

U N I X
PROGRAMMER'S MANUAL



Reference Guide

Printed by the USENIX Association as a service to the UNIX Community. This material is copyrighted by The Regents of the University of California and/or Bell Telephone Laboratories, and is reprinted by permission. Permission for the publication or other use of these materials may be granted only by the Licensors and copyright holders.

Cover design by John Lasseter, Lucasfilm, Ltd.

First Printing	July 1984
Second Printing	December 1984
Third Printing	September 1985
Fourth Printing	March 1986

UNIX PROGRAMMER'S MANUAL

Reference Guide

*4.2 Berkeley Software Distribution
Virtual VAX-11 Version*

March, 1984

Computer Science Division
Department of Electrical Engineering and Computer Science
University of California
Berkeley, California 94720

PREFACE

This manual is part of a five volume set intended for use with the 4.2 Berkeley Software Distribution for the VAX-11 computer. While the five volumes together contain virtually the same material presented in the four volume UNIX Programmer's Manual distributed with 4.2BSD, the manuals reflect a revised organization necessitated by the large quantity of information. The documentation is divided into three logically distinct *manuals*:

- UNIX User's Manual,
- UNIX Programmer's Manual, and
- UNIX System Manager's Manual.

Each of the User and Programmer manuals are two volumes: a Reference Guide, containing relevant sections from Volume 1 of the old UNIX Programmer's Manual, and a volume of Supplementary Documents, containing pertinent material from Volume 2 of the old UNIX Programmer's Manual. The System Manager's manual consists of a single volume containing information from both Volumes 1 and 2. We acknowledge those who have assisted us in putting together these manuals. In particular, we thank Tom Ferrin for pursuing the printing particulars.

M. J. Karels
S. J. Leffler

Preface to the 4.2 Berkeley distribution

This update to the 4.1 distribution of June 1981 provides support for the VAX 11/730, full networking and interprocess communication support, an entirely new file system, and many other new features. It is certainly the most ambitious release of software ever prepared here and represents many man-years of work. Bill Shannon (both at DEC and at Sun Microsystems) and Robert Elz of the University of Melbourne contributed greatly to this distribution through new device drivers and painful debugging episodes. Rob Gurwitz of BBN wrote the initial version of the code upon which the current networking support is based. Eric Allman of Britton-Lee donated countless hours to the mail system. Bill Croft (both at SRI and Sun Microsystems) aided in the debugging and development of the networking facilities. Dennis Ritchie of Bell Laboratories also contributed greatly to this distribution, providing valuable advice and guidance. Helge Skrivervik worked on the device drivers which enabled the distribution to be delivered with a TU58 console cassette and RX01 console floppy disk, and rewrote major portions of the standalone i/o system to support formatting of non-DEC peripherals.

Numerous others contributed their time and energy in organizing the user software for release, while many groups of people on campus suffered patiently through the low spots of development. As always, we are grateful to the UNIX user community for encouragement and support.

Once again, the financial support of the Defense Advanced Research Projects Agency is gratefully acknowledged.

S. J. Leffler
W. N. Joy
M. K. McKusick

NAME

intro — introduction to system calls and error numbers

SYNOPSIS

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

DESCRIPTION

This section describes all of the system calls. Most of these calls have one or more error returns. An error condition is indicated by an otherwise impossible return value. This is almost always `-1`; the individual descriptions specify the details.

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

The following is a complete list of the errors and their names as given in `<errno.h>`.

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

- 11 EAGAIN No more processes
In a *fork*, the system's process table is full or the user is not allowed to create any more processes.
- 12 ENOMEM Not enough core
During an *execve* or *break*, a program asks for more core or swap space than the system is able to supply. A lack of swap space is normally a temporary condition, however a lack of core is not a temporary condition; the maximum size of the text, data, and stack segments is a system parameter.
- 13 EACCES Permission denied
An attempt was made to access a file in a way forbidden by the protection system.
- 14 EFAULT Bad address
The system encountered a hardware fault in attempting to access the arguments of a system call.
- 15 ENOTBLK Block device required
A plain file was mentioned where a block device was required, e.g. in *mount*.
- 16 EBUSY Mount device busy
An attempt to mount a device that was already mounted or an attempt was made to dismount a device on which there is an active file directory. (open file, current directory, mounted-on file, active text segment).
- 17 EEXIST File exists
An existing file was mentioned in an inappropriate context, e.g. *link*.
- 18 EXDEV Cross-device link
A hard link to a file on another device was attempted.
- 19 ENODEV No such device
An attempt was made to apply an inappropriate system call to a device; e.g. read a write-only device.
- 20 ENOTDIR Not a directory
A non-directory was specified where a directory is required, for example in a path name or as an argument to *chdir*.
- 21 EISDIR Is a directory
An attempt to write on a directory.
- 22 EINVAL Invalid argument
Some invalid argument: dismounting a non-mounted device, mentioning an unknown signal in *signal*, reading or writing a file for which *seek* has generated a negative pointer. Also set by math functions, see *intro*(3).
- 23 ENFILE File table overflow
The system's table of open files is full, and temporarily no more *opens* can be accepted.
- 24 EMFILE Too many open files
Customary configuration limit is 20 per process.
- 25 ENOTTY Not a typewriter
The file mentioned in an *ioctl* is not a terminal or one of the other devices to which this call applies.
- 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 (about 10^9 bytes).
- 28 ENOSPC No space left on device
During a *write* to an ordinary file, there is no free space left on the device.
- 29 EPIPE Illegal seek
An *lseek* was issued to a pipe. This error may also be issued for other non-seekable devices.
- 30 EROFS Read-only file system
An attempt to modify a file or directory was made on a device mounted read-only.
- 31 EMLINK Too many links
An attempt to make more than 32767 hard links to a file.
- 32 EPIPE Broken pipe
A write on a pipe or socket for which there is no process to read the data. This condition normally generates a signal; the error is returned if the signal is ignored.
- 33 EDOM Math argument
The argument of a function in the math package (3M) is out of the domain of the function.
- 34 ERANGE Result too large
The value of a function in the math package (3M) is unrepresentable within machine precision.
- 35 EWOULDBLOCK Operation would block
An operation which would cause a process to block was attempted on a object in non-blocking mode (see *ioctl* (2)).
- 36 EINPROGRESS Operation now in progress
An operation which takes a long time to complete (such as a *connect* (2)) was attempted on a non-blocking object (see *ioctl* (2)).
- 37 EALREADY Operation already in progress
An operation was attempted on a non-blocking object which already had an operation in progress.
- 38 ENOTSOCK Socket operation on non-socket
Self-explanatory.
- 39 EDESTADDRREQ Destination address required
A required address was omitted from an operation on a socket.
- 40 EMSGSIZE Message too long
A message sent on a socket was larger than the internal message buffer.
- 41 EPROTOTYPE Protocol wrong type for socket
A protocol was specified which does not support the semantics of the socket type requested. For example you cannot use the ARPA Internet UDP protocol with type `SOCK_STREAM`.
- 42 ENOPROTOPT Bad protocol option
A bad option was specified in a *getsockopt*(2) or *setsockopt*(2) call.
- 43 EPROTONOSUPPORT Protocol not supported
The protocol has not been configured into the system or no implementation for it exists.

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

- 62 ELOOP Too many levels of symbolic links
A path name lookup involved more than 8 symbolic links.
- 63 ENAMETOOLONG File name too long
A component of a path name exceeded 255 characters, or an entire path name exceeded 1023 characters.
- 64 EHOSTDOWN Host is down
A socket operation failed because the destination host was down.
- 65 EHOSTUNREACH Host is unreachable
A socket operation was attempted to an unreachable host.
- 66 ENOTEMPTY Directory not empty
A directory with entries other than "." and ".." was supplied to a remove directory or rename call.
- 69 EDQUOT Disc quota exceeded
A file creation or write operation failed because the hard limit for that resource had been reached.

DEFINITIONS

Process ID

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

Parent process ID

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

Process Group ID

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

Tty Group ID

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

Real User ID and Real Group ID

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

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

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

Effective User Id, Effective Group Id, and Access Groups

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

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

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

Super-user

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

Special Processes

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

Descriptor

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

File Name

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

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

Note that it is generally unwise to use *, ?, [or] as part of file names because of the special meaning attached to these characters by the shell.

Path Name

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

If a path name begins with a slash, the path search begins at the *root* directory. Otherwise, the search begins from the current working directory. A slash by itself names the *root* directory. A null pathname refers to the current directory.

Directory

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

Root Directory and Current Working Directory

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

File Access Permissions

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

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

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

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

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

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

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

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

Otherwise, permission is denied.

Sockets and Address Families

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

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

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

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

SEE ALSO

intro(3), perror(3)

NAME

`accept` — accept a connection on a socket

SYNOPSIS

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

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

DESCRIPTION

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

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

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

RETURN VALUE

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

ERRORS

The *accept* will fail if:

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

SEE ALSO

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

NAME

access — determine accessibility of file

SYNOPSIS

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

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

accessible = access(path, mode)

int accessible;
char *path;
int mode;
```

DESCRIPTION

Access checks the given file *path* for accessibility according to *mode*, which is an inclusive or of the bits R_OK, W_OK and X_OK. Specifying *mode* as F_OK (i.e. 0) tests whether the directories leading to the file can be searched and the file exists.

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

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

RETURN VALUE

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

ERRORS

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

- [ENOTDIR] A component of the path prefix is not a directory.
- [ENOENT] The argument path name was too long.
- [ENOENT] Read, write, or execute (search) permission is requested for a null path name or the named file does not exist.
- [EPERM] The argument contains a byte with the high-order bit set.
- [ELOOP] Too many symbolic links were encountered in translating the pathname.
- [EROFS] Write access is requested for a file on a read-only file system.
- [ETXTBSY] Write access is requested for a pure procedure (shared text) file that is being executed.
- [EACCES] Permission bits of the file mode do not permit the requested access; or search permission is denied on a component of the path prefix. The owner of a file has permission checked with respect to the "owner" read, write, and execute mode bits, members of the file's group other than the owner have permission checked with respect to the "group" mode bits, and all others have permissions checked with respect to the "other" mode bits.
- [EFAULT] *Path* points outside the process's allocated address space.

SEE ALSO

chmod(2), stat(2)

NAME

acct — turn accounting on or off

SYNOPSIS

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

DESCRIPTION

The system is prepared to write a record in an accounting *file* for each process as it terminates. This call, with a null-terminated string naming an existing file as argument, turns on accounting; records for each terminating process are appended to *file*. An argument of 0 causes accounting to be turned off.

The accounting file format is given in *acct(5)*.

This call is permitted only to the super-user.

NOTES

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

RETURN VALUE

On error -1 is returned. The file must exist and the call may be exercised only by the super-user. It is erroneous to try to turn on accounting when it is already on.

ERRORS

Acct will fail if one of the following is true:

- | | |
|-----------|---|
| [EPERM] | The caller is not the super-user. |
| [EPERM] | The pathname contains a character with the high-order bit set. |
| [ENOTDIR] | A component of the path prefix is not a directory. |
| [ENOENT] | The named file does not exist. |
| [EISDIR] | The named file is a directory. |
| [EROFS] | The named file resides on a read-only file system. |
| [EFAULT] | <i>File</i> points outside the process's allocated address space. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |
| [EACCES] | The file is a character or block special file. |

SEE ALSO

acct(5), sa(8)

BUGS

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

NAME

`bind` — bind a name to a socket

SYNOPSIS

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

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

DESCRIPTION

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

NOTES

Binding a name in the UNIX domain creates a socket in the file system which must be deleted by the caller when it is no longer needed (using *unlink(2)*). The file created is a side-effect of the current implementation, and will not be created in future versions of the UNIX ipc domain.

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

RETURN VALUE

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

ERRORS

The *bind* call will fail if:

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

SEE ALSO

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

NAME

brk, *sbrk* — change data segment size

SYNOPSIS

```
caddr_t brk(addr)
caddr_t addr;
caddr_t sbrk(incr)
int incr;
```

DESCRIPTION

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

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

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

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

RETURN VALUE

Zero is returned if the *brk* could be set; -1 if the program requests more memory than the system limit. *Sbrk* returns -1 if the break could not be set.

ERRORS

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

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

SEE ALSO

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

BUGS

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

NAME

`chdir` — change current working directory

SYNOPSIS

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

DESCRIPTION

Path is the pathname of a directory. *Chdir* causes this directory to become the current working directory, the starting point for path names not beginning with “/”.

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

RETURN VALUE

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

ERRORS

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

- | | |
|-----------|---|
| [ENOTDIR] | A component of the pathname is not a directory. |
| [ENOENT] | The named directory does not exist. |
| [ENOENT] | The argument path name was too long. |
| [EPERM] | The argument contains a byte with the high-order bit set. |
| [EACCES] | Search permission is denied for any component of the path name. |
| [EFAULT] | <i>Path</i> points outside the process's allocated address space. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |

SEE ALSO

`chroot(2)`

NAME

chmod — change mode of file

SYNOPSIS

```

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

fchmod(fd, mode)
int fd, mode;

```

DESCRIPTION

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

```

04000 set user ID on execution
02000 set group ID on execution
01000 save text image after execution
00400 read by owner
00200 write by owner
00100 execute (search on directory) by owner
00070 read, write, execute (search) by group
00007 read, write, execute (search) by others

```

If an executable file is set up for sharing (this is the default) then mode 1000 prevents the system from abandoning the swap-space image of the program-text portion of the file when its last user terminates. Ability to set this bit is restricted to the super-user.

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

Writing or changing the owner of a file turns off the set-user-id and set-group-id bits. This makes the system somewhat more secure by protecting set-user-id (set-group-id) files from remaining set-user-id (set-group-id) if they are modified, at the expense of a degree of compatibility.

RETURN VALUE

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

ERRORS

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

[EPERM]	The argument contains a byte with the high-order bit set.
[ENOTDIR]	A component of the path prefix is not a directory.
[ENOENT]	The pathname was too long.
[ENOENT]	The named file does not exist.
[EACCES]	Search permission is denied on a component of the path prefix.
[EPERM]	The effective user ID does not match the owner of the file and the effective user ID is not the super-user.
[EROFS]	The named file resides on a read-only file system.
[EFAULT]	<i>Path</i> points outside the process's allocated address space.
[ELOOP]	Too many symbolic links were encountered in translating the pathname.

Fchmod will fail if:

[EBADF]	The descriptor is not valid.
[EINVAL]	<i>Fd</i> refers to a socket, not to a file.

[EROFS] The file resides on a read-only file system.

SEE ALSO

open(2), chown(2)

2

NAME

chown — change owner and group of a file

SYNOPSIS

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

fchown(fd, owner, group)
int fd, owner, group;
```

DESCRIPTION

The file which is named by *path* or referenced by *fd* has its *owner* and *group* changed as specified. Only the super-user may execute this call, because if users were able to give files away, they could defeat the file-space accounting procedures.

On some systems, *chown* clears the set-user-id and set-group-id bits on the file to prevent accidental creation of set-user-id and set-group-id programs owned by the super-user.

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

Only one of the owner and group id's may be set by specifying the other as *-1*.

RETURN VALUE

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

ERRORS

Chown will fail and the file will be unchanged if:

- [EINVAL] The argument *path* does not refer to a file.
- [ENOTDIR] A component of the path prefix is not a directory.
- [ENOENT] The argument pathname is too long.
- [EPERM] The argument contains a byte with the high-order bit set.
- [ENOENT] The named file does not exist.
- [EACCES] Search permission is denied on a component of the path prefix.
- [EPERM] The effective user ID does not match the owner of the file and the effective user ID is not the super-user.
- [EROFS] The named file resides on a read-only file system.
- [EFAULT] *Path* points outside the process's allocated address space.
- [ELOOP] Too many symbolic links were encountered in translating the pathname.

Fchown will fail if:

- [EBADF] *Fd* does not refer to a valid descriptor.
- [EINVAL] *Fd* refers to a socket, not a file.

SEE ALSO

chmod(2), *flock*(2)

NAME

chroot — change root directory

SYNOPSIS

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

DESCRIPTION

Dirname is the address of the pathname of a directory, terminated by a null byte. *Chroot* causes this directory to become the root directory, the starting point for path names beginning with “/”.

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

This call is restricted to the super-user.

RETURN VALUE

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

ERRORS

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

[ENOTDIR]	A component of the path name is not a directory.
[ENOENT]	The pathname was too long.
[EPERM]	The argument contains a byte with the high-order bit set.
[ENOENT]	The named directory does not exist.
[EACCES]	Search permission is denied for any component of the path name.
[EFAULT]	<i>Path</i> points outside the process's allocated address space.
[ELOOP]	Too many symbolic links were encountered in translating the pathname.

SEE ALSO

chdir(2)

NAME

close — delete a descriptor

SYNOPSIS

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

DESCRIPTION

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

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

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

RETURN VALUE

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

ERRORS

Close will fail if:

[EBADF] *D* is not an active descriptor.

SEE ALSO

accept(2), flock(2), open(2), pipe(2), socket(2), socketpair(2), execve(2), fcntl(2)

NAME

connect — initiate a connection on a socket

SYNOPSIS

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

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

DESCRIPTION

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

RETURN VALUE

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

ERRORS

The call fails if:

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

SEE ALSO

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

NAME

creat — create a new file

SYNOPSIS

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

DESCRIPTION

This interface is obsoleted by `open(2)`.

`Creat` creates a new file or prepares to rewrite an existing file called *name*, given as the address of a null-terminated string. If the file did not exist, it is given mode *mode*, as modified by the process's mode mask (see `umask(2)`). Also see `chmod(2)` for the construction of the *mode* argument.

If the file did exist, its mode and owner remain unchanged but it is truncated to 0 length.

The file is also opened for writing, and its file descriptor is returned.

NOTES

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

RETURN VALUE

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

ERRORS

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

- [EPERM] The argument contains a byte with the high-order bit set.
- [ENOTDIR] A component of the path prefix is not a directory.
- [EACCES] A needed directory does not have search permission.
- [EACCES] The file does not exist and the directory in which it is to be created is not writable.
- [EACCES] The file exists, but it is unwritable.
- [EISDIR] The file is a directory.
- [EMFILE] There are already too many files open.
- [EROFS] The named file resides on a read-only file system.
- [ENXIO] The file is a character special or block special file, and the associated device does not exist.
- [ETXTBSY] The file is a pure procedure (shared text) file that is being executed.
- [EFAULT] *Name* points outside the process's allocated address space.
- [ELOOP] Too many symbolic links were encountered in translating the pathname.
- [EOPNOTSUPP] The file was a socket (not currently implemented).

SEE ALSO

`open(2)`, `write(2)`, `close(2)`, `chmod(2)`, `umask(2)`

NAME

`dup`, `dup2` — duplicate a descriptor

SYNOPSIS

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

DESCRIPTION

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

The object referenced by the descriptor does not distinguish between references using *oldd* and *newd* in any way. Thus if *newd* and *oldd* are duplicate references to an open file, *read(2)*, *write(2)* and *lseek(2)* calls all move a single pointer into the file. If a separate pointer into the file is desired, a different object reference to the file must be obtained by issuing an additional *open(2)* call.

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

RETURN VALUE

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

ERRORS

Dup and *dup2* fail if:

- [EBADF] *Oldd* or *newd* is not a valid active descriptor
- [EMFILE] Too many descriptors are active.

SEE ALSO

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

NAME

`execve` — execute a file

SYNOPSIS

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

DESCRIPTION

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

An interpreter file begins with a line of the form “#! *interpreter*”; When an interpreter file is *execve*'d, the system *execve*'s the specified *interpreter*, giving it the name of the originally *exec*'d file as an argument, shifting over the rest of the original arguments.

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

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

The argument *envp* is also an array of character pointers to null-terminated strings. These strings pass information to the new process which are not directly arguments to the command, see *environ*(7).

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

Ignored signals remain ignored across an *execve*, but signals that are caught are reset to their default values. The signal stack is reset to be undefined; see *sigvec*(2) for more information.

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

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

process ID	see <i>getpid</i> (2)
parent process ID	see <i>getppid</i> (2)
process group ID	see <i>getpgrp</i> (2)
access groups	see <i>getgroups</i> (2)
working directory	see <i>chdir</i> (2)
root directory	see <i>chroot</i> (2)
control terminal	see <i>tty</i> (4)
resource usages	see <i>getrusage</i> (2)
interval timers	see <i>getitimer</i> (2)
resource limits	see <i>getrlimit</i> (2)
file mode mask	see <i>umask</i> (2)
signal mask	see <i>sigvec</i> (2)

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

```

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

```

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

Envp is a pointer to an array of strings that constitute the *environment* of the process. A pointer to this array is also stored in the global variable "environ". Each string consists of a name, an "=", and a null-terminated value. The array of pointers is terminated by a null pointer. The shell *sh*(1) passes an environment entry for each global shell variable defined when the program is called. See *environ*(7) for some conventionally used names.

RETURN VALUE

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

ERRORS

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

- | | |
|-----------|--|
| [ENOENT] | One or more components of the new process file's path name do not exist. |
| [ENOTDIR] | A component of the new process file is not a directory. |
| [EACCES] | Search permission is denied for a directory listed in the new process file's path prefix. |
| [EACCES] | The new process file is not an ordinary file. |
| [EACCES] | The new process file mode denies execute permission. |
| [ENOEXEC] | The new process file has the appropriate access permission, but has an invalid magic number in its header. |
| [ETXTBSY] | The new process file is a pure procedure (shared text) file that is currently open for writing or reading by some process. |
| [ENOMEM] | The new process requires more virtual memory than is allowed by the imposed maximum (<i>getrlimit</i> (2)). |
| [E2BIG] | The number of bytes in the new process's argument list is larger than the system-imposed limit of {ARG_MAX} bytes. |
| [EFAULT] | The new process file is not as long as indicated by the size values in its header. |
| [EFAULT] | <i>Path</i> , <i>argv</i> , or <i>envp</i> point to an illegal address. |

CAVEATS

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

SEE ALSO

exit(2), *fork*(2), *execl*(3), *environ*(7)

NAME

`_exit` — terminate a process

SYNOPSIS

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

DESCRIPTION

`_exit` terminates a process with the following consequences:

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

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

The parent process ID of all of the calling process's existing child processes are also set to 1. This means that the initialization process (see `intro(2)`) inherits each of these processes as well.

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

RETURN VALUE

This call never returns.

SEE ALSO

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

NAME

fcntl — file control

SYNOPSIS

```
#include <fcntl.h>
res = fcntl(fd, cmd, arg)
int res;
int fd, cmd, arg;
```

DESCRIPTION

Fcntl provides for control over descriptors. The argument *fd* is a descriptor to be operated on by *cmd* as follows:

- F_DUPFD** Return a new descriptor as follows:
 Lowest numbered available descriptor greater than or equal to *arg*.
 Same object references as the original descriptor.
 New descriptor shares the same file pointer if the object was a file.
 Same access mode (read, write or read/write).
 Same file status flags (i.e., both file descriptors share the same file status flags).
 The close-on-exec flag associated with the new file descriptor is set to remain open across *execv(2)* system calls.
- F_GETFD** Get the close-on-exec flag associated with the file descriptor *fd*. If the low-order bit is 0, the file will remain open across *exec*, otherwise the file will be closed upon execution of *exec*.
- F_SETFD** Set the close-on-exec flag associated with *fd* to the low order bit of *arg* (0 or 1 as above).
- F_GETFL** Get descriptor status flags, as described below.
- F_SETFL** Set descriptor status flags.
- 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 flags for the **F_GETFL** and **F_SETFL** flags are as follows:

- FNDELAY** Non-blocking I/O; if no data is available to a *read* call, or if a write operation would block, the call returns -1 with the error **EWOULDBLOCK**.
- FAPPEND** Force each write to append at the end of file; corresponds to the **O_APPEND** flag of *open(2)*.
- FASYNC** Enable the SIGIO signal to be sent to the process group when I/O is possible, e.g. upon availability of data to be read.

RETURN VALUE

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

- F_DUPFD** A new file descriptor.
F_GETFD Value of flag (only the low-order bit is defined).
F_GETFL Value of flags.
F_GETOWN Value of file 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] *Fildes* is not a valid open file descriptor.
- [EMFILE] *Cmd* is `F_DUPFD` and the maximum allowed number of file descriptors are currently open.
- [EINVAL] *Cmd* is `F_DUPFD` and *arg* is negative or greater the maximum allowable number (see *getdtablesize(2)*).

SEE ALSO

close(2), *execve(2)*, *getdtablesize(2)*, *open(2)*, *sigvec(2)*

BUGS

The asynchronous I/O facilities of `FNDELAY` and `FASYNC` are currently available only for tty operations. No `SIGIO` signal is sent upon draining of output sufficiently for non-blocking writes to occur.

NAME

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

SYNOPSIS

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

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

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

DESCRIPTION

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

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

The locking mechanism allows two types of locks: *shared* locks and *exclusive* locks. At any time multiple shared locks may be applied to a file, but at no time are multiple exclusive, or both shared and exclusive, locks allowed simultaneously on a file.

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

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

NOTES

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

Processes blocked awaiting a lock may be awakened by signals.

RETURN VALUE

Zero is returned if the operation was successful; on an error a `-1` is returned and an error code is left in the global location `errno`.

ERRORS

The `flock` call fails if:

- [`EWOULDBLOCK`] The file is locked and the `LOCK_NB` option was specified.
- [`EBADF`] The argument `fd` is an invalid descriptor.
- [`EINVAL`] The argument `fd` refers to an object other than a file.

SEE ALSO

`open(2)`, `close(2)`, `dup(2)`, `execve(2)`, `fork(2)`

NAME

`fork` — create a new process

SYNOPSIS

```
pid = fork(0)
int pid;
```

DESCRIPTION

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

The child process has a unique process ID.

The child process has a different parent process ID (i.e., the process ID of the parent process).

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

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

RETURN VALUE

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

ERRORS

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

- [EAGAIN] The system-imposed limit {PROC_MAX} on the total number of processes under execution would be exceeded.
- [EAGAIN] The system-imposed limit {KID_MAX} on the total number of processes under execution by a single user would be exceeded.

SEE ALSO

execve(2), *wait(2)*

NAME

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

SYNOPSIS

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

DESCRIPTION

Fsync causes all modified data and attributes of *fd* to be moved to a permanent storage device. This normally results in all in-core modified copies of buffers for the associated file to be written to a disk.

Fsync should be used by programs which require a file to be in a known state; for example in building a simple transaction facility.

RETURN VALUE

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

ERRORS

The *fsync* fails if:

[EBADF] *Fd* is not a valid descriptor.

[EINVAL] *Fd* refers to a socket, not to a file.

SEE ALSO

sync(2), *sync*(8), *update*(8)

BUGS

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

NAME

`getdtablesize` — get descriptor table size

SYNOPSIS

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

DESCRIPTION

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

SEE ALSO

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

NAME

getgid, *getegid* — get group identity

SYNOPSIS

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

DESCRIPTION

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

The real group ID is specified at login time.

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

SEE ALSO

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

NAME

`getgroups` — get group access list

SYNOPSIS

```
#include <sys/param.h>
ngroups = getgroups(gidsetlen, gidset)
int ngroups, gidsetlen, *gidset;
```

DESCRIPTION

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 which may be placed in *gidset*. *Getgroups* returns the actual number of groups returned in *gidset*. No more than NGROUPS, as defined in *<sys/param.h>*, will ever be returned.

RETURN VALUE

A successful call returns the number of groups in the group set. A value of `-1` indicates that an error occurred, and the error code is stored in the global variable *errno*.

ERRORS

The possible errors for *getgroup* are:

- [EINVAL] The argument *gidsetlen* is smaller than the number of groups in the group set.
- [EFAULT] The argument *gidset* specifies an invalid address.

SEE ALSO

`setgroups(2)`, `initgroups(3X)`

NAME

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

SYNOPSIS

```
hostid = gethostid()
int hostid;
sethostid(hostid)
int hostid;
```

DESCRIPTION

Sethostid establishes a 32-bit identifier for the current processor which is intended to be unique among all UNIX systems in existence. This is normally a DARPA Internet address for the local machine. This call is allowed only to the super-user and is normally performed at boot time.

Gethostid returns the 32-bit identifier for the current processor.

SEE ALSO

hostid(1), gethostname(2)

BUGS

32 bits for the identifier is too small.

NAME

gethostname, sethostname — get/set name of current host

SYNOPSIS

```
gethostname(name, namelen)
char *name;
int namelen;

sethostname(name, namelen)
char *name;
int namelen;
```

DESCRIPTION

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

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

RETURN VALUE

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

ERRORS

The following errors may be returned by these calls:

[EFAULT]	The <i>name</i> or <i>namelen</i> parameter gave an invalid address.
[EPERM]	The caller was not the super-user.

SEE ALSO

gethostid(2)

BUGS

Host names are limited to 255 characters.

NAME

gettimer, settimer — 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;

settimer(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* in the structure at *value*. The *settimer* call sets a timer to the specified *value* (returning the previous value of the timer if *ovalue* is nonzero).

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

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

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

Time values smaller than the resolution of the system clock are rounded up to this resolution (on the VAX, 10 milliseconds).

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

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

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

NOTES

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

RETURN VALUE

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

ERRORS

The possible errors are:

[EFAULT] The *value* structure specified a bad address.

[EINVAL] A *value* structure specified a time was too large to be handled.

SEE ALSO

sigvec(2), gettimeofday(2)

NAME

`getpagesize` — get system page size

SYNOPSIS

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

DESCRIPTION

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

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

SEE ALSO

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

NAME

`getpeername` — get name of connected peer

SYNOPSIS

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

DESCRIPTION

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

DIAGNOSTICS

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

ERRORS

The call succeeds unless:

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

SEE ALSO

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

BUGS

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

NAME

getpgrp — get process group

SYNOPSIS

```
pgrp = getpgrp(pid)  
int pgrp;  
int pid;
```

DESCRIPTION

The process group of the specified process is returned by *getpgrp*. If *pid* is zero, then the call applies to the current process.

Process groups are used for distribution of signals, and by terminals to arbitrate requests for their input: processes which have the same process group as the terminal are foreground and may read, while others will block with a signal if they attempt to read.

