **ri** may bring down non-blank lines.

Insert/Delete Character

There are two basic kinds of intelligent terminals with respect to insert/delete character which can be described using *terminfo*. The most common insert/delete character operations affect only the characters on the current line and shift characters off the end of the line rigidly. Other terminals, such as the Concept 100 and the Perkin Elmer Owl, make a distinction between typed and untyped blanks on the screen, shifting upon an insert or delete only to an untyped blank on the screen which is either eliminated, or expanded to two untyped blanks. You can determine the kind of terminal you have by clearing the screen and then typing text separated by cursor motions. Type “abc def” using local cursor motions (not spaces) between the “abc” and the “def”. Then position the cursor before the “abc” and put the terminal in insert mode. If typing characters causes the rest of the line to shift rigidly and characters to fall off the end, then your terminal does not distinguish between blanks and untyped positions. If the “abc” shifts over to the “def” which then move together around the end of the current line and onto the next as you insert, you have the second type of terminal, and should give the capability **in**, which stands for “insert null”. While these are two logically separate attributes (one line vs. multiline insert mode, and special treatment of untyped spaces) we have seen no terminals whose insert mode cannot be described with the single attribute.

Terminfo can describe both terminals which have an insert mode, and terminals which send a simple sequence to open a blank position on the current line. Give as **smir** the sequence to get into insert mode. Give as **rmir** the sequence to leave insert mode. Now give as **ich1** any sequence needed to be sent just before sending the character to be inserted. Most terminals with a true insert mode will not give **ich1**; terminals which send a sequence to open a screen position should give it here. (If your terminal has both, insert mode is usually preferable to **ich1**. Do not give both unless the terminal actually requires both to be used in combination.) If post insert padding is needed, give this as a number of milliseconds in **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.

It is occasionally necessary to move around while in insert mode to delete characters on the same line (e.g., if there is a tab after the insertion position). If your terminal allows motion while in insert mode you can give the capability **mir** to speed up inserting in this case. Omitting **mir** will affect only speed. Some terminals (notably Datamedia’s) must not have **mir** because of the way their insert mode works.

Finally, you can specify **dch1** to delete a single character, **dch** with one parameter, *n*, to delete *n* characters, and delete mode by giving **smdc** and **rmdc** to enter and exit delete mode (any mode the terminal needs to be placed in for **dch1** to work).

A command to erase *n* characters (equivalent to outputting *n* blanks without moving the cursor) can be given as **ech** with one parameter.

Highlighting, Underlining, and Visible Bells

If your terminal has one or more kinds of display attributes, these can be represented in a number of different ways. You should choose one display form as *standout mode*, representing a good, high contrast, easy-on-the-eyes, format for highlighting error messages and other attention getters. (If you have a choice, reverse video plus half-bright is good, or reverse video alone.) The sequences to enter and exit standout mode are given as **smso** and **rmso**, respectively. If the code to change into or out of standout mode leaves one or even two blank spaces on the screen, as the TVI 912 and Teleray 1061 do, then **xmc** should be given to tell how many spaces are left.

Codes to begin underlining and end underlining can be given as **smul** and **rmul** respectively. If the terminal has a code to underline the current character and move the cursor one space to the right, such as the Microterm Mime, this can be given as **uc**.

Other capabilities to enter various highlighting modes include **blink** (blinking) **bold** (bold or extra bright) **dim** (dim or half-bright) **invis** (blanking or invisible text) **prot** (protected) **rev** (reverse video) **sgr0** (turn off *all* attribute modes) **smacs** (enter alternate character set mode) and **rmacs** (exit alternate character set mode). Turning on any of these modes singly may or may not turn off other modes.

If there is a sequence to set arbitrary combinations of modes, this should be given as **sgr** (set attributes), taking 9 parameters. Each parameter is either 0 or 1, as the corresponding attribute is on or off. The 9 parameters are, in order: standout, underline, reverse, blink, dim, bold, blank, protect, alternate character set. Not all modes need be supported by **sgr**, only those for which corresponding separate attribute commands exist.

Terminals with the "magic cookie" glitch (**xmc**) deposit special "cookies" when they receive mode-setting sequences, which affect the display algorithm rather than having extra bits for each character. Some terminals, such as the HP 2621, automatically leave standout mode when they move to a new line or the cursor is addressed. Programs using standout mode should exit standout mode before moving the cursor or sending a newline, unless the **msg**r capability, asserting that it is safe to move in standout mode, is present.

If the terminal has a way of flashing the screen to indicate an error quietly (a bell replacement) then this can be given as **flash**; it must not move the cursor.

If the cursor needs to be made more visible than normal when it is not on the bottom line (to make, for example, a non-blinking underline into an easier to find block or blinking underline) give this sequence as **cvvis**. If there is a way to make the cursor completely invisible, give that as **civis**. The capability **cnorm** should be given which undoes the effects of both of these modes.

If the terminal needs to be in a special mode when running a program that uses these capabilities, the codes to enter and exit this mode can be given as **smcup** and **rmcup**. This arises, for example, from terminals like the Concept with more than one page of memory. If the terminal has only memory relative cursor addressing and not screen relative cursor addressing, a one screen-sized window must be fixed into the terminal for cursor addressing to work properly. This is also used for the Tektronix 4025, where **smcup** sets the command character to be the one used by terminfo.

If your terminal correctly generates underlined characters (with no special codes needed) even though it does not overstrike, then you should give the capability **ul**. If overstrikes are erasable with a blank, then this should be indicated by giving **eo**.

Keypad

If the terminal has a keypad that transmits codes when the keys are pressed, this information can be given. Note that it is not possible to handle terminals where the keypad only works in local (this applies, for example, to the unshifted HP 2621 keys). If the keypad can be set to transmit or not transmit, give these codes as **smkx** and **rmkx**. Otherwise the keypad is assumed to always transmit. The codes sent by the left arrow, right arrow, up arrow, down arrow, and home keys can be given as **kcub1**, **kcuf1**, **kcuu1**, **kcud1**, and **khome** respectively. If there are function keys such as **f0**, **f1**, ..., **f10**, the codes they send can be given as **kf0**, **kf1**, ..., **kf10**. If these keys have labels other than the default **f0** through **f10**, the labels can be given as **lf0**, **lf1**, ..., **lf10**. The codes transmitted by certain other special keys can be given: **kill** (home down), **kbs** (backspace), **ktbc** (clear all tabs), **kctab** (clear the tab stop in this column), **kclr** (clear screen or erase key), **kdch1** (delete character), **kdl1** (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), **knf** (next page), **kpp** (previous page), **kind** (scroll forward/down), **kri** (scroll backward/up), **khts** (set a tab stop in this column). In addition, if the keypad has a 3 by 3 array of keys including the four arrow keys, the other five keys can be given as **ka1**, **ka3**, **kb2**, **kc1**, and **kc3**. These keys are useful when the effects of a 3 by 3 directional pad are needed.

Tabs and Initialization

If the terminal has hardware tabs, the command to advance to the next tab stop can be given as **ht** (usually control I). A "backtab" command which moves leftward to the next tab stop can be given as **cbt**. By convention, if the teletype modes indicate that tabs are being expanded by the computer rather than being sent to the terminal, programs should not use **ht** or **cbt** even if they are present, since the user may not have the tab stops properly set. If the terminal has hardware tabs which are initially set every *n* spaces when the terminal is powered up, the numeric parameter *n* is given, showing the number of spaces the tabs are set to. This is normally used by the **tset** command to determine whether to set the mode for hardware tab expansion, and whether to set the tab stops. If the terminal has tab stops that can be saved in nonvolatile memory, the terminfo description can assume that they are properly set.

Other capabilities include **is1**, **is2**, and **is3**, initialization strings for the terminal, **iprog**, the path name of a program to be run to initialize the terminal, and **if**, the name of a file containing long initialization strings. These strings are expected to set the terminal into modes consistent with the rest of the terminfo description. They are normally sent to the terminal, by the **tset** program, each time the user logs in. They will be printed in the following order: **is1**; **is2**; setting tabs using **tbc** and **hts**; **if**; running the program **iprog**; and finally **is3**. Most initialization is done with **is2**. Special terminal modes can be set up without duplicating strings by putting the common sequences in **is2** and special cases in **is1** and **is3**. A pair of sequences that does a harder reset from a totally unknown state can be analogously given as **rs1**, **rs2**, **rf**, and **rs3**, analogous to **is2** and **if**. These strings are output by the **reset** program, which is used when the terminal gets into a wedged state. Commands are normally placed in **rs2** and **rf** only if they produce annoying effects on the screen and are not necessary when logging in. For example, the command to set the VT100 into 80-column mode would normally be part of **is2**, but it causes an annoying glitch of the screen and is not normally needed since the terminal is usually already in 80 column mode.

If there are commands to set and clear tab stops, they can be given as **tbc** (clear all tab stops) and **hts** (set a tab stop in the current column of every row). If a more complex sequence is needed to set the tabs than can be described by this, the sequence can be placed in **is2** or **if**.

Delays

Certain capabilities control padding in the teletype driver. These are primarily needed by hard copy terminals, and are used by the **tset** program to set teletype modes appropriately. Delays embedded in the capabilities **cr**, **ind**, **cub1**, **ff**, and **tab** will cause the appropriate delay bits to be set in the teletype driver. If **pb** (padding baud rate) is given, these values can be ignored at baud rates below the value of **pb**.

Miscellaneous

If the terminal requires other than a null (zero) character as a pad, then this can be given as **pad**. Only the first character of the **pad** string is used.

If the terminal has an extra “status line” that is not normally used by software, this fact can be indicated. If the status line is viewed as an extra line below the bottom line, into which one can cursor address normally (such as the Heathkit H19’s 25th line, or the 24th line of a VT100 which is set to a 23-line scrolling region), the capability **hs** should be given. Special strings to go to the beginning of the status line and to return from the status line can be given as **tsl** and **fsl**. (**fsl** must leave the cursor position in the same place it was before **tsl**. If necessary, the **sc** and **rc** strings can be included in **tsl** and **fsl** to get this effect.) The parameter **tsl** takes one parameter, which is the column number of the status line the cursor is to be moved to. If escape sequences and other special commands, such as **tab**, work while in the status line, the flag **eslok** can be given. A string which turns off the status line (or otherwise erases its contents) should be given as **dsl**. If the terminal has commands to save and restore the position of the cursor, give them as **sc** and **rc**. The status line is normally assumed to be the same width as the rest of the screen, e.g., **cols**. If the status line is a different width (possibly because the terminal does not allow an entire line to be loaded) the width, in columns, can be indicated with the numeric parameter **wsl**.

If the terminal can move up or down half a line, this can be indicated with **hu** (half-line up) and **hd** (half-line down). This is primarily useful for superscripts and subscripts on hardcopy terminals. If a hardcopy terminal can eject to the next page (form feed), give this as **ff** (usually control L).

If there is a command to repeat a given character a given number of times (to save time transmitting a large number of identical characters) this can be indicated with the parameterized string **rep**. The first parameter is the character to be repeated and the second is the number of times to repeat it. Thus, **tparam(repeat_char, 'x', 10)** is the same as **'xxxxxxxxxx'**.

If the terminal has a settable command character, such as the Tektronix 4025, this can be indicated with **cmdch**. A prototype command character is chosen which is used in all capabilities. This character is given in the **cmdch** capability to identify it. The following convention is supported on some UNIX systems: The environment is to be searched for a **CC** variable, and if found, all occurrences of the prototype character are replaced with the character in the environment variable.

Terminal descriptions that do not represent a specific kind of known terminal, such as *switch*, *dialup*, *patch*, and *network*, should include the **gn** (generic) capability so that programs can complain that they do not know how to talk to the terminal. (This capability does not apply to *virtual* terminal descriptions for which the escape sequences are known.)

If the terminal uses **xon/xoff** handshaking for flow control, give **xon**. Padding information should still be included so that routines can make better decisions about costs, but actual pad characters will not be transmitted.

If the terminal has a “meta key” which acts as a shift key, setting the 8th bit of any character transmitted, this fact can be indicated with **km**. Otherwise, software will assume that the 8th bit is parity and it will usually be cleared. If strings exist to turn this “meta mode” on and off, they can be given as **smm** and **rmm**.

If the terminal has more lines of memory than will fit on the screen at once, the number of lines of memory can be indicated with **lm**. A value of **lm#0** indicates that the number of lines is not fixed, but that there is still more memory than fits on the screen.

If the terminal is one of those supported by the UNIX virtual terminal protocol, the terminal number can be given as **vt**.

Media copy strings which control an auxiliary printer connected to the terminal can be given as **mc0**: print the contents of the screen, **mc4**: turn off the printer, and **mc5**: turn on the printer. When the printer is on, all text sent to the terminal will be sent to the printer. It is undefined whether the text is also displayed on the terminal screen when the printer is on. A variation **mc5p** takes one parameter, and leaves the printer on for as many characters as the value of the parameter, then turns the printer off. The parameter should not exceed 255. All text, including **mc4**, is transparently passed to the printer while an **mc5p** is in effect.

Strings to program function keys can be given as **pfkey**, **pfloc**, and **pfx**. Each of these strings takes two parameters: the function key number to program (from 0 to 10) and the string to program it with. Function key numbers out of this range may program undefined keys in a terminal dependent manner. The

difference between the capabilities is that **pfkey** causes pressing the given key to be the same as the user typing the given string; **pfloc** causes the string to be executed by the terminal in local; and **pfx** causes the string to be transmitted to the computer.

Glitches and Braindamage

Hazeltine terminals, which do not allow ‘” characters to be displayed should indicate **hz**.

Terminals which ignore a linefeed immediately after an am wrap, such as the Concept and VT100, should indicate **xenl**.

If **el** is required to get rid of standout (instead of merely writing normal text on top of it), **xhp** should be given.

Teleray terminals, where tabs turn all characters moved over to blanks, should indicate **xt** (destructive tabs). This glitch is also taken to mean that it is not possible to position the cursor on top of a “magic cookie”, that to erase standout mode it is instead necessary to use delete and insert line.

The Beehive Superbee, which is unable to correctly transmit the escape or control C characters, has **xsb**, indicating that the f1 key is used for escape and f2 for control C. (Only certain Superbees have this problem, depending on the ROM.)

Other specific terminal problems may be corrected by adding more capabilities of the form **xx**.

Similar Terminals

If there are two very similar terminals, one can be defined as being just like the other with certain exceptions. The string capability **use** can be given with the name of the similar terminal. The capabilities given before **use** override those in the terminal type invoked by **use**. A capability can be cancelled by placing **xx@** to the left of the capability definition, where **xx** is the capability. For example, the entry

```
2621-nl, smkx@, rmkx@, use=2621,
```

defines a 2621-nl that does not have the **smkx** or **rmkx** capabilities, and hence does not turn on the function key labels when in visual mode. This is useful for different modes for a terminal, or for different user preferences.

FILES

`/usr/5lib/terminfo/?/*` files containing terminal descriptions

SEE ALSO

`curses(3V)`, `printf(3S)`

NAME

terminfo – terminal capability data base

SYNOPSIS

/usr/5lib/terminfo/*/*

DESCRIPTION

terminfo is a data base describing terminals, used by *curses*(3V). Terminals are described in *terminfo* by giving a set of capabilities which they have, and by describing how operations are performed. Padding requirements and initialization sequences are included in *terminfo*.

Entries in *terminfo* consist of a number of ‘,’ separated fields. White space after each ‘,’ is ignored. The first entry for each terminal gives the names which are known for the terminal, separated by ‘|’ characters. The first name given is the most common abbreviation for the terminal, the last name given should be a long name fully identifying the terminal, and all others are understood as synonyms for the terminal name. All names but the last should be in lower case and contain no blanks; the last name may well contain upper case and blanks for readability.

Terminal names (except for the last, verbose entry) should be chosen using the following conventions. The particular piece of hardware making up the terminal should have a root name chosen, thus “hp2621”. This name should not contain hyphens, except that synonyms may be chosen that do not conflict with other names. Modes that the hardware can be in, or user preferences, should be indicated by appending a hyphen and an indicator of the mode. Thus, a VT100 in 132 column mode would be vt100-w. The following suffixes should be used where possible:

Suffix	Meaning	Example
-w	Wide mode (more than 80 columns)	vt100-w
-am	With auto. margins (usually default)	vt100-am
-nam	Without automatic margins	vt100-nam
-n	Number of lines on the screen	aaa-60
-na	No arrow keys (leave them in local)	c100-na
-np	Number of pages of memory	c100-4p
-rv	Reverse video	c100-rv

CAPABILITIES

The variable is the name by which the programmer (at the terminfo level) accesses the capability. The cap-name is the short name used in the text of the database, and is used by a person updating the database. The i.code is the two letter internal code used in the compiled database, and always corresponds to the old *termcap* capability name.

Capability names have no hard length limit, but an informal limit of 5 characters has been adopted to keep them short and to allow the tabs in the source file caps to line up nicely. Whenever possible, names are chosen to be the same as or similar to the ANSI X3.64-1979 standard. Semantics are also intended to match those of the specification.

- (P) indicates that padding may be specified
- (G) indicates that the string is passed through *tparm* with parms as given (*#i*).
- (*) indicates that padding may be based on the number of lines affected
- (#*i*.) indicates the *i*th parameter.

Variable	Cap-name	I. Code	Description
auto_left_margin,	bw	bw	cu b1 wraps from column 0 to last column
auto_right_margin,	am	am	Terminal has automatic margins
beehive_glitch,	xb	xb	Beehive (f1=escape, f2=ctrl C)
ceol_standout_glitch,	xhp	xs	Standout not erased by overwriting (hp)

eat_newline_glitch,	xenl	xn	newline ignored after 80 cols (Concept)
erase_overstrike,	eo	eo	Can erase overstrikes with a blank
generic_type,	gn	gn	Generic line type (e.g., dialup, switch).
hard_copy,	hc	hc	Hardcopy terminal
has_meta_key,	km	km	Has a meta key (shift, sets parity bit)
has_status_line,	hs	hs	Has extra "status line"
insert_null_glitch,	in	in	Insert mode distinguishes nulls
memory_above,	da	da	Display may be retained above the screen
memory_below,	db	db	Display may be retained below the screen
move_insert_mode,	mir	mi	Safe to move while in insert mode
move_standout_mode,	msgr	ms	Safe to move in standout modes
over_strike,	os	os	Terminal overstrikes
status_line_esc_ok,	eslok	es	Escape can be used on the status line
telera_y_glitch,	xt	xt	Tabs ruin, magic so char (Telera_y 1061)
tilde_glitch,	hz	hz	Hazeltine; can not print '~'s
transparent_underline,	ul	ul	underline character overstrikes
xon_xoff,	xon	xo	Terminal uses xon/xoff handshaking
Numbers:			
columns,	cols	co	Number of columns in a line
init_tabs,	it	it	Tabs initially every # spaces
lines,	lines	li	Number of lines on screen or page
lines_of_memory,	lm	lm	Lines of memory if > lines. 0 means varies
magic_cookie_glitch,	xmc	sg	Number of blank chars left by smso or rmso
padding_baud_rate,	pb	pb	Lowest baud where cr/nl padding is needed
virtual_terminal,	vt	vt	Virtual terminal number (UNIX system)
width_status_line,	wsl	ws	No. columns in status line
Strings:			
back_tab,	cbt	bt	Back tab (P)
bell,	bel	bl	Audible signal (bell) (P)
carriage_return,	cr	cr	Carriage return (P*)
change_scroll_region,	csr	cs	change to lines #1 through #2 (VT100) (PG)
clear_all_tabs,	tbc	ct	Clear all tab stops (P)
clear_screen,	clear	cl	Clear screen and home cursor (P*)
clr_eol,	el	ce	Clear to end of line (P)
clr_eos,	ed	cd	Clear to end of display (P*)
column_address,	hpa	ch	Set cursor column (PG)
command_character,	cmdch	CC	Term. settable cmd char in prototype
cursor_address,	cup	cm	Screen rel. cursor motion row #1 col #2 (PG)
cursor_down,	cudl	do	Down one line
cursor_home,	home	ho	Home cursor (if no cup)

cursor_invisible,	civis	vi	Make cursor invisible
cursor_left,	cubl	le	Move cursor left one space
cursor_mem_address,	mrcup	CM	Memory relative cursor addressing
cursor_normal,	cnorm	ve	Make cursor appear normal (undo vs/vi)
cursor_right,	cuf1	nd	Non-destructive space (cursor right)
cursor_to_ll,	ll	ll	Last line, first column (if no cup)
cursor_up,	cuul	up	Upline (cursor up)
cursor_visible,	cvvis	vs	Make cursor very visible
delete_character,	dch1	dc	Delete character (P*)
delete_line,	dll	dl	Delete line (P*)
dis_status_line,	dsl	ds	Disable status line
down_half_line,	hd	hd	Half-line down (forward 1/2 linefeed)
enter_alt_charset_mode,	smacs	as	Start alternate character set (P)
enter_blink_mode,	blink	mb	Turn on blinking
enter_bold_mode,	bold	md	Turn on bold (extra bright) mode
enter_ca_mode,	smcup	ti	String to begin programs that use cup
enter_delete_mode,	smdc	dm	Delete mode (enter)
enter_dim_mode,	dim	mh	Turn on half-bright mode
enter_insert_mode,	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 stand out mode
enter_underline_mode,	smul	us	Start underscore mode
erase_chars	ech	ec	Erase #1 characters (PG)
exit_alt_charset_mode,	rmacs	ae	End alternate character set (P)
exit_attribute_mode,	sgr0	me	Turn off all attributes
exit_ca_mode,	rmcup	te	String to end programs that use cup
exit_delete_mode,	rmdc	ed	End delete mode
exit_insert_mode,	rmir	ei	End insert mode
exit_standout_mode,	rmso	se	End stand out mode
exit_underline_mode,	rmul	ue	End underscore mode
flash_screen,	flash	vb	Visible bell (may not move cursor)
form_feed,	ff	ff	Hardcopy terminal page eject (P*)
from_status_line,	fsl	fs	Return from status line
init_1string,	is1	i1	Terminal initialization string
init_2string,	is2	i2	Terminal initialization string
init_3string,	is3	i3	Terminal initialization string
init_file,	if	if	Name of file containing is
insert_character,	ich1	ic	Insert character (P)
insert_line,	ill	al	Add new blank line (P*)
insert_padding,	ip	ip	Insert pad after character inserted (p*)
key_backspace,	kbs	kb	Sent by backspace key
key_catab,	ktbc	ka	Sent by clear-all-tabs key
key_clear,	kclr	kC	Sent by clear screen or erase key
key_ctab,	kctab	kt	Sent by clear-tab key
key_dc,	kdch1	kD	Sent by delete character key
key_dl,	kdll	kL	Sent by delete line key
key_down,	kcud1	kd	Sent by terminal down arrow key
key_eic,	krmir	kM	Sent by rmir or smir in insert mode
key_eol,	kel	kE	Sent by clear-to-end-of-line key
key_eos,	ked	kS	Sent by clear-to-end-of-screen key

key_f0,	kf0	k0	Sent by function key f0
key_f1,	kf1	k1	Sent by function key f1
key_f10,	kf10	ka	Sent by function key f10
key_f2,	kf2	k2	Sent by function key f2
key_f3,	kf3	k3	Sent by function key f3
key_f4,	kf4	k4	Sent by function key f4
key_f5,	kf5	k5	Sent by function key f5
key_f6,	kf6	k6	Sent by function key f6
key_f7,	kf7	k7	Sent by function key f7
key_f8,	kf8	k8	Sent by function key f8
key_f9,	kf9	k9	Sent by function key f9
key_home,	khome	kh	Sent by home key
key_ic,	kich1	kI	Sent by ins char/enter ins mode key
key_il,	kill	kA	Sent by insert line
key_left,	kcub1	kI	Sent by terminal left arrow key
key_ll,	kll	kH	Sent by home-down key
key_npage,	knp	kN	Sent by next-page key
key_ppage,	kpp	kP	Sent by previous-page key
key_right,	kcuf1	kr	Sent by terminal right arrow key
key_sf,	kind	kF	Sent by scroll-forward/down key
key_sr,	kri	kR	Sent by scroll-backward/up key
key_stab,	khts	kT	Sent by set-tab key
key_up,	kcuu1	ku	Sent by terminal up arrow key
keypad_local,	rmkx	ke	Out of "keypad transmit" mode
keypad_xmit,	smkx	ks	Put terminal in "keypad transmit" mode
lab_f0,	lf0	l0	Labels on function key f0 if not f0
lab_f1,	lf1	l1	Labels on function key f1 if not f1
lab_f10,	lf10	la	Labels on function key f10 if not f10
lab_f2,	lf2	l2	Labels on function key f2 if not f2
lab_f3,	lf3	l3	Labels on function key f3 if not f3
lab_f4,	lf4	l4	Labels on function key f4 if not f4
lab_f5,	lf5	l5	Labels on function key f5 if not f5
lab_f6,	lf6	l6	Labels on function key f6 if not f6
lab_f7,	lf7	l7	Labels on function key f7 if not f7
lab_f8,	lf8	l8	Labels on function key f8 if not f8
lab_f9,	lf9	l9	Labels on function key f9 if not f9
meta_on,	smm	mm	Turn on "meta mode" (8th bit)
meta_off,	rmm	mo	Turn off "meta mode"
newline,	nel	nw	Newline (behaves like cr followed by lf)
pad_char,	pad	pc	Pad character (rather than null)
parm_dch,	dch	DC	Delete #1 chars (PG*)
parm_delete_line,	dl	DL	Delete #1 lines (PG*)
parm_down_cursor,	cud	DO	Move cursor down #1 lines (PG*)
parm_ich,	ich	IC	Insert #1 blank chars (PG*)
parm_index,	indn	SF	Scroll forward #1 lines (PG)
parm_insert_line,	il	AL	Add #1 new blank lines (PG*)
parm_left_cursor,	cub	LE	Move cursor left #1 spaces (PG)
parm_right_cursor,	cuf	RI	Move cursor right #1 spaces (PG*)
parm_rindex,	rin	SR	Scroll backward #1 lines (PG)
parm_up_cursor,	cuu	UP	Move cursor up #1 lines (PG*)
pkey_key,	pfkey	pk	Prog funct key #1 to type string #2
pkey_local,	pfloc	pl	Prog funct key #1 to execute string #2

pkey_xmit,	pfx	px	Prog funct key #1 to xmit string #2
print_screen,	mc0	ps	Print contents of the screen
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. (PG*)
reset_1string,	rs1	r1	Reset terminal completely to sane modes.
reset_2string,	rs2	r2	Reset terminal completely to sane modes.
reset_3string,	rs3	r3	Reset terminal completely to sane modes.
reset_file,	rf	rf	Name of file containing reset string
restore_cursor,	rc	rc	Restore cursor to position of last sc
row_address,	vpa	cv	Vertical position absolute (set row) (PG)
save_cursor,	sc	sc	Save cursor position (P)
scroll_forward,	ind	sf	Scroll text up (P)
scroll_reverse,	ri	sr	Scroll text down (P)
set_attributes,	sgr	sa	Define the video attributes (PG9)
set_tab,	hts	st	Set a tab in all rows, current column
set_window,	wind	wi	Current window is lines #1-#2 cols #3-#4
tab,	ht	ta	Tab to next 8 space hardware tab stop
to_status_line,	tsl	ts	Go to status line, column #1
underline_char,	uc	uc	Underscore one char and move past it
up_half_line,	hu	hu	Half-line up (reverse 1/2 linefeed)
init_prog,	iprog	iP	Path name of program for init
key_a1,	ka1	K1	Upper left of keypad
key_a3,	ka3	K3	Upper right of keypad
key_b2,	kb2	K2	Center of keypad
key_c1,	kc1	K4	Lower left of keypad
key_c3,	kc3	K5	Lower right of keypad
prtr_non,	mc5p	pO	Turn on the printer for #1 bytes

A Sample Entry

The following entry, which describes the Concept 100, is among the more complex entries in the *terminfo* file as of this writing.

```
concept100|c100|concept|c104|c100-4p|concept 100,
am, bel=^G, blank=\EH, blink=\EC, clear=^L$<2*>, cnorm=\Ew,
cols#80, cr=^M$<9>, cub1=^H, cudl=^J, cufl=\E=,
cup=\Ea%p1%' '%+%c%p2%' '%+%c,
cuul=\E;, cvvis=\EW, db, dchl=\E^A$<16*>, dim=\EE, dll=\E^B$<3*>,
ed=\E^C$<16*>, el=\E^U$<16>, eo, flash=\Ek$<20>\EK, ht=\t$<8>,
ill=\E^R$<3*>, in, ind=^J, .ind=^J$<9>, ip=$<16*>,
is2=\EU\Ef\E7\E5\E8\E1\ENH\EK\E\200\Eo&\200\Eo\47\E,
kbs=^h, kcubl=\E>, kcucl=\E<, kcuf1=\E=, kcuul=\E;,
kf1=\E5, kf2=\E6, kf3=\E7, khome=\E?,
lines#24, mir, pb#9600, prot=\EI, rep=\Er%p1%c%p2%' '%+%c$<.2*>,
rev=\ED, rmcup=\Ev $<6>\Ep\r\n, rmir=\E\200, rmkx=\Ex,
rmso=\Ed\Ee, rmul=\Eg, rmul=\Eg, sgr0=\EN\200,
smcup=\EU\Ev 8p\Ep\r, smir=\E^P, smkx=\EX, smso=\EE\ED,
smul=\EG, tabs, ul, vt#8, xenl,
```

Entries may continue onto multiple lines by placing white space at the beginning of each line except the first. Comments may be included on lines beginning with “#”. Capabilities in *terminfo* are of three types: Boolean capabilities which indicate that the terminal has some particular feature, numeric capabilities giving the size of the terminal or the size of particular delays, and string capabilities, which give a sequence

which can be used to perform particular terminal operations.

Types of Capabilities

All capabilities have names. For instance, the fact that the Concept has *automatic margins* (i.e., an automatic return and linefeed when the end of a line is reached) is indicated by the capability **am**. Hence the description of the Concept includes **am**. Numeric capabilities are followed by the character '#' and then the value. Thus **cols**, which indicates the number of columns the terminal has, gives the value '80' for the Concept.

Finally, string valued capabilities, such as **el** (clear to end of line sequence) are given by the two-character code, an '=', and then a string ending at the next following ','. A delay in milliseconds may appear anywhere in such a capability, enclosed in \$<..> brackets, as in **el=\EK\$<3>**, and padding characters are supplied by *tputs* to provide this delay. The delay can be either a number, e.g., '20', or a number followed by an '*', i.e., '3*'. A '*' indicates that the padding required is proportional to the number of lines affected by the operation, and the amount given is the per-affected-unit padding required. (In the case of insert character, the factor is still the number of *lines* affected. This is always one unless the terminal has **xenl** and the software uses it.) When a '*' is specified, it is sometimes useful to give a delay of the form '3.5' to specify a delay per unit to tenths of milliseconds. (Only one decimal place is allowed.)

A number of escape sequences are provided in the string valued capabilities for easy encoding of characters there. Both **\E** and **\e** map to an ESCAPE character, **\x** maps to a control-x for any appropriate x, and the sequences **\n \r \t \b \f \s** give a newline, linefeed, return, tab, backspace, formfeed, and space. Other escapes include **\^** for ^, **** for \, **\,** for comma, **\:** for :, and **\0** for null. (**\0** will produce **\200**, which does not terminate a string but behaves as a null character on most terminals.) Finally, characters may be given as three octal digits after a ****.

Sometimes individual capabilities must be commented out. To do this, put a period before the capability name. For example, see the second **ind** in the example above.

Preparing Descriptions

We now outline how to prepare descriptions of terminals. The most effective way to prepare a terminal description is by imitating the description of a similar terminal in *terminfo* and to build up a description gradually, using partial descriptions with 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 easily test a new terminal description you can set the environment variable **TERMINFO** to a pathname of a directory containing the compiled description you are working on and programs will look there rather than in **/usr/5lib/terminfo**. To get the padding for insert line right (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 terminal messes up, more padding is usually needed. A similar test can be used for insert character.

Basic Capabilities

The number of columns on each line for the terminal is given by the **cols** numeric capability. If the terminal is a CRT, then the number of lines on the screen is given by the **lines** capability. If the terminal wraps around to the beginning of the next line when it reaches the right margin, then it should have the **am** capability. If the terminal can clear its screen, leaving the cursor in the home position, then this is given by the **clear** string capability. If the terminal overstrikes (rather than clearing a position when a character is struck over) then it should have the **os** capability. If the terminal is a printing terminal, with no soft copy unit, give it both **hc** and **os**. (**os** applies to storage scope terminals, such as Tektronix 4010 series, as well as hard copy and APL terminals.) If there is a code to move the cursor to the left edge of the current row, give this as **cr**. (Normally this will be carriage return, control M.) If there is a code to produce an audible signal (bell, beep, etc) give this as **bel**.

If there is a code to move the cursor one position to the left (such as backspace) that capability should be given as **cub1**. Similarly, codes to move to the right, up, and down should be given as **cuf1**, **cuu1**, and **cul1**. These local cursor motions should not alter the text they pass over, for example, you would not normally use 'cuf1=' because the space would erase the character moved over.

A very important point here is that the local cursor motions encoded in *terminfo* are undefined at the left and top edges of a CRT terminal. Programs should never attempt to backspace around the left edge, unless **bw** is given, and never attempt to go up locally off the top. In order to scroll text up, a program will go to the bottom left corner of the screen and send the **ind** (index) string.

To scroll text down, a program goes to the top left corner of the screen and sends the **ri** (reverse index) string. The strings **ind** and **ri** are undefined when not on their respective corners of the screen.

Parameterized versions of the scrolling sequences are **indn** and **rin** which have the same semantics as **ind** and **ri** except that they take one parameter, and scroll that many lines. They are also undefined except at the appropriate edge of the screen.

The **am** capability tells whether the cursor sticks at the right edge of the screen when text is output, but this does not necessarily apply to a **cuf1** from the last column. The only local motion which is defined from the left edge is if **bw** is given, then a **cub1** from the left edge will move to the right edge of the previous row. If **bw** is not given, the effect is undefined. This is useful for drawing a box around the edge of the screen, for example. If the terminal has switch selectable automatic margins, the *terminfo* file usually assumes that this is on; i.e., **am**. If the terminal has a command which moves to the first column of the next line, that command can be given as **nel** (newline). It does not matter if the command clears the remainder of the current line, so if the terminal has no **cr** and **lf** it may still be possible to craft a working **nel** out of one or both of them.

These capabilities suffice to describe hardcopy and "glass-tty" terminals. Thus the model 33 teletype is described as

```
33|tty33|tty|model 33 teletype,
bel=^G, cols#72, cr=^M, cudl=^J, hc, ind=^J, os,
```

while the Lear Siegler ADM-3 is described as

```
adm3|3|lsi adm3,
am, bel=^G, clear=^Z, cols#80, cr=^M, cub1=^H, cudl=^J,
ind=^J, lines#24,
```

Parameterized Strings

Cursor addressing and other strings requiring parameters in the terminal are described by a parameterized string capability, with *printf(3S)* like escapes **%x** in it. For example, to address the cursor, the **cup** capability is given, using two parameters: the row and column to address to. (Rows and columns are numbered from zero and refer to the physical screen visible to the user, not to any unseen memory.) If the terminal has memory relative cursor addressing, that can be indicated by **mrcup**.

The parameter mechanism uses a stack and special **%** codes to manipulate it. Typically a sequence will push one of the parameters onto the stack and then print it in some format. Often more complex operations are necessary.

The **%** encodings have the following meanings:

%%	outputs '%'
%d	print pop() as in printf
%2d	print pop() like %2d
%3d	print pop() like %3d
%02d	
%03d	as in printf
%c	print pop() gives %c
%s	print pop() gives %s
%p[1-9]	push ith parm
%P[a-z]	set variable [a-z] to pop()
%g[a-z]	get variable [a-z] and push it
%'c'	char constant c

<code>%{nn}</code>	integer constant nn
<code>%+ %- %* %/ %m</code>	arithmetic (<code>%m</code> is mod): push(pop() op pop())
<code>%& % %^</code>	bit operations: push(pop() op pop())
<code>%= %> %<</code>	logical operations: push(pop() op pop())
<code>%! %^-</code>	unary operations push(op pop())
<code>%i</code>	add 1 to first two parms (for ANSI terminals)
<code>%? expr %t thenpart %e elsepart %;</code>	if-then-else, <code>%e</code> elsepart is optional. else-if's are possible ala Algol 68: <code>%? c₁ %t b₁ %e c₂ %t b₂ %e c₃ %t b₃ %e c₄ %t b₄ %e %;</code> <code>c_i</code> are conditions, <code>b_i</code> are bodies.

Binary operations are in postfix form with the operands in the usual order. That is, to get x-5 one would use "`%gx%{5}%-`".

Consider the HP 2645, which, to get to row 3 and column 12, needs to be sent `\E&a12c03Y` padded for 6 milliseconds. Note that the order of the rows and columns is inverted here, and that the row and column are printed as two digits. Thus its `cup` capability is "`cup=6\E&%p2%2dc%p1%2dY`".

The Microterm ACT-IV needs the current row and column sent preceded by a `^T`, with the row and column simply encoded in binary, "`cup=^T%p1%c%p2%c`". Terminals which use "`%c`" need to be able to back-space the cursor (`cuB1`), and to move the cursor up one line on the screen (`cuu1`). This is necessary because it is not always safe to transmit `\n ^D` and `\r`, as the system may change or discard them. (The library routines dealing with terminfo set tty modes so that tabs are never expanded, so `\t` is safe to send. This turns out to be essential for the Ann Arbor 4080.)

A final example is the LSI ADM-3a, which uses row and column offset by a blank character, thus "`cup=\E=%p1% ' %+%c%p2% ' %+%c`". After sending `\E=`, this pushes the first parameter, pushes the ASCII value for a space (32), adds them (pushing the sum on the stack in place of the two previous values) and outputs that value as a character. Then the same is done for the second parameter. More complex arithmetic is possible using the stack.

If the terminal has row or column absolute cursor addressing, these can be given as single parameter capabilities `hpa` (horizontal position absolute) and `vpa` (vertical position absolute). Sometimes these are shorter than the more general two parameter sequence (as with the hp2645) and can be used in preference to `cup`. If there are parameterized local motions (e.g., move *n* spaces to the right) these can be given as `cud`, `cub`, `cuf`, and `cuu` with a single parameter indicating how many spaces to move. These are primarily useful if the terminal does not have `cup`, such as the Tektronix 4025.

Cursor Motions

If the terminal has a fast way to home the cursor (to very upper left corner of screen) then this can be given as `home`; similarly a fast way of getting to the lower left-hand corner can be given as `ll`; this may involve going up with `cuu1` from the home position, but a program should never do this itself (unless `ll` does) because it can make no assumption about the effect of moving up from the home position. Note that the home position is the same as addressing to (0,0): to the top left corner of the screen, not of memory. (Thus, the `\EH` sequence on HP terminals cannot be used for `home`.)

Area Clears

If the terminal can clear from the current position to the end of the line, leaving the cursor where it is, this should be given as `el`. If the terminal can clear from the current position to the end of the display, then this should be given as `ed`. `Ed` is only defined from the first column of a line. (Thus, it can be simulated by a request to delete a large number of lines, if a true `ed` is not available.)

Insert/delete line

If the terminal can open a new blank line before the line where the cursor is, this should be given as **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 scrolling region (like the VT100) the command to set this can be described with the **csr** capability, which takes two parameters: the top and bottom lines of the scrolling region. The cursor position is, alas, undefined after using this command. It is possible to get the effect of insert or delete line using this command – the **sc** and **rc** (save and restore cursor) commands are also useful. Inserting lines at the top or bottom of the screen can also be done using **ri** or **ind** on many terminals without a true insert/delete line, and is often faster even on terminals with those features.

If the terminal has the ability to define a window as part of memory, which all commands affect, it should be given as the parameterized string **wind**. The four parameters are the starting and ending lines in memory and the starting and ending columns in memory, in that order.

If the terminal can retain display memory above, then the **da** capability should be given; if display memory can be retained below, then **db** should be given. These indicate that deleting a line or scrolling may bring non-blank lines up from below or that scrolling back with **ri** may bring down non-blank lines.