This call is thus used by programs such as *cs(1)* to create process groups in implementing job control. The *TIOCGPRP* and *TIOCSPGRP* calls described in *ty(4)* are used to get/set the process group of the control terminal.

SEE ALSO

setpgrp(2), getuid(2), tty(4)

NAME

getpid, *getppid* — get process identification

SYNOPSIS

```
pid = getpid()  
long pid;
```

```
ppid = getppid()  
long ppid;
```

DESCRIPTION

Getpid returns the process ID of the current process. Most often it is used with the host identifier *gethostid(2)* to generate uniquely-named temporary files.

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

SEE ALSO

gethostid(2)

NAME

getpriority, setpriority — get/set program scheduling priority

SYNOPSIS

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

#define PRIO_PROCESS  0 /* process */
#define PRIO_PGRP    1 /* process group */
#define PRIO_USER    2 /* user id */

prio = getpriority(which, who)
int prio, which, who;

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

DESCRIPTION

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

The *getpriority* call returns the highest priority (lowest numerical value) enjoyed by any of the specified processes. The *setpriority* call sets the priorities of all of the specified processes to the specified value. Only the super-user may lower priorities.

RETURN VALUE

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

ERRORS

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

- [ESRCH] No process(es) were located using the *which* and *who* values specified.
- [EINVAL] *Which* was not one of PRIO_PROCESS, PRIO_PGRP, or PRIO_USER.

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

- [EACCES] A process was located, but neither its effective nor real user ID matched the effective user ID of the caller.
- [EACCES] A non super-user attempted to change a process priority to a negative value.

SEE ALSO

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

NAME

getrlimit, setrlimit — control maximum system resource consumption

SYNOPSIS

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

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

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

DESCRIPTION

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

The *resource* parameter is one of the following:

RLIMIT_CPU the maximum amount of cpu time (in milliseconds) to be used by each process.

RLIMIT_FSIZE the largest size, in bytes, of any single file which may be created.

RLIMIT_DATA the maximum size, in bytes, of the data segment for a process; this defines how far a program may extend its break with the *sbrk(2)* system call.

RLIMIT_STACK the maximum size, in bytes, of the stack segment for a process; this defines how far a program's stack segment may be extended, either automatically by the system, or explicitly by a user with the *sbrk(2)* system call.

RLIMIT_CORE the largest size, in bytes, of a *core* file which may be created.

RLIMIT_RSS the maximum size, in bytes, a process's resident set size may grow to. This imposes a limit on the amount of physical memory to be given to a process; if memory is tight, the system will prefer to take memory from processes which are exceeding their declared resident set size.

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

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

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

An "infinite" value for a limit is defined as **RLIMIT_INFINITY** (0x7fffffff).

Because this information is stored in the per-process information, this system call must be executed directly by the shell if it is to affect all future processes created by the shell; *limit* is thus a built-in command to *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 soft cpu time limit is exceeded, a signal SIGXCPU is sent to the offending process.

RETURN VALUE

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

ERRORS

The possible errors are:

- [EFAULT] The address specified for *rlp* is invalid.
- [EPERM] The limit specified to *setrlimit* would have raised the maximum limit value, and the caller is not the super-user.

SEE ALSO

csh(1), *quota*(2)

BUGS

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

NAME

getrusage — get information about resource utilization

SYNOPSIS

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

#define RUSAGE_SELF      0    /* calling process */
#define RUSAGE_CHILDREN -1   /* terminated child processes */

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

DESCRIPTION

Getrusage returns information describing the resources utilized by the current process, or all its terminated child processes. The *who* parameter is one of `RUSAGE_SELF` and `RUSAGE_CHILDREN`. If *rusage* is non-zero, the buffer it points to will be filled in with the following structure:

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

The fields are interpreted as follows:

<code>ru_utime</code>	the total amount of time spent executing in user mode.
<code>ru_stime</code>	the total amount of time spent in the system executing on behalf of the process(es).
<code>ru_maxrss</code>	the maximum resident set size utilized (in kilobytes).
<code>ru_ixrss</code>	an “integral” value indicating the amount of memory used which was also shared among other processes. This value is expressed in units of kilobytes * seconds-of-execution and is calculated by summing the number of shared memory pages in use each time the internal system clock ticks and then averaging over 1 second intervals.
<code>ru_idrss</code>	an integral value of the amount of unshared memory residing in the data segment of a process (expressed in units of kilobytes * seconds-of-execution).
<code>ru_isrss</code>	an integral value of the amount of unshared memory residing in the stack segment of a process (expressed in units of kilobytes * seconds-of-execution).
<code>ru_minflt</code>	the number of page faults serviced without any i/o activity; here i/o activity is

	avoided by "reclaiming" a page frame from the list of pages awaiting reallocation.
<i>ru_majflt</i>	the number of page faults serviced which required i/o activity.
<i>ru_nswap</i>	the number of times a process was "swapped" out of main memory.
<i>ru_inblock</i>	the number of times the file system had to perform input.
<i>ru_outblock</i>	the number of times the file system had to perform output.
<i>ru_msgsnd</i>	the number of ipc messages sent.
<i>ru_msgrcv</i>	the number of ipc messages received.
<i>ru_nsignals</i>	the number of signals delivered.
<i>ru_nvcsw</i>	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).
<i>ru_nivcsw</i>	the number of times a context switch resulted due to a higher priority process becoming runnable or because the current process exceeded its time slice.

NOTES

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

SEE ALSO

gettimeofday(2), *wait(2)*

BUGS

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

NAME

`getsockname` — get socket name

SYNOPSIS

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

DESCRIPTION

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

DIAGNOSTICS

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

ERRORS

The call succeeds unless:

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

SEE ALSO

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

BUGS

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

NAME

getsockopt, setsockopt — get and set options on sockets

SYNOPSIS

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

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

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

DESCRIPTION

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

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

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

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

RETURN VALUE

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

ERRORS

The call succeeds unless:

[EBADF]	The argument <i>s</i> is not a valid descriptor.
[ENOTSOCK]	The argument <i>s</i> is a file, not a socket.
[ENOPROTOOPT]	The option is unknown.
[EFAULT]	The options are not in a valid part of the process address space.

SEE ALSO

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

NAME

gettimeofday, settimeofday — get/set date and time

SYNOPSIS

```
#include <sys/time.h>
gettimeofday(tp, tzp)
struct timeval *tp;
struct timezone *tzp;
settimeofday(tp, tzp)
struct timeval *tp;
struct timezone *tzp;
```

DESCRIPTION

Gettimeofday returns the system's notion of the current Greenwich time and the current time zone. Time returned is expressed relative in seconds and microseconds since midnight January 1, 1970.

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

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

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

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

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

RETURN

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

ERRORS

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

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

SEE ALSO

date(1), ctime(3)

BUGS

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

NAME

`getuid`, `geteuid` — get user identity

SYNOPSIS

```
uid = getuid()
int uid;

euid = geteuid()
int euid;
```

DESCRIPTION

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

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

SEE ALSO

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

NAME

`ioctl` — control device

SYNOPSIS

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

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

DESCRIPTION

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

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

RETURN VALUE

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

ERRORS

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

- [EBADF] *D* is not a valid descriptor.
- [ENOTTY] *D* is not associated with a character special device.
- [ENOTTY] The specified request does not apply to the kind of object which the descriptor *d* references.
- [EINVAL] *Request* or *argp* is not valid.

SEE ALSO

`execve(2)`, `fcntl(2)`, `mt(4)`, `tty(4)`, `intro(4N)`

NAME

kill — send signal to a process

SYNOPSIS

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

DESCRIPTION

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

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

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

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

Processes may send signals to themselves.

RETURN VALUE

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

ERRORS

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

- [EINVAL] *Sig* is not a valid signal number.
- [ESRCH] No process can be found corresponding to that specified by *pid*.
- [EPERM] The sending process is not the super-user and its effective user id does not match the effective user-id of the receiving process.

SEE ALSO

getpid(2), getpg(2), killpg(2), sigvec(2)

NAME

killpg — send signal to a process group

SYNOPSIS

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

DESCRIPTION

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

The sending process and members of the process group must have the same effective user ID, otherwise this call is restricted to the super-user. As a single special case the continue signal SIGCONT may be sent to any process which is a descendant of the current process.

RETURN VALUE

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

ERRORS

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

- [EINVAL] *Sig* is not a valid signal number.
- [ESRCH] No process can be found corresponding to that specified by *pid*.
- [EPERM] The sending process is not the super-user and one or more of the target processes has an effective user ID different from that of the sending process.

SEE ALSO

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

NAME

link — make a hard link to a file

SYNOPSIS

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

DESCRIPTION

A hard link to *name1* is created; the link has the name *name2*. *Name1* must exist.

With hard links, both *name1* and *name2* must be in the same file system. Unless the caller is the super-user, *name1* must not be a directory. Both the old and the new *link* share equal access and rights to the underlying object.

RETURN VALUE

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

ERRORS

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

- [EPERM] Either pathname contains a byte with the high-order bit set.
- [ENOENT] Either pathname was too long.
- [ENOTDIR] A component of either path prefix is not a directory.
- [ENOENT] A component of either path prefix does not exist.
- [EACCES] A component of either path prefix denies search permission.
- [ENOENT] The file named by *name1* does not exist.
- [EEXIST] The link named by *name2* does exist.
- [EPERM] The file named by *name1* is a directory and the effective user ID is not super-user.
- [EXDEV] The link named by *name2* and the file named by *name1* are on different file systems.
- [EACCES] The requested link requires writing in a directory with a mode that denies write permission.
- [EROFS] The requested link requires writing in a directory on a read-only file system.
- [EFAULT] One of the pathnames specified is outside the process's allocated address space.
- [ELOOP] Too many symbolic links were encountered in translating the pathname.

SEE ALSO

symlink(2), unlink(2)

NAME

listen — listen for connections on a socket

SYNOPSIS

```
listen(s, backlog)  
int s, backlog;
```

DESCRIPTION

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

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

RETURN VALUE

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

ERRORS

The call fails if:

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

SEE ALSO

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

BUGS

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

NAME

`lseek` — move read/write pointer

SYNOPSIS

```
#define L_SET    0    /* set the seek pointer */
#define L_INCR  1    /* increment the seek pointer */
#define L_XTND  2    /* extend the file size */
```

```
pos = lseek(d, offset, whence)
```

```
int pos;
```

```
int d, offset, whence;
```

DESCRIPTION

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

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

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

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

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

NOTES

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

RETURN VALUE

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

ERRORS

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

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

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

[EINVAL] *Whence* is not a proper value.

[EINVAL] The resulting file pointer would be negative.

SEE ALSO

`dup(2)`, `open(2)`

BUGS

This document's use of *whence* is incorrect English, but maintained for historical reasons.

NAME

`mkdir` — make a directory file

SYNOPSIS

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

DESCRIPTION

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

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

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

RETURN VALUE

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

ERRORS

Mkdir will fail and no directory will be created if:

- | | |
|-----------|---|
| [EPERM] | The process's effective user ID is not super-user. |
| [EPERM] | The <i>path</i> argument contains a byte with the high-order bit set. |
| [ENOTDIR] | A component of the path prefix is not a directory. |
| [ENOENT] | A component of the path prefix does not exist. |
| [EROFS] | The named file resides on a read-only file system. |
| [EEXIST] | The named file exists. |
| [EFAULT] | <i>Path</i> points outside the process's allocated address space. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |
| [EIO] | An I/O error occurred while writing to the file system. |

SEE ALSO

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

NAME

`mknod` — make a special file

SYNOPSIS

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

DESCRIPTION

Mknod creates a new file whose name is *path*. The mode of the new file (including special file bits) is initialized from *mode*. (The protection part of the mode is modified by the process's mode mask; see *umask(2)*). The first block pointer of the i-node is initialized from *dev* and is used to specify which device the special file refers to.

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

Mknod may be invoked only by the super-user.

RETURN VALUE

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

ERRORS

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

- | | |
|-----------|---|
| [EPERM] | The process's effective user ID is not super-user. |
| [EPERM] | The pathname contains a character with the high-order bit set. |
| [ENOTDIR] | A component of the path prefix is not a directory. |
| [ENOENT] | A component of the path prefix does not exist. |
| [EROFS] | The named file resides on a read-only file system. |
| [EEXIST] | The named file exists. |
| [EFAULT] | <i>Path</i> points outside the process's allocated address space. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |

SEE ALSO

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

NAME

mount, umount — mount or remove file system

SYNOPSIS

```
mount(special, name, rflag)
char *special, *name;
int rflag;

umount(special)
char *special;
```

DESCRIPTION

Mount announces to the system that a removable file system has been mounted on the block-structured special file *special*; from now on, references to file *name* will refer to the root file on the newly mounted file system. *Special* and *name* are pointers to null-terminated strings containing the appropriate path names.

Name must exist already. *Name* must be a directory. Its old contents are inaccessible while the file system is mounted.

The *rflag* argument determines whether the file system can be written on; if it is 0 writing is allowed, if non-zero no writing is done. Physically write-protected and magnetic tape file systems must be mounted read-only or errors will occur when access times are updated, whether or not any explicit write is attempted.

Umount announces to the system that the *special* file is no longer to contain a removable file system. The associated file reverts to its ordinary interpretation.

RETURN VALUE

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

Umount returns 0 if the action occurred; -1 if the special file is inaccessible or does not have a mounted file system, or if there are active files in the mounted file system.

ERRORS

Mount will fail when one of the following occurs:

- | | |
|-----------|---|
| [NODEV] | The caller is not the super-user. |
| [NODEV] | <i>Special</i> does not exist. |
| [ENOTBLK] | <i>Special</i> is not a block device. |
| [ENXIO] | The major device number of <i>special</i> is out of range (this indicates no device driver exists for the associated hardware). |
| [EPERM] | The pathname contains a character with the high-order bit set. |
| [ENOTDIR] | A component of the path prefix in <i>name</i> is not a directory. |
| [EROFS] | <i>Name</i> resides on a read-only file system. |
| [EBUSY] | <i>Name</i> is not a directory, or another process currently holds a reference to it. |
| [EBUSY] | No space remains in the mount table. |
| [EBUSY] | The super block for the file system had a bad magic number or an out of range block size. |
| [EBUSY] | Not enough memory was available to read the cylinder group information for the file system. |
| [EBUSY] | An i/o error occurred while reading the super block or cylinder group information. |

Umount may fail with one of the following errors:

- [NODEV] The caller is not the super-user.
- [NODEV] *Special* does not exist.
- [ENOTBLK] *Special* is not a block device.
- [ENXIO] The major device number of *special* is out of range (this indicates no device driver exists for the associated hardware).
- [EINVAL] The requested device is not in the mount table.
- [EBUSY] A process is holding a reference to a file located on the file system.

SEE ALSO

mount(8), umount(8)

BUGS

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

NAME

`open` — open a file for reading or writing, or create a new file

SYNOPSIS

```
#include <sys/file.h>
open(path, flags, mode)
char *path;
int flags, mode;
```

DESCRIPTION

Open opens the file *path* for reading and/or writing, as specified by the *flags* argument and returns a descriptor for that file. The *flags* argument may indicate the file is to be created if it does not already exist (by specifying the `O_CREAT` flag), in which case the file is created with mode *mode* as described in *chmod(2)* and modified by the process' *umask* value (see *umask(2)*).

Path is the address of a string of ASCII characters representing a path name, terminated by a null character. The flags specified are formed by *or*'ing the following values

<code>O_RDONLY</code>	open for reading only
<code>O_WRONLY</code>	open for writing only
<code>O_RDWR</code>	open for reading and writing
<code>O_NDELAY</code>	do not block on open
<code>O_APPEND</code>	append on each write
<code>O_CREAT</code>	create file if it does not exist
<code>O_TRUNC</code>	truncate size to 0
<code>O_EXCL</code>	error if create and file exists

Opening a file with `O_APPEND` set causes each write on the file to be appended to the end. If `O_TRUNC` is specified and the file exists, the file is truncated to zero length. If `O_EXCL` is set with `O_CREAT`, then if the file already exists, the open returns an error. This can be used to implement a simple exclusive access locking mechanism. If the `O_NDELAY` flag is specified and the open call would result in the process being blocked for some reason (e.g. waiting for carrier on a dialup line), the open returns immediately. The first time the process attempts to perform i/o on the open file it will block (not currently implemented).

Upon successful completion a non-negative integer termed a file descriptor is returned. The file pointer used to mark the current position within the file is set to the beginning of the file.

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

No process may have more than `{OPEN_MAX}` file descriptors open simultaneously.

ERRORS

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

[E <code>PERM</code>]	The pathname contains a character with the high-order bit set.
[E <code>NOTDIR</code>]	A component of the path prefix is not a directory.
[E <code>NOENT</code>]	<code>O_CREAT</code> is not set and the named file does not exist.
[E <code>ACCES</code>]	A component of the path prefix denies search permission.
[E <code>ACCES</code>]	The required permissions (for reading and/or writing) are denied for the named flag.
[E <code>ISDIR</code>]	The named file is a directory, and the arguments specify it is to be opened for writing.
[E <code>ROFS</code>]	The named file resides on a read-only file system, and the file is to be modified.

- [EMFILE] (OPEN_MAX) file descriptors are currently open.
- [ENXIO] The named file is a character special or block special file, and the device associated with this special file does not exist.
- [ETXTBSY] The file is a pure procedure (shared text) file that is being executed and the *open* call requests write access.
- [EFAULT] *Path* points outside the process's allocated address space.
- [ELOOP] Too many symbolic links were encountered in translating the pathname.
- [EEXIST] O_EXCL was specified and the file exists.
- [ENXIO] The O_NDELAY flag is given, and the file is a communications device on which there is no carrier present.
- [EOPNOTSUPP] An attempt was made to open a socket (not currently implemented).

SEE ALSO

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

NAME

`pipe` — create an interprocess communication channel

SYNOPSIS

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

DESCRIPTION

The *pipe* system call creates an I/O mechanism called a pipe. The file descriptors returned can be used in read and write operations. When the pipe is written using the descriptor *fdes*[1] up to 4096 bytes of data are buffered before the writing process is suspended. A read using the descriptor *fdes*[0] will pick up the data.

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

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

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

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

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

RETURN VALUE

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

ERRORS

The *pipe* call will fail if:

[EMFILE] Too many descriptors are active.

[EFAULT] The *fdes* buffer is in an invalid area of the process's address space.

SEE ALSO

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

BUGS

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

NAME

profil — execution time profile

SYNOPSIS

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

DESCRIPTION

Buff points to an area of core whose length (in bytes) is given by *bufsiz*. After this call, the user's program counter (*pc*) is examined each clock tick (10 milliseconds); *offset* is subtracted from it, and the result multiplied by *scale*. If the resulting number corresponds to a word inside *buff*, that word is incremented.

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

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

RETURN VALUE

A 0, indicating success, is always returned.

SEE ALSO

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

NAME

ptrace -- process trace

SYNOPSIS

```
#include <signal.h>

ptrace(request, pid, addr, data)
int request, pid, *addr, data;
```

DESCRIPTION

Ptrace provides a means by which a parent process may control the execution of a child process, and examine and change its core image. Its primary use is for the implementation of breakpoint debugging. There are four arguments whose interpretation depends on a *request* argument. Generally, *pid* is the process ID of the traced process, which must be a child (no more distant descendant) of the tracing process. A process being traced behaves normally until it encounters some signal whether internally generated like "illegal instruction" or externally generated like "interrupt". See *sigvec(2)* for the list. Then the traced process enters a stopped state and its parent is notified via *wait(2)*. When the child is in the stopped state, its core image can be examined and modified using *ptrace*. If desired, another *ptrace* request can then cause the child either to terminate or to continue, possibly ignoring the signal.

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

- 0 This request is the only one used by the child process; it declares that the process is to be traced by its parent. All the other arguments are ignored. Peculiar results will ensue if the parent does not expect to trace the child.
- 1,2 The word in the child process's address space at *addr* is returned. If I and D space are separated (e.g. historically on a pdp-11), request 1 indicates I space, 2 D space. *Addr* must be even. The child must be stopped. The input *data* is ignored.
- 3 The word of the system's per-process data area corresponding to *addr* is returned. *Addr* must be even and less than 512. This space contains the registers and other information about the process; its layout corresponds to the *user* structure in the system.
- 4,5 The given *data* is written at the word in the process's address space corresponding to *addr*, which must be even. No useful value is returned. If I and D space are separated, request 4 indicates I space, 5 D space. Attempts to write in pure procedure fail if another process is executing the same file.
- 6 The process's system data is written, as it is read with request 3. Only a few locations can be written in this way: the general registers, the floating point status and registers, and certain bits of the processor status word.
- 7 The *data* argument is taken as a signal number and the child's execution continues at location *addr* as if it had incurred that signal. Normally the signal number will be either 0 to indicate that the signal that caused the stop should be ignored, or that value fetched out of the process's image indicating which signal caused the stop. If *addr* is (int *)1 then execution continues from where it stopped.
- 8 The traced process terminates.
- 9 Execution continues as in request 7; however, as soon as possible after execution of at least one instruction, execution stops again. The signal number from the stop is SIGTRAP. (On the VAX-11 the T-bit is used and just one instruction is executed.) This is part of the mechanism for implementing breakpoints.

As indicated, these calls (except for request 0) can be used only when the subject process has stopped. The *wait* call is used to determine when a process stops; in such a case the "termination" status returned by *wait* has the value 0177 to indicate stoppage rather than genuine termination.

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

On a VAX-11, "word" also means a 32-bit integer, but the "even" restriction does not apply.

RETURN VALUE

A 0 value is returned if the call succeeds. If the call fails then a -1 is returned and the global variable *errno* is set to indicate the error.

ERRORS

[EINVAL]	The request code is invalid.
[EINVAL]	The specified process does not exist.
[EINVAL]	The given signal number is invalid.
[EFAULT]	The specified address is out of bounds.
[EPERM]	The specified process cannot be traced.

SEE ALSO

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

BUGS

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

The request 0 call should be able to specify signals which are to be treated normally and not cause a stop. In this way, for example, programs with simulated floating point (which use "illegal instruction" signals at a very high rate) could be efficiently debugged.

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

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

NAME

quota — manipulate disk quotas

SYNOPSIS

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

quota(cmd, uid, arg, addr)
int cmd, uid, arg;
caddr_t addr;
```

DESCRIPTION

The *quota* call manipulates disk quotas for file systems which have had quotas enabled with *setquota(2)*. The *cmd* parameter indicates a command to be applied to the user ID *uid*. *Arg* is a command specific argument and *addr* is the address of an optional, command specific, data structure which is copied in or out of the system. The interpretation of *arg* and *addr* is given with each command below.

Q_SETDLIM

Set disc quota limits and current usage for the user with ID *uid*. *Arg* is a major-minor device indicating a particular file system. *Addr* is a pointer to a struct dqblk structure (defined in *<sys/quota.h>*). This call is restricted to the super-user.

Q_GETDLIM

Get disc quota limits and current usage for the user with ID *uid*. The remaining parameters are as for **Q_SETDLIM**.

Q_SETDUSE

Set disc usage limits for the user with ID *uid*. *Arg* is a major-minor device indicating a particular file system. *Addr* is a pointer to a struct dqusage structure (defined in *<sys/quota.h>*). This call is restricted to the super-user.

Q_SYNC

Update the on-disc copy of quota usages. The *uid*, *arg*, and *addr* parameters are ignored.

Q_SETUID

Change the calling process's quota limits to those of the user with ID *uid*. The *arg* and *addr* parameters are ignored. This call is restricted to the super-user.

Q_SETWARN

Alter the disc usage warning limits for the user with ID *uid*. *Arg* is a major-minor device indicating a particular file system. *Addr* is a pointer to a struct dqwarn structure (defined in *<sys/quota.h>*). This call is restricted to the super-user.

Q_DOWARN

Warn the user with user ID *uid* about excessive disc usage. This call causes the system to check its current disc usage information and print a message on the terminal of the caller for each file system on which the user is over quota. If the *arg* parameter is specified as **NODEV**, all file systems which have disc quotas will be checked. Otherwise, *arg* indicates a specific major-minor device to be checked. This call is restricted to the super-user.

RETURN VALUE

A successful call returns 0 and, possibly, more information specific to the *cmd* performed; when an error occurs, the value **-1** is returned and *errno* is set to indicate the reason.

ERRORS

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

[EINVAL] *Cmd* is invalid.

[ESRCH]	No disc quota is found for the indicated user.
[EPERM]	The call is privileged and the caller was not the super-user.
[EINVAL]	The <i>arg</i> parameter is being interpreted as a major-minor device and it indicates an unmounted file system.
[EFAULT]	An invalid <i>addr</i> is supplied; the associated structure could not be copied in or out of the kernel.
[EUSERS]	The quota table is full.

SEE ALSO

setquota(2), quotaon(8), quotacheck(8)

BUGS

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

The Australian spelling of *disk* is used throughout the quota facilities in honor of the implementors.

NAME

`read`, `readv` — read input

SYNOPSIS

```
cc = read(d, buf, nbytes)
int cc, d;
char *buf;
int nbytes;
#include <sys/types.h>
#include <sys/uio.h>
cc = readv(d, iov, iovcnt)
int cc, d;
struct iovec *iov;
int iovcnt;
```

DESCRIPTION

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

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

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

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

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

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

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

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

RETURN VALUE

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

ERRORS

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

- [EBADF] *Fildes* is not a valid file descriptor open for reading.
- [EFAULT] *Buf* points outside the allocated address space.
- [EINTR] A read from a slow device was interrupted before any data arrived by the delivery of a signal.

In addition, *readv* may return one of the following errors:

- [EINVAL] *Iovent* was less than or equal to 0, or greater than 16.
- [EINVAL] One of the *iov_len* values in the *iov* array was negative.

[EINVAL] The sum of the *iov_len* values in the *iov* array overflowed a 32-bit integer.

SEE ALSO

dup(2), open(2), pipe(2), socket(2), socketpair(2)

NAME

readlink — read value of a symbolic link

SYNOPSIS

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

DESCRIPTION

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

RETURN VALUE

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

ERRORS

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

- | | |
|-----------|---|
| [EPERM] | The <i>path</i> argument contained a byte with the high-order bit set. |
| [ENOENT] | The pathname was too long. |
| [ENOTDIR] | A component of the path prefix is not a directory. |
| [ENOENT] | The named file does not exist. |
| [ENXIO] | The named file is not a symbolic link. |
| [EACCES] | Search permission is denied on a component of the path prefix. |
| [EPERM] | The effective user ID does not match the owner of the file and the effective user ID is not the super-user. |
| [EINVAL] | The named file is not a symbolic link. |
| [EFAULT] | <i>Buf</i> extends outside the process's allocated address space. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |

SEE ALSO

stat(2), lstat(2), symlink(2)

NAME

reboot — reboot system or halt processor

SYNOPSIS

```
#include <sys/reboot.h>
reboot(howto)
int howto;
```

DESCRIPTION

Reboot reboots the system, and is invoked automatically in the event of unrecoverable system failures. *Howto* is a mask of options passed to the bootstrap program. The system call interface permits only `RB_HALT` or `RB_AUTOBOOT` to be passed to the reboot program; the other flags are used in scripts stored on the console storage media, or used in manual bootstrap procedures. When none of these options (e.g. `RB_AUTOBOOT`) is given, the system is rebooted from file “`vmunix`” in the root file system of unit 0 of a disk chosen in a processor specific way. An automatic consistency check of the disks is then normally performed.

The bits of *howto* are:

RB_HALT

the processor is simply halted; no reboot takes place. `RB_HALT` should be used with caution.

RB_ASKNAME

Interpreted by the bootstrap program itself, causing it to inquire as to what file should be booted. Normally, the system is booted from the file “`xx(0,0)vmunix`” without asking.

RB_SINGLE

Normally, the reboot procedure involves an automatic disk consistency check and then multi-user operations. `RB_SINGLE` prevents the consistency check, rather simply booting the system with a single-user shell on the console. `RB_SINGLE` is interpreted by the `init(8)` program in the newly booted system. This switch is not available from the system call interface.

Only the super-user may *reboot* a machine.

RETURN VALUES

If successful, this call never returns. Otherwise, a `-1` is returned and an error is returned in the global variable `errno`.

ERRORS

[`EPERM`] The caller is not the super-user.

SEE ALSO

`crash(8)`, `halt(8)`, `init(8)`, `reboot(8)`

BUGS

The notion of “console medium”, among other things, is specific to the VAX.

NAME

recv, *recvfrom*, *recvmsg* — receive a message from a socket

SYNOPSIS

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

cc = recv(s, buf, len, flags)
int cc, s;
char *buf;
int len, flags;

cc = recvfrom(s, buf, len, flags, from, fromlen)
int cc, s;
char *buf;
int len, flags;
struct sockaddr *from;
int *fromlen;

cc = recvmsg(s, msg, flags)
int cc, s;
struct msghdr msg[];
int flags;
```

DESCRIPTION

Recv, *recvfrom*, and *recvmsg* are used to receive messages from a socket.

The *recv* call may be used only on a *connected* socket (see *connect(2)*), while *recvfrom* and *recvmsg* may be used to receive data on a socket whether it is in a connected state or not.

If *from* is non-zero, the source address of the message is filled in. *Fromlen* is a value-result parameter, initialized to the size of the buffer associated with *from*, and modified on return to indicate the actual size of the address stored there. The length of the message is returned in *cc*. If a message is too long to fit in the supplied buffer, excess bytes may be discarded depending on the type of socket the message is received from; see *socket(2)*.

If no messages are available at the socket, the receive call waits for a message to arrive, unless the socket is nonblocking (see *ioctl(2)*) in which case a *cc* of *-1* is returned with the external variable *errno* set to *EWOULDBLOCK*.

The *select(2)* call may be used to determine when more data arrives.