Insert/Delete Character

There are two basic kinds of intelligent terminals with respect to insert/delete character which can be described using *terminfo*. The most common insert/delete character operations affect only the characters on the current line and shift characters off the end of the line rigidly. Other terminals, such as the Concept 100 and the Perkin Elmer Owl, make a distinction between typed and untyped blanks on the screen, shifting upon an insert or delete only to an untyped blank on the screen which is either eliminated, or expanded to two untyped blanks. You can determine the kind of terminal you have by clearing the screen and then typing text separated by cursor motions. Type “abc def” using local cursor motions (not spaces) between the “abc” and the “def”. Then position the cursor before the “abc” and put the terminal in insert mode. If typing characters causes the rest of the line to shift rigidly and characters to fall off the end, then your terminal does not distinguish between blanks and untyped positions. If the “abc” shifts over to the “def” which then move together around the end of the current line and onto the next as you insert, you have the second type of terminal, and should give the capability **in**, which stands for “insert null”. While these are two logically separate attributes (one line vs. multiline insert mode, and special treatment of untyped spaces) we have seen no terminals whose insert mode cannot be described with the single attribute.

Terminfo can describe both terminals which have an insert mode, and terminals which send a simple sequence to open a blank position on the current line. Give as **smir** the sequence to get into insert mode. Give as **rmir** the sequence to leave insert mode. Now give as **ich1** any sequence needed to be sent just before sending the character to be inserted. Most terminals with a true insert mode will not give **ich1**; terminals which send a sequence to open a screen position should give it here. (If your terminal has both, insert mode is usually preferable to **ich1**. Do not give both unless the terminal actually requires both to be used in combination.) If post insert padding is needed, give this as a number of milliseconds in **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.

It is occasionally necessary to move around while in insert mode to delete characters on the same line (e.g., if there is a tab after the insertion position). If your terminal allows motion while in insert mode you can give the capability **mir** to speed up inserting in this case. Omitting **mir** will affect only speed. Some terminals (notably Datamedia’s) must not have **mir** because of the way their insert mode works.

Finally, you can specify **dch1** to delete a single character, **dch** with one parameter, *n*, to delete *n* characters, and delete mode by giving **smdc** and **rmdc** to enter and exit delete mode (any mode the terminal needs to be placed in for **dch1** to work).

A command to erase *n* characters (equivalent to outputting *n* blanks without moving the cursor) can be given as **ech** with one parameter.

Highlighting, Underlining, and Visible Bells

If your terminal has one or more kinds of display attributes, these can be represented in a number of different ways. You should choose one display form as *standout mode*, representing a good, high contrast, easy-on-the-eyes, format for highlighting error messages and other attention getters. (If you have a choice, reverse video plus half-bright is good, or reverse video alone.) The sequences to enter and exit standout mode are given as **sms0** and **rms0**, respectively. If the code to change into or out of standout mode leaves one or even two blank spaces on the screen, as the TVI 912 and Teleray 1061 do, then **xmc** should be given to tell how many spaces are left.

Codes to begin underlining and end underlining can be given as **smul** and **rmul** respectively. If the terminal has a code to underline the current character and move the cursor one space to the right, such as the Microterm Mime, this can be given as **uc**.

Other capabilities to enter various highlighting modes include **blink** (blinking) **bold** (bold or extra bright) **dim** (dim or half-bright) **invis** (blanking or invisible text) **prot** (protected) **rev** (reverse video) **sgr0** (turn off *all* attribute modes) **smacs** (enter alternate character set mode) and **rmacs** (exit alternate character set mode). Turning on any of these modes singly may or may not turn off other modes.

If there is a sequence to set arbitrary combinations of modes, this should be given as **sgr** (set attributes), taking 9 parameters. Each parameter is either 0 or 1, as the corresponding attribute is on or off. The 9 parameters are, in order: standout, underline, reverse, blink, dim, bold, blank, protect, alternate character set. Not all modes need be supported by **sgr**, only those for which corresponding separate attribute commands exist.

Terminals with the "magic cookie" glitch (**xmc**) deposit special "cookies" when they receive mode-setting sequences, which affect the display algorithm rather than having extra bits for each character. Some terminals, such as the HP 2621, automatically leave standout mode when they move to a new line or the cursor is addressed. Programs using standout mode should exit standout mode before moving the cursor or sending a newline, unless the **msg**r capability, asserting that it is safe to move in standout mode, is present.

If the terminal has a way of flashing the screen to indicate an error quietly (a bell replacement) then this can be given as **flash**; it must not move the cursor.

If the cursor needs to be made more visible than normal when it is not on the bottom line (to make, for example, a non-blinking underline into an easier to find block or blinking underline) give this sequence as **cvvis**. If there is a way to make the cursor completely invisible, give that as **civis**. The capability **cnorm** should be given which undoes the effects of both of these modes.

If the terminal needs to be in a special mode when running a program that uses these capabilities, the codes to enter and exit this mode can be given as **smcup** and **rmcup**. This arises, for example, from terminals like the Concept with more than one page of memory. If the terminal has only memory relative cursor addressing and not screen relative cursor addressing, a one screen-sized window must be fixed into the terminal for cursor addressing to work properly. This is also used for the Tektronix 4025, where **smcup** sets the command character to be the one used by terminfo.

If your terminal correctly generates underlined characters (with no special codes needed) even though it does not overstrike, then you should give the capability **ul**. If overstrikes are erasable with a blank, then this should be indicated by giving **eo**.

Keypad

If the terminal has a keypad that transmits codes when the keys are pressed, this information can be given. Note that it is not possible to handle terminals where the keypad only works in local (this applies, for example, to the unshifted HP 2621 keys). If the keypad can be set to transmit or not transmit, give these codes as **smkx** and **rmkx**. Otherwise the keypad is assumed to always transmit. The codes sent by the left arrow, right arrow, up arrow, down arrow, and home keys can be given as **kcub1**, **kcuf1**, **kcuu1**, **kcud1**, and **khome** respectively. If there are function keys such as **f0**, **f1**, ..., **f10**, the codes they send can be given as **kf0**, **kf1**, ..., **kf10**. If these keys have labels other than the default **f0** through **f10**, the labels can be given as **lf0**, **lf1**, ..., **lf10**. The codes transmitted by certain other special keys can be given: **kll** (home down), **kbs** (backspace), **ktbc** (clear all tabs), **kctab** (clear the tab stop in this column), **kclr** (clear screen or erase key), **kdch1** (delete character), **kdll1** (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), **kill1** (insert line), **knp** (next page), **kpp** (previous page), **kind** (scroll forward/down), **kri** (scroll backward/up), **khts** (set a tab stop in this column). In addition, if the keypad has a 3 by 3 array of keys including the four arrow keys, the other five keys can be given as **ka1**, **ka3**, **kb2**, **kc1**, and **kc3**. These keys are useful when the effects of a 3 by 3 directional pad are needed.

Tabs and Initialization

If the terminal has hardware tabs, the command to advance to the next tab stop can be given as **ht** (usually control I). A "backtab" command which moves leftward to the next tab stop can be given as **cbt**. By convention, if the teletype modes indicate that tabs are being expanded by the computer rather than being sent to the terminal, programs should not use **ht** or **cbt** even if they are present, since the user may not have the tab stops properly set. If the terminal has hardware tabs which are initially set every *n* spaces when the terminal is powered up, the numeric parameter *it* is given, showing the number of spaces the tabs are set to. This is normally used by the *tset* command to determine whether to set the mode for hardware tab expansion, and whether to set the tab stops. If the terminal has tab stops that can be saved in nonvolatile memory, the terminfo description can assume that they are properly set.

Other capabilities include **is1**, **is2**, and **is3**, initialization strings for the terminal, **iprog**, the path name of a program to be run to initialize the terminal, and **if**, the name of a file containing long initialization strings. These strings are expected to set the terminal into modes consistent with the rest of the terminfo description. They are normally sent to the terminal, by the *tset* program, each time the user logs in. They will be printed in the following order: **is1**; **is2**; setting tabs using **tbc** and **hts**; **if**; running the program **iprog**; and finally **is3**. Most initialization is done with **is2**. Special terminal modes can be set up without duplicating strings by putting the common sequences in **is2** and special cases in **is1** and **is3**. A pair of sequences that does a harder reset from a totally unknown state can be analogously given as **rs1**, **rs2**, **rf**, and **rs3**, analogous to **is2** and **if**. These strings are output by the *reset* program, which is used when the terminal gets into a wedged state. Commands are normally placed in **rs2** and **rf** only if they produce annoying effects on the screen and are not necessary when logging in. For example, the command to set the VT100 into 80-column mode would normally be part of **is2**, but it causes an annoying glitch of the screen and is not normally needed since the terminal is usually already in 80 column mode.

If there are commands to set and clear tab stops, they can be given as **tbc** (clear all tab stops) and **hts** (set a tab stop in the current column of every row). If a more complex sequence is needed to set the tabs than can be described by this, the sequence can be placed in **is2** or **if**.

Delays

Certain capabilities control padding in the teletype driver. These are primarily needed by hard copy terminals, and are used by the *tset* program to set teletype modes appropriately. Delays embedded in the capabilities **cr**, **ind**, **cub1**, **ff**, and **tab** will cause the appropriate delay bits to be set in the teletype driver. If **pb** (padding baud rate) is given, these values can be ignored at baud rates below the value of **pb**.

Miscellaneous

If the terminal requires other than a null (zero) character as a pad, then this can be given as **pad**. Only the first character of the **pad** string is used.

If the terminal has an extra “status line” that is not normally used by software, this fact can be indicated. If the status line is viewed as an extra line below the bottom line, into which one can cursor address normally (such as the Heathkit H19’s 25th line, or the 24th line of a VT100 which is set to a 23-line scrolling region), the capability **hs** should be given. Special strings to go to the beginning of the status line and to return from the status line can be given as **tsl** and **fsl**. (**fsl** must leave the cursor position in the same place it was before **tsl**. If necessary, the **sc** and **rc** strings can be included in **tsl** and **fsl** to get this effect.) The parameter **tsl** takes one parameter, which is the column number of the status line the cursor is to be moved to. If escape sequences and other special commands, such as **tab**, work while in the status line, the flag **eslok** can be given. A string which turns off the status line (or otherwise erases its contents) should be given as **dsl**. If the terminal has commands to save and restore the position of the cursor, give them as **sc** and **rc**. The status line is normally assumed to be the same width as the rest of the screen, e.g., **cols**. If the status line is a different width (possibly because the terminal does not allow an entire line to be loaded) the width, in columns, can be indicated with the numeric parameter **wsl**.

If the terminal can move up or down half a line, this can be indicated with **hu** (half-line up) and **hd** (half-line down). This is primarily useful for superscripts and subscripts on hardcopy terminals. If a hardcopy terminal can eject to the next page (form feed), give this as **ff** (usually control L).

If there is a command to repeat a given character a given number of times (to save time transmitting a large number of identical characters) this can be indicated with the parameterized string **rep**. The first parameter is the character to be repeated and the second is the number of times to repeat it. Thus, **tparam(repeat_char, 'x', 10)** is the same as **'xxxxxxxxx'**.

If the terminal has a settable command character, such as the Tektronix 4025, this can be indicated with **cmdch**. A prototype command character is chosen which is used in all capabilities. This character is given in the **cmdch** capability to identify it. The following convention is supported on some UNIX systems: The environment is to be searched for a **CC** variable, and if found, all occurrences of the prototype character are replaced with the character in the environment variable.

Terminal descriptions that do not represent a specific kind of known terminal, such as *switch*, *dialup*, *patch*, and *network*, should include the **gn** (generic) capability so that programs can complain that they do not know how to talk to the terminal. (This capability does not apply to *virtual* terminal descriptions for which the escape sequences are known.)

If the terminal uses **xon/xoff** handshaking for flow control, give **xon**. Padding information should still be included so that routines can make better decisions about costs, but actual pad characters will not be transmitted.

If the terminal has a “meta key” which acts as a shift key, setting the 8th bit of any character transmitted, this fact can be indicated with **km**. Otherwise, software will assume that the 8th bit is parity and it will usually be cleared. If strings exist to turn this “meta mode” on and off, they can be given as **smm** and **rmm**.

If the terminal has more lines of memory than will fit on the screen at once, the number of lines of memory can be indicated with **lm**. A value of **lm#0** indicates that the number of lines is not fixed, but that there is still more memory than fits on the screen.

If the terminal is one of those supported by the UNIX virtual terminal protocol, the terminal number can be given as **vt**.

Media copy strings which control an auxiliary printer connected to the terminal can be given as **mc0**: print the contents of the screen, **mc4**: turn off the printer, and **mc5**: turn on the printer. When the printer is on, all text sent to the terminal will be sent to the printer. It is undefined whether the text is also displayed on the terminal screen when the printer is on. A variation **mc5p** takes one parameter, and leaves the printer on for as many characters as the value of the parameter, then turns the printer off. The parameter should not exceed 255. All text, including **mc4**, is transparently passed to the printer while an **mc5p** is in effect.

Strings to program function keys can be given as **pfkey**, **pfloc**, and **pxf**. Each of these strings takes two parameters: the function key number to program (from 0 to 10) and the string to program it with. Function key numbers out of this range may program undefined keys in a terminal dependent manner. The

difference between the capabilities is that **pfkey** causes pressing the given key to be the same as the user typing the given string; **pfloc** causes the string to be executed by the terminal in local; and **pfx** causes the string to be transmitted to the computer.

Glitches and Braindamage

Hazeltine terminals, which do not allow “~” characters to be displayed should indicate **hz**.

Terminals which ignore a linefeed immediately after an am wrap, such as the Concept and VT100, should indicate **xenl**.

If **el** is required to get rid of standout (instead of merely writing normal text on top of it), **xhp** should be given.

Telera terminals, where tabs turn all characters moved over to blanks, should indicate **xt** (destructive tabs). This glitch is also taken to mean that it is not possible to position the cursor on top of a “magic cookie”, that to erase standout mode it is instead necessary to use delete and insert line.

The Beehive Superbee, which is unable to correctly transmit the escape or control C characters, has **xsb**, indicating that the f1 key is used for escape and f2 for control C. (Only certain Superbees have this problem, depending on the ROM.)

Other specific terminal problems may be corrected by adding more capabilities of the form **xx**.

Similar Terminals

If there are two very similar terminals, one can be defined as being just like the other with certain exceptions. The string capability **use** can be given with the name of the similar terminal. The capabilities given before **use** override those in the terminal type invoked by **use**. A capability can be cancelled by placing **xx@** to the left of the capability definition, where **xx** is the capability. For example, the entry

```
2621-nl, smkx@, rmkx@, use=2621,
```

defines a 2621-nl that does not have the **smkx** or **rmkx** capabilities, and hence does not turn on the function key labels when in visual mode. This is useful for different modes for a terminal, or for different user preferences.

FILES

`/usr/5lib/terminfo/?/*` files containing terminal descriptions

SEE ALSO

`curses(3V)`, `printf(3S)`

NAME

tp – DEC/mag tape formats

DESCRIPTION

Tp dumps files to and extracts files from DECtape and magtape. The formats of these tapes are the same except that magtapes have larger directories.

Block zero contains a copy of a stand-alone bootstrap program. See *reboot*(8).

Blocks 1 through 24 for DECtape (1 through 62 for magtape) contain a directory of the tape. There are 192 (resp. 496) entries in the directory; 8 entries per block; 64 bytes per entry. Each entry has the following format:

```

struct {
    char        pathname[32];
    unsigned short mode;
    char        uid;
    char        gid;
    char        unused1;
    char        size[3];
    long        modtime;
    unsigned short tapeaddr;
    char        unused2[16];
    unsigned short checksum;
};

```

The path name entry is the path name of the file when put on the tape. If the pathname starts with a zero word, the entry is empty. It is at most 32 bytes long and ends in a null byte. Mode, uid, gid, size and time modified are the same as described under i-nodes (see file system *fs*(5)). The tape address is the tape block number of the start of the contents of the file. Every file starts on a block boundary. The file occupies $(size+511)/512$ blocks of continuous tape. The checksum entry has a value such that the sum of the 32 words of the directory entry is zero.

Blocks above 25 (resp. 63) are available for file storage.

A fake entry has a size of zero.

SEE ALSO

fs(5)

BUGS

The *pathname*, *uid*, *gid*, and *size* fields are too small.

NAME

ttys – terminal initialization data

DESCRIPTION

The *ttys* file is read by the *init* program and specifies which terminal special files are to have a process created for them so that people can log in. There is one line in the *ttys* file per special file associated with a terminal.

The first character of a line in the *ttys* file is either '0' or '1'. If the first character on the line is a '0', the *init* program ignores that line. If the first character on the line is a '1', the *init* program creates a login process for that line.

The second character on each line is used as an argument to *getty*(8), which performs such tasks as baud-rate recognition, reading the login name, and calling *login*. For normal lines, the second character is '0'; other characters can be used, for example, with hard-wired terminals where speed recognition is unnecessary or which have special characteristics. The remainder of the line is the terminal's entry in the device directory, */dev*.

Getty uses the second character in the *ttys* file to look up the characteristics of the terminal in the */etc/gettytab* file. Consult the *gettytab*(5) manual page for an explanation of the layout of */etc/gettytab*.

FILES

/etc/ttys

SEE ALSO

init(8), *getty*(8), *login*(1), *gettytab*(5)

NAME

ttytype – data base of terminal types by port

SYNOPSIS

/etc/ttytype

DESCRIPTION

Ttytype is a database containing, for each tty port on the system, the kind of terminal that is attached to it. There is one line per port, containing the terminal kind (as a name listed in *termcap* (5)), a space, and the name of the tty, minus */dev/*.

This information is read by *tset*(1) and by *login*(1) to initialize the TERM variable at login time.

SEE ALSO

tset(1), *login*(1)

BUGS

Some lines are merely known as “dialup” or “plugboard”.

NAME

types – primitive system data types

SYNOPSIS

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

DESCRIPTION

The data types defined in the include file are used in UNIX system code; some data of these types are accessible to user code:

```
/*      @(#)types.h 1.2 86/08/01 SMI; from UCB 4.11 83/07/01*/
```

```
/*
 * Basic system types and major/minor device constructing/busting macros.
 */
```

```
#ifndef _TYPES_
#define _TYPES_
```

```
#ifndef KERNEL
#include<sys/sysmacros.h>
#else
#include"./h/sysmacros.h"
#endif
```

```
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
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 int     off_t;
typedef long    key_t;
```

```
typedef struct  fd_set { int fds_bits[1]; } fd_set;
```

#endif

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

fs(5), *time(3C)*, *lseek(2)*, *adb(1)*

NAME

utmp, *wtmp*, *lastlog*, *usracct* – login records

SYNOPSIS

```
#include <utmp.h>
```

DESCRIPTION

The *utmp* file records information about who is currently using the system. The file is a sequence of entries with the following structure declared in the include file:

```
/*      @(#)utmp.h 1.1 86/07/07 SMI; from UCB 4.2 83/05/22      */

/*
 * Structure of utmp and wtmp files.
 *
 * Assuming the number 8 is unwise.
 */
struct utmp {
    char    ut_line[8];           /* tty name */
    char    ut_name[8];          /* user id */
    char    ut_host[16];         /* host name, if remote */
    long    ut_time;            /* time on */
};

/*
 * This is a utmp entry that does not correspond to a genuine user
 */
#define nonuser(ut) ((ut).ut_host[0] == 0 &&      strcmp((ut).ut_line, "tty", 3) == 0 && ((ut).ut_line[3] ==
```

This structure gives the name of the special file associated with the user's terminal, the user's login name, and the time of the login in the form of *time(3C)*.

The *wtmp* file records all logins and logouts. A null user name indicates a logout on the associated terminal. Furthermore, the terminal name “” indicates that the system was rebooted at the indicated time; the adjacent pair of entries with terminal names ‘|’ and ‘}’ indicate the system-maintained time just before and just after a *date* command has changed the system's idea of the time.

wtmp is maintained by *login(1)* and *init(8)*. Neither of these programs creates the file, so if it is removed, record-keeping is turned off. It is summarized by *ac(8)*.

/usr/adm/wtmp is appended to whenever a user logs in or out, and should be truncated periodically.

The *lastlog* file records the most recent login-date for every user logged in. When reporting (and updating) the most recent login date, *login(1)* performs an to a byte-offset in */usr/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* 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.