The *flags* argument to a send call is formed by *or*'ing one or more of the values,

```
#define MSG_PEEK    0x1    /* peek at incoming message */
#define MSG_OOB    0x2    /* process out-of-band data */
```

The *recvmsg* call uses a *msghdr* structure to minimize the number of directly supplied parameters. This structure has the following form, as defined in *<sys/socket.h>*:

```
struct msghdr {
    caddr_t msg_name;          /* optional address */
    int msg_namelen;          /* size of address */
    struct iov *msg_iov;      /* scatter/gather array */
    int msg_iovlen;           /* # elements in msg_iov */
    caddr_t msg_accrightr;    /* access rights sent/received */
    int msg_accrightrlen;
};
```



Here *msg_name* and *msg_namelen* specify the destination address if the socket is unconnected; *msg_name* may be given as a null pointer if no names are desired or required. The *msg_iov* and *msg_iovlen* describe the scatter gather locations, as described in *read(2)*. Access rights to be sent along with the message are specified in *msg_accrights*, which has length *msg_accrightslen*.

RETURN VALUE

These calls return the number of bytes received, or -1 if an error occurred.

ERRORS

The calls fail if:

- | | |
|---------------|---|
| [EBADF] | The argument <i>s</i> is an invalid descriptor. |
| [ENOTSOCK] | The argument <i>s</i> is not a socket. |
| [EWOULDBLOCK] | The socket is marked non-blocking and the receive operation would block. |
| [EINTR] | The receive was interrupted by delivery of a signal before any data was available for the receive. |
| [EFAULT] | The data was specified to be received into a non-existent or protected part of the process address space. |

SEE ALSO

read(2), *send(2)*, *socket(2)*

NAME

rename — change the name of a file

SYNOPSIS

```
rename(from, to)
char *from, *to;
```

DESCRIPTION

Rename causes the link named *from* to be renamed as *to*. If *to* exists, then it is first removed. Both *from* and *to* must be of the same type (that is, both directories or both non-directories), and must reside on the same file system.

Rename guarantees that an instance of *to* will always exist, even if the system should crash in the middle of the operation.

CAVEAT

The system can deadlock if a loop in the file system graph is present. This loop takes the form of an entry in directory "a", say "a/foo", being a hard link to directory "b", and an entry in directory "b", say "b/bar", being a hard link to directory "a". When such a loop exists and two separate processes attempt to perform "rename a/foo b/bar" and "rename b/bar a/foo", respectively, the system may deadlock attempting to lock both directories for modification. Hard links to directories should be replaced by symbolic links by the system administrator.

RETURN VALUE

A 0 value is returned if the operation succeeds, otherwise *rename* returns -1 and the global variable *errno* indicates the reason for the failure.

ERRORS

Rename will fail and neither of the argument files will be affected if any of the following are true:

- [ENOTDIR] A component of either path prefix is not a directory.
- [ENOENT] A component of either path prefix does not exist.
- [EACCES] A component of either path prefix denies search permission.
- [ENOENT] The file named by *from* does not exist.
- [EPERM] The file named by *from* is a directory and the effective user ID is not super-user.
- [EXDEV] The link named by *to* and the file named by *from* are on different logical devices (file systems). Note that this error code will not be returned if the implementation permits cross-device links.
- [EACCES] The requested link requires writing in a directory with a mode that denies write permission.
- [EROFS] The requested link requires writing in a directory on a read-only file system.
- [EFAULT] *Path* points outside the process's allocated address space.
- [EINVAL] *From* is a parent directory of *to*.

SEE ALSO

open(2)

NAME

`rmdir` — remove a directory file

SYNOPSIS

```
rmdir(path)
char *path;
```

DESCRIPTION

Rmdir removes a directory file whose name is given by *path*. The directory must not have any entries other than “.” and “..”.

RETURN VALUE

A 0 is returned if the remove succeeds; otherwise a -1 is returned and an error code is stored in the global location *errno*.

ERRORS

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

- [ENOTEMPTY] The named directory contains files other than “.” and “..” in it.
- [EPERM] The pathname contains a character with the high-order bit set.
- [ENOENT] The pathname was too long.
- [ENOTDIR] A component of the path prefix is not a directory.
- [ENOENT] The named file does not exist.
- [EACCES] A component of the path prefix denies search permission.
- [EACCES] Write permission is denied on the directory containing the link to be removed.
- [EBUSY] The directory to be removed is the mount point for a mounted file system.
- [EROFS] The directory entry to be removed resides on a read-only file system.
- [EFAULT] *Path* points outside the process's allocated address space.
- [ELOOP] Too many symbolic links were encountered in translating the pathname.

SEE ALSO

`mkdir(2)`, `unlink(2)`

NAME

select — synchronous i/o multiplexing

SYNOPSIS

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

nfound = select(nfds, readfds, writefds, exceptfds, timeout)
int nfound, nfds, *readfds, *writefds, *exceptfds;
struct timeval *timeout;
```

DESCRIPTION

Select examines the i/o descriptors specified by the bit masks *readfds*, *writefds*, and *exceptfds* to see if they are ready for reading, writing, or have an exceptional condition pending, respectively. File descriptor *f* is represented by the bit “1<<*f*” in the mask. *Nfds* descriptors are checked, i.e. the bits from 0 through *nfds*-1 in the masks are examined. *Select* returns, in place, a mask of those descriptors which are ready. The total number of ready descriptors is returned in *nfound*.

If *timeout* is a non-zero pointer, it specifies a maximum interval to wait for the selection to complete. If *timeout* is a zero pointer, the *select* blocks indefinitely. To affect a poll, the *timeout* argument should be non-zero, pointing to a zero valued *timeval* structure.

Any of *readfds*, *writefds*, and *exceptfds* may be given as 0 if no descriptors are of interest.

RETURN VALUE

Select returns the number of descriptors which are contained in the bit masks, or -1 if an error occurred. If the time limit expires then *select* returns 0.

ERRORS

An error return from *select* indicates:

- | | |
|---------|---|
| [EBADF] | One of the bit masks specified an invalid descriptor. |
| [EINTR] | An signal was delivered before any of the selected for events occurred or the time limit expired. |

SEE ALSO

accept(2), *connect(2)*, *read(2)*, *write(2)*, *recv(2)*, *send(2)*

BUGS

The descriptor masks are always modified on return, even if the call returns as the result of the *timeout*.

NAME

send, sendto, sendmsg — send a message from a socket

SYNOPSIS

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

cc = send(s, msg, len, flags)
int cc, s;
char *msg;
int len, flags;

cc = sendto(s, msg, len, flags, to, tolen)
int cc, s;
char *msg;
int len, flags;
struct sockaddr *to;
int tolen;

cc = sendmsg(s, msg, flags)
int cc, s;
struct msghdr msg[];
int flags;
```

DESCRIPTION

Send, *sendto*, and *sendmsg* are used to transmit a message to another socket. *Send* may be used only when the socket is in a *connected* state, while *sendto* and *sendmsg* may be used at any time.

The address of the target is given by *to* with *tolen* specifying its size. The length of the message is given by *len*. If the message is too long to pass atomically through the underlying protocol, then the error EMSGSIZE is returned, and the message is not transmitted.

No indication of failure to deliver is implicit in a *send*. Return values of -1 indicate some locally detected errors.

If no messages space is available at the socket to hold the message to be transmitted, then *send* normally blocks, unless the socket has been placed in non-blocking *i/o* mode. The *select(2)* call may be used to determine when it is possible to send more data.

The *flags* parameter may be set to 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

[EBADF]	An invalid descriptor was specified.
[ENOTSOCK]	The argument <i>s</i> is not a socket.
[EFAULT]	An invalid user space address was specified for a parameter.
[EMSGSIZE]	The socket requires that message be sent atomically, and the size of the message to be sent made this impossible.
[EWOULDBLOCK]	The socket is marked non-blocking and the requested operation would block.

SEE ALSO

recv(2), *socket(2)*

NAME

setgroups — set group access list

SYNOPSIS

```
#include <sys/param.h>
setgroups(ngroups, gidset)
int ngroups, *gidset;
```

DESCRIPTION

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

Only the super-user may set new groups.

RETURN VALUE

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

ERRORS

The *setgroups* call will fail if:

- [EPERM] The caller is not the super-user.
- [EFAULT] The address specified for *gidset* is outside the process address space.

SEE ALSO

getgroups(2), initgroups(3X)

NAME

`setpgrp` — set process group

SYNOPSIS

```
setpgrp(pid, pgrp)  
int pid, pgrp;
```

DESCRIPTION

Setpgrp sets the process group of the specified process *pid* to the specified *pgrp*. If *pid* is zero, then the call applies to the current process.

If the invoker is not the super-user, then the affected process must have the same effective user-id as the invoker or be a descendant of the invoking process.

RETURN VALUE

Setpgrp returns when the operation was successful. If the request failed, `-1` is returned and the global variable *errno* indicates the reason.

ERRORS

Setpgrp will fail and the process group will not be altered if one of the following occur:

- | | |
|---------|---|
| [ESRCH] | The requested process does not exist. |
| [EPERM] | The effective user ID of the requested process is different from that of the caller and the process is not a descendent of the calling process. |

SEE ALSO

`getpgrp(2)`

NAME

setquota -- enable/disable quotas on a file system

SYNOPSIS

setquota(*special*, *file*)
char **special*, **file*;

DESCRIPTION

Disc quotas are enabled or disabled with the *setquota* call. *Special* indicates a block special device on which a mounted file system exists. If *file* is nonzero, it specifies a file in that file system from which to take the quotas. If *file* is 0, then quotas are disabled on the file system. The quota file must exist; it is normally created with the *checkquota*(8) program.

Only the super-user may turn quotas on or off.

SEE ALSO

quota(2), quotacheck(8), quotaon(8)

RETURN VALUE

A 0 return value indicates a successful call. A value of -1 is returned when an error occurs and *errno* is set to indicate the reason for failure.

ERRORS

Setquota will fail when one of the following occurs:

- | | |
|-----------|---|
| [NODEV] | The caller is not the super-user. |
| [NODEV] | <i>Special</i> does not exist. |
| [ENOTBLK] | <i>Special</i> is not a block device. |
| [ENXIO] | The major device number of <i>special</i> is out of range (this indicates no device driver exists for the associated hardware). |
| [EPERM] | The pathname contains a character with the high-order bit set. |
| [ENOTDIR] | A component of the path prefix in <i>file</i> is not a directory. |
| [EROFS] | <i>File</i> resides on a read-only file system. |
| [EACCES] | <i>File</i> resides on a file system different from <i>special</i> . |
| [EACCES] | <i>File</i> is not a plain file. |

BUGS

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

NAME

setregid — set real and effective group ID

SYNOPSIS

setregid (rgid, egid)
int rgid, egid;

DESCRIPTION

The real and effective group ID's of the current process are set to the arguments. Unprivileged users may change the real group ID to the effective group ID and vice-versa; only the super-user may make other changes.

Supplying a value of -1 for either the real or effective group ID forces the system to substitute the current ID in place of the -1 parameter.

RETURN VALUE

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

ERRORS

[EPERM] The current process is not the super-user and a change other than changing the effective group-id to the real group-id was specified.

SEE ALSO

getgid(2), setreuid(2), setgid(3)

NAME

setreuid — set real and effective user ID's

SYNOPSIS

```
setreuid(ruid, euid)
int ruid, euid;
```

DESCRIPTION

The real and effective user ID's of the current process are set according to the arguments. If *ruid* or *euid* is -1 , the current uid is filled in by the system. Unprivileged users may change the real user ID to the effective user ID and vice-versa; only the super-user may make other changes.

RETURN VALUE

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

ERRORS

[EPERM] The current process is not the super-user and a change other than changing the effective user-id to the real user-id was specified.

SEE ALSO

getuid(2), setregid(2), setuid(3)

NAME

shutdown — shut down part of a full-duplex connection

SYNOPSIS

```
shutdown(s, how)
int s, how;
```

DESCRIPTION

The *shutdown* call causes all or part of a full-duplex connection on the socket associated with *s* to be shut down. If *how* is 0, then further receives will be disallowed. If *how* is 1, then further sends will be disallowed. If *how* is 2, then further sends and receives will be disallowed.

DIAGNOSTICS

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

ERRORS

The call succeeds unless:

[EBADF] *S* is not a valid descriptor.

[ENOTSOCK] *S* is a file, not a socket.

[ENOTCONN] The specified socket is not connected.

SEE ALSO

connect(2), socket(2)

NAME

sigblock — block signals

SYNOPSIS

```
sigblock(mask);  
int mask;
```

DESCRIPTION

Sigblock causes the signals specified in *mask* to be added to the set of signals currently being blocked from delivery. Signal *i* is blocked if the *i*-th bit in *mask* is a 1.

It is not possible to block SIGKILL, SIGSTOP, or SIGCONT; this restriction is silently imposed by the system.

RETURN VALUE

The previous set of masked signals is returned.

SEE ALSO

kill(2), sigvec(2), sigsetmask(2),

NAME

`sigpause` — atomically release blocked signals and wait for interrupt

SYNOPSIS

```
sigpause(sigmask)  
int sigmask;
```

DESCRIPTION

Sigpause assigns *sigmask* to the set of masked signals and then waits for a signal to arrive; on return the set of masked signals is restored. *Sigmask* is usually 0 to indicate that no signals are now to be blocked. *Sigpause* always terminates by being interrupted, returning `EINTR`.

In normal usage, a signal is blocked using *sigblock(2)*, to begin a critical section, variables modified on the occurrence of the signal are examined to determine that there is no work to be done, and the process pauses awaiting work by using *sigpause* with the mask returned by *sigblock*.

SEE ALSO

`sigblock(2)`, `sigvec(2)`

NAME

sigsetmask — set current signal mask

SYNOPSIS

```
sigsetmask(mask);  
int mask;
```

DESCRIPTION

Sigsetmask sets the current signal mask (those signals which are blocked from delivery). Signal *i* is blocked if the *i*-th bit in *mask* is a 1.

The system quietly disallows SIGKILL, SIGSTOP, or SIGCONT to be blocked.

RETURN VALUE

The previous set of masked signals is returned.

SEE ALSO

kill(2), sigvec(2), sigblock(2), sigpause(2)

NAME

sigstack — set and/or get signal stack context

SYNOPSIS

```
#include <signal.h>
struct sigstack {
    caddr_t  ss_sp;
    int     ss_onstack;
};
sigstack(ss, oss);
struct sigstack *ss, *oss;
```

DESCRIPTION

Sigstack allows users to define an alternate stack on which signals are to be processed. If *ss* is non-zero, it specifies a *signal stack* on which to deliver signals and tells the system if the process is currently executing on that stack. When a signal's action indicates its handler should execute on the signal stack (specified with a *sigvec(2)* call), the system checks to see if the process is currently executing on that stack. If the process is not currently executing on the signal stack, the system arranges a switch to the signal stack for the duration of the signal handler's execution. If *oss* is non-zero, the current signal stack state is returned.

NOTES

Signal stacks are not “grown” automatically, as is done for the normal stack. If the stack overflows unpredictable results may occur.

RETURN VALUE

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

ERRORS

Sigstack will fail and the signal stack context will remain unchanged if one of the following occurs.

[EFAULT] Either *ss* or *oss* points to memory which is not a valid part of the process address space.

SEE ALSO

sigvec(2), *setjmp(3)*

NAME

sigvec — software signal facilities

SYNOPSIS

```
#include <signal.h>

struct sigvec {
    int      (*sv_handler) 0;
    int      sv_mask;
    int      sv_onstack;
};

sigvec(sig, vec, ovec)
int sig;
struct sigvec *vec, *ovec;
```

DESCRIPTION

The system defines a set of signals that may be delivered to a process. Signal delivery resembles the occurrence of a hardware interrupt: the signal is blocked from further occurrence, the current process context is saved, and a new one is built. A process may specify a *handler* to which a signal is delivered, or specify that a signal is to be *blocked* or *ignored*. A process may also specify that a default action is to be taken by the system when a signal occurs. Normally, signal handlers execute on the current stack of the process. This may be changed, on a per-handler basis, so that signals are taken on a special *signal stack*.

All signals have the same *priority*. Signal routines execute with the signal that caused their invocation *blocked*, but other signals may yet occur. A global *signal mask* defines the set of signals currently blocked from delivery to a process. The signal mask for a process is initialized from that of its parent (normally 0). It may be changed with a *sigblock(2)* or *sigsetmask(2)* call, or when a signal is delivered to the process.

When a signal condition arises for a process, the signal is added to a set of signals pending for the process. If the signal is not currently *blocked* by the process then it is delivered to the process. When a signal is delivered, the current state of the process is saved, a new signal mask is calculated (as described below), and the signal handler is invoked. The call to the handler is arranged so that if the signal handling routine returns normally the process will resume execution in the context from before the signal's delivery. If the process wishes to resume in a different context, then it must arrange to restore the previous context itself.

When a signal is delivered to a process a new signal mask is installed for the duration of the process' signal handler (or until a *sigblock* or *sigsetmask* call is made). This mask is formed by taking the current signal mask, adding the signal to be delivered, and *or'ing* in the signal mask associated with the handler to be invoked.

Sigvec assigns a handler for a specific signal. If *vec* is non-zero, it specifies a handler routine and mask to be used when delivering the specified signal. Further, if *sv_onstack* is 1, the system will deliver the signal to the process on a *signal stack*, specified with *sigstack(2)*. If *ovec* is non-zero, the previous handling information for the signal is returned to the user.

The following is a list of all signals with names as in the include file *<signal.h>*:

SIGHUP	1	hangup
SIGINT	2	interrupt
SIGQUIT	3*	quit
SIGILL	4*	illegal instruction
SIGTRAP	5*	trace trap
SIGIOT	6*	IOT instruction
SIGEMT	7*	EMT instruction
SIGFPE	8*	floating point exception

SIGKILL	9	kill (cannot be caught, blocked, or ignored)
SIGBUS	10•	bus error
SIGSEGV	11•	segmentation violation
SIGSYS	12•	bad argument to system call
SIGPIPE	13	write on a pipe with no one to read it
SIGALRM	14	alarm clock
SIGTERM	15	software termination signal
SIGURG	16•	urgent condition present on socket
SIGSTOP	17†	stop (cannot be caught, blocked, or ignored)
SIGTSTP	18†	stop signal generated from keyboard
SIGCONT	19•	continue after stop (cannot be blocked)
SIGCHLD	20•	child status has changed
SIGTTIN	21†	background read attempted from control terminal
SIGTTOU	22†	background write attempted to control terminal
SIGIO	23•	i/o is possible on a descriptor (see <i>fcntl(2)</i>)
SIGXCPU	24	cpu time limit exceeded (see <i>setrlimit(2)</i>)
SIGXFSZ	25	file size limit exceeded (see <i>setrlimit(2)</i>)
SIGVTALRM	26	virtual time alarm (see <i>setitimer(2)</i>)
SIGPROF	27	profiling timer alarm (see <i>setitimer(2)</i>)

The starred signals in the list above cause a core image if not caught or ignored.

Once a signal handler is installed, it remains installed until another *sigvec* call is made, or an *execve(2)* is performed. The default action for a signal may be reinstated by setting *sv_handler* to *SIG_DFL*; this default is termination (with a core image for starred signals) except for signals marked with • or †. Signals marked with • are discarded if the action is *SIG_DFL*; signals marked with † cause the process to stop. If *sv_handler* is *SIG_IGN* the signal is subsequently ignored, and pending instances of the signal are discarded.

If a caught signal occurs during certain system calls, causing the call to terminate prematurely, the call is automatically restarted. In particular this can occur during a *read* or *write(2)* on a slow device (such as a terminal; but not a file) and during a *wait(2)*.

After a *fork(2)* or *vfork(2)* the child inherits all signals, the signal mask, and the signal stack.

Execve(2) resets all caught signals to default action; ignored signals remain ignored; the signal mask remains the same; the signal stack state is reset.

NOTES

The mask specified in *vec* is not allowed to block *SIGKILL*, *SIGSTOP*, or *SIGCONT*. This is done silently by the system.

RETURN VALUE

A 0 value indicated that the call succeeded. A -1 return value indicates an error occurred and *errno* is set to indicated the reason.

ERRORS

Sigvec will fail and no new signal handler will be installed if one of the following occurs:

[EFAULT]	Either <i>vec</i> or <i>ovec</i> points to memory which 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 <i>SIGKILL</i> or <i>SIGSTOP</i> .
[EINVAL]	An attempt is made to ignore <i>SIGCONT</i> (by default <i>SIGCONT</i> is ignored).

SEE ALSO

kill(1), *ptrace(2)*, *kill(2)*, *sigblock(2)*, *sigsetmask(2)*, *sigpause(2)*, *sigstack(2)*, *sigvec(2)*, *setjmp(3)*, *tty(4)*

NOTES (VAX-11)

The handler routine can be declared:

```
handler(sig, code, scp)
int sig, code;
struct sigcontext *scp;
```

Here *sig* is the signal number, into which the hardware faults and traps are mapped as defined below. *Code* is a parameter which is either a constant as given below or, for compatibility mode faults, the code provided by the hardware (Compatibility mode faults are distinguished from the other SIGILL traps by having PSL_CM set in the psl). *Scp* is a pointer to the *sigcontext* structure (defined in `<signal.h>`), used to restore the context from before the signal.

The following defines the mapping of hardware traps to signals and codes. All of these symbols are defined in `<signal.h>`:

Hardware condition	Signal	Code
Arithmetic traps:		
Integer overflow	SIGFPE	FPE_INTOVF_TRAP
Integer division by zero	SIGFPE	FPE_INTDIV_TRAP
Floating overflow trap	SIGFPE	FPE_FLTOVF_TRAP
Floating/decimal division by zero	SIGFPE	FPE_FLTDIV_TRAP
Floating underflow trap	SIGFPE	FPE_FLTUND_TRAP
Decimal overflow trap	SIGFPE	FPE_DECOVF_TRAP
Subscript-range	SIGFPE	FPE_SUBRNG_TRAP
Floating overflow fault	SIGFPE	FPE_FLTOVF_FAULT
Floating divide by zero fault	SIGFPE	FPE_FLTDIV_FAULT
Floating underflow fault	SIGFPE	FPE_FLTUND_FAULT
Length access control	SIGSEGV	
Protection violation	SIGBUS	
Reserved instruction	SIGILL	ILL_RESAD_FAULT
Customer-reserved instr.	SIGEMT	
Reserved operand	SIGILL	ILL_PRIVIN_FAULT
Reserved addressing	SIGILL	ILL_RESOP_FAULT
Trace pending	SIGTRAP	
Bpt instruction	SIGTRAP	
Compatibility-mode	SIGILL	hardware supplied code
Chme	SIGSEGV	
Chms	SIGSEGV	
Chmu	SIGSEGV	

BUGS

This manual page is confusing.

NAME

socket — create an endpoint for communication

SYNOPSIS

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

s = socket(af, type, protocol)
int s, af, type, protocol;
```

DESCRIPTION

Socket creates an endpoint for communication and returns a descriptor.

The *af* parameter specifies an address format with which addresses specified in later operations using the socket should be interpreted. These formats are defined in the include file `<sys/socket.h>`. The currently understood formats are

AF_UNIX	(UNIX path names),
AF_INET	(ARPA Internet addresses),
AF_PUP	(Xerox PUP-I Internet addresses), and
AF_IMPLINK	(IMP “host at IMP” addresses).

The socket has the indicated *type* which specifies the semantics of communication. Currently defined types are:

```
SOCK_STREAM
SOCK_DGRAM
SOCK_RAW
SOCK_SEQPACKET
SOCK_RDM
```

A `SOCK_STREAM` type provides sequenced, reliable, two-way connection based byte streams with an out-of-band data transmission mechanism. A `SOCK_DGRAM` socket supports datagrams (connectionless, unreliable messages of a fixed (typically small) maximum length). `SOCK_RAW` sockets provide access to internal network interfaces. The types `SOCK_RAW`, which is available only to the super-user, and `SOCK_SEQPACKET` and `SOCK_RDM`, which are planned, but not yet implemented, are not described here.

The *protocol* specifies a particular protocol to be used with the socket. Normally only a single protocol exists to support a particular socket type using a given address format. However, it is possible that many protocols may exist in which case a particular protocol must be specified in this manner. The protocol number to use is particular to the “communication domain” in which communication is to take place; see *services(3N)* and *protocols(3N)*.

Sockets of type `SOCK_STREAM` are full-duplex byte streams, similar to pipes. A stream socket must be in a *connected* state before any data may be sent or received on it. A connection to another socket is created with a *connect(2)* call. Once connected, data may be transferred using *read(2)* and *write(2)* calls or some variant of the *send(2)* and *recv(2)* calls. When a session has been completed a *close(2)* may be performed. Out-of-band data may also be transmitted as described in *send(2)* and received as described in *recv(2)*.

The communications protocols used to implement a `SOCK_STREAM` insure that data is not lost or duplicated. If a piece of data for which the peer protocol has buffer space cannot be successfully transmitted within a reasonable length of time, then the connection is considered broken and calls will indicate an error with `-1` returns and with `ETIMEDOUT` as the specific code in the global variable `errno`. The protocols optionally keep sockets “warm” by forcing transmissions roughly every minute in the absence of other activity. An error is then indicated if no response can be elicited on an otherwise idle connection for an extended period (e.g. 5 minutes). A `SIGPIPE` signal is raised if a process sends on a broken stream; this causes naive

processes, which do not handle the signal, to exit.

SOCK_DGRAM and SOCK_RAW sockets allow sending of datagrams to correspondents named in *send(2)* calls. It is also possible to receive datagrams at such a socket with *recv(2)*.

An *fcntl(2)* call can be used to specify a process group to receive a SIGURG signal when the out-of-band data arrives.

The operation of sockets is controlled by socket level *options*. These options are defined in the file *<sys/socket.h>* and explained below. *Setsockopt* and *getsockopt(2)* are used to set and get options, respectively.

SO_DEBUG	turn on recording of debugging information
SO_REUSEADDR	allow local address reuse
SO_KEEPALIVE	keep connections alive
SO_DONTROUTE	do not apply routing on outgoing messages
SO_LINGER	linger on close if data present
SO_DONTLINGER	do not linger on close

SO_DEBUG enables debugging in the underlying protocol modules. SO_REUSEADDR indicates the rules used in validating addresses supplied in a *bind(2)* call should allow reuse of local addresses. SO_KEEPALIVE enables the periodic transmission of messages on a connected socket. Should the connected party fail to respond to these messages, the connection is considered broken and processes using the socket are notified via a SIGPIPE signal. SO_DONTROUTE indicates that outgoing messages should bypass the standard routing facilities. Instead, messages are directed to the appropriate network interface according to the network portion of the destination address. SO_LINGER and SO_DONTLINGER control the actions taken when unsent messages are queued on socket and a *close(2)* is performed. If the socket promises reliable delivery of data and SO_LINGER is set, the system will block the process on the *close* attempt until it is able to transmit the data or until it decides it is unable to deliver the information (a timeout period, termed the linger interval, is specified in the *setsockopt* call when SO_LINGER is requested). If SO_DONTLINGER is specified and a *close* is issued, the system will process the close in a manner which allows the process to continue as quickly as possible.

RETURN VALUE

A -1 is returned if an error occurs, otherwise the return value is a descriptor referencing the socket.

ERRORS

The *socket* call fails if:

[EAFNOSUPPORT]	The specified address family is not supported in this version of the system.
[ESOCKTNOSUPPORT]	The specified socket type is not supported in this address family.
[EPROTONOSUPPORT]	The specified protocol is not supported.
[EMFILE]	The per-process descriptor table is full.
[ENOBUFS]	No buffer space is available. The socket cannot be created.

SEE ALSO

accept(2), *bind(2)*, *connect(2)*, *getsockname(2)*, *getsockopt(2)*, *ioctl(2)*, *listen(2)*, *recv(2)*, *select(2)*, *send(2)*, *shutdown(2)*, *socketpair(2)*
 "A 4.2BSD Interprocess Communication Primer".

BUGS

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

2

NAME

socketpair — create a pair of connected sockets

SYNOPSIS

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

socketpair(d, type, protocol, sv)
int d, type, protocol;
int sv[2];
```

DESCRIPTION

The *socketpair* call creates an unnamed pair of connected sockets in the specified domain *d*, of the specified *type*, and using the optionally specified *protocol*. The descriptors used in referencing the new sockets are returned in *sv*[0] and *sv*[1]. The two sockets are indistinguishable.

DIAGNOSTICS

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

ERRORS

The call succeeds unless:

- [EMFILE] Too many descriptors are in use by this process.
- [EAFNOSUPPORT] The specified address family is not supported on this machine.
- [EPROTONOSUPPORT] The specified protocol is not supported on this machine.
- [EOPNOSUPPORT] The specified protocol does not support creation of socket pairs.
- [EFAULT] The address *sv* does not specify a valid part of the process address space.

SEE ALSO

read(2), write(2), pipe(2)

BUGS

This call is currently implemented only for the UNIX domain.

NAME

stat, lstat, fstat — get file status

SYNOPSIS

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

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

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

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

DESCRIPTION

Stat obtains information about the file *path*. Read, write or execute permission of the named file is not required, but all directories listed in the path name leading to the file must be reachable.

Lstat is like *stat* except in the case where the named file is a symbolic link, in which case *lstat* returns information about the link, while *stat* returns information about the file the link references.

Fstat obtains the same information about an open file referenced by the argument descriptor, such as would be obtained by an *open* call.

Buf is a pointer to a *stat* structure into which information is placed concerning the file. The contents of the structure pointed to by *buf*

```
struct stat {
    dev_t    st_dev;    /* device inode resides on */
    ino_t    st_ino;    /* this inode's number */
    u_short  st_mode;   /* protection */
    short    st_nlink;  /* number 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 */
    time_t   st_atime;  /* file last access time */
    int      st_spare1;
    time_t   st_mtime;  /* file last modify time */
    int      st_spare2;
    time_t   st_ctime;  /* file last status change time */
    int      st_spare3;
    long     st_blksize; /* optimal blocksize for file system i/o ops */
    long     st_blocks;  /* actual number of blocks allocated */
    long     st_spare4[2];
};
```

st_atime Time when file data was last read or modified. Changed by the following system calls: *mknod(2)*, *utimes(2)*, *read(2)*, and *write(2)*. For reasons of efficiency, *st_atime* is not set when a directory is searched, although this would be more logical.