The *usracct* file keeps login records in an older format.

FILES

```
/etc/utmp
/usr/adm/wtmp
/usr/adm/lastlog
/usr/adm/usracct
```

SEE ALSO

login(1), *init(8)*, *who(1)*, *ac(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* 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

uuencode(1C), uudecode(1C), usend(1C), uucp(1C), mail(1)

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 short size;       /* 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 short addr;       /* &(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 doesn't 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), pti(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 <i>lex</i>)
\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 yellowpages database and directory structure

DESCRIPTION

The yellow pages (YP) network lookup service uses a database of *dbm* files in the directory hierarchy at */etc/yp*. A *dbm* database consists of two files, created by calls to the *dbm(3X)* 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 named set of YP maps. Each YP domain is implemented as a subdirectory of */etc/yp* containing the map. 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 */usr/lib/crontab* 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 UNIX 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* 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. 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 */etc/yp/Makefile*. All Sun-supplied maps have entries in */etc/yp/Makefile*; if you add a YP map, edit this file to support the new map. The makefile uses *makedbm* to generate the YP map on the master, and *yppush* 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)*.

SEE ALSO

makedbm(8), *ypinit(8)*, *ypmake(8)*, *ypxfr(8)*, *yppush(8)*, *yppoll(8)*, *ypserv(8)*, *rpcinfo(8)*,

Index

Special Characters

ltolower — convert character to lower-case, System V, 360
uoupper — convert character to upper-case, System V, 360

1

1/2-inch tape drive
 tm — tapemaster, 494
 xt — Xylogics 472, 510
1/4-inch tape drive
 ar — Archive 1/4-inch Streaming Tape Drive, 409
 st — Sysgen SC 4000 (Archive) Tape Drive, 482 *thru* 483
10 Mb/s 3Com Ethernet interface — ec, 427
10 Mb/s Sun Ethernet interface — ie, 435
10 Mb/s Sun-3/50 Ethernet interface — le, 460 *thru* 461

2

2180 SMD Disk driver — ip, 442 *thru* 443

3

3Com 10 Mb/s Ethernet interface — ec, 427

4

450 SMD Disk driver — xy, 511 *thru* 512
472 1/2-inch tape drive — xt, 510

8

8530 SCC serial communications driver — zs, 513

A

a.out — assembler and link editor output, 517
a64l — convert long integer to base-64 ASCII, 161
abort — generate fault, 162
abs — integer absolute value, 163
absolute value — abs, 163
absolute value function — fabs, 278
accept — connection on socket, 14
access, 15
access times of file, change — utimes, 145
accounting
 process accounting, turn on or off — acct, 17
accounting file — acct, 521
acct
 acct — execution accounting file, 521
 acct — process accounting on or off, 17
acos — trigonometric arccosine, 286

acosh — inverse hyperbolic function, 275
Adaptec ST-506 disk driver — sd, 480 *thru* 481
add password file entry — putpwent, 223
add route ioctl — SIOCADDRT, 479
addmntent — get filesystem descriptor file entry, 192
adjtime — adjust time, 18
advise paging system — vadvise, 146
alarm — schedule signal, 260
aliases — sendmail aliases file, 522
alloca — allocate on stack, 211
allocate aligned memory — memalign, 210
allocate aligned memory — valloc, 211
allocate memory — calloc, 210
allocate memory — malloc, 210
allocate on stack — alloca, 211
alphasort — sort directory, 231
ANSI standard terminal emulation, 418 *thru* 422
ANSI terminal emulation — console, 418 *thru* 422
ar — Archive 1/4-inch Streaming Tape Drive, 409
ar — archive file format, 525
arc — plot arc, 402
archive file format — ar, 525
argument lists, varying length — varargs, 256
arp — Address Resolution Protocol, 410 *thru* 411
arp ioctl
 SIOCDAARP — delete arp entry, 410
 SIOCGARP — get arp entry, 410
 SIOCSARP — set arp entry, 410

ASCII

string to double — strtod, 243
string to long integer — strtol, 244
to float — atof, 243
to integer — atoi, 244
to long — atol, 244
ASCII to Ethernet address — ether_aton, 292
asctime — date and time conversion, 169
asctime — date and time conversion, System V, 358
asin — trigonometric arcsine, 286
asinh — inverse hyperbolic function, 275
assembler output — a.out, 517
assert — program verification, 164
assert — program verification,, 357
assign buffering to stream
 setbuf — assign buffering, System V, 384
 setbuf — assign buffering, 348

assign buffering to stream, *continued*
 setbuffer — assign buffering, System V, 384
 setbuffer — assign buffering, 348
 setlinebuf — assign buffering, System V, 384
 setlinebuf — assign buffering, 348
 setvbuf — assign buffering, System V, 384
 setvbuf — assign buffering, 348

assign to memory characters, 213

async_daemon, 84

at OOB mark? `ioctl` — `SIOCATMARK`, 485

atan — trigonometric arctangent, 286

atan2 — trigonometric arctangent, 286

atanh — inverse hyperbolic function, 275

atof — ASCII to float, 243

atoi — ASCII to integer, 244

atol — ASCII to long, 244

attributes of file `fstat`, 132

attributes of file `lstat`, 132

attributes of file `stat`, 132

B

bcmp — compare byte strings, 167

bcopy — copy byte strings, 167

Bessel functions

- j0, 281
- j1, 281
- jn, 281
- y0, 281
- y1, 281
- yn, 281

binary I/O, buffered

- `fread` — read from stream, 336
- `fread` — read from stream, System V, 371
- `fwrite` — write to stream, 336
- `fwrite` — write to stream, System V, 371

binary search of sorted table — `bsearch`, 165

binary tree routines, 248

bind, 19

bit clear local mode bits `ioctl` — `TIOCLBIC`, 503

bit set local mode bits `ioctl` — `TIOCLBIS`, 503

bit string functions — `ffs`, 167

bk — machine-machine communication line discipline, 412

bk `ioctl`'s

- `TIOCGETD` — get line discipline, 412
- `TIOCSETD` — set line discipline, 412

block signals, 121

blocked signals, release — `sigpause`, 122

both real and effective group ID, set — `setgid`, 234

both real and effective group ID, set, System V — `setgid`, 386

both real and effective user ID, set — `setuid`, 234

both real and effective user ID, set, System V — `setuid`, 386

brk — set data segment break, 21

`bsearch` — binary search of a sorted table, 165

buffered binary I/O

- `fread` — read from stream, 336
- `fread` — read from stream, System V, 371
- `fwrite` — write to stream, 336
- `fwrite` — write to stream, System V, 371

buffered I/O library functions, introduction to, 329

buffering

buffering, *continued*

- assign to stream — `setbuf`, 348
- assign to stream, System V — `setbuf`, 384
- assign to stream — `setbuffer`, 348
- assign to stream, System V — `setbuffer`, 384
- assign to stream — `setlinebuf`, 348
- assign to stream, System V — `setlinebuf`, 384
- assign to stream — `setvbuf`, 348
- assign to stream, System V — `setvbuf`, 384

bwone — Sun-1 black and white frame buffer, 413

bwtwo — Sun-3/Sun-2 black and white frame buffer, 414

byte order, functions to convert between host and network, 291

byte string functions

- bcmp, 167
- bcopy, 167
- bzero, 167

bzero — zero byte strings, 167

C

C library functions, introduction to, 153 *thru* 160

cabs — Euclidean distance, 279

calloc — allocate memory, 210

cbirt — cube root function, 288

ceil — ceiling of, 278

cfree — free memory, 210

cgfour — Sun-3 color graphics interface, 415

cgone — Sun-1 color graphics interface, 416

cgtwo — Sun-3/Sun-2 color graphics interface, 417

change

- current working directory, 22
- data segment size — `sbrk`, 21
- file access times — `utimes`, 145
- file mode — `chmod`, 23
- file name — `rename`, 101
- owner and group of file — `chown`, 25
- root directory — `chroot`, 27

change translation table entry `ioctl` — `KIOCKETKEY`, 444

character

- get from `stdin` — `getchar`, 338
- get from `stdin`, System V — `getchar`, 372
- get from stream — `fgetc`, 338
- get from stream, System V — `fgetc`, 372
- get from stream — `getc`, 338
- get from stream, System V — `getc`, 372
- push back to stream — `ungetc`, 352
- put to `stdin` — `putchar`, 344
- put to stream — `fputc`, 344
- put to stream — `putc`, 344

character classification

- isalnum, 171
- isalpha, 171
- isascii, 171
- iscntrl, 171
- isdigit, 171
- isgraph, 171
- islower, 171
- isprint, 171
- ispunct, 171
- isspace, 171
- isupper, 171
- isxdigit, 171

character classification, System V

- character classification, System V, *continued*
 - isalnum, 360
 - isalpha, 360
 - isascii, 360
 - isctrl, 360
 - isdigit, 360
 - isgraph, 360
 - islower, 360
 - isprint, 360
 - ispunct, 360
 - isspace, 360
 - isupper, 360
 - isxdigit, 360
- character conversion
 - toascii, 171
 - tolower, 171
 - toupper, 171
- character conversion, System V
 - _tolower, 360
 - _toupper, 360
 - toascii, 360
 - tolower, 360
 - toupper, 360
- chdir, 22
- check buffer state `ioctl` — `GPIO_GET_GBUFFER_STATE`, 431
- check heap — `malloc_verify`, 211
- chmod, 23
- chown, 25
- chroot — change root directory, 27
- circle — plot circle, 402
- clear break bit `ioctl` — `TIOCCBRK`, 502
- clear DTR `ioctl` — `TIOCCDTR`, 502
- clear user table `ioctl` — `NDIOCCLEAR`, 471
- clearerr — clear error on stream, 334
- clearerr — clear error on stream, System V, 368
- clock, 261
- close, 28
- close directory stream — `closedir`, 173
- close stream — `fclose`, 333
- closedir — close directory stream, 173
- closelog — close system log file, 246
- closepl — close plot device, 402
- color graphics interface
 - `cgfour` — Sun-3 color graphics interface, 415
 - `cgone` — Sun-1 color graphics interface, 416
 - `cgtwo` — Sun-3/Sun-2 color graphics interface, 417
- command
 - return stream to remote — `rcmd`, 303
 - return stream to remote — `rexec`, 304
- compare
 - byte strings — `bcmp`, 167
 - memory characters — `memcmp`, 213
 - strings — `strcmp`, 241
 - strings — `strncmp`, 241
- compatibility library functions, introduction to, 259
- compile regular expression — `re_comp`, 227
- concatenate strings
 - `strcat`, 241
 - `strncat`, 241
- connect, 29
- connected peer, get name of, 54
- connection
 - accept on socket — `accept`, 14
 - listen for on socket — `listen`, 69
- console — console driver/terminal emulator, 418 *thru* 422
- console I/O `ioctl`, `TIOCCONS`, 418
- cont — continue line, 402
- control devices — `ioctl`, 65
- control resource consumption — `vlimit`, 270
- control system log
 - close system log — `closelog`, 246
 - start system log — `openlog`, 246
 - write to system log — `syslog`, 246
- control terminal, hangup — `vhangup`, 148
- convert
 - functions to between host and network byte order, 291
 - host to network long — `htonl`, 291
 - host to network short — `htons`, 291
 - network to host long — `ntohl`, 291
 - network to host short — `ntohs`, 291
- convert base-64 ASCII to long integer — 164a, 161
- convert character
 - to ASCII — `toascii`, 171
 - to ASCII, System V — `toascii`, 360
 - to lower-case — `tolower`, 171
 - to lower-case, System V — `_tolower`, 360
 - to lower-case, System V — `tolower`, 360
 - to upper-case — `toupper`, 171
 - to upper-case, System V — `_toupper`, 360
 - to upper-case, System V — `toupper`, 360
- convert long integer to base-64 ASCII — 164a, 161
- convert numbers to strings
 - `ecvt`, 177
 - `fcvt`, 177
 - `fprintf`, 341
 - `gcvt`, 177
 - `printf`, 341
 - `sprintf`, 341
- convert numbers to strings, System V
 - `fprintf`, 377
 - `printf`, 377
 - `sprintf`, 377
- convert strings to numbers
 - `atof`, 243
 - `atoi`, 244
 - `atol`, 244
 - `fscanf`, 346
 - `scanf`, 346
 - `sscanf`, 346
 - `strtod`, 243
 - `strtol`, 244
- convert strings to numbers, System V
 - `fscanf`, 381
 - `scanf`, 381
 - `sscanf`, 381
- convert time and date
 - `asctime`, 169
 - `ctime`, 169
 - `dysize`, 170
 - `gmtime`, 169
 - `localtime`, 169
 - `timezone`, 169

- convert time and date, System V
 - asctime, 358
 - ctime, 358
 - gmtime, 358
 - localtime, 358
 - copy
 - byte strings — bcopy, 167
 - memory character fields — memcpy, 213
 - memory character strings — memccpy, 213
 - strings — strcpy, 241
 - strings — strncpy, 241
 - copysign — IEEE floating-point function, 280
 - core — memory image file format, 526
 - cos — trigonometric cosine, 286
 - cosh — hyperbolic cosine, 287
 - cpio — cpio archive format, 527
 - creat, 31
 - create
 - file — open, 85
 - hash table — hcreate, 201
 - interprocess communication channel — pipe, 88
 - interprocess communication endpoint — socket, 129
 - name for temporary file — tmpnam, 351
 - pair of connected sockets — socketpair, 131
 - special file, 73
 - symbolic link — symlink, 137
 - unique file name — mktemp, 214
 - create directory, 71
 - create new process, 42
 - crontab — periodic jobs table, 528
 - crypt — encryption, 168
 - ctermid — generate filename for terminal, 331
 - ctime — date and time conversion, 169
 - ctime — date and time conversion, System V, 358
 - current directory
 - change, 22
 - get pathname — getwd, 200
 - current host, get identifier of — gethostid, 49
 - current working directory — getcwd, 186
 - cursor functions, System V, 362
 - cuserid — get user name, 332
- ## D
- daemons
 - network file system, 84
 - data segment size, change — sbrk, 21
 - data types — types, 619
 - database functions — dbm
 - dbm_init, 396
 - delete, 396
 - fetch, 396
 - firstkey, 396
 - nextkey, 396
 - store, 396
 - database functions — ndbm, *continued*
 - dbm_firstkey, 400
 - dbm_nextkey, 400
 - dbm_open, 400
 - dbm_store, 400
 - database library
 - ldb option to cc, 396
 - ndbm, 400
 - date and time
 - get — time, 266
 - get — gettimeofday, 63
 - get — ftime, 266
 - set — settimeofday, 63
 - date and time conversion
 - asctime, 169
 - ctime, 169
 - dysize, 170
 - gmtime, 169
 - localtime, 169
 - timezone, 169
 - date and time conversion, System V
 - asctime, 358
 - ctime, 358
 - gmtime, 358
 - localtime, 358
 - date and time display — fdate, 182
 - dbm_clearerr — clear ndbm database error condition, 400
 - dbm_close — close ndbm routine, 400
 - dbm_delete — remove data from ndbm database, 400
 - dbm_err — ndbm database routine, 400
 - dbm_error — return ndbm database error condition, 400
 - dbm_fetch — fetch ndbm database data, 400
 - dbm_firstkey — access ndbm database, 400
 - dbm_nextkey — access ndbm database, 400
 - dbm_open — open ndbm database, 400
 - dbm_store — add data to ndbm database, 400
 - dbm_init — open database, 396
 - debugging memory management, 211 *thru* 212
 - malloc_debug — set debug level, 211
 - malloc_verify — verify heap, 211
 - debugging support — assert, 164
 - debugging support, System V — assert, 357
 - delete
 - directory — rmdir, 103
 - directory entry — unlink, 143
 - delete arp entry ioctl — SIOCDELR, 410
 - delete datum and key — delete, 396
 - delete descriptor, 28
 - delete — delete datum and key, 396
 - delete route ioctl — SIOCDELRT, 479
 - demount file system — unmount, 144
 - des — DES encryption chip interface, 423
 - descriptors
 - close, 28
 - delete, 28
 - dup, 33
 - dup2, 33
 - fcntl, 38
 - flock, 41
 - getdtablesize, 46
 - lockf, 206

descriptors, *continued*

- select, 104
- DESIOCBLOCK — process block, 423
- DESIOCQUICK — process quickly, 423
- destroy hash table — `hdestroy`, 201
- device controls — `ioctl`, 65
- devices, introduction to, 407 *thru* 408
- `dir` — directory format, 529
- directory
 - change current, 22
 - change root — `chroot`, 27
 - delete — `rmdir`, 103
 - erase — `rmdir`, 103
 - get entries, 44
 - make, 71
 - remove — `rmdir`, 103
 - scan, 231
- directory operations
 - `closedir`, 173
 - `opendir`, 173
 - `readdir`, 173
 - `rewinddir`, 173
 - `seekdir`, 173
 - `tellldir`, 173
- disk driver
 - `sd` — Adaptec ST-506, 480 *thru* 481
 - `ip` — Interphase, 442 *thru* 443
 - `si` — Sun SCSI, 480 *thru* 481
 - `xy` — Xylogics, 511 *thru* 512
- disk quotas — `quotactl`, 93
- `dkio` — disk control operations, 424 *thru* 425
- DKIOCGGEO — get disk geometry, 425
- DKIOCGPART — get disk partition info, 425
- DKIOCINFO — get disk info, 425
- DKIOCSGEO — set disk geometry, 425
- DKIOCSPART — set disk partition info, 425
- domain
 - get name of current — `getdomainname`, 45
 - set name of current — `setdomainname`, 45
- `drand48` — generate uniformly distributed random numbers, 175
- `drem` — IEEE floating-point function, 280
- `drum` — paging device, 426
- `dump` — incremental dump format, 531
- `dup`, 33
- `dup2`, 33
- duplicate descriptor, 33
- `dysize` — date and time conversion, 170

E

- E2BIG error number, 1
- EACCES error number, 2
- EADDRINUSE error number, 4
- EADDRNOTAVAIL error number, 4
- EAFNOSUPPORT error number, 4
- EAGAIN error number, 2
- EALREADY error number, 3
- EBADF error number, 1
- EBUSY error number, 2
- `ec` — 3Com 10 Mb/s Ethernet interface, 427
- ECHILD error number, 2
- ECONNABORTED error number, 4
- ECONNREFUSED error number, 4
- ECONNRESET error number, 4
- `ecvt` — convert number to ASCII, 177
- `edata` — end of program data, 178
- EDESTADDRREQ error number, 3
- EDOM error number, 3
- EDQUOT error number, 5
- EEXIST error number, 2
- EFAULT error number, 2
- EFBIG error number, 3
- effective group ID
 - get, 47
 - set, 113
- effective group ID, set — `setegid`, 234
- effective group ID, set, System V — `setegid`, 386
- effective user ID
 - get, 64
 - set — `setreuid`, 114
- effective user ID, set — `seteuid`, 234
- effective user ID, set, System V — `seteuid`, 386
- EHOSTDOWN error number, 5
- EHOSTUNREACH error number, 5
- EIDRM error number, 5
- EINPROGRESS error number, 3
- EINTR error number, 1
- EINVAL error number, 2
- EIO error number, 1
- EISCONN error number, 4
- EISDIR error number, 2
- ELOOP error number, 4
- EMFILE error number, 2
- EMLINK error number, 3
- EMSGSIZE error number, 3
- ENAMETOOLONG error number, 5
- encrypt — encryption, 168
- encryption
 - `crypt`, 168
 - `encrypt`, 168
 - `setkey`, 168
- encryption chip — `des`, 423
- `end` — end of program, 178
- end locations in program, 178
- `endfsent` — get file system descriptor file entry, 188
- `endgrent` — get group file entry, 189
- `endhostent` — get network host entry, 293
- `endmntent` — get filesystem descriptor file entry, 192
- `endnetent` — get network entry, 295
- `endnetgrent` — get network group entry, 297
- `endprotoent` — get protocol entry, 298
- `endpwent` — get password file entry, 198
- `endpwent` — get password file entry, System V, 374
- `endrpcent` — get RPC entry, 299
- `endservent` — get service entry, 300
- ENETDOWN error number, 4
- ENETRESET error number, 4
- ENETUNREACH error number, 4
- ENFILE error number, 2
- ENOBUFS error number, 4

- ENODEV error number, 2
 - ENOENT error number, 1
 - ENOEXEC error number, 1
 - ENOMEM error number, 2
 - ENOMSG error number, 5
 - ENOPROTOPT error number, 3
 - ENOSPC error number, 3
 - ENOTBLK error number, 2
 - ENOTCONN error number, 4
 - ENOTDIR error number, 2
 - ENOTEMPTY error number, 5
 - ENOTSOCK error number, 3
 - ENOTTY error number, 2
 - enquire stream status
 - clearerr — clear error on stream, 334
 - clearerr — clear error on stream, *System V*, 368
 - feof — enquire EOF on stream, 334
 - feof — enquire EOF on stream, *System V*, 368
 - ferror — inquire error on stream, 334
 - ferror — inquire error on stream, *System V*, 368
 - fileno — get stream descriptor number, 334
 - fileno — get stream descriptor number, *System V*, 368
 - environ — user environment, 533
 - environ — execute file, 179
 - environment
 - get value — *getenv*, 187
 - set value — *putenv*, 222
 - ENXIO error number, 1
 - EOPNOTSUPP error number, 4
 - EPERM error number, 1
 - EPFNOSUPPORT error number, 4
 - EPIPE error number, 3
 - EPROTONOSUPPORT error number, 3
 - EPROTOTYPE error number, 3
 - erand48 — generate uniformly distributed random numbers, 175
 - ERANGE error number, 3
 - erase
 - directory — *rmdir*, 103
 - directory entry — *unlink*, 143
 - erase — start new plot frame, 402
 - EREMOTE error number, 5
 - erf — error functions, 276
 - erfc — error functions, 276
 - EROFS error number, 3
 - errno — system error messages, 219
 - error messages, 219
 - ESHUTDOWN error number, 4
 - ESOCKTNOSUPPORT error number, 4
 - ESPIPE error number, 3
 - ESRCH error number, 1
 - ESTALE error number, 5
 - etext — end of program text, 178
 - Ethernet address mapping, 292
 - Ethernet address to ASCII — *ether_ntoa*, 292
 - Ethernet address to hostname — *ether_ntohost*, 292
 - Ethernet controller
 - ec — 10 Mb/s 3Com Ethernet interface, 427
 - ie — Sun Ethernet interface, 435
 - le — 10 Mb/s LANCE Ethernet interface, 460 *thru* 461
 - ethers file — Ethernet addresses, 534
 - ETIMEDOUT error number, 4
 - ETXTBSY error number, 3
 - Euclidean distance functions
 - cabs, 279
 - hypot, 279
 - EWOULDBLOCK error number, 3
 - EXDEV error number, 2
 - execl — execute file, 179
 - execle — execute file, 179
 - execlp — execute file, 179
 - execute file, 34, 179
 - environ, 179
 - execl, 179
 - execle, 179
 - execlp, 179
 - execv, 179
 - execvp, 179
 - execute regular expression — *re_exec*, 227
 - execution
 - suspend for interval, 239, 390
 - suspend for interval in microseconds, 254
 - execution accounting file — *acct*, 521
 - execution profile, prepare — *monitor*, 215
 - execv — execute file, 179
 - execve, 34
 - execvp — execute file, 179
 - exit, 37
 - exit — terminate process, 181
 - exp — exponential function, 277
 - exponent and mantissa, split into — *frexp*, 183
 - exponential function — *exp*, 277
 - external data representation routines, 307
- ## F
- fabs — absolute value, 278
 - fb — Sun console frame buffer driver, 428
 - fbio — frame buffers general properties, 429 *thru* 430
 - fchmod, 23
 - fchown, 25
 - fclose — close stream, 333
 - fcntl — file control system call, 38
 - fcntl — file control options, 536
 - fcvt — convert number to ASCII, 177
 - fdate — return date and time in ASCII format, 182
 - fdopen — associate descriptor, 335
 - fdopen — associate descriptor, *System V*, 369
 - feof — enquire EOF on stream, 334
 - feof — enquire EOF on stream, *System V*, 368
 - ferror — inquire error on stream, 334
 - ferror — inquire error on stream, *System V*, 368
 - fetch — retrieve datum under key, 396
 - fflush — flush stream, 333
 - ffs — find first one bit, 167
 - fgetc — get character from stream, 338
 - fgetc — get character from stream, *System V*, 372
 - fgetgrent — get group file entry, 189
 - fgetpwent — get password file entry, 198
 - fgetpwent — get password file entry, *System V*, 374

- `fgets` — get string from stream, 339
- file**
- `ftw` — traverse file tree, 185
 - create new, 31
 - create temporary name — `tmpnam`, 351
 - determine accessibility of, 15
 - execute, 34
 - make hard link to, 68
 - synchronize state — `fsync`, 43
- file attributes**
- `fstat`, 132
 - `lstat`, 132
 - `stat`, 132
- file control**
- options header file — `fcntl`, 536
 - system call — `fcntl`, 38
- file formats**, 515
- file position, move** — `lseek`, 70
- file system**
- access, 15
 - `chdir`, 22
 - `chmod`, 23
 - `chown`, 25
 - create file — `open`, 85
 - delete directory entry — `unlink`, 143
 - delete directory — `rmdir`, 103
 - unmount — demount file system, 144
 - erase directory entry — `unlink`, 143
 - erase directory — `rmdir`, 103
 - `fchmod`, 23
 - `fchown`, 25
 - format — `fs`, 537
 - `ftruncate`, 140
 - get file descriptor entry, 188
 - `getdirentries`, 44
 - link, 68
 - `lseek`, 70
 - `mkdir`, 71
 - `mknod`, 73
 - `mntent` — static information, 550
 - mount, 77
 - mounted table — `mtab`, 552
 - `open`, 85
 - `quotactl` — disk quotas, 93
 - `readlink`, 97
 - remove directory entry — `unlink`, 143
 - remove directory — `rmdir`, 103
 - rename file — `rename`, 101
 - statistics — `fstatfs`, 134
 - statistics — `statfs`, 134
 - `symlink`, 137
 - `tell`, 70
 - `truncate`, 140
 - `umask`, 141
 - unmount — demount file system, 144
 - `utimes` — set file times, 145
- file times, set** — `utime`, 269
- filename, change** — `rename`, 101
- fileno** — get stream descriptor number, 334
- fileno** — get stream descriptor number, Sysem V, 368
- files used by programs**
- `/etc/dumpdates` — dump record, 532
 - `/etc/ethers` — host ethernet map, 534
- files used by programs, continued**
- `/etc/exports` — list of filesystems accessible to clients, 535
 - `/etc/fstab` — table of filesystems to mount at boot, 551
 - `/etc/gettytab` — terminal characteristics for `getty`, 545
 - `/etc/group` — local group file, 546
 - `/etc/hosts.equiv` — list of trusted clients, 548, 553
 - `/etc/hosts` — host ID map, 547
 - `/etc/mntab` — table of mounted filesystems, 552
 - `/etc/netgroup` — network groups, 553
 - `/etc/networks` — DARPA Internet known networks, 555
 - `/etc/passwd` — password file, 557
 - `/etc/phones` — remote host phone numbers for `tip`, 558
 - `/etc/protocols` — DARPA Internet known protocols, 562
 - `/etc/remote` — remote host description file for `tip`, 565
 - `/etc/resolv.conf` — configuration file for name server, 566
 - `/etc/rmtab` — list of hosts with local filesystems mounted, 567
 - `/etc/servers` — list of Internet server processes, 572
 - `/etc/termcap` — terminal capabilities file, 589
 - `/etc/ttys` — list of terminals to start at boot, 617
 - `/etc/utmp` — login accounting, 621
 - `/etc/yp/domain/netgroup*` — list of network groups for YP domain, 553
 - `/etc/yp/group` — group YP map, 546
 - `/usr/5lib/terminfo` — directory of Sytem V terminal-description files, 602, 615
 - `/usr/adm/lastlog` — login accounting, 621
 - `/usr/adm/usracct` — login accounting, 621
 - `/usr/adm/wtmp` — login accounting, 621
 - `/usr/lib/crontab` — table of timed events, 528
 - `/usr/lib/fonts/fixedwidthfonts` — directory of fixed width (screen) font files, 623
 - `/usr/lib/term` — directory of `nroff` terminal-support files, 579
 - `/usr/lib/vgrindefs` — `vgrind` code formatting specifications, 625
 - `~/.netrc` — ftp remote login data, 554
- filesystem descriptor, get file entry**, 192
- find**
- first key in `dbm` database — `firstkey`, 396
 - first one bit — `ffs`, 167
 - name of terminal — `ttyname`, 251
 - next key in `dbm` database — `nextkey`, 396
- finite** — IEEE floating-point function, 280
- FIOASYNC** — set/clear async I/O, 408
- FIOCLEX** — set close-on-exec flag for `fd`, 408
- FIOGETOWN** — get file owner, 408
- FIONBIO** — set/clear non-blocking I/O, 408
- FIONCLEX** — remove close-on-exec flag, 408
- FIONREAD** — get # bytes to read, 408
- FIOSETOWN** — set file owner, 408
- firstkey** — find first key, 396
- floating point**
- `isinf` — test infinite value, 205
 - `isnan` — test not a number, 205
- flock**, 41
- floor** — floor of, 278
- flush buffers** `ioctl` — `TIOCFLUSH`, 502

flush stream — `fflush`, 333
 fopen — open stream, 335
 fopen — open stream, System V, 369
 fork a new process — `fork`, 42
 format of memory image file — `core`, 526
 formatted input conversion
 `fscanf` — convert from stream, 346
 `fscanf` — convert from stream, System V, 381
 `scanf` — convert from stdin, 346
 `scanf` — convert from stdin, System V, 381
 `sscanf` — convert from string, 346
 `sscanf` — convert from string, System V, 381
 .forward — mail forwarding file, 522
 fprintf — formatted output conversion, 341
 fprintf — format to stream, System V, 377
 fputc — put character on stream, 344
 fputs — put string to stream, 345
 frame buffer
 bwone — Sun-1 black and white frame buffer, 413
 bwtwo — Sun-3/Sun-2 black and white frame buffer, 414
 fread — read from stream, 336
 fread — read from stream, System V, 371
 free — free memory, 210
 free memory — `cfree`, 210
 free memory — `free`, 210
 free static block `ioctl` — `GP1IO_FREE_STATIC_BLOCK`, 431
 freopen — reopen stream, 335
 freopen — reopen stream, System V, 369
 frexp — split into mantissa and exponent, 183
 fs — file system format, 537
 fscanf — convert from stream, 346
 fscanf — convert from stream, System V, 381
 fseek — seek on stream, 337
 fspec text file tabstop specifications, 540
 fstat — obtain file attributes, 132
 fstatfs — obtain file system statistics, 134
 fsync — synchronize disk file with core image, 43
 ftell — get stream position, 337
 ftime — get date and time, 266
 ftok — interprocess communication routine, 184
 ftp — remote login data — `.netrc` file, 554
 ftpusers — ftp prohibited users list, 542
 ftruncate, 140
 ftw — traverse file tree, 185
 full-duplex connection, shut down — `shutdown`, 120
 fwrite — write to stream, 336
 fwrite — write to stream, System V, 371

G

gamma — log gamma, 282
 gather write — `writev`, 151
 gcd — multiple precision GCD, 398
 gcvt — convert number to ASCII, 177
 generate
 fault — `abort`, 162
 generate random numbers
 `initstate`, 225
 `rand`, 264

generate random numbers, *continued*

`random`, 225
 `setstate`, 225
 `srand`, 264
 `srandom`, 225
 `drand48`, 175
 `erand48`, 175
 `jrand48`, 175
 `lcong48`, 175
 `lrand48`, 175
 `mrnd48`, 175
 `nrnd48`, 175
 `seed48`, 175
 `srand48`, 175

generate random numbers, System V

`rand`, 380
 `srand`, 380

generic disk control operations — `dkio`, 424 *thru* 425

generic operations

 gather write — `writev`, 151
 `ioctl`, 65
 read, 95
 scatter read — `readv`, 95
 write, 151

get

 arp entry `ioctl` — `SIOCGARP`, 410
 character from stream — `fgetc`, 338
 character from stream — `getc`, 338
 console I/O `ioctl` — `TIOCCONS`, 418
 count of bytes to read `ioctl` — `FIONREAD`, 408
 current working directory pathname — `getwd`, 200
 date and time — `ftime`, 266
 date and time — `time`, 266
 disk geometry `ioctl` — `DKIOCGGEO`, 425
 disk info `ioctl` — `DKIOCINFO`, 425
 disk partition info `ioctl` — `DKIOCGPART`, 425
 entries from name list — `nlist`, 217
 environment value — `getenv`, 187
 file owner `ioctl` — `FIOGETOWN`, 408
 file system descriptor file entry, 188
 high water mark `ioctl` — `SIOCGHIWAT`, 485
 ifnet address `ioctl` — `SIOCGIFADDR`, 436
 ifnet flags `ioctl` — `SIOCGIFFLAGS`, 436
 ifnet list `ioctl` — `SIOCGIFCONF`, 436
 info on resource usage — `vtimes`, 271
 line discipline `ioctl` — `TIOCGTD`, 412, 495, 502
 local mode bits `ioctl` — `TIOCLGET`, 503
 local special chars `ioctl` — `TIOCGLTC`, 504
 login name — `getlogin`, 191
 low water mark `ioctl` — `SIOCGLOWAT`, 485
 network entry — `getnetent`, 295
 network group entry — `getnetgrent`, 297
 network host entry — `gethostent`, 293
 network service entry — `getservent`, 300
 number of characters in output queue `ioctl` — `TIOCOUTQ`, 502
 options on sockets — `getsockopt`, 62
 p-p address `ioctl` — `SIOCGIFDSTADDR`, 436
 parameters — `gty` `ioctl` — `TIOCGETP`, 501
 parent process identification — `getppid`, 55
 pathname of current working directory — `getcwd`, 186
 position of stream — `ftell`, 337
 process domain name — `getdomainname`, 45

get, continued

- process group of tty `ioctl` — `TIOCGGRP`, 502
- process identification — `getpid`, 55
- process times — `times`, 267
- protocol entry — `getprotoent`, 298
- requested minor device `ioctl` — `GPIO_GET_REQDEV`, 431
- restart count `ioctl` — `GPIO_GET_RESTART_COUNT`, 431
- RPC program entry — `getrpcent`, 299
- scheduling priority — `getpriority`, 56
- signal stack context — `sigstack`, 124
- special characters `ioctl` — `TIOCGETC`, 503
- static block `ioctl` — `GPIO_GET_STATIC_BLOCK`, 431
- string from `stdin` — `gets`, 339
- string from stream — `fgets`, 339
- tape status `ioctl` — `MTIOCGGET`, 468
- terminal state — `gtty`, 265
- true minor device `ioctl` — `GPIO_GET_TRUINORDEV`, 432
- user limits — `ulimit`, 268
- word from stream — `getw`, 338
- get character from stream, System V — `fgetc`, 372
- get character from stream, System V — `getc`, 372
- get date and time, 63
- get filesystem descriptor file entry
 - `admntent`, 192
 - `endmntent`, 192
 - `getmntent`, 192
 - `hasmntopt`, 192
 - `setmntent`, 192
- get group file entry
 - `endgrent`, 189
 - `fgetgrent`, 189
 - `getgrent`, 189
 - `getgrgid`, 189
 - `getgrnam`, 189
 - `setgrent`, 189
- get high water mark `ioctl` — `SIOCGHIWAT`, 506
- get keyboard “direct input” state `ioctl` — `KIOCGDIRECT`, 445
- get keyboard translation `ioctl` — `KIOCGTRANS`, 444
- get keyboard type `ioctl` — `KIOCTYPE`, 445
- get low water mark `ioctl` — `SIOCGLOWAT`, 506
- get option letter from argument vector — `getopt`, 194
- get password file entry
 - `endpwent`, 198
 - `fgetpwent`, 198
 - `getpwent`, 198
 - `getpwnam`, 198
 - `getpwuid`, 198
 - `setpwent`, 198
- get password file entry, System V
 - `endpwent`, 374
 - `fgetpwent`, 374
 - `getpwent`, 374
 - `getpwnam`, 374
 - `getpwuid`, 374
 - `setpwent`, 374
- get process times, System V — `times`, 391
- get translation table entry `ioctl` — `KIOCGETKEY`, 445
- get user name — `cuserid`, 332
- get word from stream, System V — `getw`, 372
- `getc` — get character from stream, 338
- `getc` — get character from stream, System V, 372
- `getchar` — get character from `stdin`, 338
- `getchar` — get character from `stdin`, System V, 372
- `getcwd` — get pathname of current directory, 186
- `getdirent`, 44
- `getdomainname` — get process domain, 45
- `getdtablesize`, 46
- `getegid` — get effective group ID, 47
- `getenv` — get value from environment, 187
- `geteuid` — get effective user ID, 64
- `getfsent` — get file system descriptor file entry, 188
- `getfsfile` — get file system descriptor file entry, 188
- `getfssize` — get file system descriptor file entry, 188
- `getfstype` — get file system descriptor file entry, 188
- `getgid` — get group ID, 47
- `getgrent` — get group file entry, 189
- `getgrgid` — get group file entry, 189
- `getgrnam` — get group file entry, 189
- `getgroups`, 48
- `gethostbyaddr` — get network host entry, 293
- `gethostbyname` — get network host entry, 293
- `gethostent` — get network host entry, 293
- `gethostid`, 49
- `gethostname`, 50
- `getitimer`, 51
- `getlogin` — get login name, 191
- `getmntent` — get filesystem descriptor file entry, 192
- `getnetbyaddr` — get network entry, 295
- `getnetbyname` — get network entry, 295
- `getnetent` — get network entry, 295
- `getnetgrent` — get network group entry, 297
- `getopt` — get option letter, 194
- `getpagesize`, 53
- `getpass` — read password, 196
- `getpass` — read password, System V, 373
- `getpeername`, 54
- `getpgrp`, 112
- `getpid`, 55
- `getppid`, 55
- `getpriority`, 56
- `getprotobyname` — get protocol entry, 298
- `getprotoent` — get protocol entry, 298
- `getpw` — get name from uid, 197
- `getpwent` — get password file entry, 198
- `getpwent` — get password file entry, System V, 374
- `getpwnam` — get password file entry, 198
- `getpwnam` — get password file entry, System V, 374
- `getpwuid` — get password file entry, 198
- `getpwuid` — get password file entry, System V, 374
- `getrlimit`, 57
- `getrpcbyname` — get RPC entry, 299
- `getrpcbynumber` — get RPC entry, 299
- `getrpcent` — get RPC entry, 299
- `getrpcport` — get RPC port number, 316
- `getrusage`, 59

gets — get string from stdin, 339
 getservbyname — get service entry, 300
 getservbyport — get service entry, 300
 getservent — get service entry, 300
 getsockname, 61
 getsockopt, 62
 gettimeofday, 63
 gettytab — terminal configuration data base, 543
 getuid — get user ID, 64
 getw — get word from stream, 338
 getw — get word from stream, System V, 372
 getwd — get current working directory pathname, 200
 gmtime — date and time conversion, 169
 gmtime — date and time conversion, System V, 358
 GP1IO_CHK_GP — restart GP, 431
 GP1IO_FREE_STATIC_BLOCK — free static block, 431
 GP1IO_GET_GBUFFER_STATE — check buffer state, 431
 GP1IO_GET_REQDEV — get requested minor device, 431
 GP1IO_GET_RESTART_COUNT — get restart count, 431
 GP1IO_GET_STATIC_BLOCK — get static block, 431
 GP1IO_GET_TRUMINORDEV — get true minor device, 432
 GP1IO_PUT_INFO — pass framebuffer info, 431
 GP1IO_REDIRECT_DEVFB — reconfigure fb, 431
 gpone — graphics processor interface, 431 *thru* 432
 graphics interface
 arc, 402
 circle, 402
 closepl, 402
 cont, 402
 erase, 402
 label, 402
 line, 402
 linemod, 402
 move, 402
 openpl, 402
 point, 402
 space, 402
 graphics interface files — plot, 559
 graphics processor interface — gpone, 431 *thru* 432
 group access list
 initialize — initgroups, 203
 group entry, network — getnetgrent, 297
 group — group file format, 546
 group file entry — getgrent, 189
 group ID
 get, 47
 get effective, 47
 set real and effective, 113
 groups access list, get — getgroups, 48
 groups access list, set — setgroups, 48
 gtty — get terminal state, 265