- st_mtime** Time when data was last modified. It is not set by changes of owner, group, link count, or mode. Changed by the following system calls: *mknod(2)*, *utimes(2)*, *write(2)*.
- st_ctime** Time when file status was last changed. It is set both both by writing and changing the i-node. Changed by the following system calls: *chmod(2)* *chown(2)*, *link(2)*, *mknod(2)*, *unlink(2)*, *utimes(2)*, *write(2)*.

The status information word *st_mode* has bits:

```
#define S_IFMT      0170000    /* type of file */
#define S_IFDIR    0040000    /* directory */
#define S_IFCHR    0020000    /* character special */
#define S_IFBLK    0060000    /* block special */
#define S_IFREG    0100000    /* regular */
#define S_IFLNK    0120000    /* symbolic link */
#define S_IFSOCK   0140000    /* socket */
#define S_ISUID    0004000    /* set user id on execution */
#define S_ISGID    0002000    /* set group id on execution */
#define S_ISVTX    0001000    /* save swapped text even after use */
#define S_IRREAD   0000400    /* read permission, owner */
#define S_IWWRITE  0000200    /* write permission, owner */
#define S_IXEXEC   0000100    /* execute/search permission, owner */
```

The mode bits 0000070 and 0000007 encode group and others permissions (see *chmod(2)*).

When *fd* is associated with a pipe, *fstat* reports an ordinary file with an i-node number, restricted permissions, and a not necessarily meaningful length.

RETURN VALUE

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

ERRORS

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

- [ENOTDIR] A component of the path prefix is not a directory.
- [EPERM] The pathname contains a character with the high-order bit set.
- [ENOENT] The pathname was too long.
- [ENOENT] The named file does not exist.
- [EACCESS] Search permission is denied for a component of the path prefix.
- [EFAULT] *Buf* or *name* points to an invalid address.

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

- [EBADF] *Fildes* is not a valid open file descriptor.
- [EFAULT] *Buf* points to an invalid address.
- [ELOOP] Too many symbolic links were encountered in translating the pathname.

CAVEAT

The fields in the *stat* structure currently marked *st_spare1*, *st_spare2*, and *st_spare3* are present in preparation for inode time stamps expanding to 64 bits. This, however, can break certain programs which depend on the time stamps being contiguous (in calls to *utimes(2)*).

SEE ALSO

chmod(2), *chown(2)*, *utimes(2)*

BUGS

Applying *fstat* to a socket returns a zero'd buffer.

The list of calls which modify the various fields should be carefully checked with reality.

NAME

swapon — 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)

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:

- | | |
|-----------|---|
| [EPERM] | Either <i>name1</i> or <i>name2</i> contains a character with the high-order bit set. |
| [ENOENT] | One of the pathnames specified was too long. |
| [ENOTDIR] | A component of the <i>name2</i> prefix is not a directory. |
| [EEXIST] | <i>Name2</i> already exists. |
| [EACCES] | A component of the <i>name2</i> path prefix denies search permission. |
| [EROFS] | The file <i>name2</i> would reside on a read-only file system. |
| [EFAULT] | <i>Name1</i> or <i>name2</i> points outside the process's allocated address space. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |

SEE ALSO

link(2), ln(1), unlink(2)

NAME

`sync` — update super-block

SYNOPSIS

`sync()`

DESCRIPTION

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

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

SEE ALSO

`fsync(2)`, `sync(8)`, `update(8)`

BUGS

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

NAME

syscall — indirect system call

SYNOPSIS

syscall(*number*, *arg*, ...) (VAX-11)

DESCRIPTION

Syscall performs the system call whose assembly language interface has the specified *number*, register arguments *r0* and *r1* and further arguments *arg*.

The *r0* value of the system call is returned.

DIAGNOSTICS

When the C-bit is set, *syscall* returns *-1* and sets the external variable *errno* (see *intro(2)*).

BUGS

There is no way to simulate system calls such as *pipe(2)*, which return values in register *r1*.

NAME

truncate — truncate a file to a specified length

SYNOPSIS

truncate(*path*, *length*)

char *path;

int length;

ftruncate(*fd*, *length*)

int fd, length;

DESCRIPTION

Truncate causes the file named by *path* or referenced by *fd* to be truncated to at most *length* bytes in size. If the file previously was larger than this size, the extra data is lost. With *ftruncate*, the file must be open for writing.

RETURN VALUES

A value of 0 is returned if the call succeeds. If the call fails a -1 is returned, and the global variable *errno* specifies the error.

ERRORS

Truncate succeeds unless:

- [EPERM] The pathname contains a character with the high-order bit set.
- [ENOENT] The pathname was too long.
- [ENOTDIR] A component of the path prefix of *path* is not a directory.
- [ENOENT] The named file does not exist.
- [EACCES] A component of the *path* prefix denies search permission.
- [EISDIR] The named file is a directory.
- [EROFS] The named file resides on a read-only file system.
- [ETXTBSY] The file is a pure procedure (shared text) file that is being executed.
- [EFAULT] *Name* points outside the process's allocated address space.

Ftruncate succeeds unless:

- [EBADF] The *fd* is not a valid descriptor.
- [EINVAL] The *fd* references a socket, not a file.

SEE ALSO

open(2)

BUGS

Partial blocks discarded as the result of truncation are not zero filled; this can result in holes in files which do not read as zero.

These calls should be generalized to allow ranges of bytes in a file to be discarded.

NAME

`umask` — set file creation mode mask

SYNOPSIS

```
oumask = umask(numask)  
int oumask, numask;
```

DESCRIPTION

Umask sets the process's file mode creation mask to *numask* and returns the previous value of the mask. The low-order 9 bits of *numask* are used whenever a file is created, clearing corresponding bits in the file mode (see *chmod(2)*). This clearing allows each user to restrict the default access to his files.

The value is initially 022 (write access for owner only). The mask is inherited by child processes.

RETURN VALUE

The previous value of the file mode mask is returned by the call.

SEE ALSO

chmod(2), *mknod(2)*, *open(2)*

NAME

`unlink` — remove directory entry

SYNOPSIS

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

DESCRIPTION

Unlink removes the entry for the file *path* from its directory. If this entry was the last link to the file, and no process has the file open, then all resources associated with the file are reclaimed. If, however, the file was open in any process, the actual resource reclamation is delayed until it is closed, even though the directory entry has disappeared.

RETURN VALUE

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

ERRORS

The *unlink* succeeds unless:

- | | |
|-----------|---|
| [EPERM] | The path contains a character with the high-order bit set. |
| [ENOENT] | The path name is too long. |
| [ENOTDIR] | A component of the path prefix is not a directory. |
| [ENOENT] | The named file does not exist. |
| [EACCES] | Search permission is denied for a component of the path prefix. |
| [EACCES] | Write permission is denied on the directory containing the link to be removed. |
| [EPERM] | The named file is a directory and the effective user ID of the process is not the super-user. |
| [EBUSY] | The entry to be unlinked is the mount point for a mounted file system. |
| [EROFS] | The named file resides on a read-only file system. |
| [EFAULT] | <i>Path</i> points outside the process's allocated address space. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |

SEE ALSO

`close(2)`, `link(2)`, `rmdir(2)`

NAME

utimes — set file times

SYNOPSIS

```
#include <sys/time.h>
utimes(file, tvp)
char *file;
struct timeval *tvp[2];
```

DESCRIPTION

The *utimes* call uses the “accessed” and “updated” times in that order from the *tvp* vector to set the corresponding recorded times for *file*.

The caller must be the owner of the file or the super-user. The “inode-changed” time of the file is set to the current time.

RETURN VALUE

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

ERRORS

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

- | | |
|-----------|--|
| [EPERM] | The pathname contained a character with the high-order bit set. |
| [ENOENT] | The pathname was too long. |
| [ENOENT] | The named file does not exist. |
| [ENOTDIR] | A component of the path prefix is not a directory. |
| [EACCES] | A component of the path prefix denies search permission. |
| [EPERM] | The process is not super-user and not the owner of the file. |
| [EACCES] | The effective user ID is not super-user and not the owner of the file and <i>times</i> is NULL and write access is denied. |
| [EROFS] | The file system containing the file is mounted read-only. |
| [EFAULT] | <i>Tvp</i> points outside the process's allocated address space. |
| [ELOOP] | Too many symbolic links were encountered in translating the pathname. |

SEE ALSO

stat(2)

NAME

vfork — spawn new process in a virtual memory efficient way

SYNOPSIS

```
pid = vfork()
int pid;
```

DESCRIPTION

Vfork can be used to create new processes without fully copying the address space of the old process, which is horrendously inefficient in a paged environment. It is useful when the purpose of *fork(2)* would have been to create a new system context for an *execve*. *Vfork* differs from *fork* in that the child borrows the parent's memory and thread of control until a call to *execve(2)* or an exit (either by a call to *exit(2)* or abnormally.) The parent process is suspended while the child is using its resources.

Vfork returns 0 in the child's context and (later) the pid of the child in the parent's context.

Vfork can normally be used just like *fork*. It does not work, however, to return while running in the child's context from the procedure which called *vfork* since the eventual return from *vfork* would then return to a no longer existent stack frame. Be careful, also, to call *_exit* rather than *exit* if you can't *execve*, since *exit* will flush and close standard I/O channels, and thereby mess up the parent 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 which are children in the middle of a *vfork* are never sent SIGTTOU or SIGTTIN signals; rather, output or *ioctl*s are allowed and input attempts result in an end-of-file indication.

NAME

`vhangup` — virtually “hangup” the current control terminal

SYNOPSIS

`vhangup()`

DESCRIPTION

Vhangup is used by the initialization process *init*(8) (among others) to arrange that users are given “clean” terminals at login, by revoking access of the previous users' processes to the terminal. To effect this, *vhangup* searches the system tables for references to the control terminal of the invoking process, revoking access permissions on each instance of the terminal which it finds. Further attempts to access the terminal by the affected processes will yield i/o errors (EBADF). Finally, a hangup signal (SIGHUP) is sent to the process group of the control terminal.

SEE ALSO

`init` (8)

BUGS

Access to the control terminal via `/dev/tty` is still possible.

This call should be replaced by an automatic mechanism which takes place on process exit.

NAME

wait, wait3 — wait for process to terminate

SYNOPSIS

```
#include <sys/wait.h>
pid = wait(status)
int pid;
union wait *status;

pid = wait(0)
int pid;

#include <sys/time.h>
#include <sys/resource.h>
pid = wait3(status, options, rusage)
int pid;
union wait *status;
int options;
struct rusage *rusage;
```

DESCRIPTION

Wait causes its caller to delay until a signal is received or one of its child processes terminates. If any child has died since the last *wait*, return is immediate, returning the process id and exit status of one of the terminated children. If there are no children, return is immediate with the value -1 returned.

On return from a successful *wait* call, *status* is nonzero, and the high byte of *status* contains the low byte of the argument to *exit* supplied by the child process; the low byte of *status* contains the termination status of the process. A more precise definition of the *status* word is given in *<sys/wait.h>*.

Wait3 provides an alternate interface for programs which must not block when collecting the status of child processes. The *status* parameter is defined as above. The *options* parameter is used to indicate the call should not block if there are no processes which wish to report status (WNOHANG), and/or that children of the current process which are stopped due to a SIGTTIN, SIGTTOU, SIGTSTP, or SIGSTOP signal should also have their status reported (WUNTRACED). If *rusage* is non-zero, a summary of the resources used by the terminated process and all its children is returned (this information is currently not available for stopped processes).

When the WNOHANG option is specified and no processes wish to report status, *wait3* returns a *pid* of 0. The WNOHANG and WUNTRACED options may be combined by *or*'ing the two values.

NOTES

See *sigvec(2)* for a list of termination statuses (signals); 0 status indicates normal termination. A special status (0177) is returned for a stopped process which has not terminated and can be restarted; see *ptrace(2)*. If the 0200 bit of the termination status is set, a core image of the process was produced by the system.

If the parent process terminates without waiting on its children, the initialization process (process ID = 1) inherits the children.

Wait and *wait3* are automatically restarted when a process receives a signal while awaiting termination of a child process.

RETURN VALUE

If *wait* returns due to a stopped or terminated child process, the process ID of the child is returned to the calling process. Otherwise, a value of -1 is returned and *errno* is set to

indicate the error.

Wait3 returns -1 if there are no children not previously waited for; 0 is returned if *WNOHANG* is specified and there are no stopped or exited children.

ERRORS

Wait will fail and return immediately if one or more of the following are true:

- [ECHILD] The calling process has no existing unwaited-for child processes.
- [EFAULT] The *status* or *rusage* arguments point to an illegal address.

SEE ALSO

exit(2)

NAME

write, writev — write on a file

SYNOPSIS

```
write(d, buf, nbytes)
int d;
char *buf;
int nbytes;

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

writev(d, iov, ioveclen)
int d;
struct iovec *iov;
int ioveclen;
```

DESCRIPTION

Write attempts to write *nbytes* of data to the object referenced by the descriptor *d* from the buffer pointed to by *buf*. *Writev* performs the same action, but gathers the output data from the *iovlen* buffers specified by the members of the *iovec* array: *iov*[0], *iov*[1], etc.

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

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

If the real user is not the super-user, then *write* clears the set-user-id bit on a file. This prevents penetration of system security by a user who “captures” a writable set-user-id file owned by the super-user.

RETURN VALUE

Upon successful completion the number of bytes actually written is returned. Otherwise a *-1* is returned and *errno* is set to indicate the error.

ERRORS

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

- [EBADF] *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.
- [EPIPE] An attempt is made to write to a socket of type SOCK_STREAM which is not connected to a peer socket.
- [EFBIG] An attempt was made to write a file that exceeds the process's file size limit or the maximum file size.
- [EFAULT] Part of *iov* or data to be written to the file points outside the process's allocated address space.

SEE ALSO

lseek(2), *open*(2), *pipe*(2)

NAME

intro — introduction to library functions

DESCRIPTION

This section describes functions that may be found in various libraries. The library functions are those other than the functions which directly invoke UNIX system primitives, described in section 2. This section has the libraries physically grouped together. This is a departure from older versions of the UNIX Programmer's Reference Manual, which did not group functions by library. The functions described in this section are grouped into various libraries:

(3) and (3S)

The straight "3" functions are the standard C library functions. The C library also includes all the functions described in section 2. The 3S functions comprise the standard I/O library. Together with the (3N), (3X), and (3C) routines, these functions constitute library *libc*, which is automatically loaded by the C compiler *cc*(1), the Pascal compiler *pc*(1), and the Fortran compiler *f77*(1). The link editor *ld*(1) searches this library under the '-lc' option. Declarations for some of these functions may be obtained from include files indicated on the appropriate pages.

(3F) The 3F functions are all functions callable from FORTRAN. These functions perform the same jobs as do the straight "3" functions.

(3M) These functions constitute the math library, *libm*. They are automatically loaded as needed by the Pascal compiler *pc*(1) and the Fortran compiler *f77*(1). The link editor searches this library under the '-lm' option. Declarations for these functions may be obtained from the include file *<math.h>*.

(3N) These functions constitute the internet network library,

(3S) These functions constitute the 'standard I/O package', see *intro*(3S). These functions are in the library *libc* already mentioned. Declarations for these functions may be obtained from the include file *<stdio.h>*.

(3X) Various specialized libraries have not been given distinctive captions. Files in which such libraries are found are named on appropriate pages.

(3C) Routines included for compatibility with other systems. In particular, a number of system call interfaces provided in previous releases of 4BSD have been included for source code compatibility. The manual page entry for each compatibility routine indicates the proper interface to use.

FILES

/lib/libc.a
/usr/lib/libm.a
/usr/lib/libc_p.a
/usr/lib/libm_p.a

SEE ALSO

intro(3C), *intro*(3S), *intro*(3F), *intro*(3M), *intro*(3N), *nm*(1), *ld*(1), *cc*(1), *f77*(1), *intro*(2)

DIAGNOSTICS

Functions in the math library (3M) may return conventional values when the function is undefined for the given arguments or when the value is not representable. In these cases the external variable *errno* (see *intro*(2)) is set to the value EDOM (domain error) or ERANGE (range error). The values of EDOM and ERANGE are defined in the include file *<math.h>*.

LIST OF FUNCTIONS

<i>Name</i>	<i>Appears on Page</i>	<i>Description</i>
abort	abort.3	generate a fault
abort	abort.3f	terminate abruptly with memory image

abs	abs.3	integer absolute value
access	access.3f	determine accessibility of a file
acos	sin.3m	trigonometric functions
alarm	alarm.3c	schedule signal after specified time
alarm	alarm.3f	execute a subroutine after a specified time
alloca	malloc.3	memory allocator
arc	plot.3x	graphics interface
asctime	ctime.3	convert date and time to ASCII
asin	sin.3m	trigonometric functions
assert	assert.3x	program verification
atan	sin.3m	trigonometric functions
atan2	sin.3m	trigonometric functions
atof	atof.3	convert ASCII to numbers
atoi	atof.3	convert ASCII to numbers
atol	atof.3	convert ASCII to numbers
bcmp	bstring.3	bit and byte string operations
bcopy	bstring.3	bit and byte string operations
bessel	bessel.3f	of two kinds for integer orders
bit	bit.3f	and, or, xor, not, rshift, lshift bitwise functions
bzero	bstring.3	bit and byte string operations
cabs	hypot.3m	Euclidean distance
calloc	malloc.3	memory allocator
ceil	floor.3m	absolute value, floor, ceiling functions
chdir	chdir.3f	change default directory
chmod	chmod.3f	change mode of a file
circle	plot.3x	graphics interface
clearerr	ferorr.3s	stream status inquiries
closedir	directory.3	directory operations
closelog	syslog.3	control system log
closepl	plot.3x	graphics interface
cont	plot.3x	graphics interface
cos	sin.3m	trigonometric functions
cosh	sinh.3m	hyperbolic functions
crypt	crypt.3	DES encryption
ctime	ctime.3	convert date and time to ASCII
ctime	time.3f	return system time
curses	curses.3x	screen functions with "optimal" cursor motion
dbmunit	dbm.3x	data base subroutines
delete	dbm.3x	data base subroutines
dfrac	fmin.3f	return extreme values
dflmax	fmin.3f	return extreme values
dflmax	range.3f	return extreme values
dflmin	fmin.3f	return extreme values
dflmin	range.3f	return extreme values
drand	rand.3f	return random values
dtime	etime.3f	return elapsed execution time
ecvt	ecvt.3	output conversion
edata	end.3	last locations in program
encrypt	crypt.3	DES encryption
end	end.3	last locations in program
endsent	getfsent.3x	get file system descriptor file entry
endgrent	getgrent.3	get group file entry

endhostent	gethostent.3n	get network host entry
endnetent	getnetent.3n	get network entry
endprotoent	getprotoent.3n	get protocol entry
endpwent	getpwent.3	get password file entry
endservent	getservent.3n	get service entry
environ	execl.3	execute a file
erase	plot.3x	graphics interface
etext	end.3	last locations in program
etime	etime.3f	return elapsed execution time
exec	execl.3	execute a file
exece	execl.3	execute a file
execl	execl.3	execute a file
execle	execl.3	execute a file
execlp	execl.3	execute a file
execst	execl.3	execute a file
execv	execl.3	execute a file
execvp	execl.3	execute a file
exit	exit.3	terminate a process after flushing any pending output
exit	exit.3f	terminate process with status
exp	exp.3m	exponential, logarithm, power, square root
fabs	floor.3m	absolute value, floor, ceiling functions
fclose	fclose.3s	close or flush a stream
fcvt	ecvt.3	output conversion
fdate	fdate.3f	return date and time in an ASCII string
feof	ferror.3s	stream status inquiries
ferror	ferror.3s	stream status inquiries
fetch	dbm.3x	data base subroutines
fflush	fclose.3s	close or flush a stream
ffrac	fmin.3f	return extreme values
ffs	bstring.3	bit and byte string operations
fgetc	getc.3f	get a character from a logical unit
fgetc	getc.3s	get character or word from stream
fgets	gets.3s	get a string from a stream
fileno	ferror.3s	stream status inquiries
firstkey	dbm.3x	data base subroutines
fimax	fmin.3f	return extreme values
fimax	range.3f	return extreme values
fimin	fmin.3f	return extreme values
fimin	range.3f	return extreme values
floor	floor.3m	absolute value, floor, ceiling functions
flush	flush.3f	flush output to a logical unit
fopen	fopen.3s	open a stream
fork	fork.3f	create a copy of this process
fpecont	trpfpe.3f	trap and repair floating point faults
fprintf	printf.3s	formatted output conversion
fputc	putc.3f	write a character to a fortran logical unit
fputc	putc.3s	put character or word on a stream
fputs	puts.3s	put a string on a stream
fread	fread.3s	buffered binary input/output
free	malloc.3	memory allocator
frexp	frexp.3	split into mantissa and exponent
fscanf	scanf.3s	formatted input conversion

fseek	fseek.3f	reposition a file on a logical unit
fseek	fseek.3s	reposition a stream
fstat	stat.3f	get file status
ftell	fseek.3f	reposition a file on a logical unit
ftell	fseek.3s	reposition a stream
ftime	time.3c	get date and time
fwrite	fread.3s	buffered binary input/output
gamma	gamma.3m	log gamma function
gcvt	ecvt.3	output conversion
gerror	perror.3f	get system error messages
getarg	getarg.3f	return command line arguments
getc	getc.3f	get a character from a logical unit
getc	getc.3s	get character or word from stream
getchar	getc.3s	get character or word from stream
getcwd	getcwd.3f	get pathname of current working directory
getdiskbyname	getdisk.3x	get disk description by its name
getenv	getenv.3	value for environment name
getenv	getenv.3f	get value of environment variables
getfsent	getfsent.3x	get file system descriptor file entry
getfsfile	getfsent.3x	get file system descriptor file entry
getfsspec	getfsent.3x	get file system descriptor file entry
getfstype	getfsent.3x	get file system descriptor file entry
getgid	getuid.3f	get user or group ID of the caller
getgrent	getgrent.3	get group file entry
getgrgid	getgrent.3	get group file entry
getgrnam	getgrent.3	get group file entry
gethostbyaddr	gethostent.3n	get network host entry
gethostbyname	gethostent.3n	get network host entry
gethostent	gethostent.3n	get network host entry
getlog	getlog.3f	get user's login name
getlogin	getlogin.3	get login name
getnetbyaddr	getnetent.3n	get network entry
getnetbyname	getnetent.3n	get network entry
getnetent	getnetent.3n	get network entry
getpass	getpass.3	read a password
getpid	getpid.3f	get process id
getprotobyname	getprotoent.3n	get protocol entry
getprotobyname	getprotoent.3n	get protocol entry
getprotoent	getprotoent.3n	get protocol entry
getpw	getpw.3	get name from uid
getpwent	getpwent.3	get password file entry
getpwnam	getpwent.3	get password file entry
getpwuid	getpwent.3	get password file entry
gets	gets.3s	get a string from a stream
getservbyname	getservent.3n	get service entry
getservbyname	getservent.3n	get service entry
getservent	getservent.3n	get service entry
getuid	getuid.3f	get user or group ID of the caller
getw	getc.3s	get character or word from stream
getwd	getwd.3	get current working directory pathname
gmtime	ctime.3	convert date and time to ASCII
gmtime	time.3f	return system time

gtty	stty.3c	set and get terminal state (defunct)
hostnm	hostnm.3f	get name of current host
htonl	byteorder.3n	convert values between host and network byte order
htons	byteorder.3n	convert values between host and network byte order
hypot	hypot.3m	Euclidean distance
iargc	getarg.3f	return command line arguments
idate	idate.3f	return date or time in numerical form
ierrno	perror.3f	get system error messages
index	index.3f	tell about character objects
index	string.3	string operations
inet_addr	inet.3n	Internet address manipulation routines
inet_lnaof	inet.3n	Internet address manipulation routines
inet_makeaddr	inet.3n	Internet address manipulation routines
inet_netof	inet.3n	Internet address manipulation routines
inet_network	inet.3n	Internet address manipulation routines
initgroups	initgroups.3x	initialize group access list
initstate	random.3	better random number generator
inmax	flmin.3f	return extreme values
inmax	range.3f	return extreme values
insque	insque.3	insert/remove element from a queue
ioinit	ioinit.3f	change f77 I/O initialization
irand	rand.3f	return random values
isalnum	ctype.3	character classification macros
isalpha	ctype.3	character classification macros
isascii	ctype.3	character classification macros
isatty	ttynam.3f	find name of a terminal port
isatty	ttynam.3	find name of a terminal
iscntrl	ctype.3	character classification macros
isdigit	ctype.3	character classification macros
islower	ctype.3	character classification macros
isprint	ctype.3	character classification macros
ispunct	ctype.3	character classification macros
isspace	ctype.3	character classification macros
isupper	ctype.3	character classification macros
itime	idate.3f	return date or time in numerical form
j0	j0.3m	bessel functions
j1	j0.3m	bessel functions
jn	j0.3m	bessel functions
kill	kill.3f	send a signal to a process
label	plot.3x	graphics interface
ldexp	frexp.3	split into mantissa and exponent
len	index.3f	tell about character objects
lib2648	lib2648.3x	subroutines for the HP 2648 graphics terminal
line	plot.3x	graphics interface
linemod	plot.3x	graphics interface
link	link.3f	make a link to an existing file
lnblk	index.3f	tell about character objects
loc	loc.3f	return the address of an object
localtime	ctime.3	convert date and time to ASCII
log	exp.3m	exponential, logarithm, power, square root
log10	exp.3m	exponential, logarithm, power, square root
long	long.3f	integer object conversion

longjmp	setjmp.3	non-local goto
lstat	stat.3f	get file status
ltime	time.3f	return system time
malloc	malloc.3	memory allocator
mktemp	mktemp.3	make a unique file name
modf	frexp.3	split into mantissa and exponent
moncontrol	monitor.3	prepare execution profile
monitor	monitor.3	prepare execution profile
monstartup	monitor.3	prepare execution profile
move	plot.3x	graphics interface
nextkey	dbm.3x	data base subroutines
nice	nice.3c	set program priority
nlist	nlist.3	get entries from name list
ntohl	byteorder.3n	convert values between host and network byte order
ntohs	byteorder.3n	convert values between host and network byte order
opendir	directory.3	directory operations
openlog	syslog.3	control system log
openpl	plot.3x	graphics interface
pause	pause.3c	stop until signal
pclose	popen.3	initiate I/O to/from a process
perror	perror.3	system error messages
perror	perror.3f	get system error messages
point	plot.3x	graphics interface
popen	popen.3	initiate I/O to/from a process
pow	exp.3m	exponential, logarithm, power, square root
printf	printf.3s	formatted output conversion
psignal	psignal.3	system signal messages
putc	putc.3f	write a character to a fortran logical unit
putc	putc.3s	put character or word on a stream
putchar	putc.3s	put character or word on a stream
puts	puts.3s	put a string on a stream
putw	putc.3s	put character or word on a stream
qsort	qsort.3	quicker sort
qsort	qsort.3f	quick sort
rand	rand.3c	random number generator
rand	rand.3f	return random values
random	random.3	better random number generator
rcmd	rcmd.3x	routines for returning a stream to a remote command
re_comp	regex.3	regular expression handler
re_exec	regex.3	regular expression handler
readdir	directory.3	directory operations
realloc	malloc.3	memory allocator
remque	insque.3	insert/remove element from a queue
rename	rename.3f	rename a file
rewind	fseek.3s	reposition a stream
rewinddir	directory.3	directory operations
rexec	rexec.3x	return stream to a remote command
rindex	index.3f	tell about character objects
rindex	string.3	string operations
resvport	rcmd.3x	routines for returning a stream to a remote command
ruserok	rcmd.3x	routines for returning a stream to a remote command
scandir	scandir.3	scan a directory

scanf	scanf.3s	formatted input conversion
seekdir	directory.3	directory operations
setbuf	setbuf.3s	assign buffering to a stream
setbuffer	setbuf.3s	assign buffering to a stream
setegid	setuid.3	set user and group ID
seteuid	setuid.3	set user and group ID
setfsent	getfsent.3x	get file system descriptor file entry
setgid	setuid.3	set user and group ID
setgrent	getgrent.3	get group file entry
sethostent	gethostent.3n	get network host entry
setjmp	setjmp.3	non-local goto
setkey	crypt.3	DES encryption
setlinebuf	setbuf.3s	assign buffering to a stream
setnetent	getnetent.3n	get network entry
setprotoent	getprotoent.3n	get protocol entry
setpwent	getpwent.3	get password file entry
setrgid	setuid.3	set user and group ID
setruid	setuid.3	set user and group ID
setservent	getservent.3n	get service entry
setstate	random.3	better random number generator
setuid	setuid.3	set user and group ID
short	long.3f	integer object conversion
signal	signal.3	simplified software signal facilities
signal	signal.3f	change the action for a signal
sin	sin.3m	trigonometric functions
sinh	sinh.3m	hyperbolic functions
sleep	sleep.3	suspend execution for interval
sleep	sleep.3f	suspend execution for an interval
space	plot.3x	graphics interface
sprintf	printf.3s	formatted output conversion
sqrt	exp.3m	exponential, logarithm, power, square root
srand	rand.3c	random number generator
srandom	random.3	better random number generator
sscanf	scanf.3s	formatted input conversion
stat	stat.3f	get file status
stdio	intro.3s	standard buffered input/output package
store	dbm.3x	data base subroutines
strcat	string.3	string operations
strcmp	string.3	string operations
strcpy	string.3	string operations
strlen	string.3	string operations
strncat	string.3	string operations
strncmp	string.3	string operations
strncpy	string.3	string operations
stty	stty.3c	set and get terminal state (defunct)
swab	swab.3	swap bytes
sys_errlist	perror.3	system error messages
sys_nerr	perror.3	system error messages
sys_siglist	psignal.3	system signal messages
syslog	syslog.3	control system log
system	system.3	issue a shell command
system	system.3f	execute a UNIX command

tan	sin.3m	trigonometric functions
tanh	sinh.3m	hyperbolic functions
tclose	topen.3f	f77 tape I/O
telldir	directory.3	directory operations
tgetent	termcap.3x	terminal independent operation routines
tgetflag	termcap.3x	terminal independent operation routines
tgetnum	termcap.3x	terminal independent operation routines
tgetstr	termcap.3x	terminal independent operation routines
tgoto	termcap.3x	terminal independent operation routines
time	time.3c	get date and time
time	time.3f	return system time
times	times.3c	get process times
timezone	ctime.3	convert date and time to ASCII
topen	topen.3f	f77 tape I/O
tputs	termcap.3x	terminal independent operation routines
traper	traper.3f	trap arithmetic errors
trapov	trapov.3f	trap and repair floating point overflow
tread	topen.3f	f77 tape I/O
trewin	topen.3f	f77 tape I/O
trpfpe	trpfpe.3f	trap and repair floating point faults
tskipf	topen.3f	f77 tape I/O
tstate	topen.3f	f77 tape I/O
ttynam	ttynam.3f	find name of a terminal port
ttynam	ttynam.3	find name of a terminal
ttyslot	ttynam.3	find name of a terminal
twrite	topen.3f	f77 tape I/O
ungetc	ungetc.3s	push character back into input stream
unlink	unlink.3f	remove a directory entry
utime	utime.3c	set file times
valloc	valloc.3	aligned memory allocator
varargs	varargs.3	variable argument list
vlimit	vlimit.3c	control maximum system resource consumption
vtimes	vtimes.3c	get information about resource utilization
wait	wait.3f	wait for a process to terminate
y0	j0.3m	bessel functions
y1	j0.3m	bessel functions
yn	j0.3m	bessel functions

NAME

abort — generate a fault

DESCRIPTION

Abort executes an instruction which is illegal in user mode. This causes a signal that normally terminates the process with a core dump, which may be used for debugging.

SEE ALSO

adb(1), sigvec(2), exit(2)

DIAGNOSTICS

Usually 'IOT trap — core dumped' from the shell.

BUGS

The abort() function does not flush standard I/O buffers. Use *flush*(3S).

NAME

abs — integer absolute value

SYNOPSIS

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

DESCRIPTION

Abs returns the absolute value of its integer operand.

SEE ALSO

floor(3M) for *fabs*

BUGS

Applying the *abs* function to the most negative integer generates a result which is the most negative integer. That is,

```
abs(0x80000000)
```

returns 0x80000000 as a result.

NAME

atof, *atoi*, *atol* — convert ASCII to numbers

SYNOPSIS

double *atof*(*nptr*)

char **nptr*;

atol(*nptr*)

char **nptr*;

long *atol*(*nptr*)

char **nptr*;

DESCRIPTION

These functions convert a string pointed to by *nptr* to floating, integer, and long integer representation respectively. The first unrecognized character ends the string.

Atof recognizes an optional string of spaces, then an optional sign, then a string of digits optionally containing a decimal point, then an optional 'e' or 'E' followed by an optionally signed integer.

Atoi and *atol* recognize an optional string of spaces, then an optional sign, then a string of digits.

SEE ALSO

scanf(3S)

BUGS

There are no provisions for overflow.

NAME

bcopy, *bcmp*, *bzero*, *ffs* — bit and byte string operations

SYNOPSIS

***bcopy*(*b1*, *b2*, *length*)**

char **b1*, **b2*;

int *length*;

***bcmp*(*b1*, *b2*, *length*)**

char **b1*, **b2*;

int *length*;

***bzero*(*b*, *length*)**

char **b*;

int *length*;

***ffs*(*i*)**

int *i*;

DESCRIPTION

The functions *bcopy*, *bcmp*, and *bzero* operate on variable length strings of bytes. They do not check for null bytes as the routines in *string(3)* do.

Bcopy copies *length* bytes from string *b1* to the string *b2*.

Bcmp compares byte string *b1* against byte string *b2*, returning zero if they are identical, non-zero otherwise. Both strings are assumed to be *length* bytes long.

Bzero places *length* 0 bytes in the string *b1*.

Ffs find the first bit set in the argument passed it and returns the index of that bit. Bits are numbered starting at 1. A return value of -1 indicates the value passed is zero.

BUGS

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

NAME

`crypt`, `setkey`, `encrypt` — DES encryption

SYNOPSIS

```
char *crypt(key, salt)
char *key, *salt;

setkey(key)
char *key;

encrypt(block, edflag)
char *block;
```

DESCRIPTION

Crypt is the password encryption routine. It is based on the NBS Data Encryption Standard, with variations intended (among other things) to frustrate use of hardware implementations of the DES for key search.

The first argument to *crypt* is normally a user's typed password. The second is a 2-character string chosen from the set [a-zA-Z0-9./]. The *salt* string is used to perturb the DES algorithm in one of 4096 different ways, after which the password is used as the key to encrypt repeatedly a constant string. The returned value points to the encrypted password, in the same alphabet as the salt. The first two characters are the salt itself.

The other entries provide (rather primitive) access to the actual DES algorithm. The argument of *setkey* is a character array of length 64 containing only the characters with numerical value 0 and 1. If this string is divided into groups of 8, the low-order bit in each group is ignored, leading to a 56-bit key which is set into the machine.

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

SEE ALSO

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

BUGS

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

NAME

time, localtime, gmtime, asctime, timezone — convert date and time to ASCII

SYNOPSIS

```
char *ctime(clock)
long *clock;
#include <sys/time.h>
struct tm *localtime(clock)
long *clock;
struct tm *gmtime(clock)
long *clock;
char *asctime(tm)
struct tm *tm;
char *timezone(zone, dst)
```

DESCRIPTION

Ctime converts a time pointed to by *clock* such as returned by *time(2)* into ASCII and returns a pointer to a 26-character string in the following form. All the fields have constant width.

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

Localtime and *gmtime* return pointers to structures containing the broken-down time. *Localtime* corrects for the time zone and possible daylight savings time; *gmtime* converts directly to GMT, which is the time UNIX uses. *Asctime* converts a broken-down time to ASCII and returns a pointer to a 26-character string.

The structure declaration from the include file is:

```
struct tm {
    int    tm_sec;
    int    tm_min;
    int    tm_hour;
    int    tm_mday;
    int    tm_mon;
    int    tm_year;
    int    tm_wday;
    int    tm_yday;
    int    tm_isdst;
};
```

These quantities give the time on a 24-hour clock, day of month (1-31), month of year (0-11), day of week (Sunday = 0), year — 1900, day of year (0-365), and a flag that is nonzero if daylight saving time is in effect.

When local time is called for, the program consults the system to determine the time zone and whether the U.S.A., Australian, Eastern European, Middle European, or Western European daylight saving time adjustment is appropriate. The program knows about various peculiarities in time conversion over the past 10-20 years; if necessary, this understanding can be extended.

Timezone returns the name of the time zone associated with its first argument, which is measured in minutes westward from Greenwich. If the second argument is 0, the standard name is used, otherwise the Daylight Saving version. If the required name does not appear in a table built into the routine, the difference from GMT is produced; e.g. in Afghanistan *timezone(-60*4+30)*, 0) is appropriate because it is 4:30 ahead of GMT and the string **GMT+4:30** is produced.

SEE ALSO

gettimeofday(2), time(3)

BUGS

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

NAME

isalpha, *isupper*, *islower*, *isdigit*, *isalnum*, *isspace*, *ispunct*, *isprint*, *isctrnl*, *isascii* — character classification macros

SYNOPSIS

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

```
isalpha(c)
```

```
...
```

DESCRIPTION

These macros classify ASCII-coded integer values by table lookup. Each is a predicate returning nonzero for true, zero for false. *isascii* is defined on all integer values; the rest are defined only where *isascii* is true and on the single non-ASCII value EOF (see *stdio(3S)*).

<i>isalpha</i>	<i>c</i> is a letter
<i>isupper</i>	<i>c</i> is an upper case letter
<i>islower</i>	<i>c</i> is a lower case letter
<i>isdigit</i>	<i>c</i> is a digit
<i>isalnum</i>	<i>c</i> is an alphanumeric character
<i>isspace</i>	<i>c</i> is a space, tab, carriage return, newline, or formfeed
<i>ispunct</i>	<i>c</i> is a punctuation character (neither control nor alphanumeric)
<i>isprint</i>	<i>c</i> is a printing character, code 040(8) (space) through 0176 (tilde)
<i>isctrnl</i>	<i>c</i> is a delete character (0177) or ordinary control character (less than 040).
<i>isascii</i>	<i>c</i> is an ASCII character, code less than 0200

SEE ALSO

ascii(7)

NAME

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

SYNOPSIS

```
#include <sys/dir.h>
DIR *opendir(filename)
char *filename;
struct direct *readdir(dirp)
DIR *dirp;
long telldir(dirp)
DIR *dirp;
seekdir(dirp, loc)
DIR *dirp;
long loc;
rewinddir(dirp)
DIR *dirp;
closedir(dirp)
DIR *dirp;
```

DESCRIPTION

Opendir opens the directory named by *filename* and associates a *directory stream* with it. *Opendir* returns a pointer to be used to identify the *directory stream* in subsequent operations. The pointer NULL is returned if *filename* cannot be accessed, or if it cannot *malloc(3)* enough memory to hold the whole thing.

Readdir returns a pointer to the next directory entry. It returns NULL upon reaching the end of the directory or detecting an invalid *seekdir* operation.

Telldir returns the current location associated with the named *directory stream*.

Seekdir sets the position of the next *readdir* operation on the *directory stream*. The new position reverts to the one associated with the *directory stream* when the *telldir* operation was performed. Values returned by *telldir* are good only for the lifetime of the DIR pointer from which they are derived. If the directory is closed and then reopened, the *telldir* value may be invalidated due to undetected directory compaction. It is safe to use a previous *telldir* value immediately after a call to *opendir* and before any calls to *readdir*.

Rewinddir resets the position of the named *directory stream* to the beginning of the directory.

Closedir closes the named *directory stream* and frees the structure associated with the DIR pointer.

Sample code which searches a directory for entry "name" is:

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

SEE ALSO

open(2), close(2), read(2), lseek(2), dir(5)

NAME

ecvt, *fcvt*, *gcvt* — output conversion

SYNOPSIS

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

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

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

DESCRIPTION

Ecvt converts the *value* to a null-terminated string of *ndigit* ASCII digits and returns a pointer thereto. The position of the decimal point relative to the beginning of the string is stored indirectly through *decpt* (negative means to the left of the returned digits). If the sign of the result is negative, the word pointed to by *sign* is non-zero, otherwise it is zero. The low-order digit is rounded.

Fcvt is identical to *ecvt*, except that the correct digit has been rounded for Fortran F-format output of the number of digits specified by *ndigits*.

Gcvt converts the *value* to a null-terminated ASCII string in *buf* and returns a pointer to *buf*. It attempts to produce *ndigit* significant digits in Fortran F format if possible, otherwise E format, ready for printing. Trailing zeros may be suppressed.

SEE ALSO

`printf(3)`

BUGS

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

NAME

end, *etext*, *edata* — last locations in program

SYNOPSIS

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

DESCRIPTION

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

When execution begins, the program break coincides with *end*, but it is reset by the routines *brk(2)*, *malloc(3)*, standard input/output (*stdio(3)*), the profile (**-p**) option of *cc(1)*, etc. The current value of the program break is reliably returned by 'sbrk(0)', see *brk(2)*.

SEE ALSO

brk(2), *malloc(3)*

NAME

`execl`, `execv`, `execle`, `execlp`, `execvp`, `exec`, `exece`, `exec`, `environ` — execute a file

SYNOPSIS

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

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

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

exec(name, argv, envp)
char *name, *argv[], *envp[];

extern char **environ;
```

DESCRIPTION

These routines provide various interfaces to the *execve* system call. Refer to *execve(2)* for a description of their properties; only brief descriptions are provided here.

Exec in all its forms overlays the calling process with the named file, then transfers to the entry point of the core image of the file. There can be no return from a successful *exec*; the calling core image is lost.

The *name* argument is a pointer to the name of the file to be executed. The pointers *arg[0]*, *arg[1]* ... address null-terminated strings. Conventionally *arg[0]* is the name of the file.

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

The *execv* version is useful when the number of arguments is unknown in advance; the arguments to *execv* are the name of the file to be executed and a vector of strings containing the arguments. The last argument string must be followed by a 0 pointer.

The *exec* version is used when the executed file is to be manipulated with *ptrace(2)*. The program is forced to single step a single instruction giving the parent an opportunity to manipulate its state. On the VAX-11 this is done by setting the trace bit in the process status longword.

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

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

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

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

Envp is a pointer to an array of strings that constitute the *environment* of the process. Each string consists of a name, an "=", and a null-terminated value. The array of pointers is terminated by a null pointer. The shell *sh(1)* passes an environment entry for each global shell variable defined when the program is called. See *environ(7)* for some conventionally used names. The C run-time start-off routine places a copy of *envp* in the global cell *environ*, which is used by *execv* and *execl* to pass the environment to any subprograms executed by the current program.

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

FILES

/bin/sh shell, invoked if command file found by *execlp* or *execvp*

SEE ALSO

execve(2), *fork(2)*, *environ(7)*, *csh(1)*

DIAGNOSTICS

If the file cannot be found, if it is not executable, if it does not start with a valid magic number (see *a.out(5)*), if maximum memory is exceeded, or if the arguments require too much space, a return constitutes the diagnostic; the return value is -1 . Even for the super-user, at least one of the execute-permission bits must be set for a file to be executed.

BUGS

If *execvp* is called to execute a file that turns out to be a shell command file, and if it is impossible to execute the shell, the values of *argv[0]* and *argv[-1]* will be modified before return.

NAME

`exit` — terminate a process after flushing any pending output

SYNOPSIS

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

DESCRIPTION

Exit terminates a process after calling the Standard I/O library function `_cleanup` to flush any buffered output. *Exit* never returns.

SEE ALSO

`exit(2)`, `intro(3S)`

NAME

frexp, *ldexp*, *modf* — split into mantissa and exponent

SYNOPSIS

```
double frexp(value, eptr)
double value;
int *eptr;

double ldexp(value, exp)
double value;

double modf(value, iptr)
double value, *iptr;
```

DESCRIPTION

Frexp returns the mantissa of a double *value* as a double quantity, *x*, of magnitude less than 1 and stores an integer *n* such that $value = x \cdot 2^n$ indirectly through *eptr*.

Ldexp returns the quantity $value \cdot 2^{exp}$.

Modf returns the positive fractional part of *value* and stores the integer part indirectly through *iptr*.

NAME

`getenv` — value for environment name

SYNOPSIS

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

DESCRIPTION

Getenv searches the environment list (see *environ(7)*) for a string of the form *name=value* and returns a pointer to the string *value* if such a string is present, otherwise *getenv* returns the value 0 (NULL).

SEE ALSO

environ(7), *execve(2)*

NAME

getgrent, *getgrgid*, *getgrnam*, *setgrent*, *endgrent* — get group file entry

SYNOPSIS

```
#include <grp.h>

struct group *getgrent()
struct group *getgrgid(gid)
int gid;
struct group *getgrnam(name)
char *name;
setgrent()
endgrent()
```

DESCRIPTION

Getgrent, *getgrgid* and *getgrnam* each return pointers to an object with the following structure containing the broken-out fields of a line in the group file.

```
struct group { /* see getgrent(3) */
    char    *gr_name;
    char    *gr_passwd;
    int     gr_gid;
    char    **gr_mem;
};

struct group *getgrent(), *getgrgid(), *getgrnam();
```

The members of this structure are:

gr_name The name of the group.
gr_passwd The encrypted password of the group.
gr_gid The numerical group-ID.
gr_mem Null-terminated vector of pointers to the individual member names.

Getgrent simply reads the next line while *getgrgid* and *getgrnam* search until a matching *gid* or *name* is found (or until EOF is encountered). Each routine picks up where the others leave off so successive calls may be used to search the entire file.

A call to *setgrent* has the effect of rewinding the group file to allow repeated searches. *Endgrent* may be called to close the group file when processing is complete.

FILES

/etc/group

SEE ALSO

getlogin(3), *getpwent(3)*, *group(5)*

DIAGNOSTICS

A null pointer (0) is returned on EOF or error.

BUGS

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

NAME

`getlogin` — get login name

SYNOPSIS

`char *getlogin()`

DESCRIPTION

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

If *getlogin* is called within a process that is not attached to a typewriter, it returns NULL. The correct procedure for determining the login name is to first call *getlogin* and if it fails, to call *getpw(getuid())*.

FILES

/etc/utmp

SEE ALSO

getpwent(3), *getgrent(3)*, *utmp(5)*, *getpw(3)*

DIAGNOSTICS

Returns NULL (0) if name not found.

BUGS

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

NAME

`getpass` — read a password

SYNOPSIS

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

DESCRIPTION

Getpass reads a password from the file */dev/tty*, or if that cannot be opened, from the standard input, after prompting with the null-terminated string *prompt* and disabling echoing. A pointer is returned to a null-terminated string of at most 8 characters.

FILES

/dev/tty

SEE ALSO

`crypt(3)`

BUGS

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

NAME

`getpwent`, `getpwuid`, `getpwnam`, `setpwent`, `endpwent` — get password file entry

SYNOPSIS

```
#include <pwd.h>

struct passwd *getpwent()
struct passwd *getpwuid(uid)
int uid;
struct passwd *getpwnam(name)
char *name;
int setpwent()
int endpwent()
```

DESCRIPTION

`Getpwent`, `getpwuid` and `getpwnam` each return a pointer to an object with the following structure containing the broken-out fields of a line in the password file.

```
struct passwd { /* see getpwent(3) */
    char *pw_name;
    char *pw_passwd;
    int pw_uid;
    int pw_gid;
    int pw_quota;
    char *pw_comment;
    char *pw_gecos;
    char *pw_dir;
    char *pw_shell;
};
```

```
struct passwd *getpwent(), *getpwuid(), *getpwnam();
```

The fields `pw_quota` and `pw_comment` are unused; the others have meanings described in `passwd(5)`.

`Getpwent` reads the next line (opening the file if necessary); `setpwent` rewinds the file; `endpwent` closes it.

`Getpwuid` and `getpwnam` search from the beginning until a matching `uid` or `name` is found (or until EOF is encountered).

FILES

`/etc/passwd`

SEE ALSO

`getlogin(3)`, `getgrent(3)`, `passwd(5)`

DIAGNOSTICS

Null pointer (0) returned on EOF or error.

BUGS

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

NAME

`getwd` — get current working directory pathname

SYNOPSIS

```
char *getwd(pathname)
char *pathname;
```

DESCRIPTION

Getwd copies the absolute pathname of the current working directory to *pathname* and returns a pointer to the result.

LIMITATIONS

Maximum pathname length is `MAXPATHLEN` characters (1024).

DIAGNOSTICS

Getwd returns zero and places a message in *pathname* if an error occurs.

BUGS

Getwd may fail to return to the current directory if an error occurs.

NAME

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.

SEE ALSO

"VAX Architecture Handbook", pp. 228-235.

NAME

`malloc`, `free`, `realloc`, `calloc`, `alloca` — memory allocator

SYNOPSIS

```
char *malloc(size)
unsigned size;

free(ptr)
char *ptr;

char *realloc(ptr, size)
char *ptr;
unsigned size;

char *calloc(nelem, elsize)
unsigned nelem, elsize;

char *alloca(size)
int size;
```

DESCRIPTION

Malloc and *free* provide a simple general-purpose memory allocation package. *Malloc* returns a pointer to a block of at least *size* bytes beginning on a word boundary.

The argument to *free* is a pointer to a block previously allocated by *malloc*; this space is made available for further allocation, but its contents are left undisturbed.

Needless to say, grave disorder will result if the space assigned by *malloc* is overrun or if some random number is handed to *free*.

Malloc maintains multiple lists of free blocks according to size, allocating space from the appropriate list. It calls *sbrk* (see *brk(2)*) to get more memory from the system when there is no suitable space already free.

Realloc changes the size of the block pointed to by *ptr* to *size* bytes and returns a pointer to the (possibly moved) block. The contents will be unchanged up to the lesser of the new and old sizes.

In order to be compatible with older versions, *realloc* also works if *ptr* points to a block freed since the last call of *malloc*, *realloc* or *calloc*; sequences of *free*, *malloc* and *realloc* were previously used to attempt storage compaction. This procedure is no longer recommended.

Calloc allocates space for an array of *nelem* elements of size *elsize*. The space is initialized to zeros.

Alloca allocates *size* bytes of space in the stack frame of the caller. This temporary space is automatically freed on return.

Each of the allocation routines returns a pointer to space suitably aligned (after possible pointer coercion) for storage of any type of object.

DIAGNOSTICS

Malloc, *realloc* and *calloc* return a null pointer (0) if there is no available memory or if the arena has been detectably corrupted by storing outside the bounds of a block. *Malloc* may be recompiled to check the arena very stringently on every transaction; those sites with a source code license may check the source code to see how this can be done.

BUGS

When *realloc* returns 0, the block pointed to by *ptr* may be destroyed.

Alloca is machine dependent; it's use is discouraged.

NAME

mktemp — make a unique file name

SYNOPSIS

char *mktemp(template)

char *template;

DESCRIPTION

Mktemp replaces *template* by a unique file name, and returns the address of the template. The template should look like a file name with six trailing X's, which will be replaced with the current process id and a unique letter.

SEE ALSO

getpid(2)

NAME

monitor, monstartup, moncontrol — prepare execution profile

SYNOPSIS

```
monitor(lowpc, highpc, buffer, bufsize, nfunc)
int (*lowpc)(), (*highpc)();
short buffer[];
```

```
monstartup(lowpc, highpc)
int (*lowpc)(), (*highpc)();
```

```
moncontrol(mode)
```

DESCRIPTION

There are two different forms of monitoring available: An executable program created by:

```
cc -p . . .
```

automatically includes calls for the *profil*(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 *gprofil*(1) monitor.

Monstartup is a high level interface to *profil*(2). *Lowpc* and *highpc* specify the address range that is to be sampled; the lowest address sampled is that of *lowpc* and the highest is just below *highpc*. *Monstartup* allocates space using *sbrk*(2) and passes it to *monitor* (see below) to record a histogram of periodically sampled values of the program counter, and of counts of calls of certain functions, in the buffer. Only calls of functions compiled with the profiling option *-p* of *cc*(1) are recorded.

To profile the entire program, it is sufficient to use

```
extern etext();
. . .
monstartup((int) 2, etext);
```

Etect lies just above all the program text, see *end*(3).

To stop execution monitoring and write the results on the file *mon.out*, use

```
monitor(0);
```

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

Moncontrol is used to selectively control profiling within a program. This works with either *profil*(1) or *gprofil*(1) type profiling. When the program starts, profiling begins. To stop the collection of histogram ticks and call counts use *moncontrol*(0); to resume the collection of histogram ticks and call counts use *moncontrol*(1). This allows the cost of particular operations to be measured. Note that an output file will be produced upon program exit irregardless of the state of *moncontrol*.

Monitor is a low level interface to *profil*(2). *Lowpc* and *highpc* are the addresses of two functions; *buffer* is the address of a (user supplied) array of *bufsize* short integers. At most *nfunc* call counts can be kept. For the results to be significant, especially where there are small, heavily used routines, it is suggested that the buffer be no more than a few times smaller than the range of locations sampled. *Monitor* divides the buffer into space to record the histogram of program counter samples over the range *lowpc* to *highpc*, and space to record call counts of functions compiled with the *-p* option to *cc*(1).

To profile the entire program, it is sufficient to use

```
extern etext();
...
monitor((int) 2, etext, buf, bufsize, nfunc);
```

FILES

mon.out

SEE ALSO

cc(1), prof(1), gprof(1), profil(2), sbrk(2)

3

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 given executable output file and selectively extracts a list of values. The name list consists of an array of structures containing names, types and values. The list is terminated with a null name. Each name is looked up in the name list of the file. If the name is found, the type and value of the name are inserted in the next two fields. If the name is not found, both entries are set to 0. See *a.out(5)* for the structure declaration.

This subroutine is useful for examining the system name list kept in the file `/vmunix`. In this way programs can obtain system addresses that are up to date.

SEE ALSO

a.out(5)

DIAGNOSTICS

All type entries are set to 0 if the file cannot be found or if it is not a valid namelist.

NAME

`perror`, `sys_errlist`, `sys_nerr` — system error messages

SYNOPSIS

```
perror(s)
char *s;
int sys_nerr;
char *sys_errlist[];
```

DESCRIPTION

Perror produces a short error message on the standard error file describing the last error encountered during a call to the system from a C program. First the argument string *s* is printed, then a colon, then the message and a new-line. Most usefully, the argument string is the name of the program which incurred the error. The error number is taken from the external variable *errno* (see *intro(2)*), which is set when errors occur but not cleared when non-erroneous calls are made.

To simplify variant formatting of messages, the vector of message strings *sys_errlist* is provided; *errno* can be used as an index in this table to get the message string without the newline. *sys_nerr* is the number of messages provided for in the table; it should be checked because new error codes may be added to the system before they are added to the table.

SEE ALSO

intro(2), *psignal(3)*

NAME

popen, *pclose* — initiate I/O to/from a process

SYNOPSIS

```
#include <stdio.h>
FILE *popen (command, type)
char *command, *type;
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. It creates a pipe between the calling process and the command to be executed. The value returned is a stream pointer that can be used (as appropriate) to write to the standard input of the command or read from its standard output.

A stream opened by *popen* should be closed by *pclose*, which waits for the associated process to terminate and returns the exit status of the command.

Because open files are shared, a type "r" command may be used as an input filter, and a type "w" as an output filter.

SEE ALSO

pipe(2), *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 'popened' command.

BUGS

Buffered reading before opening an input filter may leave the standard input of that filter mispositioned. Similar problems with an output filter may be forestalled by careful buffer flushing, for instance, with *flush*, see *fclose(3S)*.

Popen always calls *sh*, never calls *csh*.

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

`sigvec`(2), `perror`(3)

NAME

qsort — quicker sort

SYNOPSIS

```
qsort(base, nel, width, compar)
char *base;
int (*compar)();
```

DESCRIPTION

Qsort is an implementation of the quicker-sort algorithm. The first argument is a pointer to the base of the data; the second is the number of elements; the third is the width of an element in bytes; the last is the name of the comparison routine to be called with two arguments which are pointers to the elements being compared. The routine must return an integer less than, equal to, or greater than 0 according as the first argument is to be considered less than, equal to, or greater than the second.

SEE ALSO

sort(1)

NAME

random, srandom, initstate, setstate — better random number generator; routines for changing generators

SYNOPSIS

```
long random()
srandom(seed)
int seed;
char *initstate(seed, state, n)
unsigned seed;
char *state;
int n;
char *setstate(state)
char *state;
```

DESCRIPTION

Random uses a non-linear additive feedback random number generator employing a default table of size 31 long integers to return successive pseudo-random numbers in the range from 0 to $2^{31}-1$. The period of this random number generator is very large, approximately $16 \cdot (2^{31}-1)$.

Random/srandom have (almost) the same calling sequence and initialization properties as *rand/srand*. The difference is that *rand(3)* produces a much less random sequence -- in fact, the low dozen bits generated by *rand* go through a cyclic pattern. All the bits generated by *random* are usable. For example, "random(&01)" will produce a random binary value.

Unlike *srand*, *srandom* does not return the old seed; the reason for this is that the amount of state information used is much more than a single word. (Two other routines are provided to deal with restarting/changing random number generators). Like *rand(3)*, however, *random* will by default produce a sequence of numbers that can be duplicated by calling *srandom* with 1 as the seed.

The *initstate* routine allows a state array, passed in as an argument, to be initialized for future use. The size of the state array (in bytes) is used by *initstate* to decide how sophisticated a random number generator it should use -- the more state, the better the random numbers will be. (Current "optimal" values for the amount of state information are 8, 32, 64, 128, and 256 bytes; other amounts will be rounded down to the nearest known amount. Using less than 8 bytes will cause an error). The seed for the initialization (which specifies a starting point for the random number sequence, and provides for restarting at the same point) is also an argument. *Initstate* returns a pointer to the previous state information array.

Once a state has been initialized, the *setstate* routine provides for rapid switching between states. *Setstate* returns a pointer to the argument state array is used for further random number generation until the next call to *initstate* or *setstate*.

Once a state array has been initialized, it may be restarted at a different point either by calling *initstate* (with the desired seed, the state array, and its size) or by calling both *setstate* (with the state array) and *srandom* (with the desired seed). The advantage of calling both *setstate* and *srandom* is that the size of the state array does not have to be remembered after it is initialized.

With 256 bytes of state information, the period of the random number generator is greater than 2^{69} , which should be sufficient for most purposes.

AUTHOR

Earl T. Cohen

DIAGNOSTICS

If *initstate* is called with less than 8 bytes of state information, or if *setstate* detects that the state information has been garbled, error messages are printed on the standard error output.

SEE ALSO

rand(3)

BUGS

About 2/3 the speed of *rand(3C)*.

NAME

`re_comp`, `re_exec` — regular expression handler

SYNOPSIS

```
char *re_comp(s)
char *s;
re_exec(s)
char *s;
```

DESCRIPTION

Re_comp compiles a string into an internal form suitable for pattern matching. *Re_exec* checks the argument string against the last string passed to *re_comp*.

Re_comp returns 0 if the string *s* was compiled successfully; otherwise a string containing an error message is returned. If *re_comp* is passed 0 or a null string, it returns without changing the currently compiled regular expression.

Re_exec returns 1 if the string *s* matches the last compiled regular expression, 0 if the string *s* failed to match the last compiled regular expression, and -1 if the compiled regular expression was invalid (indicating an internal error).

The strings passed to both *re_comp* and *re_exec* may have trailing or embedded newline characters; they are terminated by nulls. The regular expressions recognized are described in the manual entry for *ed*(1), given the above difference.

SEE ALSO

ed(1), *ex*(1), *egrep*(1), *fgrep*(1), *grep*(1)

DIAGNOSTICS

Re_exec returns -1 for an internal error.