H

halt processor, 98
 hang up on last close ioctl — TIOCHPCL, 499, 502
 hangup, control terminal — vhangup, 148
 hard link to file — link, 68
 hardware support, introduction to, 407 *thru* 408
 hash table search routine — hsearch, 201

hasmntopt — get filesystem descriptor file entry, 192
 havedisk — disk inquiry of remote kernel, 323
 hcreate — create hash table, 201
 hdestroy — destroy hash table, 201
 host
 functions to convert to network byte order, 291
 get identifier of, 49
 get network entry — gethostent, 293
 get/set name — gethostname, 50
 phone numbers file — phones, 558
 hostname to Ethernet address — ether_hostton, 292
 hosts — host name data base, 547
 hosts.equiv — trusted hosts list, 548
 hsearch — hash table search routine, 201
 htonl — convert network to host long, 291
 htons — convert host to network short, 291
 hyperbolic functions
 cosh, 287
 sinh, 287
 tanh, 287
 hypot — Euclidean distance, 279

I

I/O, buffered binary
 fread — read from stream, 336
 fread — read from stream, System V, 371
 fwrite — write to stream, 336
 fwrite — write to stream, System V, 371
 icmp — Internet Control Message Protocol, 433 *thru* 434
 identifier of current host, get — gethostid, 49
 ie — Sun 10 Mb/s Ethernet interface, 435
 if — network interface general properties, 436 *thru* 437
 Ikon 10071-5 printer interface — vp, 507
 incremental dump format — dump, 531
 indeterminate floating point values, test for — isinf, 205
 index — find character in string, 241
 index memory characters — memchr, 213
 index strings — index, 241
 index strings — rindex, 241
 indirect system call, 139
 inet — Internet protocol family, 438 *thru* 439
 inet server database — servers, 572
 inet_addr — Internet address manipulation, 301
 inet_lnaof — Internet address manipulation, 301
 inet_makeaddr — Internet address manipulation, 301
 inet_netof — Internet address manipulation, 301
 inet_network — Internet address manipulation, 301
 inet_ntoa — Internet address manipulation, 301
 initgroups — initialize group access list, 203
 initialize group access list — initgroups, 203
 initiate
 connection on socket — connect, 29
 I/O to/from process — popen, 340
 initstate — random number routines, 225
 innetgr — get network group entry, 297
 input conversion
 fscanf — convert from stream, 346
 fscanf — convert from stream, System V, 381
 scanf — convert from stdin, 346
 scanf — convert from stdin, System V, 381

- input conversion, *continued*
 - sscanf — convert from string, 346
 - sscanf — convert from string, System V, 381
- input stream, push character back to — ungetc, 352
- inquire stream status
 - clearerr — clear error on stream, 334
 - clearerr — clear error on stream, System V, 368
 - feof — enquire EOF on stream, 334
 - feof — enquire EOF on stream, System V, 368
 - ferror — inquire error on stream, 334
 - ferror — inquire error on stream, System V, 368
 - fileno — get stream descriptor number, 334
 - fileno — get stream descriptor number, System V, 368
- insert element in queue — insque, 204
- insque — insert element in queue, 204
- integer absolute value — abs, 163
- Internet
 - control message protocol — icmp, 433 *thru* 434
 - protocol family — inet, 438 *thru* 439
 - Protocol — ip, 440 *thru* 441
 - to Ethernet address resolution — arp, 410 *thru* 411
 - Transmission Control Protocol — tcp, 484, 485
 - User Datagram Protocol — udp, 505, 506
- Internet address manipulation functions, 301
- Interphase SMD Disk driver — ip, 442 *thru* 443
- interprocess communication
 - accept connection — accept, 14
 - bind, 19
 - connect, 29
 - ftok, 184
 - getsockname, 61
 - getsockopt, 62
 - listen, 69
 - pipe, 88
 - recv, 99
 - recvfrom, 99
 - recvmsg, 99
 - send, 111
 - sendmsg, 111
 - sendto, 111
 - setsockopt, 62
 - shutdown, 120
 - socket, 129
 - socketpair, 131
- interrupts, release blocked signals — sigpause, 122
- interval timers
 - clock, 261
 - get, 51
 - set, 51
 - timerclear — macro, 51
 - timercmp — macro, 51
 - timerisset — macro, 51
- introduction
 - C library functions, 153 *thru* 160
 - compatibility library functions, 259
 - devices, 407 *thru* 408
 - file formats, 515
 - hardware support, 407 *thru* 408
 - mathematical library functions, 273
 - miscellaneous library functions, 393
 - network library functions, 289
 - RPC library functions, 313
 - special files, 407 *thru* 408
- introduction, *continued*
 - standard I/O library functions, 329
 - system calls, 1 *thru* 10
 - system error numbers, 1 *thru* 5
 - System V library functions, 355
- ioctl, 65
- ioctl's for des chip
 - DESIOCBLOCK — process block, 423
 - DESIOCFIX — process quickly, 423
- ioctls for disks
 - DKIOCGGEO — get disk geometry, 425
 - DKIOCGPART — get disk partition info, 425
 - DKIOCINFO — get disk info, 425
 - DKIOCSGEO — set disk geometry, 425
 - DKIOCSPART — set disk partition info, 425
- ioctl's for files
 - FIOASYNC — set/clear async I/O, 408
 - FIOCLEX — set close-on-exec for fd, 408
 - FIOGETOWN — get owner, 408
 - FIONBIO — set/clear non-blocking I/O, 408
 - FIONCLEX — remove close-on-exec flag, 408
 - FIONREAD — get # bytes to read, 408
 - FIOSETOWN — set owner, 408
- ioctl's for graphics processor
 - GP1IO_CHK_GP — restart GP, 431
 - GP1IO_FREE_STATIC_BLOCK — free static block, 431
 - GP1IO_GET_GBUFFER_STATE — check buffer state, 431
 - GP1IO_GET_REQDEV — get requested minor device, 431
 - GP1IO_GET_RESTART_COUNT — get restart count, 431
 - GP1IO_GET_STATIC_BLOCK — get static block, 431
 - GP1IO_GET_TRUMINORDEV — get true minor device, 432
 - GP1IO_PUT_INFO — pass framebuffer info, 431
 - GP1IO_REDIRECT_DEVFB — reconfigure fb, 431
- ioctl's for keyboards
 - KIOCCMD — send a keyboard command, 445
 - KIOCGDIRECT — get keyboard "direct input" state, 445
 - KIOCGKEY — get translation table entry, 445
 - KIOCGTRANS — get keyboard translation, 444
 - KIOCSDIRECT — set keyboard "direct input" state, 445
 - KIOCSETKEY — change translation table entry, 444
 - KIOCTRANS — set keyboard translation, 444
 - KIOCTYPE — get keyboard type, 445
- ioctl's for network disks
 - NDIOCCLEAR — clear user table, 471
 - NDIOCETHER — set ether address, 471
 - NDIOCSAT — server at ipaddress, 471
 - NDIOCSOFF — server off, 471
 - NDIOCSON — server on, 471
 - NDIOCUSER — set user parameters, 471
 - NDIOCVER — version number, 471
- ioctl's for sockets
 - SIOCADDRT — add route, 479
 - SIOCATMARK — at OOB mark?, 485
 - SIOCARP — delete arp entry, 410
 - SIOCDELRT — delete route, 479
 - SIOCGARP — get arp entry, 410
 - SIOCGHIWAT — get high water mark, 485, 506
 - SIOCGIFADDR — get ifnet address, 436
 - SIOCGIFCONF — get ifnet list, 436
 - SIOCGIFDSTADDR — get p-p address, 436
 - SIOCGIFFLAGS — get ifnet flags, 436
 - SIOCGLOWAT — get low water mark, 485, 506
 - SIOCSARP — set arp entry, 410

ioctl's for sockets, continued

- SIOCSHIWAT — set high water mark, 485, 506
- SIOCSIFADDR — set ifnet address, 436
- SIOCSIFDSTADDR — set p-p address, 436
- SIOCSIFFLAGS — set ifnet flags, 436
- SIOCSLOWAT — set low water mark, 485, 506

ioctl's for tapes

- MTIOCGET — get tape status, 468
- MTIOCTOP — tape operation, 467

ioctl's for terminals

- TIOCCBRK — clear break bit, 502
- TIOCCDTR — clear DTR, 502
- TIOCCONS — get console I/O, 418
- TIOCEXCL — set exclusive use of tty, 502
- TIOCFLUSH — flush buffers, 502
- TIOCGETC — get special characters, 503
- TIOCGETD — get line discipline, 412, 495, 502
- TIOCGETP — get parameters — gtty, 501
- TIOCG LTC — get local special chars, 504
- TIOCGPGRP — get process group of tty, 502
- TIOCHPCL — hang up on last close, 499, 502
- TIOCLBIC — bit clear local mode bits, 503
- TIOCLBIS — bit set local mode bits, 503
- TIOCLGET — get local mode bits, 503
- TIOCLSET — set local mode bits, 503
- TIOCNOTTY — void tty association, 495
- TIOCNXCL — remove exclusive use of tty, 502
- TIOCOUTQ — get number of characters in output queue, 502
- TIOCPKT — set/clear packet mode (pty), 477
- TIOCREMOTE — remote input editing, 477
- TIOCSBRK — set break bit, 502
- TIOCS DTR — set DTR, 502
- TIOCSETC — set special characters, 503
- TIOCSETD — set line discipline, 412, 495, 502
- TIOCSETN — set parameters, 501
- TIOCSETP — set parameters — gtty, 501
- TIOCS LTC — set local special chars, 504
- TIOCS PGRP — set process group of tty, 502
- TIOCSTART — start output (like control-Q), 477, 502
- TIOCSI — simulate terminal input, 502
- TIOCS STOP — stop output (like control-S), 477, 502

IP raw sockets, 441

- ip — Internet Protocol, 440 *thru* 441
- ip — Interphase SMD Disk driver, 442 *thru* 443
- isalnum — is character alphanumeric, 171
- isalnum — is character alphanumeric, System V, 360
- isalpha — is character letter, 171
- isalpha — is character letter, System V, 360
- isascii — is character ASCII, 171
- isascii — is character ASCII, System V, 360
- isatty — test if device is terminal, 251
- isctrl — is character control, 171
- isctrl — is character control, System V, 360
- isdigit — is character digit, 171
- isdigit — is character digit, System V, 360
- isgraph — is character graphic, 171
- isgraph — is character graphic, System V, 360
- isinf — test infinite value, 205
- islower — is character lower-case, 171
- islower — is character lower-case, System V, 360
- isnan — test not a number, 205
- isprint — is character printable, 171

- isprint — is character printable, System V, 360
- ispunct — is character punctuation, 171
- ispunct — is character punctuation, System V, 360
- isspace — is character whitespace, 171
- isspace — is character whitespace, System V, 360
- issue shell command — system, 247
- isupper — is character upper-case, 171
- isupper — is character upper-case, System V, 360
- isxdigit — is character hex digit, 171
- isxdigit — is character hex digit, System V, 360
- itom — integer to multiple precision, 398

J

- j0 — Bessel function, 281
- j1 — Bessel function, 281
- jn — Bessel function, 281
- jrand48 — generate uniformly distributed random numbers, 175

K

- kb — Sun keyboard
- kbd — Sun keyboard
- kill — send signal to process, 66
- killpg — send signal to process group, 67
- KIOCCMD — send a keyboard command, 445
- KIOCGDIRECT — get keyboard “direct input” state, 445
- KIOCGKEY — get translation table entry, 445
- KIOCGTRANS — get keyboard translation, 444
- KIOCSDIRECT — set keyboard “direct input” state, 445
- KIOCSKEY — change translation table entry, 444
- KIOCTRANS — set keyboard translation, 444
- KIOCTYPE — get keyboard type, 445
- kmem — kernel memory space, 463

L

- l64a — convert base-64 ASCII, 161
- label — plot label, 402
- LANCE 10 Mb/s Ethernet interface — le, 460 *thru* 461
- last locations in program, 178
- lastlog — login records, 621
- lcong48 — generate uniformly distributed random numbers, 175
- ldexp — split into mantissa and exponent, 183
- le — Sun-3/50 10 Mb/s Ethernet interface, 460 *thru* 461
- lfind — linear search routine, 208
- library file format — ar, 525
- library functions
 - introduction to C, 153 *thru* 160
 - introduction to compatibility, 259
 - introduction to mathematical, 273
 - introduction to miscellaneous, 393
 - introduction to network, 289
 - introduction to RPC, 313
 - introduction to standard I/O, 329
 - introduction to System V, 355
- limits
 - get for user — ulimit, 268
 - set for user — ulimit, 268
- line discipline — bk, 412
- line discipline *ioctl's*
 - TIOCGETD — get line discipline, 412

- line discipline *ioctl's, continued*
 TIOCSETD — set line discipline, 412
- line — plot line, 402
- line to Ethernet address — *ether_line*, 292
- linear search and update routine — *lsearch*, 208
- linear search routine — *lfind*, 208
- linemod — set line style, 402
- link, 68
 make symbolic, 137
 read value of symbolic, 97
- link editor output — *a.out*, 517
- listen, 69
- lo — software loopback network interface, 462
- localtime — date and time conversion, 169
- localtime — date and time conversion, 358
- lock
 file — *flock*, 41
 record — *fcntl*, 38, 206
- lockf, 206
- log — natural logarithm, 277
- log gamma function — *gamma*, 282
- log10 — logarithm, base 10, 277
- logarithm, base 10 — *log10*, 277
- logarithm, natural — *log*, 277
- logb — IEEE floating-point function, 280
- environ", 533
- login name, get — *getlogin*, 191
- login records
 lastlog file, 621
 utmp file, 621
 wtmp file, 621
- long jmp — non-local goto, 232
- lrand48 — generate uniformly distributed random numbers, 175
- lsearch* — linear search and update routine, 208
- lseek* — move file position, 70
- lstat* — obtain file attributes, 132
- ## M
- machine-dependent values — *values*, 255
- machine-machine communication line discipline — *bk*, 412
- madd — multiple precision add, 398
- magic file — file command's magic numbers table, 549
- magnetic tape *ioctl's*
 MTIOCGET — get tape status, 468
 MTIOCTOP — tape operation, 467
- make
 special file, 73
- make directory, 71
- make hard link to file, 68
- malloc — allocate memory, 210
- malloc_debug — set debug level, 211
- malloc_verify — verify heap, 211
- manipulate Internet addresses, 301
- mantissa and exponent, split into — *frexp*, 183
- map memory pages — *mmap*, 75
- mask, set current signal — *sigsetmask*, 123
- mathematical functions
 acos, 286
- mathematical functions, *continued*
 asin, 286
 atan, 286
 atan2, 286
 cabs, 279
 ceil — ceiling of, 278
 cos, 286
 cosh, 287
 exp — exponential, 277
 fabs — absolute value, 278
 floor — floor of, 278
 gamma, 282
 hypot, 279
 j0, 281
 j1, 281
 jn, 281
 log — natural logarithm, 277
 log10 — logarithm, base 10, 277
 pow — raise to power, 277
 sin, 286
 sinh, 287
 tan, 286
 tanh, 287
 y0, 281
 y1, 281
 yn, 281
- mathematical library functions, introduction to, 273
- matherr — math library error-handline routine, 283
- mbio — Multibus I/O space, 463
- mbmem — Multibus memory space, 463
- mdiv — multiple precision divide, 398
- mem — main memory space, 463
- memalign — allocate aligned memory, 210
- memcpy — copy memory character strings, 213
- memchr — index memory characters, 213
- memcmp compare memory characters, 213
- memcpy copy memory character fields, 213
- memory allocation debugging, 211 *thru* 212
- memory image file format — *core*, 526
- memory images
 kmem — kernel memory space, 463
 mbio — Multibus I/O space, 463
 mbmem — Multibus memory space, 463
 mem — main memory space, 463
 virtual — virtual address space, 463
 vme16 — VMEbus 16-bit space, 463
 vme16d16 — VMEbus address space, 463
 vme16d32 — VMEbus address space, 463
 vme24 — VMEbus 24-bit space, 463
 vme24d16 — VMEbus address space, 463
 vme24d32 — VMEbus address space, 463
 vme32d16 — VMEbus address space, 463
 vme32d32 — VMEbus address space, 463
- memory management, 210 *thru* 212
 alloca — allocate on stack, 211
 brk — set data segment break, 21
 calloc — allocate memory, 210
 cfree — free memory, 210
 free — free memory, 210
 getpagesize, 53
 malloc — allocate memory, 210
 malloc_debug — set debug level, 211

- memory management, *continued*
 - `malloc_verify` — verify heap, 211
 - `memalign` — allocate aligned memory, 210
 - `mmap`, 75
 - `realloc` — reallocate memory, 210
 - `sbrk` — change data segment size, 21
 - `valloc` — allocate aligned memory, 211
 - memory management debugging, 211 *thru* 212
 - memory operations, 213
 - `memset` assign to memory characters, 213
 - message
 - receive from socket — `recv`, 99
 - send from socket — `send`, 111
 - message control operations
 - `msgctl`, 79
 - `msgget`, 80
 - `msgsnd`, 81
 - messages
 - system error, 219
 - system signal, 221
 - `mfree` — release multiple precision storage, 398
 - `min` — multiple precision decimal input, 398
 - miscellaneous library functions, introduction to, 393
 - `mkdir`, 71
 - `mknod`, 73
 - `mktemp` — make unique file name, 214
 - `mmap`, 75
 - `mntent` — file system static information, 550
 - `modf` — split into mantissa and exponent, 183
 - `moncontrol` — make execution profile, 215
 - `monitor` — make execution profile, 215
 - monitor traffic on the Ethernet, 314
 - monochrome frame buffer — `bwone`, 413
 - monochrome frame buffer — `bwtwo`, 414
 - `monstartup` — make execution profile, 215
 - `mount`, 77
 - mounted file system table — `mtab`, 552
 - mouse — Sun mouse, 464
 - `mout` — multiple precision decimal output, 398
 - move file position — `lseek`, 70
 - move — move current point, 402
 - `rand48` — generate uniformly distributed random numbers, 175
 - `msgctl`, 79
 - `msgget`, 80
 - `msgsnd`, 81
 - `msqrt` — multiple precision exponential, 398
 - `msub` — multiple precision subtract, 398
 - `mtab` — mounted file system table, 552
 - `mti` — Systech MTI-800/1600 multi-terminal interface, 465 *thru* 466
 - `mtio` — UNIX magnetic tape interface, 467 *thru* 468
 - `MTIOCGET` — get tape status, 468
 - `MTIOCTOP` — tape operation, 467
 - `mtox` — multiple precision to hexadecimal string, 398
 - `mult` — multiple precision multiply, 398
 - multiple precision integer arithmetic
 - `gcd`, 398
 - `itom`, 398
 - `madd`, 398
 - `mdiv`, 398
 - multiple precision integer arithmetic, *continued*
 - `mfree`, 398
 - `min`, 398
 - `mout`, 398
 - `msqrt`, 398
 - `msub`, 398
 - `mtox`, 398
 - `mult`, 398
 - `pow`, 398
 - `rpow`, 398
 - `sdiv`, 398
 - `xtom`, 398
- ## N
- name list, get entries from — `nlist`, 217
 - name of terminal, find — `ttyname`, 251
 - name termination handler — `on_exit`, 218
 - natural logarithm — `log`, 277
 - `nd` — network disk driver, 469 *thru* 471
 - `NDIOCCLEAR` — clear user table, 471
 - `NDIOCETHER` — set ether address, 471
 - `NDIOCSAT` — server at ipaddress, 471
 - `NDIOCSOFF` — server off, 471
 - `NDIOCSON` — server on, 471
 - `NDIOCUSER` — set user parameters, 471
 - `NDIOCVER` — version number, 471
 - `netgroup` — network groups list, 553
 - `.netrc` — ftp remote login data file, 554
 - network byte order
 - function to convert to host, 291
 - network disk `ioctl`'s
 - `NDIOCCLEAR` — clear user table, 471
 - `NDIOCETHER` — set ether address, 471
 - `NDIOCSAT` — server at ipaddress, 471
 - `NDIOCSOFF` — server off, 471
 - `NDIOCSON` — server on, 471
 - `NDIOCUSER` — set user parameters, 471
 - `NDIOCVER` — version number, 471
 - network entry, get — `getnetent`, 295
 - network file system daemons, 84
 - network group entry
 - `get`, 297
 - network host entry, get — `gethostent`, 293
 - network interface `ioctl`'s
 - `SIOCGIFADDR` — get ifnet address, 436
 - `SIOCGIFCONF` — get ifnet list, 436
 - `SIOCGIFDSTADDR` — get p-p address, 436
 - `SIOCGIFFLAGS` — get ifnet flags, 436
 - `SIOCSIFADDR` — set ifnet address, 436
 - `SIOCSIFDSTADDR` — set p-p address, 436
 - `SIOCSIFFLAGS` — set ifnet flags, 436
 - network interface tap protocol — `nit`, 473 *thru* 475
 - network library functions, introduction to, 289
 - network loopback interface — `lo`, 462
 - network packet routing device — `routing`, 479
 - network service entry, get — `getservent`, 300
 - network services status monitor files, 575
 - networks — network name data base, 555
 - `nextkey` — find next key, 396
 - NFS exported file systems — `exports`, 535
 - NFS, network file system protocol, 472