Re_comp returns 'one of the following strings if an error occurs:

```
No previous regular expression,
Regular expression too long,
unmatched \,
missing |,
too many \(\) pairs,
unmatched \.
```


NAME

`scandir` — scan a directory

SYNOPSIS

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

scandir(dirname, namelist, select, compar)
char *dirname;
struct direct *(*namelist[]);
int (*select)();
int (*compar)();
alphasort(d1, d2)
struct direct **d1, **d2;
```

DESCRIPTION

Scandir reads the directory *dirname* and builds an array of pointers to directory entries using *malloc(3)*. It returns the number of entries in the array and a pointer to the array through *namelist*.

The *select* parameter is a pointer to a user supplied subroutine which is called by *scandir* to select which entries are to be included in the array. The *select* routine is passed a pointer to a directory entry and should return a non-zero value if the directory entry is to be included in the array. If *select* is null, then all the directory entries will be included.

The *compar* parameter is a pointer to a user supplied subroutine which is passed to *qsort(3)* to sort the completed array. If this pointer is null, the array is not sorted. *Alphasort* is a routine which can be used for the *compar* parameter to sort the array alphabetically.

The memory allocated for the array can be deallocated with *free* (see *malloc(3)*) by freeing each pointer in the array and the array itself.

SEE ALSO

directory(3), *malloc(3)*, *qsort(3)*, *dir(5)*

DIAGNOSTICS

Returns `-1` if the directory cannot be opened for reading or if *malloc(3)* cannot allocate enough memory to hold all the data structures.

NAME

setjmp, longjmp — non-local goto

SYNOPSIS

```
#include <setjmp.h>
setjmp(env)
jmp_buf env;
longjmp(env, val)
jmp_buf env;
_setjmp(env)
jmp_buf env;
_longjmp(env, val)
jmp_buf env;
```

DESCRIPTION

These routines are useful for dealing with errors and interrupts encountered in a low-level sub-routine of a program.

Setjmp saves its stack environment in *env* for later use by *longjmp*. It returns value 0.

Longjmp restores the environment saved by the last call of *setjmp*. It then returns in such a way that execution continues as if the call of *setjmp* had just returned the value *val* to the function that invoked *setjmp*, which must not itself have returned in the interim. All accessible data have values as of the time *longjmp* was called.

Setjmp and *longjmp* save and restore the signal mask *sigmask(2)*, while *_setjmp* and *_longjmp* manipulate only the C stack and registers.

SEE ALSO

sigvec(2), sigstack(2), signal(3)

BUGS

Setjmp does not save current notion of whether the process is executing on the signal stack. The result is that a *longjmp* to some place on the signal stack leaves the signal stack state incorrect.

NAME

setuid, seteuid, setruid, setgid, setegid, setrgid — set user and group ID

SYNOPSIS

setuid (uid)
seteuid (euid)
setruid (ruid)
setgid (gid)
setegid (egid)
setrgid (rgid)

DESCRIPTION

Setuid (setgid) sets both the real and effective user ID (group ID) of the current process to as specified.

Seteuid (setegid) sets the effective user ID (group ID) of the current process.

Setruid (setruid) sets the real user ID (group ID) of the current process.

These calls are only permitted to the super-user or if the argument is the real or effective ID.

SEE ALSO

setreuid(2), setregid(2), getuid(2), getgid(2)

DIAGNOSTICS

Zero is returned if the user (group) ID is set; -1 is returned otherwise.

NAME

sleep — suspend execution for interval

SYNOPSIS

sleep(seconds)
unsigned seconds;

DESCRIPTION

The current process is suspended from execution for the number of seconds specified by the argument. The actual suspension time may be up to 1 second less than that requested, because scheduled wakeups occur at fixed 1-second intervals, and an arbitrary amount longer because of other activity in the system.

The routine is implemented by setting an interval timer and pausing until it occurs. The previous state of this timer is saved and restored. If the sleep time exceeds the time to the expiration of the previous timer, the process sleeps only until the signal would have occurred, and the signal is sent 1 second later.

SEE ALSO

setitimer(2), sigpause(2)

BUGS

An interface with finer resolution is needed.

NAME

strcat, strncat, strcmp, strncmp, strcpy, strncpy, strlen, index, rindex — string operations

SYNOPSIS

```
#include <strings.h>
char *strcat(s1, s2)
char *s1, *s2;
char *strncat(s1, s2, n)
char *s1, *s2;
strcmp(s1, s2)
char *s1, *s2;
strncmp(s1, s2, n)
char *s1, *s2;
char *strcpy(s1, s2)
char *s1, *s2;
char *strncpy(s1, s2, n)
char *s1, *s2;
strlen(s)
char *s;
char *index(s, c)
char *s, c;
char *rindex(s, c)
char *s, c;
```

DESCRIPTION

These functions operate on null-terminated strings. They do not check for overflow of any receiving string.

Strcat appends a copy of string *s2* to the end of string *s1*. *Strncat* copies at most *n* characters. Both return a pointer to the null-terminated result.

Strcmp compares its arguments and returns an integer greater than, equal to, or less than 0, according as *s1* is lexicographically greater than, equal to, or less than *s2*. *Strncmp* makes the same comparison but looks at at most *n* characters.

Strcpy copies string *s2* to *s1*, stopping after the null character has been moved. *Strncpy* copies exactly *n* characters, truncating or null-padding *s2*; the target may not be null-terminated if the length of *s2* is *n* or more. Both return *s1*.

Strlen returns the number of non-null characters in *s*.

Index (*rindex*) returns a pointer to the first (last) occurrence of character *c* in string *s*, or zero if *c* does not occur in the string.

NAME

`swab` -- swap bytes

SYNOPSIS

`swab`(*from*, *to*, *nbytes*)
`char` **from*, **to*;

DESCRIPTION

Swab copies *nbytes* bytes pointed to by *from* to the position pointed to by *to*, exchanging adjacent even and odd bytes. It is useful for carrying binary data between PDP11's and other machines. *Nbytes* should be even.

NAME

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*(3) 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 it cannot open the file *ldevlog*, it writes on *ldevconsole* instead and returns -1.

Closelog can be used to close the log file.

EXAMPLES

```
syslog(LOG_ALERT, "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 indicates the shell couldn't be executed.

3

NAME

ttyname, *isatty*, *ttyslot* — find name of a terminal

SYNOPSIS

char **ttyname*(*filedes*)

***isatty*(*filedes*)**

***ttyslot*()**

DESCRIPTION

Ttyname returns a pointer to the null-terminated path name of the terminal device associated with file descriptor *filedes* (this is a system file descriptor and has nothing to do with the standard I/O FILE typedef).

Isatty returns 1 if *filedes* is associated with a terminal device, 0 otherwise.

Ttyslot returns the number of the entry in the *ttys*(5) file for the control terminal of the current process.

FILES

*/dev/**

/etc/ttys

SEE ALSO

ioctl(2), *ttys*(5)

DIAGNOSTICS

Ttyname returns a null pointer (0) if *filedes* does not describe a terminal device in directory *'/dev'*.

Ttyslot returns 0 if *'/etc/ttys'* is inaccessible or if it cannot determine the control terminal.

BUGS

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

NAME

`valloc` — aligned memory allocator

SYNOPSIS

```
char *valloc(size)  
unsigned size;
```

DESCRIPTION

Valloc allocates *size* bytes aligned on a page boundary. It is implemented by calling *malloc(3)* with a slightly larger request, saving the true beginning of the block allocated, and returning a properly aligned pointer.

DIAGNOSTICS

Valloc returns a null pointer (0) if there is no available memory or if the arena has been detectably corrupted by storing outside the bounds of a block.

BUGS

Vfree isn't implemented.

3

NAME

`varargs` — variable argument list

SYNOPSIS

```
#include <varargs.h>
function(va_allst)
va_dcl
va_list pvar;
va_start(pvar);
f = va_arg(pvar, type);
va_end(pvar);
```

DESCRIPTION

This set of macros provides a means of writing portable procedures that accept variable argument lists. Routines having variable argument lists (such as *printf(3)*) that do not use `varargs` are inherently nonportable, since different machines use different argument passing conventions.

`va_allst` is used in a function header to declare a variable argument list.

`va_dcl` is a declaration for `va_allst`. Note that there is no semicolon after `va_dcl`.

`va_list` is a type which can be used for the variable `pvar`, which is used to traverse the list. One such variable must always be declared.

`va_start(pvar)` is called to initialize `pvar` to the beginning of the list.

`va_arg(pvar, type)` will return the next argument in the list pointed to by `pvar`. *Type* is the type the argument is expected to be. Different types can be mixed, but it is up to the routine to know what type of argument is expected, since it cannot be determined at runtime.

`va_end(pvar)` is used to finish up.

Multiple traversals, each bracketed by `va_start ... va_end`, are possible.

EXAMPLE

```
#include <varargs.h>
execl(va_allst)
va_dcl
{
    va_list ap;
    char *file;
    char *args[100];
    int argno = 0;

    va_start(ap);
    file = va_arg(ap, char *);
    while (args[argno++] = va_arg(ap, char *))
        ;
    va_end(ap);
    return execv(file, args);
}
```

BUGS

It is up to the calling routine to determine how many arguments there are, since it is not possible to determine this from the stack frame. For example, *execl* passes a 0 to signal the end of the list. *Printf* can tell how many arguments are supposed to be there by the format.

NAME

intro — introduction to FORTRAN library functions

DESCRIPTION

This section describes those functions that are in the FORTRAN run time library. The functions listed here provide an interface from *f77* programs to the system in the same manner as the C library does for C programs. They are automatically loaded as needed by the Fortran compiler *f77*(1).

Most of these functions are in libU77.a. Some are in libF77.a or libI77.a. A few intrinsic functions are described for the sake of completeness.

For efficiency, the SCCS ID strings are not normally included in the *a.out* file. To include them, simply declare

```
external f77lid
```

in any *f77* module.

LIST OF FUNCTIONS

<i>Name</i>	<i>Appears on Page</i>	<i>Description</i>
abort	abort.3f	terminate abruptly with memory image
access	access.3f	determine accessibility of a file
alarm	alarm.3f	execute a subroutine after a specified time
bessel	bessel.3f	of two kinds for integer orders
bit	bit.3f	and, or, xor, not, rshift, lshift bitwise functions
chdir	chdir.3f	change default directory
chmod	chmod.3f	change mode of a file
ctime	time.3f	return system time
dfrac	flmin.3f	return extreme values
dflmax	flmin.3f	return extreme values
dflmin	flmin.3f	return extreme values
drand	rand.3f	return random values
dtime	etime.3f	return elapsed execution time
etime	etime.3f	return elapsed execution time
exit	exit.3f	terminate process with status
fdate	fdate.3f	return date and time in an ASCII string
frac	flmin.3f	return extreme values
fgetc	getc.3f	get a character from a logical unit
flmax	flmin.3f	return extreme values
flmin	flmin.3f	return extreme values
flush	flush.3f	flush output to a logical unit
fork	fork.3f	create a copy of this process
fpecnt	trfppe.3f	trap and repair floating point faults
fputc	putc.3f	write a character to a fortran logical unit
fseek	fseek.3f	reposition a file on a logical unit
fstat	stat.3f	get file status
ftell	fseek.3f	reposition a file on a logical unit
gerror	perror.3f	get system error messages
getarg	getarg.3f	return command line arguments
getc	getc.3f	get a character from a logical unit
getcwd	getcwd.3f	get pathname of current working directory
getenv	getenv.3f	get value of environment variables
getgid	getuid.3f	get user or group ID of the caller
getlog	getlog.3f	get user's login name

getpid	getpid.3f	get process id
getuid	getuid.3f	get user or group ID of the caller
gmtime	time.3f	return system time
hostnm	hostnm.3f	get name of current host
iargc	getarg.3f	return command line arguments
idate	idate.3f	return date or time in numerical form
ierrno	perror.3f	get system error messages
index	index.3f	tell about character objects
inmax	flmin.3f	return extreme values
intro	intro.3f	introduction to FORTRAN library functions
ioinit	ioinit.3f	change f77 I/O initialization
irand	rand.3f	return random values
isatty	ttynam.3f	find name of a terminal port
itime	idate.3f	return date or time in numerical form
kill	kill.3f	send a signal to a process
len	index.3f	tell about character objects
link	link.3f	make a link to an existing file
lnblnk	index.3f	tell about character objects
loc	loc.3f	return the address of an object
long	long.3f	integer object conversion
lstat	stat.3f	get file status
ltime	time.3f	return system time
perror	perror.3f	get system error messages
putc	putc.3f	write a character to a fortran logical unit
qsort	qsort.3f	quick sort
rand	rand.3f	return random values
rename	rename.3f	rename a file
rindex	index.3f	tell about character objects
short	long.3f	integer object conversion
signal	signal.3f	change the action for a signal
sleep	sleep.3f	suspend execution for an interval
stat	stat.3f	get file status
system	system.3f	execute a UNIX command
tclose	topen.3f	f77 tape I/O
time	time.3f	return system time
topen	topen.3f	f77 tape I/O
traper	traper.3f	trap arithmetic errors
trapov	trapov.3f	trap and repair floating point overflow
tread	topen.3f	f77 tape I/O
trewin	topen.3f	f77 tape I/O
trpfe	trpfe.3f	trap and repair floating point faults
tskipf	topen.3f	f77 tape I/O
tstate	topen.3f	f77 tape I/O
ttynam	ttynam.3f	find name of a terminal port
twrite	topen.3f	f77 tape I/O
unlink	unlink.3f	remove a directory entry
wait	wait.3f	wait for a process to terminate

NAME

abort — terminate abruptly with memory image

SYNOPSIS

subroutine abort (string)
character*(*) string

DESCRIPTION

Abort cleans up the I/O buffers and then aborts producing a *core* file in the current directory. If *string* is given, it is written to logical unit 0 preceeded by "abort:".

FILES

/usr/lib/libF77.a

SEE ALSO

abort(3)

BUGS

String is ignored on the PDP11.

NAME

access — determine accessibility of a file

SYNOPSIS

integer function access (**name**, **mode**)
character*(*) name, **mode**

DESCRIPTION

Access checks the given file, *name*, for accessibility with respect to the caller according to *mode*. *Mode* may include in any order and in any combination one or more of:

r	test for read permission
w	test for write permission
x	test for execute permission
(blank)	test for existence

An error code is returned if either argument is illegal, or if the file can not be accessed in all of the specified modes. 0 is returned if the specified access would be successful.

FILES

/usr/lib/libU77.a

SEE ALSO

access(2), perror(3F)

BUGS

Pathnames can be no longer than MAXPATHLEN as defined in *<sys/param.h>*.

NAME

alarm — execute a subroutine after a specified time

SYNOPSIS

integer function alarm (*time*, *proc*)
integer *time*
external *proc*

DESCRIPTION

This routine arranges for subroutine *proc* to be called after *time* seconds. If *time* is "0", the alarm is turned off and no routine will be called. The returned value will be the time remaining on the last alarm.

FILES

/usr/lib/libU77.a

SEE ALSO

alarm(3C), sleep(3F), signal(3F)

BUGS

Alarm and *sleep* interact. If *sleep* is called after *alarm*, the *alarm* process will never be called. SIGALRM will occur at the lesser of the remaining *alarm* time or the *sleep* time.

NAME

bessel functions — of two kinds for integer orders

SYNOPSIS

function *besj0* (*x*)

function *besj1* (*x*)

function *besjn* (*n*, *x*)

function *besy0* (*x*)

function *besy1* (*x*)

function *besyn* (*n*, *x*)

double precision function *dbesj0* (*x*)
double precision *x*

double precision function *dbesj1* (*x*)
double precision *x*

double precision function *dbesjn* (*n*, *x*)
double precision *x*

double precision function *dbesy0* (*x*)
double precision *x*

double precision function *dbesy1* (*x*)
double precision *x*

double precision function *dbesyn* (*n*, *x*)
double precision *x*

DESCRIPTION

These functions calculate Bessel functions of the first and second kinds for real arguments and integer orders.

DIAGNOSTICS

Negative arguments cause *besy0*, *besy1*, and *besyn* to return a huge negative value. The system error code will be set to EDOM (33).

FILES

/usr/lib/libF77.a

SEE ALSO

j0(3M), *perror*(3F)

NAME

bit — and, or, xor, not, rshift, lshift bitwise functions

SYNOPSIS

(intrinsic) function **and** (**word1**, **word2**)

(intrinsic) function **or** (**word1**, **word2**)

(intrinsic) function **xor** (**word1**, **word2**)

(intrinsic) function **not** (**word**)

(intrinsic) function **rshift** (**word**, **nbits**)

(intrinsic) function **lshift** (**word**, **nbits**)

DESCRIPTION

These bitwise functions are built into the compiler and return the data type of their argument(s). It is recommended that their arguments be **integer** values; inappropriate manipulation of **real** objects may cause unexpected results.

The bitwise combinatorial functions return the bitwise “and” (**and**), “or” (**or**), or “exclusive or” (**xor**) of two operands. **Not** returns the bitwise complement of its operand.

Lshift, or *rshift* with a negative *nbits*, is a logical left shift with no end around carry. *Rshift*, or *lshift* with a negative *nbits*, is an arithmetic right shift with sign extension. No test is made for a reasonable value of *nbits*.

FILES

These functions are generated in-line by the *f77* compiler.

NAME

chdir — change default directory

SYNOPSIS

integer function chdir (**dirname**)
character*(*) dirname

DESCRIPTION

The default directory for creating and locating files will be changed to *dirname*. Zero is returned if successful; an error code otherwise.

FILES

/usr/lib/libU77.a

SEE ALSO

chdir(2), cd(1), perror(3F)

BUGS

Pathnames can be no longer than MAXPATHLEN as defined in *<sys/param.h>*.

Use of this function may cause **inquire** by unit to fail.

NAME

chmod — change mode of a file

SYNOPSIS

integer function chmod (**name, mode**)
character*(*) name, mode

DESCRIPTION

This function changes the filesystem *mode* of file *name*. *Mode* can be any specification recognized by *chmod(1)*. *Name* must be a single pathname.

The normal returned value is 0. Any other value will be a system error number.

FILES

/usr/lib/libU77.a
/bin/chmod exec'ed to change the mode.

SEE ALSO

chmod(1)

BUGS

Pathnames can be no longer than MAXPATHLEN as defined in *<sys/param.h>*.

NAME

etime, dtime — return elapsed execution time

SYNOPSIS

function etime (tarray)
real tarray(2)

function dtime (tarray)
real tarray(2)

DESCRIPTION

These two routines return elapsed runtime in seconds for the calling process. *Dtime* returns the elapsed time since the last call to *dtime*, or the start of execution on the first call.

The argument array returns user time in the first element and system time in the second element. The function value is the sum of user and system time.

The resolution of all timing is 1/HZ sec. where HZ is currently 60.

FILES

/usr/lib/libU77.a

SEE ALSO

times(2)

NAME

`exit` — terminate process with status

SYNOPSIS

subroutine `exit` (*status*)
integer *status*

DESCRIPTION

Exit flushes and closes all the process's files, and notifies the parent process if it is executing a *wait*. The low-order 8 bits of *status* are available to the parent process. (Therefore *status* should be in the range 0 — 255)

This call will never return.

The C function *exit* may cause cleanup actions before the final 'sys exit'.

FILES

/usr/lib/libF77.a

SEE ALSO

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

NAME

`fdate` — return date and time in an ASCII string

SYNOPSIS

subroutine `fdate` (string)

character*(*) string

character*(*) 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
external    fdate
```

```
write(*,*) fdate()
```

FILES

/usr/lib/libU77.a

SEE ALSO

ctime(3), *time*(3F), *itime*(3F), *idate*(3F), *ltime*(3F)

NAME

*f*lmin, *f*lmax, *f*frac, *d*flmin, *d*flmax, *d*frac, *i*nmax — return extreme values

SYNOPSIS

function flmin()

function flmax()

function ffrac()

double precision function dflmin()

double precision function dflmax()

double precision function dfrac()

function inmax()

DESCRIPTION

Functions *f*lmin and *f*lmax return the minimum and maximum positive floating point values respectively. Functions *d*flmin and *d*flmax return the minimum and maximum positive double precision floating point values. Function *i*nmax returns the maximum positive integer value.

The functions *f*frac and *d*frac return the fractional accuracy of single and double precision floating point numbers respectively. These are the smallest numbers that can be added to 1.0 without being lost.

These functions can be used by programs that must scale algorithms to the numerical range of the processor.

FILES

/usr/lib/libF77.a

NAME

flush — flush output to a logical unit

SYNOPSIS

subroutine flush (lunit)

DESCRIPTION

Flush causes the contents of the buffer for logical unit *lunit* to be flushed to the associated file. This is most useful for logical units 0 and 6 when they are both associated with the control terminal.

FILES

/usr/lib/libI77.a

SEE ALSO

fclose(3S)

NAME

fork — create a copy of this process

SYNOPSIS

integer function fork()

DESCRIPTION

Fork creates a copy of the calling process. The only distinction between the 2 processes is that the value returned to one of them (referred to as the 'parent' process) will be the process id if the copy. The copy is usually referred to as the 'child' process. The value returned to the 'child' process will be zero.

All logical units open for writing are flushed before the fork to avoid duplication of the contents of I/O buffers in the external file(s).

If the returned value is negative, it indicates an error and will be the negation of the system error code. See *perror*(3F).

A corresponding *exec* routine has not been provided because there is no satisfactory way to retain open logical units across the *exec*. However, the usual function of *forklexec* can be performed using *system*(3F).

FILES

/usr/lib/libU77.a

SEE ALSO

fork(2), wait(3F), kill(3F), system(3F), perror(3F)

NAME

fseek, *ftell* — reposition a file on a logical unit

SYNOPSIS

integer function *fseek* (*lunit*, *offset*, *from*)
integer *offset*, *from*

integer function *ftell* (*lunit*)

DESCRIPTION

lunit must refer to an open logical unit. *offset* is an offset in bytes relative to the position specified by *from*. Valid values for *from* are:

- 0 meaning 'beginning of the file'
- 1 meaning 'the current position'
- 2 meaning 'the end of the file'

The value returned by *fseek* will be 0 if successful, a system error code otherwise. (See *perror*(3F))

Ftell returns the current position of the file associated with the specified logical unit. The value is an offset, in bytes, from the beginning of the file. If the value returned is negative, it indicates an error and will be the negation of the system error code. (See *perror*(3F))

FILES

/usr/lib/libU77.a

SEE ALSO

fseek(3S), *perror*(3F)

NAME

getarg, *iargc* — return command line arguments

SYNOPSIS

subroutine *getarg* (**k**, **arg**)
character*(*) **arg**

function *iargc* ()

DESCRIPTION

A call to *getarg* will return the *k*th command line argument in character string *arg*. The 0th argument is the command name.

Iargc returns the index of the last command line argument.

FILES

/usr/lib/libU77.a

SEE ALSO

getenv(3F), *execve*(2)

NAME

getc, fgetc — get a character from a logical unit

SYNOPSIS

integer function **getc (char)**
character **char**

integer function **fgetc (lunit, char)**
character **char**

DESCRIPTION

These routines return the next character from a file associated with a fortran logical unit, bypassing normal fortran I/O. *Getc* reads from logical unit 5, normally connected to the control terminal input.

The value of each function is a system status code. Zero indicates no error occurred on the read; -1 indicates end of file was detected. A positive value will be either a UNIX system error code or an I77 I/O error code. See *perror*(3F).

FILES

/usr/lib/libU77.a

SEE ALSO

getc(3S), intro(2), perror(3F)

NAME

`getcwd` — get pathname of current working directory

SYNOPSIS

integer function `getcwd` (*dirname*)
character*(*) *dirname*

DESCRIPTION

The pathname of the default directory for creating and locating files will be returned in *dirname*. The value of the function will be zero if successful; an error code otherwise.

FILES

`/usr/lib/libU77.a`

SEE ALSO

`chdir(3F)`, `perror(3F)`

BUGS

Pathnames can be no longer than `MAXPATHLEN` as defined in `<sysparam.h>`.

NAME

getenv — get value of environment variables

SYNOPSIS

subroutine *getenv* (**ename**, **eval**)
character*(*) **ename**, **eval**

DESCRIPTION

Getenv searches the environment list (see *environ(7)*) for a string of the form *ename=value* and returns *value* in *eval* if such a string is present, otherwise fills *eval* with blanks.

FILES

/usr/lib/libU77.a

SEE ALSO

environ(7), *execve(2)*

NAME

getlog — get user's login name

SYNOPSIS

subroutine **getlog** (**name**)

character*(*) **name**

character*(*) function **getlog**()

DESCRIPTION

Getlog will return the user's login name or all blanks if the process is running detached from a terminal.

FILES

/usr/lib/libU77.a

SEE ALSO

getlogin(3)

NAME

getpid — get process id

SYNOPSIS

integer function **getpid**(0)

DESCRIPTION

Getpid returns the process ID number of the current process.

FILES

/usr/lib/libU77.a

SEE ALSO

getpid(2)

NAME

getuid, getgid — get user or group ID of the caller

SYNOPSIS

integer function getuid()

integer function getgid()

DESCRIPTION

These functions return the real user or group ID of the user of the process.

FILES

/usr/lib/libU77.a

SEE ALSO

getuid(2)

NAME

hostnm — get name of current host

SYNOPSIS

integer function hostnm (name)
character*(*) name

DESCRIPTION

This function puts the name of the current host into character string *name*. The return value should be 0; any other value indicates an error.

FILES

/usr/lib/libU77.a

SEE ALSO

gethostname(2)

3

NAME

idate, *itime* — return date or time in numerical form

SYNOPSIS

subroutine *idate* (*iarray*)
integer *iarray*(3)

subroutine *itime* (*iarray*)
integer *iarray*(3)

DESCRIPTION

Idate returns the current date in *iarray*. The order is: day, mon, year. Month will be in the range 1-12. Year will be \geq 1969.

Itime returns the current time in *iarray*. The order is: hour, minute, second.

FILES

/usr/lib/libU77.a

SEE ALSO

ctime(3F), *fdate*(3F)

NAME

index, *rindex*, *lnblk*, *len* — tell about character objects

SYNOPSIS

(intrinsic) function *index* (*string*, *substr*)
character*(*) *string*, *substr*

integer function *rindex* (*string*, *substr*)
character*(*) *string*, *substr*

function *lnblk* (*string*)
character*(*) *string*

(intrinsic) function *len* (*string*)
character*(*) *string*

DESCRIPTION

Index (*rindex*) returns the index of the first (last) occurrence of the substring *substr* in *string*, or zero if it does not occur. *Index* is an f77 intrinsic function; *rindex* is a library routine.

Lnblk returns the index of the last non-blank character in *string*. This is useful since all f77 character objects are fixed length, blank padded. Intrinsic function *len* returns the size of the character object argument.

FILES

/usr/lib/libF77.a

NAME

`ioinit` — change *f77* I/O initialization

SYNOPSIS

logical function `ioinit` (`cctl`, `bzro`, `apnd`, `prefix`, `vrbose`)
logical `cctl`, `bzro`, `apnd`, `vrbose`
character*(*) `prefix`

DESCRIPTION

This routine will initialize several global parameters in the *f77* I/O system, and attach externally defined files to logical units at run time. The effect of the flag arguments applies to logical units opened after *ioinit* is called. The exception is the preassigned units, 5 and 6, to which *cctl* and *bzro* will apply at any time. *ioinit* is written in Fortran-77.

By default, carriage control is not recognized on any logical unit. If *cctl* is **.true.**, then carriage control will be recognized on formatted output to all logical units except unit 0, the diagnostic channel. Otherwise the default will be restored.

By default, trailing and embedded blanks in input data fields are ignored. If *bzro* is **.true.**, then such blanks will be treated as zero's. Otherwise the default will be restored.

By default, all files opened for sequential access are positioned at their beginning. It is sometimes necessary or convenient to open at the END-OF-FILE so that a write will append to the existing data. If *apnd* is **.true.**, then files opened subsequently on any logical unit will be positioned at their end upon opening. A value of **.false.** will restore the default behavior.

Many systems provide an automatic association of global names with forttran logical units when a program is run. There is no such automatic association in *f77*. However, if the argument *prefix* is a non-blank string, then names of the form `prefixNN` will be sought in the program environment. The value associated with each such name found will be used to open logical unit NN for formatted sequential access. For example, if *f77* program *myprogram* included the call

```
call ioinit (.true., .false., .false., 'FORT', .false.)
```

then when the following sequence

```
% setenv FORT01 mydata
% setenv FORT12 myresults
% myprogram
```

would result in logical unit 1 opened to file *mydata* and logical unit 12 opened to file *myresults*. Both files would be positioned at their beginning. Any formatted output would have column 1 removed and interpreted as carriage control. Embedded and trailing blanks would be ignored on input.

If the argument *vrbose* is **.true.**, then *ioinit* will report on its activity.

The effect of

```
call ioinit (.true., .true., .false., ", .false.)
```

can be achieved without the actual call by including “-II66” on the *f77* command line. This gives carriage control on all logical units except 0, causes files to be opened at their beginning, and causes blanks to be interpreted as zero's.

The internal flags are stored in a labeled common block with the following definition:

```
integer*2 ieof, ictl, ibzr
```

common /ioiflg/ ieof, ictl, ibzr

FILES

/usr/lib/libI77.a f77 I/O library
/usr/lib/libI66.a sets older fortran I/O modes

SEE ALSO

getarg(3F), getenv(3F), "Introduction to the f77 I/O Library"

BUGS

Prefix can be no longer than 30 characters. A pathname associated with an environment name can be no longer than 255 characters.

The "+" carriage control does not work.

NAME

kill — send a signal to a process

SYNOPSIS

function kill (pid, signum)
integer pid, signum

DESCRIPTION

Pid must be the process id of one of the user's processes. *Signum* must be a valid signal number (see `sigvec(2)`). The returned value will be 0 if successful; an error code otherwise.

FILES

/usr/lib/libU77.a

SEE ALSO

kill(2), sigvec(2), signal(3F), fork(3F), perror(3F)

NAME

link — make a link to an existing file

SYNOPSIS

function link (name1, name2)
character*(*) name1, name2

integer function symlink (name1, name2)
character*(*) name1, name2

DESCRIPTION

Name1 must be the pathname of an existing file. *Name2* is a pathname to be linked to file *name1*. *Name2* must not already exist. The returned value will be 0 if successful; a system error code otherwise.

Symlink creates a symbolic link to *name1*.

FILES

/usr/lib/libU77.a

SEE ALSO

link(2), symlink(2), perror(3F), unlink(3F)

BUGS

Pathnames can be no longer than MAXPATHLEN as defined in *<sys/param.h>*.

NAME

`loc` — return the address of an object

SYNOPSIS

function `loc (arg)`

DESCRIPTION

The returned value will be the address of *arg*.

FILES

`/usr/lib/libU77.a`

NAME

long, short — integer object conversion

SYNOPSIS

integer*4 function long (int2)
integer*2 int2

integer*2 function short (int4)
integer*4 int4

DESCRIPTION

These functions provide conversion between short and long integer objects. *Long* is useful when constants are used in calls to library routines and the code is to be compiled with “-i2”. *Short* is useful in similar context when an otherwise long object must be passed as a short integer.

FILES

/usr/lib/libF77.a

NAME

`perror`, `gerror`, `ierrno` — get system error messages

SYNOPSIS

**subroutine `perror` (string)
character*(*) string**

**subroutine `gerror` (string)
character*(*) string**

character*(*) function `gerror`()

function `ierrno`()

DESCRIPTION

Perror will write a message to fortran logical unit 0 appropriate to the last detected system error. *String* will be written preceding the standard error message.

Gerror returns the system error message in character variable *string*. *Gerror* may be called either as a subroutine or as a function.

Ierrno will return the error number of the last detected system error. This number is updated only when an error actually occurs. Most routines and I/O statements that might generate such errors return an error code after the call; that value is a more reliable indicator of what caused the error condition.

FILES

/usr/lib/libU77.a

SEE ALSO

`intro(2)`, `perror(3)`

D. L. Wasley, *Introduction to the f77 I/O Library*

BUGS

String in the call to *perror* can be no longer than 127 characters.

The length of the string returned by *gerror* is determined by the calling program.

NOTES

UNIX system error codes are described in `intro(2)`. The f77 I/O error codes and their meanings are:

100	“error in format”
101	“illegal unit number”
102	“formatted io not allowed”
103	“unformatted io not allowed”
104	“direct io not allowed”
105	“sequential io not allowed”
106	“can't backspace file”
107	“off beginning of record”
108	“can't stat file”
109	“no * after repeat count”
110	“off end of record”
111	“truncation failed”
112	“incomprehensible list input”
113	“out of free space”
114	“unit not connected”
115	“read unexpected character”

- 116 "blank logical input field"
- 117 "'new' file exists"
- 118 "can't find 'old' file"
- 119 "unknown system error"
- 120 "requires seek ability"
- 121 "illegal argument"
- 122 "negative repeat count"
- 123 "illegal operation for unit"

3

NAME

putc, fputc — write a character to a fortran logical unit

SYNOPSIS

integer function putc (**char**)
character char

integer function fputc (**lunit, char**)
character char

DESCRIPTION

These functions write a character to the file associated with a fortran logical unit bypassing normal fortran I/O. *Putc* writes to logical unit 6, normally connected to the control terminal output.

The value of each function will be zero unless some error occurred; a system error code otherwise. See *perror*(3F).

FILES

/usr/lib/libU77.a

SEE ALSO

putc(3S), intro(2), perror(3F)

NAME

qsort — quick sort

SYNOPSIS

subroutine **qsort** (**array**, **len**, **isize**, **compar**)
 external **compar**
 integer*2 **compar**

DESCRIPTION

One dimensional *array* contains the elements to be sorted. *len* is the number of elements in the array. *isize* is the size of an element, typically -

4 for **integer** and **real**
 8 for **double precision** or **complex**
 16 for **double complex**
 (length of character object) for **character** arrays

Compar is the name of a user supplied integer*2 function that will determine the sorting order. This function will be called with 2 arguments that will be elements of *array*. The function must return -

negative if arg 1 is considered to precede arg 2
 zero if arg 1 is equivalent to arg 2
 positive if arg 1 is considered to follow arg 2

On return, the elements of *array* will be sorted.

FILES

/usr/lib/libU77.a

SEE ALSO

qsort(3)

NAME

rand, drand, irand — return random values

SYNOPSIS

function irand (iflag)

function rand (iflag)

double precision function drand (iflag)

DESCRIPTION

These functions use *rand(3C)* to generate sequences of random numbers. If *iflag* is '1', the generator is restarted and the first random value is returned. If *iflag* is otherwise non-zero, it is used as a new seed for the random number generator, and the first new random value is returned.

Irand returns positive integers in the range 0 through 2147483647. *Rand* and *drand* return values in the range 0. through 1.0 .

FILES

/usr/lib/libF77.a

SEE ALSO

rand(3C)

BUGS

The algorithm returns a 15 bit quantity on the PDP11; a 31 bit quantity on the VAX. *Irand* on the PDP11 calls *rand(3C)* twice to form a 31 bit quantity, but bit 15 will always be 0.

NAME

fmin, *fmax*, *dfmin*, *dfmax*, *imax* — return extreme values

SYNOPSIS

function *fmin*()

function *fmax*()

double precision function *dfmin*()

double precision function *dfmax*()

function *imax*()

DESCRIPTION

Functions *fmin* and *fmax* return the minimum and maximum positive floating point values respectively. Functions *dfmin* and *dfmax* return the minimum and maximum positive double precision floating point values. Function *imax* returns the maximum positive integer value.

These functions can be used by programs that must scale algorithms to the numerical range of the processor.

FILES

/usr/lib/libF77.a

3

NAME

rename — rename a file

SYNOPSIS

integer function rename (*from*, *to*)
character*(*) *from*, *to*

DESCRIPTION

From must be the pathname of an existing file. *To* will become the new pathname for the file. If *to* exists, then both *from* and *to* must be the same type of file, and must reside on the same filesystem. If *to* exists, it will be removed first.

The returned value will be 0 if successful; a system error code otherwise.

FILES

/usr/lib/libU77.a

SEE ALSO

rename(2), perror(3F)

BUGS

Pathnames can be no longer than MAXPATHLEN as defined in `<sys/param.h>`.

NAME

signal — change the action for a signal

SYNOPSIS

integer function **signal**(**signum**, **proc**, **flag**)
integer **signum**, **flag**
external proc

DESCRIPTION

When a process incurs a signal (see *signal(3C)*) the default action is usually to clean up and abort. The user may choose to write an alternative signal handling routine. A call to *signal* is the way this alternate action is specified to the system.

Signum is the signal number (see *signal(3C)*). If *flag* is negative, then *proc* must be the name of the user signal handling routine. If *flag* is zero or positive, then *proc* is ignored and the value of *flag* is passed to the system as the signal action definition. In particular, this is how previously saved signal actions can be restored. Two possible values for *flag* have specific meanings: 0 means "use the default action" (See NOTES below), 1 means "ignore this signal".

A positive returned value is the previous action definition. A value greater than 1 is the address of a routine that was to have been called on occurrence of the given signal. The returned value can be used in subsequent calls to *signal* in order to restore a previous action definition. A negative returned value is the negation of a system error code. (See *perror(3F)*)

FILES

/usr/lib/libU77.a

SEE ALSO

signal(3C), *kill(3F)*, *kill(1)*

NOTES

f77 arranges to trap certain signals when a process is started. The only way to restore the default *f77* action is to save the returned value from the first call to *signal*.

If the user signal handler is called, it will be passed the signal number as an integer argument.

NAME

`sleep` — suspend execution for an interval

SYNOPSIS

`subroutine sleep (itime)`

DESCRIPTION

Sleep causes the calling process to be suspended for *itime* seconds. The actual time can be up to 1 second less than *itime* due to granularity in system timekeeping.

FILES

`/usr/lib/libU77.a`

SEE ALSO

`sleep(3)`



NAME

stat, *lstat*, *fstat* — get file status

SYNOPSIS

integer function *stat* (*name*, *statb*)
character*(*) *name*
integer *statb*(12)

integer function *lstat* (*name*, *statb*)
character*(*) *name*
integer *statb*(12)

integer function *fstat* (*lunit*, *statb*)
integer *statb*(12)

DESCRIPTION

These routines return detailed information about a file. *Stat* and *lstat* return information about file *name*; *fstat* returns information about the file associated with fortran logical unit *lunit*. The order and meaning of the information returned in array *statb* is as described for the structure *stat* under *stat*(2). The "spare" values are not included.

The value of either function will be zero if successful; an error code otherwise.

FILES

/usr/lib/libU77.a

SEE ALSO

stat(2), *access*(3F), *pererr*(3F), *time*(3F)

BUGS

Pathnames can be no longer than MAXPATHLEN as defined in *<sys/param.h>*.

NAME

system — execute a UNIX command

SYNOPSIS

integer function system (string)
character*(*) string

DESCRIPTION

System causes *string* to be given to your shell as input as if the string had been typed as a command. If environment variable **SHELL** is found, its value will be used as the command interpreter (shell); otherwise *sh*(1) is used.

The current process waits until the command terminates. The returned value will be the exit status of the shell. See *wait*(2) for an explanation of this value.

FILES

/usr/lib/libU77.a

SEE ALSO

exec(2), *wait*(2), *system*(3)

BUGS

String can not be longer than **NCARGS**—50 characters, as defined in *<sys/param.h>*.

NAME

time, *ctime*, *ltime*, *gmtime* — return system time

SYNOPSIS

integer function *time*()

character*(*) function *ctime* (*stime*)
integer *stime*

subroutine *ltime* (*stime*, *tarray*)
integer *stime*, *tarray*(9)

subroutine *gmtime* (*stime*, *tarray*)
integer *stime*, *tarray*(9)

DESCRIPTION

Time returns the time since 00:00:00 GMT, Jan. 1, 1970, measured in seconds. This is the value of the UNIX system clock.

Ctime converts a system time to a 24 character ASCII string. The format is described under *ctime*(3). No 'newline' or NULL will be included.

Ltime and *gmtime* dissect a UNIX time into month, day, etc., either for the local time zone or as GMT. The order and meaning of each element returned in *tarray* is described under *ctime*(3).

FILES

/usr/lib/libU77.a

SEE ALSO

ctime(3), *itime*(3F), *idate*(3F), *fdate*(3F)

NAME

topen, *tclose*, *tread*, *twrite*, *trewin*, *tskipf*, *tstate* — *f77* tape I/O

SYNOPSIS

integer function *topen* (*tlu*, *devnam*, *label*)
integer *tlu*
character*(*) *devnam*
logical *label*

integer function *tclose* (*tlu*)
integer *tlu*

integer function *tread* (*tlu*, *buffer*)
integer *tlu*
character*(*) *buffer*

integer function *twrite* (*tlu*, *buffer*)
integer *tlu*
character*(*) *buffer*

integer function *trewin* (*tlu*)
integer *tlu*

integer function *tskipf* (*tlu*, *nfiles*, *nrecs*)
integer *tlu*, *nfiles*, *nrecs*

integer function *tstate* (*tlu*, *fileno*, *recno*, *errf*, *eof*, *eotf*, *tcsr*)
integer *tlu*, *fileno*, *recno*, *tcsr*
logical *errf*, *eof*, *eotf*

DESCRIPTION

These functions provide a simple interface between *f77* and magnetic tape devices. A “tape logical unit”, *tlu*, is “*topen*”ed in much the same way as a normal *f77* logical unit is “*open*”ed. All other operations are performed via the *tlu*. The *tlu* has no relationship at all to any normal *f77* logical unit.

Topen associates a device name with a *tlu*. *Tlu* must be in the range 0 to 3. The logical argument *label* should indicate whether the tape includes a tape label. This is used by *trewin* below. *Topen* does not move the tape. The normal returned value is 0. If the value of the function is negative, an error has occurred. See *perorr*(3F) for details.

Tclose closes the tape device channel and removes its association with *tlu*. The normal returned value is 0. A negative value indicates an error.

Tread reads the next physical record from tape to *buffer*. *Buffer* must be of type **character**. The size of *buffer* should be large enough to hold the largest physical record to be read. The actual number of bytes read will be returned as the value of the function. If the value is 0, the end-of-file has been detected. A negative value indicates an error.

Twrite writes a physical record to tape from *buffer*. The physical record length will be the size of *buffer*. *Buffer* must be of type **character**. The number of bytes written will be returned. A value of 0 or negative indicates an error.

Trewin rewinds the tape associated with *tlu* to the beginning of the first data file. If the tape is a labelled tape (see *topen* above) then the label is skipped over after rewinding. The normal returned value is 0. A negative value indicates an error.

Tskipf allows the user to skip over files and/or records. First, *nfiles* end-of-file marks are skipped. If the current file is at EOF, this counts as 1 file to skip. (Note: This is the way to reset the EOF status for a *thu*.) Next, *nrecs* physical records are skipped over. The normal returned value is 0. A negative value indicates an error.

Finally, *tstate* allows the user to determine the logical state of the tape I/O channel and to see the tape drive control status register. The values of *fileno* and *recno* will be returned and indicate the current file and record number. The logical values *errf*, *eof*, and *eotf* indicate an error has occurred, the current file is at EOF, or the tape has reached logical end-of-tape. End-of-tape (EOT) is indicated by an empty file, often referred to as a double EOF mark. It is not allowed to read past EOT although it is allowed to write. The value of *tsr* will reflect the tape drive control status register. See *ht(4)* for details.

FILES

/usr/lib/libU77.a

SEE ALSO

ht(4), *perror(3F)*, *rewind(1)*

NAME

traper — trap arithmetic errors

SYNOPSIS

integer function traper (mask)

DESCRIPTION

NOTE: This routine applies only to the VAX. It is ignored on the PDP11.

Integer overflow and floating point underflow are not normally trapped during execution. This routine enables these traps by setting status bits in the process status word. These bits are reset on entry to a subprogram, and the previous state is restored on return. Therefore, this routine must be called *inside* each subprogram in which these conditions should be trapped. If the condition occurs and trapping is enabled, signal SIGFPE is sent to the process. (See *signal(3C)*)

The argument has the following meaning:

value	meaning
0	do not trap either condition
1	trap integer overflow only
2	trap floating underflow only
3	trap both the above

The previous value of these bits is returned.

FILES

/usr/lib/libF77.a

SEE ALSO

signal(3C), signal(3F)

NAME

trapov — trap and repair floating point overflow

SYNOPSIS

subroutine trapov (numesg, rtnval)
double precision rtnval

DESCRIPTION

NOTE: This routine applies only to the older VAX 11/780's. VAX computers made or upgraded since spring 1983 handle errors differently. See *trpfpe*(3F) for the newer error handler. This routine has always been ineffective on the VAX 11/750. It is a null routine on the PDP11.

This call sets up signal handlers to trap arithmetic exceptions and the use of illegal operands. Trapping arithmetic exceptions allows the user's program to proceed from instances of floating point overflow or divide by zero. The result of such operations will be an illegal floating point value. The subsequent use of the illegal operand will be trapped and the operand replaced by the specified value.

The first *numesg* occurrences of a floating point arithmetic error will cause a message to be written to the standard error file. If the resulting value is used, the value given for *rtnval* will replace the illegal operand generated by the arithmetic error. *Rtnval* must be a double precision value. For example, "0d0" or "dflmax()".

FILES

/usr/lib/libF77.a

SEE ALSO

trpfpe(3F), *signal*(3F), *range*(3F)

BUGS

Other arithmetic exceptions can be trapped but not repaired.

There is no way to distinguish between an integer value of 32768 and the illegal floating point form. Therefore such an integer value may get replaced while repairing the use of an illegal operand.

NAME

trpfpe, *fpcent* — trap and repair floating point faults

SYNOPSIS

subroutine *trpfpe* (*numesg*, *rtnval*)
double precision *rtnval*

integer function *fpcent* ()

common /*fpflt*/ *fperr*
logical *fperr*

DESCRIPTION

NOTE: This routine applies only to Vax computers. It is a null routine on the PDP11.

Trpfpe sets up a signal handler to trap arithmetic exceptions. If the exception is due to a floating point arithmetic fault, the result of the operation is replaced with the *rtnval* specified. *Rtnval* must be a double precision value. For example, "0d0" or "dflmax()".

The first *numesg* occurrences of a floating point arithmetic error will cause a message to be written to the standard error file. Any exception that can't be repaired will result in the default action, typically an abort with core image.

Fpcent returns the number of faults since the last call to *trpfpe*.

The logical value in the common block labelled *fpflt* will be set to **.true.** each time a fault occurs.

FILES

/usr/lib/libF77.a

SEE ALSO

signal(3F), range(3F)

BUGS

This routine works only for *faults*, not *traps*. This is primarily due to the Vax architecture.

If the operation involves changing the stack pointer, it can't be repaired. This seldom should be a problem with the *f77* compiler, but such an operation might be produced by the optimizer.

The POLY and EMOD opcodes are not dealt with.

NAME

`ttynam`, `isatty` — find name of a terminal port

SYNOPSIS

character*(*) function `ttynam (lunit)`

logical function `isatty (lunit)`

DESCRIPTION

Ttynam returns a blank padded path name of the terminal device associated with logical unit *lunit*.

Isatty returns **.true.** if *lunit* is associated with a terminal device, **.false.** otherwise.

FILES

`/dev/*`
`/usr/lib/libU77.a`

DIAGNOSTICS

Ttynam returns an empty string (all blanks) if *lunit* is not associated with a terminal device in directory `'/dev'`.

NAME

unlink — remove a directory entry

SYNOPSIS

integer function unlink (**name**)
character*(*) name

DESCRIPTION

Unlink causes the directory entry specified by pathname *name* to be removed. If this was the last link to the file, the contents of the file are lost. The returned value will be zero if successful; a system error code otherwise.

FILES

/usr/lib/libU77.a

SEE ALSO

unlink(2), link(3F), filesys(5), perror(3F)

BUGS

Pathnames can be no longer than MAXPATHLEN as defined in `<sys/param.h>`.

NAME

`wait` — wait for a process to terminate

SYNOPSIS

integer function `wait (status)`

integer status

DESCRIPTION

Wait causes its caller to be suspended until a signal is received or one of its child processes terminates. If any child has terminated since the last *wait*, return is immediate; if there are no children, return is immediate with an error code.

If the returned value is positive, it is the process ID of the child and *status* is its termination status (see *wait(2)*). If the returned value is negative, it is the negation of a system error code.

FILES

`/usr/lib/libU77.a`

SEE ALSO

`wait(2)`, `signal(3F)`, `kill(3F)`, `perror(3F)`

NAME

intro — introduction to mathematical library functions

DESCRIPTION

These functions constitute the math library, *libm*. They are automatically loaded as needed by the Fortran compiler *f77(1)*. The link editor searches this library under the “-lm” option. Declarations for these functions may be obtained from the include file *<math.h>*.

LIST OF FUNCTIONS

<i>Name</i>	<i>Appears on Page</i>	<i>Description</i>
acos	sin.3m	trigonometric functions
asin	sin.3m	trigonometric functions
atan	sin.3m	trigonometric functions
atan2	sin.3m	trigonometric functions
cabs	hypot.3m	Euclidean distance
ceil	floor.3m	absolute value, floor, ceiling functions
cos	sin.3m	trigonometric functions
cosh	sinh.3m	hyperbolic functions
exp	exp.3m	exponential, logarithm, power, square root
fabs	floor.3m	absolute value, floor, ceiling functions
floor	floor.3m	absolute value, floor, ceiling functions
gamma	gamma.3m	log gamma function
hypot	hypot.3m	Euclidean distance
j0	j0.3m	bessel functions
j1	j0.3m	bessel functions
jn	j0.3m	bessel functions
log	exp.3m	exponential, logarithm, power, square root
log10	exp.3m	exponential, logarithm, power, square root
pow	exp.3m	exponential, logarithm, power, square root
sin	sin.3m	trigonometric functions
sinh	sinh.3m	hyperbolic functions
sqrt	exp.3m	exponential, logarithm, power, square root
tan	sin.3m	trigonometric functions
tanh	sinh.3m	hyperbolic functions
y0	j0.3m	bessel functions
y1	j0.3m	bessel functions
yn	j0.3m	bessel functions

NAME

`exp`, `log`, `log10`, `pow`, `sqrt` — exponential, logarithm, power, square root

SYNOPSIS

```
#include <math.h>  
double exp(x)  
double x;  
double log(x)  
double x;  
double log10(x)  
double x;  
double pow(x, y)  
double x, y;  
double sqrt(x)  
double x;
```

DESCRIPTION

Exp returns the exponential function of x .

Log returns the natural logarithm of x ; *log10* returns the base 10 logarithm.

Pow returns x^y .

Sqrt returns the square root of x .

SEE ALSO

`hypot(3M)`, `sinh(3M)`, `intro(3M)`

DIAGNOSTICS

Exp and *pow* return a huge value when the correct value would overflow; *errno* is set to ERANGE. *Pow* returns 0 and sets *errno* to EDOM when the first argument is negative and the second is non-integral or when first argument is 0 and the second is less than or equal to 0.

Log returns 0 when x is zero or negative; *errno* is set to EDOM.

Sqrt returns 0 when x is negative; *errno* is set to EDOM.

NAME

fabs, *floor*, *ceil* — absolute value, floor, ceiling functions

SYNOPSIS

```
#include <math.h>
double floor(x)
double x;
double ceil(x)
double x;
double fabs(x)
double x;
```

DESCRIPTION

Fabs returns the absolute value $|x|$.

Floor returns the largest integer not greater than x .

Ceil returns the smallest integer not less than x .

SEE ALSO

[abs\(3\)](#)

NAME

gamma — log gamma function

SYNOPSIS

```
#include <math.h>
double gamma(x)
double x;
```

DESCRIPTION

Gamma returns $\ln |\Gamma(|x|)|$. The sign of $\Gamma(|x|)$ is returned in the external integer *signgam*. The following C program might be used to calculate Γ :

```
y = gamma(x);
if (y > 88.0)
    error();
y = exp(y);
if (signgam)
    y = -y;
```

DIAGNOSTICS

A huge value is returned for negative integer arguments.

BUGS

There should be a positive indication of error.

NAME

hypot, *cabs* — Euclidean distance

SYNOPSIS

```
#include <math.h>
double hypot(x, y)
double x, y;
double cabs(z)
struct { double x, y;} z;
```

DESCRIPTION

Hypot and *cabs* return
 $\text{sqrt}(x*x + y*y)$,
taking precautions against unwarranted overflows.

SEE ALSO

exp(3M) for *sqrt*

NAME

j_0 , j_1 , j_n , y_0 , y_1 , y_n — Bessel functions

SYNOPSIS

```
#include <math.h>
```

```
double j0(x)
```

```
double x;
```

```
double j1(x)
```

```
double x;
```

```
double jn(n, x)
```

```
double x;
```

```
double y0(x)
```

```
double x;
```

```
double y1(x)
```

```
double x;
```

```
double yn(n, x)
```

```
double x;
```

DESCRIPTION

These functions calculate Bessel functions of the first and second kinds for real arguments and integer orders.

DIAGNOSTICS

Negative arguments cause y_0 , y_1 , and y_n to return a huge negative value and set *errno* to EDOM.

NAME

sin, *cos*, *tan*, *asin*, *acos*, *atan*, *atan2* — trigonometric functions

SYNOPSIS

```
#include <math.h>

double sin(x)
double x;

double cos(x)
double x;

double asin(x)
double x;

double acos(x)
double x;

double atan(x)
double x;

double atan2(x, y)
double x, y;
```

DESCRIPTION

Sin, *cos* and *tan* return trigonometric functions of radian arguments. The magnitude of the argument should be checked by the caller to make sure the result is meaningful.

Asin returns the arc sin in the range $-\pi/2$ to $\pi/2$.

Acos returns the arc cosine in the range 0 to π .

Atan returns the arc tangent of *x* in the range $-\pi/2$ to $\pi/2$.

Atan2 returns the arc tangent of *xy* in the range $-\pi$ to π .

DIAGNOSTICS

Arguments of magnitude greater than 1 cause *asin* and *acos* to return value 0; *errno* is set to EDOM. The value of *tan* at its singular points is a huge number, and *errno* is set to ERANGE.

BUGS

The value of *tan* for arguments greater than about $2 \cdot 10^{31}$ is garbage.

NAME

sinh, cosh, tanh — hyperbolic functions

SYNOPSIS

```
#include <math.h>
```

```
double sinh(x)
```

```
double cosh(x)
```

```
double x;
```

```
double tanh(x)
```

```
double x;
```

DESCRIPTION

These functions compute the designated hyperbolic functions for real arguments.

DIAGNOSTICS

Sinh and *cosh* return a huge value of appropriate sign when the correct value would overflow.

NAME

intro — introduction to network library functions

DESCRIPTION

This section describes functions that are applicable to the DARPA Internet network.

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
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 routines
inet_lnaof	inet.3n	Internet address manipulation routines
inet_makeaddr	inet.3n	Internet address manipulation routines
inet_netof	inet.3n	Internet address manipulation routines
inet_network	inet.3n	Internet address manipulation routines
ntohl	byteorder.3n	convert values between host and network byte order
ntohs	byteorder.3n	convert values between host and network byte order
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

NAME

htonl, htons, ntohl, ntohs — convert values between host and network byte order

SYNOPSIS

```
#include <sys/types.h>
#include <netinet/in.h>
netlong = htonl(hostlong);
u_long netlong, hostlong;
netshort = htons(hostshort);
u_short netshort, hostshort;
hostlong = ntohl(netlong);
u_long hostlong, netlong;
hostshort = ntohs(netshort);
u_short hostshort, netshort;
```

DESCRIPTION

These routines convert 16 and 32 bit quantities between network byte order and host byte order. On machines such as the SUN these routines are defined as null macros in the include file *<netinet/in.h>*.

These routines are most often used in conjunction with Internet addresses and ports as returned by *gethostent(3N)* and *getservent(3N)*.

SEE ALSO

gethostent(3N), *getservent(3N)*

BUGS

The VAX handles bytes backwards from most everyone else in the world. This is not expected to be fixed in the near future.

NAME

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

SYNOPSIS

```
#include <netdb.h>

struct hostent *gethostent()

struct hostent *gethostbyname(name)
char *name;

struct hostent *gethostbyaddr(addr, len, type)
char *addr; int len, type;

sethostent(stayopen)
int stayopen

endhostent()
```

DESCRIPTION

Gethostent, *gethostbyname*, and *gethostbyaddr* each return a pointer to an object with the following structure containing the broken-out fields of a line in the network host data base, */etc/hosts*.

```
struct hostent {
    char *h_name; /* official name of host */
    char **h_aliases; /* alias list */
    int h_addrtype; /* address type */
    int h_length; /* length of address */
    char *h_addr; /* address */
};
```

The members of this structure are:

h_name Official name of the host.

h_aliases A zero terminated array of alternate names for the host.

h_addrtype The type of address being returned; currently always AF_INET.

h_length The length, in bytes, of the address.

h_addr A pointer to the network address for the host. Host addresses are returned in network byte order.

Gethostent reads the next line of the file, opening the file if necessary.

Sethostent opens and rewinds the file. If the *stayopen* flag is non-zero, the host data base will not be closed after each call to *gethostent* (either directly, or indirectly through one of the other "gethost" calls).

Endhostent closes the file.

Gethostbyname and *gethostbyaddr* sequentially search from the beginning of the file until a matching host name or host address is found, or until EOF is encountered. Host addresses are supplied in network order.

FILES

/etc/hosts

SEE ALSO

hosts(5)

DIAGNOSTICS

Null pointer (0) returned on EOF or error.

BUGS

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

NAME

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

SYNOPSIS

```
#include <netdb.h>

struct netent *getnetent()

struct netent *getnetbyname(name)
char *name;

struct netent *getnetbyaddr(net)
long net;

setnetent(stayopen)
int stayopen

endnetent()
```

DESCRIPTION

Getnetent, *getnetbyname*, and *getnetbyaddr* each return a pointer to an object with the following structure containing the broken-out fields of a line in the network data base, */etc/networks*.

```
struct netent {
    char   *n_name;      /* official name of net */
    char  **n_aliases;   /* alias list */
    int    n_addrtype;   /* net number type */
    long   n_net;        /* net number */
};
```

The members of this structure are:

- n_name** The official name of the network.
- n_aliases** A zero terminated list of alternate names for the network.
- n_addrtype** The type of the network number returned; currently only AF_INET.
- n_net** The network number. Network numbers are returned in machine byte order.

Getnetent reads the next line of the file, opening the file if necessary.

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

Endnetent closes the file.

Getnetbyname and *getnetbyaddr* sequentially search from the beginning of the file until a matching net name or net address is found, or until EOF is encountered. Network numbers are supplied in host order.

FILES

/etc/networks

SEE ALSO

networks(5)

DIAGNOSTICS

Null pointer (0) returned on EOF or error.

BUGS

All information is contained in a static area so it must be copied if it is to be saved. Only Internet network numbers are currently understood. Expecting network numbers to fit in no more than 32 bits is probably naive.

NAME

getprotoent, getprotobynumber, getprotobynname, setprotoent, endprotoent — get protocol entry

SYNOPSIS

```
#include <netdb.h>

struct protoent *getprotoent()
struct protoent *getprotobynname(name)
char *name;
struct protoent *getprotobynumber(proto)
int proto;
setprotoent(stayopen)
int stayopen
endprotoent()
```

DESCRIPTION

Getprotoent, *getprotobynname*, and *getprotobynumber* each return a pointer to an object with the following structure containing the broken-out fields of a line in the network protocol data base, */etc/protocols*.

```
struct protoent {
    char *p_name;      /* official name of protocol */
    char **p_aliases; /* alias list */
    long p_proto;     /* protocol number */
};
```

The members of this structure are:

p_name The official name of the protocol.
p_aliases A zero terminated list of alternate names for the protocol.
p_proto The protocol number.

Getprotoent reads the next line of the file, opening the file if necessary.

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

Endprotoent closes the file.

Getprotobynname and *getprotobynumber* sequentially search from the beginning of the file until a matching protocol name or protocol number is found, or until EOF is encountered.

FILES

/etc/protocols

SEE ALSO

protocols(5)

DIAGNOSTICS

Null pointer (0) returned on EOF or error.

BUGS

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

NAME

getservent, getservbyport, getservbyname, setservent, endservent — get service entry

SYNOPSIS

```
#include <netdb.h>

struct servent *getservent()

struct servent *getservbyname(name, proto)
char *name, *proto;

struct servent *getservbyport(port, proto)
int port; char *proto;

setservent(stayopen)
int stayopen

endservent()
```

DESCRIPTION

Getservent, *getservbyname*, and *getservbyport* each return a pointer to an object with the following structure containing the broken-out fields of a line in the network services data base, */etc/services*.