- nfssvc, 84
 - nice — change priority of a process, 262
 - nice — change priority of a process, System V, 376
 - nit — network interface tap protocol, 473 *thru* 475
 - nlist — get entries from name list, 217
 - non-local goto
 - non-local goto — longjmp, 232
 - non-local goto — setjmp, 232
 - nrnd48 — generate uniformly distributed random numbers, 175
 - ntohl — convert network to host long, 291
 - ntohs — convert host to network short, 291
 - null — null device, 476
 - null-terminated strings
 - compare — strcmp, 241
 - compare — strncmp, 241
 - concatenate — strcat, 241
 - concatenate — strncat, 241
 - copy — strcpy, 241
 - copy — strncpy, 241
 - index — index, 241
 - index — rindex, 241
 - reverse index — rindex, 241
 - numbers, convert to strings — ecvt, 177
- O**
- on_exit — name termination handler, 218
 - open, 85
 - open database — dbmopen, 396
 - open directory stream — opendir, 173
 - open stream — fopen, 335
 - open stream, System V — fopen, 369
 - opendir — open directory stream, 173
 - openlog — initialize system log file, 246
 - openpl — open plot device, 402
 - optarg — get option letter, 194
 - optind — get option letter, 194
 - option letter, get from argument vector — getopt, 194
 - options on sockets
 - get, 62
 - set, 62
 - output conversion
 - fprintf — convert to stream, 341
 - fprintf — convert to stream, System V, 377
 - printf — convert to stdout, 341
 - printf — convert to stdout, System V, 377
 - sprintf — convert to string, 341
 - sprintf — convert to string, System V, 377
- P**
- packet routing device — routing, 479
 - packet routing ioctl's
 - SIOCADDRT — add route, 479
 - SIOCDELRT — delete route, 479
 - page size, get — getpagesize, 53
 - paging device — swapon, 136, 426
 - paging system, advise — vadvise, 146
 - parent process identification, get — getppid, 55
 - pass framebuffer info ioctl — GP1IO_PUT_INFO, 431
 - passwd — password file, 556
 - password
 - password, *continued*
 - read — getpass, 196
 - read, System V — getpass, 373
 - password file
 - add entry — putpwent, 223
 - get entry — endpwent, 198
 - get entry, System V — endpwent, 374
 - get entry — fgetpwent, 198
 - get entry, System V — fgetpwent, 374
 - get entry — getpwent, 198
 - get entry, System V — getpwent, 374
 - get entry — getpwnam, 198
 - get entry, System V — getpwnam, 374
 - get entry — getpwuid, 198
 - get entry, System V — getpwuid, 374
 - get entry — setpwent, 198
 - get entry, System V — setpwent, 374
 - pause — stop until signal, 263
 - pclose — close stream to process, 340
 - peer name, get — getpeername, 54
 - periodic jobs table — crontab, 528
 - perror — system error messages, 219
 - phones — remote host phone numbers, 558
 - pipe, 88
 - plot — graphics interface files, 559
 - point — plot point, 402
 - popen — open stream to process, 340
 - position of directory stream — telldir, 173
 - pow — raise to power, 277, 398
 - power function — pow, 277
 - prepare execution profile
 - moncontrol — make execution profile, 215
 - monitor — make execution profile, 215
 - monstartup — make execution profile, 215
 - primitive system data types — types, 619
 - printcap — printer capability data base, 560
 - printer interface
 - vp — Ikon 10071-5 Versatec parallel printer interface, 507
 - vpc — Systech VPC-2200 Versatec/Centronics interface, 508
 - printf — formatted output conversion, 341
 - printf — format to stdout, System V, 377
 - priority
 - get, 56
 - set, 56
 - priority of process — nice, 262
 - priority of process, System V — nice, 376
 - process
 - and child process times, System V — times, 391
 - create, 42
 - get identification — getpid, 55
 - get times — times, 267
 - initiate I/O to/from, 340
 - priority — nice, 262
 - priority, System V — nice, 376
 - send signal to — kill, 66
 - software signals — sigvec, 125 *thru* 127
 - terminate, 37
 - terminate and cleanup — exit, 181
 - process block ioctl — DESIOCBLOCK, 423
 - process group
 - get — getpgrp, 112

process group, *continued*
 send signal to — `killpg`, 67
 set — `setpgrp`, 112

process quickly `ioctl` — `DESIQCQUICK`, 423

process tracing — `ptrace`

processes and protection

- `execve`, 34
- `exit`, 37
- `fork`, 42
- `getdomainname`, 45
- `getegid`, 47
- `geteuid`, 64
- `getgid`, 47
- `getgroups`, 48
- `gethostid`, 49
- `gethostname`, 50
- `getpgrp`, 112
- `getpid`, 55
- `getppid`, 55
- `getuid`, 64
- `ptrace`, 90 *thru* 92
- `setdomainname`, 45
- `setgroups`, 48
- `sethostname`, 50
- `setpgrp`, 112
- `setregid`, 113
- `setreuid`, 114
- `vfork`, 147
- `vhangup`, 148
- `wait`, 149
- `wait3`, 149

`prof` — profile within a function, 220

`profil`, 89

profile, execution — `monitor`, 215

`prof`, 220

program verification — `assert`, 164

program verification, System V — `assert`, 357

protocol entry
 `get`, 298

protocols — protocol name data base, 562

`psignal` — system signal messages, 221

`ptrace`, 90 *thru* 92

`pty` — pseudo terminal driver, 477 *thru* 478

push character back to input stream — `ungetc`, 352

put character to stdout — `putchar`, 344

put character to stream — `fputc`, 344

put character to stream — `putc`, 344

put string to stdout — `puts`, 345

put string to stream — `fputs`, 345

put word to stream — `putw`, 344

`putc` — put character on stream, 344

`putchar` — put character on stdout, 344

`putenv` — set environment value, 222

`putpwent` — add password file entry, 223

`puts` — put string to stdout, 345

`putw` — put word on stream, 344

Q

`qsort` — quicker sort, 224

queue
 insert element in — `insque`, 204

queue, *continued*
 remove element from — `remque`, 204

quicker sort — `qsort`, 224

`quotact1` — disk quotas, 93

R

`rand` — generate random numbers, 264

`rand` — generate random numbers, System V, 380

random — generate random number, 225

random number generator

- `drand48`, 175
- `erand48`, 175
- `initstate`, 225
- `jrand48`, 175
- `lcong48`, 175
- `lrand48`, 175
- `mrnd48`, 175
- `nrnd48`, 175
- `rand`, 264
- random, 225
- `seed48`, 175
- `setstate`, 225
- `srnd`, 264
- `srnd48`, 175
- `srandom`, 225

random number generator, System V

- `rand`, 380
- `srnd`, 380

rasterfile, 563

raw IP sockets, 441

`rcmd` — execute command remotely, 303

`re_comp` — compile regular expression, 227

`re_exec` — execute regular expression, 227

`read`, 95

read directory stream — `readdir`, 173

read formatted

- `fscanf` — convert from stream, 346
- `fscanf` — convert from stream, System V, 381
- `scanf` — convert from stdin, 346
- `scanf` — convert from stdin, System V, 381
- `sscanf` — convert from string, 346
- `sscanf` — convert from string, System V, 381

read from stream — `fread`, 336

read from stream, System V — `fread`, 371

read password — `getpass`, 196

read password, System V — `getpass`, 373

read scattered — `readv`, 95

read/write pointer, move — `lseek`, 70

`readdir` — read directory stream, 173

`readlink`, 97

real group ID
 set, 113

real group ID, set — `setrgid`, 234

real group ID, set, System V — `setrgid`, 386

real user ID
 `get` — `getuid`, 64
 set — `setreuid`, 114

real user ID, set — `setruid`, 234

real user ID, set, System V — `setruid`, 386

`realloc` — reallocate memory, 210

reallocate memory — `realloc`, 210

- reboot — halt processor, 98
 - receive message from socket, 99
 - reconfigure fb `ioctl` — `GP1IO_REDIRECT_DEVFB`, 431
 - recv — receive message from socket, 99
 - recvfrom, 99
 - recvmsg, 99
 - regex — regular expression compile and match routines, 228
 - regular expressions
 - compile — `re_comp`, 227
 - execute — `re_exec`, 227
 - release blocked signals — `sigpause`, 122
 - remote command, return stream to — `rcmd`, 303
 - remote command, return stream to — `rexec`, 304
 - remote execution protocol — `rex`, 317
 - remote — remote host descriptions, 564
 - remote host
 - number of users — `rusers`, 319
 - phone numbers — `phones`, 558
 - remote input editing `ioctl` — `TIOCREMOTE`, 477
 - remote kernel performance, 323
 - remote procedure calls, 305
 - remote users, number of — `rnusers`, 319
 - remove
 - close-on-exec flag `ioctl` — `FIONCLEX`, 408
 - directory — `rmdir`, 103
 - directory entry — `unlink`, 143
 - element from queue — `remque`, 204
 - exclusive use of tty `ioctl` — `TIOCNXCL`, 502
 - file system — `unmount`, 144
 - `remque` — remove element from queue, 204
 - rename file — `rename`, 101
 - reopen stream — `freopen`, 335
 - reopen stream, System V — `freopen`, 369
 - reposition stream
 - `fseek`, 337
 - `ftell`, 337
 - `rewind`, 337
 - resolver file — name server initialization info, 566
 - resource consumption, control — `vlimit`, 270
 - resource control
 - `getrlimit`, 57
 - `getrusage`, 59
 - `setrlimit`, 57
 - resource controls
 - `getpriority`, 56
 - `setpriority`, 56
 - resource usage, get information about — `vtimes`, 271
 - resource utilization, get information about — `getrusage`, 59
 - restart GP `ioctl` — `GP1IO_CHK_GP`, 431
 - restart output `ioctl` — `TIOCSTART`, 502
 - retrieve datum under key — `fetch`, 396
 - return stream to remote command — `rcmd`, 303
 - return stream to remote command — `rexec`, 304
 - return to saved environment — `longjmp`, 232
 - reverse index strings — `rindex`, 241
 - rewind directory stream — `rewinddir`, 173
 - rewind — rewind stream, 337
 - rewind stream — `rewind`, 337
 - `rewinddir` — rewind directory stream, 173
 - `rexec` — return stream to remote command, 304
 - `rindex` — find character in string, 241
 - `rmdir` — remove directory, 103
 - `rmtab` — remote mounted file system table, 567
 - root directory, change — `chroot`, 27
 - routing — local network packet routing, 479
 - routing `ioctl`'s
 - `SIOCADDRT` — add route, 479
 - `SIOCDELRT` — delete route, 479
 - RPC routines, 305
 - RPC library functions, introduction to, 313
 - RPC program entry, get — `getrpcent`, 299
 - `rpc` — rpc name data base, 568
 - `rpow` — multiple precision exponential, 398
 - `rresvport` — get privileged socket, 303
 - `rstat` — performance data from remote kernel, 323
 - `ruserok` — authenticate user, 303
 - `rwall` — write to specified remote machines, 325
- ## S
- save stack environment — `setjmp`, 232
 - `sbrk` — change data segment size, 21
 - `scalb` — IEEE floating-point function, 280
 - scan directory — `alphasort`, 231
 - scan directory — `scandir`, 231
 - `scandir` — scan directory, 231
 - `scanf` — convert from stdin, 346
 - `scanf` — convert from stdin, System V, 381
 - scatter read — `readv`, 95
 - `scsfile` — SCCS file format, 569
 - schedule signal in microsecond precision — `ualarm`, 253, 260
 - scheduling priority
 - `get`, 56
 - `set`, 56
 - `sd` — Adaptec ST-506 Disk driver, 480 *thru* 481
 - `sdiv` — multiple precision divide, 398
 - search functions
 - `bsearch` binary search, 165
 - `hsearch` — hash table search, 201
 - `lsearch` — linear search and update, 208
 - `seed48` — generate uniformly distributed random numbers, 175
 - seek in directory stream — `seekdir`, 173
 - seek on stream — `fseek`, 337
 - `seekdir` — seek in directory stream, 173
 - `select`, 104
 - semaphore
 - control — `semctl`, 105
 - get set of — `semget`, 107
 - operations — `semop`, 108
 - `semctl` — semaphore controls, 105
 - `semget` — get semaphore set, 107
 - `semop` — semaphore operations, 108
 - send
 - message from socket — `send`, 111
 - signal to process — `kill`, 66
 - signal to process group — `killpg`, 67
 - send a keyboard command `ioctl` — `KIOCCMD`, 445
 - sendmail aliases file — `aliases`, 522
 - sendmail aliases file — `.forward`, 522
 - `sendmsg` — send message over socket, 111

- sendto — send message to socket, 111
 PAGE, 513
 server at ipaddress ioctl — NDIOSAT, 471
 server off ioctl — NDIOSOFF, 471
 server on ioctl — NDIOSON, 471
 servers — inet server database, 572
 service entry, get — getservent, 300
 inet server database — services, 573
 set
 arp entry ioctl — SIOCSARP, 410
 break bit ioctl — TIOCSBRK, 502
 close-on-exec for fd ioctl — FIOCLEX, 408
 current signal mask — sigsetmask, 123
 date and time — gettimeofday, 63
 disk geometry ioctl — DKIOCSGEM, 425
 disk partition info ioctl — DKIOCSPART, 425
 DTR ioctl — TIOCSDTR, 502
 environment value — putenv, 222
 ether address ioctl — NDIOCETHER, 471
 file creation mode mask — umask, 141
 file owner ioctl — FIOSETOWN, 408
 file times — utime, 269
 high water mark ioctl — SIOCSEHIWAT, 485
 ifnet address ioctl — SIOCSIFADDR, 436
 ifnet flags ioctl — SIOCSIFFLAGS, 436
 line discipline ioctl — TIOCSETD, 412
 low water mark ioctl — SIOCSLOWAT, 485
 network group entry — setnetgrent, 297
 network service entry — getservent, 300
 p-p address ioctl — SIOCSIFDSTADDR, 436
 process domain name — setdomainname, 45
 RPC program entry — setrpcnt, 299
 setpriority", 56
 signal stack context — sigstack, 124
 terminal state — stty, 265
 user limits — ulimit, 268
 user mask — umask, 141
 user parameters ioctl — NDIOCUSER, 471
 set exclusive use of tty ioctl — TIOCEXCL, 502
 set high water mark ioctl — SIOCSEHIWAT, 506
 set keyboard "direct input" state ioctl — KIOCSDIRECT, 445
 set keyboard translation ioctl — KIOCTRANS, 444
 set line discipline ioctl — TIOCSETD, 495, 502
 set local mode bits ioctl — TIOCLSET, 503
 set local special chars ioctl — TIOCSLTC, 504
 set low water mark ioctl — SIOCSLOWAT, 506
 set options sockets, 62
 set parameters — gtty ioctl — TIOCSETP, 501
 set parameters ioctl — TIOCSETN, 501
 set process group of tty ioctl — TIOCSPGRP, 502
 set special characters ioctl — TIOCSETC, 503
 set/clear
 async I/O ioctl — FIOASYNC, 408
 non-blocking I/O ioctl — FIONBIO, 408
 packet mode (pty) ioctl — TIOCPKT, 477
 setbuf — assign buffering, 348
 setbuf — assign buffering, System V, 384
 setbuffer — assign buffering, 348
 setbuffer — assign buffering, System V, 384
 setdomainname — set process domain, 45
 setegid — set effective group ID, 234
 setegid — set effective group ID, System V, 386
 seteuid — set effective user ID, 234
 seteuid — set effective user ID, System V, 386
 setfsent — get file system descriptor file entry, 188
 setgid — set group ID, 234
 setgid — set group ID, System V, 386
 setgrent — get group file entry, 189
 setgroups, 48
 sethostent — get network host entry, 293
 sethostname, 50
 setitimer, 51
 setjmp — save stack environment, 232
 setjmp — non-local goto, 232
 setkey — encryption, 168
 setlinebuf — assign buffering, 348
 setlinebuf — assign buffering, System V, 384
 setmntent — get filesystem descriptor file entry, 192
 setnetent — get network entry, 295
 setnetgrent — get network group entry, 297
 setpgrp, 112
 setpriority, 56
 setprotoent — get protocol entry, 298
 setpwent — get password file entry, 198
 setpwent — get password file entry, System V, 374
 setregid, 113
 setreuid, 114
 setrgid — set real group ID, 234
 setrgid — set real group ID, System V, 386
 setrlimit, 57
 setrpcnt — get RPC entry, 299
 setruid — set real user ID, 234
 setruid — set real user ID, System V, 386
 setservent — get service entry, 300
 setsockopt, 62
 setstate — random number routines, 225
 settimeofday, 63
 setuid — set user ID, 234
 setuid — set user ID, System V, 386
 setvbuf — assign buffering, 348
 setvbuf — assign buffering, System V, 384
 shared memory
 control — shmctl, 115
 get segment — shmget, 116
 operation — shmop, 118
 shell command, issuing — system, 247
 shmctl — shared memory control, 115
 shmget — get shared memory segment, 116
 shmop — get shared memory operations, 118
 shutdown, 120
 si — Sun SCSI Disk driver, 480 *thru* 481
 sigblock, 121
 siginterrupt — interrupt system calls with software signal, 235
 signal
 schedule in microsecond precision — ualarm, 253, 260
 stop until — pause, 263
 signal — software signals, 236, 240

- signal — software signals, System V, 387
- signal messages
 - psignal, 221
 - sys_siglist, 221
- signals
 - kill, 66
 - killpg — send to process group, 67
 - sigblock, 121
 - sigpause, 122
 - sigsetmask, 123
 - sigstack — signal stack context, 124
 - sigvec, 125 *thru* 127
- sigpause, 122
- sigsetmask, 123
- sigstack — signal stack context, 124
- sigvec — software signals, 125 *thru* 127
- simulate terminal input `ioctl` — TIOCSSTI, 502
- sin — trigonometric sine, 286
- sinh — hyperbolic sine, 287
- SIOCADDRT — add route, 479
- SIOCATMARK — at OOB mark?, 485
- SIOCDELR — delete arp entry, 410
- SIOCDELRT — delete route, 479
- SIOCGARP — get arp entry, 410
- SIOCGHIWAT — get high water mark, 485, 506
- SIOCGIFADDR — get ifnet address, 436
- SIOCGIFCONF — get ifnet list, 436
- SIOCGIFDSTADDR — get p-p address, 436
- SIOCGIFFLAGS — get ifnet flags, 436
- SIOCGLOWAT — get low water mark, 485, 506
- SIOCSARP — set arp entry, 410
- SIOCSHIWAT — set high water mark, 485, 506
- SIOCSIFADDR — set ifnet address, 436
- SIOCSIFDSTADDR — set p-p address, 436
- SIOCSIFFLAGS — set ifnet flags, 436
- SIOCSLOWAT — set low water mark, 485, 506
- sleep — suspend execution, 239, 390
- SMD disk controller
 - ip — Interphase 2180, 442 *thru* 443
 - xy — Xylogics 450, 511 *thru* 512
- socket, 129
- socket operations
 - async_daemon, 84
 - bind, 19
 - connect, 29
 - getpeername, 54
 - getsockname, 61
 - getsockopt, 62
 - listen, 69
 - nfssvc, 84
 - recv, 99
 - recvfrom, 99
 - recvmsg, 99
 - send, 111
 - sendmsg, 111
 - sendto, 111
 - setsockopt, 62
 - shutdown, 120
 - socket, 129
 - socketpair, 131
- socket operations, accept connection
 - socketpair create connected socket pair, 131
 - sockets, raw IP, 441
 - interrupt system calls with software signal — `siginterrupt`, 235, 236, 240
 - software signal — signal, System V, 387
 - software signals — `sigvec` PAGE END, 127
 - software signals — `sigvec` PAGE START, 125
 - sort quicker — `qsort`, 224
 - space — specify plot space, 402
 - spawn process, 147
 - special file
 - make, 73
 - special files, introduction to, 407 *thru* 408
 - specify paging/swapping device — `swapon`, 136
 - split into mantissa and exponent — `frexp`, 183
 - spray — scatter data to check network, 326
 - sprintf — formatted output conversion, 341
 - sprintf — format to string, System V, 377
 - sqrt — square root function, 288
 - srand — generate random numbers, 264
 - srand — generate random numbers, System V, 380
 - srandom — generate random number, 225
 - sscanf — convert from string, 346
 - sscanf — convert from string, System V, 381
 - st — Sysgen SC 4000 (Archive) Tape Driver, 482 *thru* 483
 - standard I/O library functions, introduction to, 329
 - start output (like control-Q) `ioctl` — TIOCSTART, 477
 - stat — obtain file attributes, 132
 - state of terminal
 - get — `gtty`, 265
 - set — `stty`, 265
 - statfs — obtain file system statistics, 134
 - static file system information — `mntent`, 550
 - statistics
 - of file system — `fstatfs`, 134
 - of file system — `statfs`, 134
 - profil, 89
 - status monitor files for network services, 575
 - stdin
 - get character — `getchar`, 338
 - get character, System V — `getchar`, 372
 - get string from — `gets`, 339
 - input conversion — `scanf`, 346
 - input conversion, System V — `scanf`, 381
 - stdout
 - output conversion, System V — `printf`, 377
 - put character to — `putchar`, 344
 - sticky bit — `chmod`, 23
 - stop output (like control-S) `ioctl` — TIOCSTOP, 477
 - stop output `ioctl` — TIOCSTOP, 502
 - stop processor, 98
 - stop until signal — `pause`, 263
 - storage allocation, 210 *thru* 212
 - `alloca` — allocate on stack, 211
 - `calloc` — allocate memory, 210

storage allocation, *continued*

- `cfree` — free memory, 210
- `free` — free memory, 210
- `malloc` — allocate memory, 210
- `malloc_debug` — set debug level, 211
- `malloc_verify` — verify heap, 211
- `memalign` — allocate aligned memory, 210
- `realloc` — reallocate memory, 210
- `valloc` — allocate aligned memory, 211

storage management, 210 *thru* 212

storage management debugging, 211 *thru* 212

`store datum under key` — `store`, 396

`store` — store datum under key, 396

`strcat` — concatenate strings, 241

`index` — find character in string, 241

`strcmp` — compare strings, 241

`strcpy` — copy strings, 241

stream

- `fopen` — open stream, System V, 369
- assign buffering — `setbuf`, 348
- assign buffering, System V — `setbuf`, 384
- assign buffering — `setbuffer`, 348
- assign buffering, System V — `setbuffer`, 384
- assign buffering — `setlinebuf`, 348
- assign buffering, System V — `setlinebuf`, 384
- assign buffering — `setvbuf`, 348
- assign buffering, System V — `setvbuf`, 384
- associate descriptor — `fdopen`, 335
- associate descriptor, System V — `fdopen`, 369
- `close` — `fclose`, 333
- `flush` — `fflush`, 333
- `fprintf` — format to stream, System V, 377
- read from stream, System V — `fread`, 371
- write to stream, System V — `fwrite`, 371
- get character — `fgetc`, 338
- get character, System V — `fgetc`, 372
- get character — `getc`, 338
- get character, System V — `getc`, 372
- get character — `getchar`, 338
- get character, System V — `getchar`, 372
- get position of — `ftell`, 337
- get string from — `fgets`, 339
- get word — `getw`, 338
- get word, System V — `getw`, 372
- input conversion — `scanf`, 346
- input conversion, System V — `scanf`, 381
- `open` — `fopen`, 335
- output conversion, System V — `printf`, 377
- `printf` — format to stdout, System V, 377
- push character back to — `ungetc`, 352
- put character to — `fputc`, 344
- put character to — `putc`, 344
- put string to — `puts`, 345
- put string to — `fputs`, 345
- put word to — `putw`, 344
- read from stream — `fread`, 336
- `reopen` — `freopen`, 335
- `reopen`, System V — `freopen`, 369
- reposition — `rewind`, 337
- return to remote command — `rcmd`, 303
- return to remote command — `rexec`, 304
- `rewind` — `rewind`, 337
- write to stream — `fwrite`, 336

stream, *continued*

- `seek` — `fseek`, 337
- `sprintf` — format to string, System V, 377

stream status enquiries

- `clearerr` — clear error on stream, 334
- `clearerr` — clear error on stream, System V, 368
- `feof` — enquire EOF on stream, 334
- `feof` — enquire EOF on stream, System V, 368
- `ferror` — inquire error on stream, 334
- `ferror` — inquire error on stream, System V, 368
- `fileno` — get stream descriptor number, 334
- `fileno` — get stream descriptor number, System V, 368

stream, formatted output

- `fprintf` — format to stream, System V, 377
- `printf` — format to stdout, System V, 377
- `sprintf` — format to string, System V, 377

streaming 1/4-inch tape drive — `ar`, 409

string

- number conversion — `scanf`, 346
- number conversion, System V — `printf`, 377, 381

string operations

- compare — `strcmp`, 241
- compare — `strncmp`, 241
- concatenate — `strcat`, 241
- concatenate — `strncat`, 241
- convert from numbers — `ecvt`, 177
- copy — `strcpy`, 241
- copy — `strncpy`, 241
- get from stdin — `gets`, 339
- get from stream — `fgets`, 339
- `index` — `nndex`, 241
- put to stdout — `puts`, 345
- put to stream — `fputs`, 345
- reverse index — `rindex`, 241
- reverse index — `rindex`, 241

`strlen` — get length of string, 241

`strncat` — concatenate strings, 241

`strncmp` — compare strings, 241

`strncpy` — copy strings, 241

`rindex` — find character in string, 241

`strtod` — ASCII string to double, 243

`strtol` — ASCII string to long integer, 244

`stty` — set terminal state, 265

Sun 10 Mb/s Ethernet interface — `ie`, 435

Sun mouse device — `mouse`, 464

Sun SCSI disk driver — `si`, 480 *thru* 481

Sun-3/50 10 Mb/s Ethernet interface — `le`, 460 *thru* 461

super block, update — `sync`, 138

suspend execution — `sleep`, 239, 390

suspend execution for interval in microseconds — `usleep`, 254

`swab` — swap bytes, 245

swap bytes — `swab`, 245

`swapon` — specify paging device, 136

swapping device — `swapon`, 136

symbolic link

- create, 137
- read value of, 97

`symlink`, 137

`sync` — update super block, 138

synchronize file state — `fsync`, 43

synchronous I/O multiplexing, 104

- sys_errlist — system error messages, 219
 - sys_nerr — system error messages, 219
 - sys_siglist — system signal messages, 221
 - syscall, 139
 - Sysgen SC 4000 (Archive) Tape Driver — *st*, 482 *thru* 483
 - syslog — write message to system log, 246
 - System V VPC-2200 interface — *vpc*, 508
 - system calls, introduction to, 1 *thru* 10
 - system data types — *types*, 619
 - system error messages
 - errno* — system error messages, 219
 - perror* — system error messages, 219
 - sys_errlist* — system error messages, 219
 - sys_nerr* — system error messages, 219
 - system error numbers, introduction to, 1 *thru* 5
 - system* — issue shell command, 247
 - system log, control — *syslog*, 246
 - system operation support
 - mount*, 77
 - process accounting — *acct*, 17
 - reboot*, 98
 - swapon* — specify paging device, 136
 - sync*, 138
 - vadvise*, 146
 - system page size, get — *getpagesize*, 53
 - system resource consumption
 - control — *vlimit*, 270
 - system signal messages
 - psignal*, 221
 - sys_siglist*, 221
 - System V library functions, introduction to, 355
 - System V library, system call versions
 - open*, 85
 - read*, 95
 - setpgrp*, 112
 - uname*, 142
 - write*, 151
- ## T
- tabstop specifications in text files — *fspec*, 540
 - tan* — trigonometric tangent, 286
 - tanh* — hyperbolic tangent, 287
 - tap protocol for network — *nit*, 473 *thru* 475
 - tape drive, 1/2-inch
 - tm* — tapemaster, 494
 - xt* — Xylogics 472, 510
 - tape drive, 1/4-inch
 - ar* — Archive 1/4-inch Streaming Tape Drive, 409
 - Sysgen SC 4000 (Archive) Tape Driver — *st*, 482 *thru* 483
 - tape interface — *mtio*, 467 *thru* 468
 - tape operation *ioctl* — *MTIOCTOP*, 467
 - tapemaster 1/2-inch tape drive — *tm*, 494
 - tar* — tape archive file format, 576
 - tcp* — Internet Transmission Control Protocol, 484 *thru* 485
 - TCP *ioctl*'s
 - SIOCATMARK* — at OOB mark?, 485
 - SIOCGHIWAT* — get high water mark, 485, 506
 - SIOCGLOWAT* — get low water mark, 485, 506
 - SIOCSDHIWAT* — set high water mark, 485, 506
 - SIOCSLOWAT* — set low water mark, 485, 506
 - tdelete* — delete binary tree node, 248
 - tell*, 70
 - tellidir* — position of directory stream, 173
 - temporary file
 - create name for — *tmpnam*, 351
 - term* — terminal driving tables, 578
 - termcap* — terminal capability data base, 582
 - terminal
 - configuration data base — *gettytab*, 543
 - find name of — *ttynam*, 251
 - terminal description, 486
 - terminal emulation, ANSI, 418 *thru* 422
 - terminal emulator — *console*, 418 *thru* 422
 - terminal independent operations
 - tgetent*, 404
 - tgetflag*, 404
 - tgetnum*, 404
 - tgetstr*, 404
 - tgoto*, 404
 - tputs*, 404
 - terminal interface — *tty*, 495 *thru* 504
 - terminal state
 - get — *gtty*, 265
 - set — *stty*, 265
 - terminal types — *ttytype*, 618
 - terminate process, 37, 181
 - terminate program — *abort*, 162
 - termination handler, name — *on_exit*, 218
 - terminfo* — System V terminal capability data base, 590, 603
 - termio* — terminal description, 486
 - test for indeterminate floating values
 - isinf* — test infinite value, 205
 - isnan* — test not a number, 205
 - tfind* — search binary tree, 248
 - tgetent* — get entry for terminal, 404
 - tgetflag* — get Boolean capability, 404
 - tgetnum* — get numeric capability, 404
 - tgetstr* — get string capability, 404
 - tgoto* — go to position, 404
 - time
 - adjust — *adjtime*, 18
 - time and date
 - get — *time*, 266
 - get — *gettimeofday* day, 63
 - get — *ftime*, 266
 - set — *settimeofday* day, 63
 - time and date conversion
 - asctime*, 169
 - ctime*, 169
 - dysize*, 170
 - gmtime*, 169
 - localtime*, 169
 - timezone*, 169
 - time and date conversion, System V
 - asctime*, 358
 - ctime*, 358
 - gmtime*, 358
 - localtime*, 358
 - time* — get date and time, 266
 - timed event jobs table — *crontab*, 528
 - timerclear* — macro, 51
 - timercmp* — macro, 51

timerisset — macro, 51
times — get process times, 267
times — get process and child process times, System V, 391
timezone — date and time conversion, 169
timing and statistics
 clock, 261
 getitimer, 51
 gettimeofday, 63
 profil, 89
 setitimer, 51
 settimeofday, 63
 timerclear — macro, 51
 timercmp — macro, 51
 timerisset — macro, 51
TIOCCBRK — clear break bit, 502
TIOCCDTR — clear DTR, 502
TIOCCONS — get console I/O, 418
TIOCEXCL — set exclusive use of tty, 502
TIOCFLUSH — flush buffers, 502
TIOCGETC — get special characters, 503
TIOCGETD — get line discipline, 412, 495, 502
TIOCGETP — get parameters — gtty, 501
TIOCGLTC — get local special chars, 504
TIOCGPGRP — get process group of tty, 502
TIOCHPCL — hang up on last close, 499, 502
TIOCLBIC — bit clear local mode bits, 503
TIOCLBIS — bit set local mode bits, 503
TIOCLGET — get local mode bits, 503
TIOCLSET — set local mode bits, 503
TIOCNOTTY — void tty association, 495
TIOCNXCL — remove exclusive use of tty, 502
TIOCOUNTQ — get number of characters in output queue, 502
TIOCPKT — set/clear packet mode (pty), 477
TIOCREMOTE — remote input editing, 477
TIOCSBRK — set break bit, 502
TIOCSDTR — set DTR, 502
TIOCSETC — set special characters, 503
TIOCSETD — set line discipline, 412, 495, 502
TIOCSETN — set parameters, 501
TIOCSETP — set parameters — gtty, 501
TIOCSLTC — set local special chars, 504
TIOCSGRP — set process group of tty, 502
TIOCSTART — restart output, 502
TIOCSTART — start output (like control-Q), 477
TIOCSTI — simulate terminal input, 502
TIOCSTOP — stop output, 502
TIOCSTOP — stop output (like control-S), 477
tm — tapemaster 1/2-inch tape drive, 494
tmpfile — create temporary file, 350
tmpnam — make temporary file name, 351
toascii — convert character to ASCII, System V, 360
toascii — convert character to ASCII, 171
tolower — convert character to lower-case, System V, 360
tolower — convert character to lower-case, 171
toupper — convert character to upper-case, System V, 360
toupper — convert character to upper-case, 171
tp — DEC/mag tape formats, 616
tputs — decode padding information, 404
trace process — **ptrace**, 90 *thru* 92

trigonometric functions, 286
 acos, 286
 asin, 286
 atan, 286
 atan2, 286
 cos, 286
 sin, 286
 tan, 286
truncate, 140
trusted hosts list — **hosts.equiv**, 548, 553
tsearch — build and search binary tree, 248
tty — general terminal interface, 495 *thru* 504
ttynam — find terminal name, 251
ttys — terminal initialization data, 617
ttyslot — get utmp slot number, 252
ttyslot — get utmp slot number, System V, 392
ttytype — connected terminal types, 618
twalk — traverse binary tree, 248
types — primitive system data types, 619

U

ualarm — schedule signal in microsecond precision, 253
udp — Internet User Datagram Protocol, 505 *thru* 506
ulimit — get and set user limits, 268
umask, 141
uname — get system name, 142
ungetc — push character back to stream, 352
unique file name
 create — **mktemp**, 214
UNIX magnetic tape interface — **mtio**, 467 *thru* 468
unlink — remove directory entry, 143
unmount — demount file system, 144
update super block — **sync**, 138
user ID
 get, 64
 set real and effective — **setreuid**, 114
user limits
 get — **ulimit**, 268
 set — **ulimit**, 268
user mask, set — **umask**, 141
user name, get — **cuserid**, 332
usleep — suspend execution, 254
usracct — login records, 621
utime — set file times, 269
utimes — set file times, 145
utmp — login records, 621
uuencode — UUCP encoded file format, 622

V

va_arg — next argument in variable list, 256
va_dcl — variable argument declarations, 256
va_end — finish variable argument list, 256
va_list — variable argument declarations, 256
va_start — initialize varargs, 256
vadvise — advise paging system, 146
valloc — allocate aligned memory, 211
values — machine-dependent values, 255
varargs — variable argument list, 256
variable argument list, — **varargs**, 256

verify heap — `malloc_verify`, 211
 version number `ioctl` — `NDIOCVER`, 471
 vfont — font formats, 623
 vfork, 147
`vfprintf` — format and print variable argument list, 353
`vgrind` — `vgrind` language definitions, 624
 vhangup, 148
 virtual — virtual address space, 463
 vlimit — control consumption, 270
 vme16 — VMEbus 16-bit space, 463
 vme16d16 — VMEbus address space, 463
 vme16d32 — VMEbus address space, 463
 vme24 — VMEbus 24-bit space, 463
 vme24d16 — VMEbus address space, 463
 vme24d32 — VMEbus address space, 463
 vme32d16 — VMEbus address space, 463
 vme32d32 — VMEbus address space, 463
 void tty association `ioctl` — `TIOCNOTTY`, 495
 vp — Ikon 10071-5 Versatec parallel printer interface, 507
 vpc — Systech VPC-2200 Versatec/Centronics interface, 508
`vprintf` — format and print variable argument list, 353
`vsprintf` — format and print variable argument list, 353
 vtimes — resource use information, 271

W

wait, 149
 wait3, 149
 win — Sun window system, 509
 word
 get from stream — `getw`, 338
 get from stream, System V — `getw`, 372
 put to stream — `putw`, 344
 working directory
 change, 22
 get pathname — `getwd`, 200
 write, 151
 write formatted
 `fprintf` — convert to stream, System V, 377
 `printf` — convert to stdout, System V, 377
 `sprintf` — convert to string, System V, 377
 write gathered — `writev`, 151
 write to stream — `fwrite`, 336
 write to stream, System V — `fwrite`, 371
 wttmp — login records, 621

X

XDR routines, 307
 xt — Xylogics 472 1/2-inch tape drive, 510
 xtom — hexadecimal string to multiple precision, 398
 xy — Xylogics SMD Disk driver, 511 *thru* 512
 Xylogics 472 1/2-inch tape drive — `xt`, 510
 Xylogics SMD Disk driver — `xy`, 511 *thru* 512

Y

y0 — Bessel function, 281
 y1 — Bessel function, 281
 yellow pages client interface, 309
 yn — Bessel function, 281
 yp_all — yellow pages client interface, 309

yp_bind — yellow pages client interface, 309
 yp_first — yellow pages client interface, 309
 yp_get_default_domain — yellow pages client interface,
 309
 yp_master — yellow pages client interface, 309
 yp_match — yellow pages client interface, 309
 yp_next — yellow pages client interface, 309
 yp_order — yellow pages client interface, 309
 yp_unbind — yellow pages client interface, 309
 yperr_string — yellow pages client interface, 309
 ypfiles — yellowpages database and directory, 626
 yppasswd — update YP password entry, 327
 ypprot_err — yellow pages client interface,
 309

Z

zero byte strings — `bzero`, 167
 zs — zilog 8530 SCC serial communications
 driver, 513

Revision History

Version	Date	Comments
A	23 February 1983	First edition of this manual under the title <i>System Interface Manual for the Sun Workstation</i> .
B	15 April 1983	Second edition of this manual with corrections to numerous manual pages.
C	1 August 1983	Third edition of this manual with corrections to numerous manual pages. Added a glossary of system calls and system error responses.
D	1 November 1983	Fourth edition of this manual with numerous corrections. Corrected numerous incorrect cross-references. Added a <i>System Interface Overview</i> and the <i>Interprocess Communication Primer</i> .
E	7 January 1984	Fifth edition of this manual with numerous corrections.
F	15 May 1985	Sixth edition with numerous corrections. The <i>Interprocess Communication Primer</i> made a part of the manual, <i>Networking on the Sun Workstation</i> . Made page numbering contiguous throughout, and replaced the <i>Permuted Index</i> with a conventional one.

— *Continued*

Version	Date	Comments
G	1 January 1986	Formerly the <i>System Interface Manual for the Sun Workstation</i> , this seventh edition contains many corrections to manual pages. The former section entitled <i>System Interface Overview</i> is now a separate manual entitled <i>UNIX Interface Overview</i> . The index has been upgraded to refer to ioctl's and system error numbers.
H	15 October 1986	Eighth (draft) edition, for Sun 3.2 Release. Includes numerous additions for System V compatibility, and updates from U.C. Berkeley 4.3 BSD, as well as numerous corrections to manual pages.

Notes

Notes