```
struct servent {
    char   *s_name;      /* official name of service */
    char   **s_aliases; /* alias list */
    long   s_port;      /* port service resides at */
    char   *s_proto;    /* protocol to use */
};
```

The members of this structure are:

- s_name** The official name of the service.
- s_aliases** A zero terminated list of alternate names for the service.
- s_port** The port number at which the service resides. Port numbers are returned in network byte order.
- s_proto** The name of the protocol to use when contacting the service.

Getservent reads the next line of the file, opening the file if necessary.

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

Endservent closes the file.

Getservbyname and *getservbyport* sequentially search from the beginning of the file until a matching protocol name or port number is found, or until EOF is encountered. If a protocol name is also supplied (non-NULL), searches must also match the protocol.

FILES

/etc/services

SEE ALSO

getprotoent(3N), services(5)

DIAGNOSTICS

Null pointer (0) returned on EOF or error.

BUGS

All information is contained in a static area so it must be copied if it is to be saved. Expecting port numbers to fit in a 32 bit quantity is probably naive.

NAME

`inet_addr`, `inet_network`, `inet_ntoa`, `inet_makeaddr`, `inet_lnaof`, `inet_netof` — Internet address manipulation routines

SYNOPSIS

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

struct in_addr inet_addr(cp)
char *cp;

int inet_network(cp)
char *cp;

char *inet_ntoa(in)
struct in_addr in;

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

int inet_lnaof(in)
struct in_addr in;

int inet_netof(in)
struct in_addr in;
```

DESCRIPTION

The routines `inet_addr` and `inet_network` each interpret character strings representing numbers expressed in the Internet standard “.” notation, returning numbers suitable for use as Internet addresses and Internet network numbers, respectively. The routine `inet_ntoa` takes an Internet address and returns an ASCII string representing the address in “.” notation. The routine `inet_makeaddr` takes an Internet network number and a local network address and constructs an Internet address from it. The routines `inet_netof` and `inet_lnaof` break apart Internet host addresses, returning the network number and local network address part, respectively.

All Internet address are returned in network order (bytes ordered from left to right). All network numbers and local address parts are returned as machine format integer values.

INTERNET ADDRESSES

Values specified using the “.” notation take one of the following forms:

```
a.b.c.d
a.b.c
a.b
a
```

When four parts are specified, each is interpreted as a byte of data and assigned, from left to right, to the four bytes of an Internet address. Note that when an Internet address is viewed as a 32-bit integer quantity on the VAX the bytes referred to above appear as “d.c.b.a”. That is, VAX bytes are ordered from right to left.

When a three part address is specified, the last part is interpreted as a 16-bit quantity and placed in the right most two bytes of the network address. This makes the three part address format convenient for specifying Class B network addresses as “128.net.host”.

When a two part address is supplied, the last part is interpreted as a 24-bit quantity and placed in the right most three bytes of the network address. This makes the two part address format convenient for specifying Class A network addresses as “net.host”.

When only one part is given, the value is stored directly in the network address without any byte rearrangement.

All numbers supplied as "parts" in a "." notation may be decimal, octal, or hexadecimal, as specified in the C language (i.e. a leading 0x or 0X implies hexadecimal; otherwise, a leading 0 implies octal; otherwise, the number is interpreted as decimal).

SEE ALSO

gethostent(3N), getnetent(3N), hosts(5), networks(5),

DIAGNOSTICS

The value -1 is returned by *inet_addr* and *inet_network* for malformed requests.

BUGS

The problem of host byte ordering versus network byte ordering is confusing. A simple way to specify Class C network addresses in a manner similar to that for Class B and Class A is needed. The string returned by *inet_ntoa* resides in a static memory area.

NAME

stdio — standard buffered input/output package

SYNOPSIS

```
#include <stdio.h>
FILE *stdin;
FILE *stdout;
FILE *stderr;
```

DESCRIPTION

The functions described in section 3S constitute a user-level buffering scheme. The in-line macros *getc* and *putc*(3S) handle characters quickly. The higher level routines *gets*, *fgets*, *scanf*, *fscanf*, *fread*, *puts*, *fputs*, *printf*, *fprintf*, *fwrite* all use *getc* and *putc*; they can be freely intermixed.

A file with associated buffering is called a *stream*, and is declared to be a pointer to a defined type **FILE**. *Fopen*(3S) creates certain descriptive data for a stream and returns a pointer to designate the stream in all further *transactions*. There are three normally open streams with constant pointers declared in the include file and associated with the standard open files:

```
stdin    standard input file
stdout   standard output file
stderr   standard error file
```

A constant 'pointer' **NULL** (0) designates no stream at all.

An integer constant **EOF** (-1) is returned upon end of file or error by integer functions that deal with streams.

Any routine that uses the standard input/output package must include the header file *<stdio.h>* of pertinent macro definitions. The functions and constants mentioned in sections labeled 3S are declared in the include file and need no further declaration. The constants, and the following 'functions' are implemented as macros; redeclaration of these names is perilous: *getc*, *getchar*, *putc*, *putchar*, *feof*, *error*, *fileno*.

SEE ALSO

open(2), *close*(2), *read*(2), *write*(2), *fread*(3S), *fseek*(3S), *f**(3S)

DIAGNOSTICS

The value **EOF** is returned uniformly to indicate that a **FILE** pointer has not been initialized with *fopen*, input (output) has been attempted on an output (input) stream, or a **FILE** pointer designates corrupt or otherwise unintelligible **FILE** data.

For purposes of efficiency, this implementation of the standard library has been changed to line buffer output to a terminal by default and attempts to do this transparently by flushing the output whenever a *read*(2) from the standard input is necessary. This is almost always transparent, but may cause confusion or malfunctioning of programs which use standard i/o routines but use *read*(2) themselves to read from the standard input.

In cases where a large amount of computation is done after printing part of a line on an output terminal, it is necessary to *fflush*(3S) the standard output before going off and computing so that the output will appear.

BUGS

The standard buffered functions do not interact well with certain other library and system functions, especially *vfork* and *abort*.

LIST OF FUNCTIONS

<i>Name</i>	<i>Appears on Page</i>	<i>Description</i>
clearerr	error.3s	stream status inquiries
fclose	fclose.3s	close or flush a stream

feof	error.3s	stream status inquiries
ferror	error.3s	stream status inquiries
fflush	fclose.3s	close or flush a stream
fgetc	getc.3s	get character or word from stream
fgets	gets.3s	get a string from a stream
fileno	error.3s	stream status inquiries
fprintf	printf.3s	formatted output conversion
fputc	putc.3s	put character or word on a stream
fputs	puts.3s	put a string on a stream
fread	fread.3s	buffered binary input/output
fscanf	scanf.3s	formatted input conversion
fseek	fseek.3s	reposition a stream
ftell	fseek.3s	reposition a stream
fwrite	fread.3s	buffered binary input/output
getc	getc.3s	get character or word from stream
getchar	getc.3s	get character or word from stream
gets	gets.3s	get a string from a stream
getw	getc.3s	get character or word from stream
printf	printf.3s	formatted output conversion
putc	putc.3s	put character or word on a stream
putchar	putc.3s	put character or word on a stream
puts	puts.3s	put a string on a stream
putw	putc.3s	put character or word on a stream
rewind	fseek.3s	reposition a stream
scanf	scanf.3s	formatted input conversion
setbuf	setbuf.3s	assign buffering to a stream
setbuffer	setbuf.3s	assign buffering to a stream
setlinebuf	setbuf.3s	assign buffering to a stream
sprintf	printf.3s	formatted output conversion
sscanf	scanf.3s	formatted input conversion
ungetc	ungetc.3s	push character back into input stream

NAME

fclose, *fflush* — close or flush a stream

SYNOPSIS

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

```
fclose(stream)
```

```
FILE *stream;
```

```
fflush(stream)
```

```
FILE *stream;
```

DESCRIPTION

Fclose causes any buffers for the named *stream* to be emptied, and the file to be closed. Buffers allocated by the standard input/output system are freed.

Fclose is performed automatically upon calling *exit*(3).

Fflush causes any buffered data for the named output *stream* to be written to that file. The stream remains open.

SEE ALSO

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

DIAGNOSTICS

These routines return EOF if *stream* is not associated with an output file, or if buffered data cannot be transferred to that file.

NAME

feof, *feof*, *clearerr*, *fileno* -- stream status inquiries

SYNOPSIS

```
#include <stdio.h>

feof(stream)
FILE *stream;

ferror(stream)
FILE *stream

clearerr(stream)
FILE *stream

fileno(stream)
FILE *stream;
```

DESCRIPTION

Feof returns non-zero when end of file is read on the named input *stream*, otherwise zero.

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

Cleerr resets the error indication on the named *stream*.

Fileno returns the integer file descriptor associated with the *stream*, see *open(2)*.

These functions are implemented as macros; they cannot be redeclared.

SEE ALSO

fopen(3S), *open(2)*

3

NAME

fopen, *freopen*, *fdopen* — open a stream

SYNOPSIS

```
#include <stdio.h>

FILE *fopen(filename, type)
char *filename, *type;

FILE *freopen(filename, type, stream)
char *filename, *type;
FILE *stream;

FILE *fdopen(fildev, type)
char *type;
```

DESCRIPTION

Fopen opens the file named by *filename* and associates a stream with it. *Fopen* returns a pointer to be used to identify the stream in subsequent operations.

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

"r" open for reading

"w" create for writing

"a" append: open for writing at end of file, or create for writing

In addition, each *type* may be followed by a '+' to have the file opened for reading and writing. "r+" positions the stream at the beginning of the file, "w+" creates or truncates it, and "a+" positions it at the end. Both reads and writes may be used on read/write streams, with the limitation that an *fseek*, *rewind*, or reading an end-of-file must be used between a read and a write or vice-versa.

Freopen substitutes the named file in place of the open *stream*. It returns the original value of *stream*. The original stream is closed.

Freopen is typically used to attach the preopened constant names, *stdin*, *stdout*, *stderr*, to specified files.

Fdopen associates a stream with a file descriptor obtained from *open*, *dup*, *creat*, or *pipe(2)*. The *type* of the stream must agree with the mode of the open file.

SEE ALSO

open(2), *fclose(3)*

DIAGNOSTICS

Fopen and *freopen* return the pointer NULL if *filename* cannot be accessed.

BUGS

Fdopen is not portable to systems other than UNIX.

The read/write *types* do not exist on all systems. Those systems without read/write modes will probably treat the *type* as if the '+' was not present. These are unreliable in any event.

NAME

fread, *fwrite* — buffered binary input/output

SYNOPSIS

```
#include <stdio.h>

fread(ptr, sizeof(*ptr), nitems, stream)
FILE *stream;

fwrite(ptr, sizeof(*ptr), nitems, stream)
FILE *stream;
```

DESCRIPTION

Fread reads, into a block beginning at *ptr*, *nitems* of data of the type of **ptr* from the named input *stream*. It returns the number of items actually read.

If *stream* is *stdin* and the standard output is line buffered, then any partial output line will be flushed before any call to *read(2)* to satisfy the *fread*.

Fwrite appends at most *nitems* of data of the type of **ptr* beginning at *ptr* to the named output *stream*. It returns the number of items actually written.

SEE ALSO

read(2), *write(2)*, *fopen(3S)*, *getc(3S)*, *putc(3S)*, *gets(3S)*, *puts(3S)*, *printf(3S)*, *scanf(3S)*

DIAGNOSTICS

Fread and *fwrite* return 0 upon end of file or error.

NAME

fseek, *ftell*, *rewind* — reposition a stream

SYNOPSIS

```
#include <stdio.h>
fseek(stream, offset, ptrname)
FILE *stream;
long offset;
long ftell(stream)
FILE *stream;
rewind(stream)
```

DESCRIPTION

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

Fseek undoes any effects of *ungetc*(3S).

Ftell returns the current value of the offset relative to the beginning of the file associated with the named *stream*. It is measured in bytes on UNIX; on some other systems it is a magic cookie, and the only foolproof way to obtain an *offset* for *fseek*.

Rewind(*stream*) is equivalent to *fseek*(*stream*, 0L, 0).

SEE ALSO

lseek(2), *fopen*(3S)

DIAGNOSTICS

Fseek returns -1 for improper seeks.

NAME

`getc`, `getchar`, `fgetc`, `getw` — get character or word from stream

SYNOPSIS

```
#include <stdio.h>

int getc(stream)
FILE *stream;

int getchar()

int fgetc(stream)
FILE *stream;

int getw(stream)
FILE *stream;
```

DESCRIPTION

getc returns the next character from the named input *stream*.

Getchar() is identical to *getc(stdin)*.

Fgetc behaves like *getc*, but is a genuine function, not a macro; it may be used to save object text.

Getw returns the next word (in a 32-bit integer on a VAX-11) from the named input *stream*. It returns the constant **EOF** upon end of file or error, but since that is a good integer value, *feof* and *feof(3S)* should be used to check the success of *getw*. *Getw* assumes no special alignment in the file.

SEE ALSO

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

DIAGNOSTICS

These functions return the integer constant **EOF** at end of file or upon read error.

A stop with message, 'Reading bad file', means an attempt has been made to read from a stream that has not been opened for reading by *fopen*.

BUGS

The end-of-file return from *getchar* is incompatible with that in UNIX editions 1-6.

Because it is implemented as a macro, *getc* treats a *stream* argument with side effects incorrectly. In particular, '`getc(*f++);`' doesn't work sensibly.

NAME

gets, *fgets* — get a string from a stream

SYNOPSIS

```
#include <stdio.h>
char *gets(s)
char *s;
char *fgets(s, n, stream)
char *s;
FILE *stream;
```

DESCRIPTION

Gets reads a string into *s* from the standard input stream **stdin**. The string is terminated by a newline character, which is replaced in *s* by a null character. *Gets* returns its argument.

Fgets reads *n* - 1 characters, or up to a newline character, whichever comes first, from the *stream* into the string *s*. The last character read into *s* is followed by a null character. *Fgets* returns its first argument.

SEE ALSO

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

DIAGNOSTICS

Gets and *fgets* return the constant pointer **NULL** upon end of file or error.

BUGS

Gets deletes a newline, *fgets* keeps it, all in the name of backward compatibility.

NAME

printf, fprintf, sprintf – formatted output conversion

SYNOPSIS

```
#include <stdio.h>
printf(format [, arg ] ... )
char *format;
fprintf(stream, format [, arg ] ... )
FILE *stream;
char *format;
sprintf(s, format [, arg ] ... )
char *s, format;
#include <varargs.h>
_doprnt(format, args, stream)
char *format;
va_list *args;
FILE *stream;
```

DESCRIPTION

Printf places output on the standard output stream **stdout**. *Fprintf* places output on the named output *stream*. *Sprintf* places ‘output’ in the string *s*, followed by the character ‘\0’. All of these routines work by calling the internal routine **_doprnt**, using the variable-length argument facilities of *varargs*(3).

Each of these functions converts, formats, and prints its arguments after the first under control of the first argument. The first argument is a character string which contains two types of objects: plain characters, which are simply copied to the output stream, and conversion specifications, each of which causes conversion and printing of the next successive *arg printf*.

Each conversion specification is introduced by the character **%**. Following the **%**, there may be

- an optional minus sign ‘-’ which specifies *left adjustment* of the converted value in the indicated field;
- an optional digit string specifying a *field width*; if the converted value has fewer characters than the field width it will be blank-padded on the left (or right, if the left-adjustment indicator has been given) to make up the field width; if the field width begins with a zero, zero-padding will be done instead of blank-padding;
- an optional period ‘.’ which serves to separate the field width from the next digit string;
- an optional digit string specifying a *precision* which specifies the number of digits to appear after the decimal point, for *e*- and *f*-conversion, or the maximum number of characters to be printed from a string;
- an optional ‘#’ character specifying that the value should be converted to an “alternate form”. For *c*, *d*, *s*, and *u*, conversions, this option has no effect. For *o* conversions, the precision of the number is increased to force the first character of the output string to a zero. For *x*(*X*) conversion, a non-zero result has the string **0x**(**0X**) prepended to it. For *e*, *E*, *f*, *g*, and *G*, conversions, the result will always contain a decimal point, even if no digits follow the point (normally, a decimal point only appears in the results of those conversions if a digit follows the decimal point). For *g* and *G* conversions, trailing zeros are not removed from the result as they would otherwise be.
- the character *l* specifying that a following *d*, *o*, *x*, or *u* corresponds to a long integer *arg*.
- a character which indicates the type of conversion to be applied.

A field width or precision may be '*' instead of a digit string. In this case an integer *arg* supplies the field width or precision.

The conversion characters and their meanings are

- d** **ox** The integer *arg* is converted to decimal, octal, or hexadecimal notation respectively.
- f** The float or double *arg* is converted to decimal notation in the style '[−]ddd.ddd' where the number of d's after the decimal point is equal to the precision specification for the argument. If the precision is missing, 6 digits are given; if the precision is explicitly 0, no digits and no decimal point are printed.
- e** The float or double *arg* is converted in the style '[−]d.ddde±dd' where there is one digit before the decimal point and the number after is equal to the precision specification for the argument; when the precision is missing, 6 digits are produced.
- g** The float or double *arg* is printed in style **d**, in style **f**, or in style **e**, whichever gives full precision in minimum space.
- c** The character *arg* is printed.
- s** *Arg* is taken to be a string (character pointer) and characters from the string are printed until a null character or until the number of characters indicated by the precision specification is reached; however if the precision is 0 or missing all characters up to a null are printed.
- u** The unsigned integer *arg* is converted to decimal and printed (the result will be in the range 0 through MAXUINT, where MAXUINT equals 4294967295 on a VAX-11 and 65535 on a PDP-11).
- %** Print a '%'; no argument is converted.

In no case does a non-existent or small field width cause truncation of a field; padding takes place only if the specified field width exceeds the actual width. Characters generated by *printf* are printed by *putc*(3S).

Examples

To print a date and time in the form 'Sunday, July 3, 10:02', where *weekday* and *month* are pointers to null-terminated strings:

```
printf("%s, %s %d, %02d:%02d", weekday, month, day, hour, min);
```

To print π to 5 decimals:

```
printf("pi = %.5f", 4*atan(1.0));
```

SEE ALSO

putc(3S), *scanf*(3S), *ecvt*(3)

BUGS

Very wide fields (>128 characters) fail.

NAME

`putc`, `putchar`, `fputc`, `putw` — put character or word on a stream

SYNOPSIS

```
#include <stdio.h>

int putc(c, stream)
char c;
FILE *stream;

putc(c)

fputc(c, stream)
FILE *stream;

putw(w, stream)
FILE *stream;
```

DESCRIPTION

Putc appends the character *c* to the named output *stream*. It returns the character written.

Putchar(c) is defined as *putc(c, stdout)*.

Fputc behaves like *putc*, but is a genuine function rather than a macro.

Putw appends word (that is, `int`) *w* to the output *stream*. It returns the word written. *Putw* neither assumes nor causes special alignment in the file.

SEE ALSO

`fopen(3S)`, `fclose(3S)`, `getc(3S)`, `puts(3S)`, `printf(3S)`, `fread(3S)`

DIAGNOSTICS

These functions return the constant `EOF` upon error. Since this is a good integer, `feof(3S)` should be used to detect *putw* errors.

BUGS

Because it is implemented as a macro, *putc* treats a *stream* argument with side effects improperly. In particular

```
putc(c, *f++);
```

doesn't work sensibly.

Errors can occur long after the call to *putc*.

NAME

`puts`, `fputs` — put a string on a stream

SYNOPSIS

```
#include <stdio.h>

puts(s)
char *s;

fputs(s, stream)
char *s;
FILE *stream;
```

DESCRIPTION

Puts copies the null-terminated string *s* to the standard output stream `stdout` and appends a newline character.

Fputs copies the null-terminated string *s* to the named output *stream*.

Neither routine copies the terminal null character.

SEE ALSO

`fopen(3S)`, `gets(3S)`, `putc(3S)`, `printf(3S)`, `ferror(3S)`
`fread(3S)` for *fwrite*

BUGS

Puts appends a newline, *fputs* does not, all in the name of backward compatibility.

NAME

scanf, fscanf, sscanf — formatted input conversion

SYNOPSIS

```
#include <stdio.h>

scanf(format [ , pointer ] . . . )
char *format;

fscanf(stream, format [ , pointer ] . . . )
FILE *stream;
char *format;

sscanf(s, format [ , pointer ] . . . )
char *s, *format;
```

DESCRIPTION

Scanf reads from the standard input stream *stdin*. *Fscanf* reads from the named input *stream*. *Sscanf* reads from the character string *s*. Each function reads characters, interprets them according to a format, and stores the results in its arguments. Each expects as arguments a control string *format*, described below, and a set of *pointer* arguments indicating where the converted input should be stored.

The control string usually contains conversion specifications, which are used to direct interpretation of input sequences. The control string may contain:

1. Blanks, tabs or newlines, which match optional white space in the input.
2. An ordinary character (not %) which must match the next character of the input stream.
3. Conversion specifications, consisting of the character %, an optional assignment suppressing character *, an optional numerical maximum field width, and a conversion character.

A conversion specification directs the conversion of the next input field; the result is placed in the variable pointed to by the corresponding argument, unless assignment suppression was indicated by *. An input field is defined as a string of non-space characters; it extends to the next inappropriate character or until the field width, if specified, is exhausted.

The conversion character indicates the interpretation of the input field; the corresponding pointer argument must usually be of a restricted type. The following conversion characters are legal:

- %** a single '%' is expected in the input at this point; no assignment is done.
- d** a decimal integer is expected; the corresponding argument should be an integer pointer.
- o** an octal integer is expected; the corresponding argument should be a integer pointer.
- x** a hexadecimal integer is expected; the corresponding argument should be an integer pointer.
- s** a character string is expected; the corresponding argument should be a character pointer pointing to an array of characters large enough to accept the string and a terminating '\0', which will be added. The input field is terminated by a space character or a newline.
- c** a character is expected; the corresponding argument should be a character pointer. The normal skip over space characters is suppressed in this case; to read the next non-space character, try '%1s'. If a field width is given, the corresponding argument should refer to a character array, and the indicated number of characters is read.
- e** a floating point number is expected; the next field is converted accordingly and stored through the corresponding argument, which should be a pointer to a *float*. The input format for floating point numbers is an optionally signed string of digits possibly containing a decimal point, followed by an optional exponent field consisting of an E or e followed by

an optionally signed integer.

- [indicates a string not to be delimited by space characters. The left bracket is followed by a set of characters and a right bracket; the characters between the brackets define a set of characters making up the string. If the first character is not circumflex (^), the input field is all characters until the first character not in the set between the brackets; if the first character after the left bracket is ^, the input field is all characters until the first character which is in the remaining set of characters between the brackets. The corresponding argument must point to a character array.

The conversion characters **d**, **o** and **x** may be capitalized or preceded by **l** to indicate that a pointer to **long** rather than to **int** is in the argument list. Similarly, the conversion characters **e** or **f** may be capitalized or preceded by **l** to indicate a pointer to **double** rather than to **float**. The conversion characters **d**, **o** and **x** may be preceded by **h** to indicate a pointer to **short** rather than to **int**.

The *scanf* functions return the number of successfully matched and assigned input items. This can be used to decide how many input items were found. The constant **EOF** is returned upon end of input; note that this is different from 0, which means that no conversion was done; if conversion was intended, it was frustrated by an inappropriate character in the input.

For example, the call

```
int i; float x; char name[50];
scanf("%d%f%s", &i, &x, name);
```

with the input line

```
25 54.32E-1 thompson
```

will assign to *i* the value 25, *x* the value 5.432, and *name* will contain 'thompson\0'. Or,

```
int i; float x; char name[50];
scanf("%2d%f%*d%[1234567890]", &i, &x, name);
```

with input

```
56789 0123 56a72
```

will assign 56 to *i*, 789.0 to *x*, skip '0123', and place the string '56\0' in *name*. The next call to *getchar* will return 'a'.

SEE ALSO

atof(3), *getc*(3S), *printf*(3S)

DIAGNOSTICS

The *scanf* functions return **EOF** on end of input, and a short count for missing or illegal data items.

BUGS

The success of literal matches and suppressed assignments is not directly determinable.

NAME

`setbuf`, `setbuffer`, `setlinebuf` — assign buffering to a stream

SYNOPSIS

```
#include <stdio.h>
setbuf(stream, buf)
FILE *stream;
char *buf;

setbuffer(stream, buf, size)
FILE *stream;
char *buf;
int size;

setlinebuf(stream)
FILE *stream;
```

DESCRIPTION

The three types of buffering available are unbuffered, block buffered, and line buffered. When an output stream is unbuffered, information appears on the destination file or terminal as soon as written; when it is block buffered many characters are saved up and written as a block; when it is line buffered characters are saved up until a newline is encountered or input is read from `stdin`. *Eflush* (see `fclose(3S)`) may be used to force the block out early. Normally all files are block buffered. A buffer is obtained from `malloc(3)` upon the first `getc` or `putc(3S)` on the file. If the standard stream `stdout` refers to a terminal it is line buffered. The standard stream `stderr` is always unbuffered.

Setbuf is used after a stream has been opened but before it is read or written. The character array *buf* is used instead of an automatically allocated buffer. If *buf* is the constant pointer `NULL`, input/output will be completely unbuffered. A manifest constant `BUFSIZ` tells how big an array is needed:

```
char buf[BUFSIZ];
```

Setbuffer, an alternate form of *setbuf*, is used after a stream has been opened but before it is read or written. The character array *buf* whose size is determined by the *size* argument is used instead of an automatically allocated buffer. If *buf* is the constant pointer `NULL`, input/output will be completely unbuffered.

Setlinebuf is used to change `stdout` or `stderr` from block buffered or unbuffered to line buffered. Unlike *setbuf* and *setbuffer* it can be used at any time that the file descriptor is active.

A file can be changed from unbuffered or line buffered to block buffered by using *freopen* (see `fopen(3S)`). A file can be changed from block buffered or line buffered to unbuffered by using *freopen* followed by *setbuf* with a buffer argument of `NULL`.

SEE ALSO

`fopen(3S)`, `getc(3S)`, `putc(3S)`, `malloc(3)`, `fclose(3S)`, `puts(3S)`, `printf(3S)`, `fread(3S)`

BUGS

The standard error stream should be line buffered by default.

The *setbuffer* and *setlinebuf* functions are not portable to non 4.2 BSD versions of UNIX.

NAME

`ungetc` — push character back into input stream

SYNOPSIS

```
#include <stdio.h>
ungetc(c, stream)
FILE *stream;
```

DESCRIPTION

Ungetc pushes the character *c* back on an input stream. That character will be returned by the next *getc* call on that stream. *Ungetc* returns *c*.

One character of pushback is guaranteed provided something has been read from the stream and the stream is actually buffered. Attempts to push EOF are rejected.

Fseek(3S) erases all memory of pushed back characters.

SEE ALSO

getc(3S), *setbuf*(3S), *fseek*(3S)

DIAGNOSTICS

Ungetc returns EOF if it can't push a character back.

NAME

intro — introduction to miscellaneous library functions

DESCRIPTION

These functions constitute minor libraries and other miscellaneous run-time facilities. Most are available only when programming in C. The list below includes libraries which provide device independent plotting functions, terminal independent screen management routines for two dimensional non-bitmap display terminals, functions for managing data bases with inverted indexes, and sundry routines used in executing commands on remote machines. The routines *getdiskbyname*, *rcmd*, *rresvport*, *ruserok*, and *rexec* reside in the standard C run-time library “-lc”. All other functions are located in separate libraries indicated in each manual entry.

FILES

/lib/libc.a
 /usr/lib/libdbm.a
 /usr/lib/libtermcap.a
 /usr/lib/libcurses.a
 /usr/lib/lib2648.a
 /usr/lib/libplot.a

LIST OF FUNCTIONS

<i>Name</i>	<i>Appears on Page</i>	<i>Description</i>
arc	plot.3x	graphics interface
assert	assert.3x	program verification
circle	plot.3x	graphics interface
closepl	plot.3x	graphics interface
cont	plot.3x	graphics interface
curses	curses.3x	screen functions with “optimal” cursor motion
dbminit	dbm.3x	data base subroutines
delete	dbm.3x	data base subroutines
endsent	getfsent.3x	get file system descriptor file entry
erase	plot.3x	graphics interface
fetch	dbm.3x	data base subroutines
firstkey	dbm.3x	data base subroutines
getdiskbyname	getdisk.3x	get disk description by its name
getfsent	getfsent.3x	get file system descriptor file entry
getfsfile	getfsent.3x	get file system descriptor file entry
getfsspec	getfsent.3x	get file system descriptor file entry
getfstype	getfsent.3x	get file system descriptor file entry
initgroups	initgroups.3x	initialize group access list
label	plot.3x	graphics interface
lib2648	lib2648.3x	subroutines for the HP 2648 graphics terminal
line	plot.3x	graphics interface
linemod	plot.3x	graphics interface
move	plot.3x	graphics interface
nextkey	dbm.3x	data base subroutines
plot: openpl	plot.3x	graphics interface
point	plot.3x	graphics interface
rcmd	rcmd.3x	routines for returning a stream to a remote command
rexec	rexec.3x	return stream to a remote command
rresvport	rcmd.3x	routines for returning a stream to a remote command
ruserok	rcmd.3x	routines for returning a stream to a remote command
setfsent	getfsent.3x	get file system descriptor file entry
space	plot.3x	graphics interface

store	dbm.3x	data base subroutines
tgetent	termcap.3x	terminal independent operation routines
tgetflag	termcap.3x	terminal independent operation routines
tgetnum	termcap.3x	terminal independent operation routines
tgetstr	termcap.3x	terminal independent operation routines
tgoto	termcap.3x	terminal independent operation routines
tputs	termcap.3x	terminal independent operation routines

NAME

`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

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

Screen Updating and Cursor Movement Optimization: A Library Package, Ken Arnold, ioctl(2), getenv(3), tty(4), termcap(5)

AUTHOR

Ken Arnold

FUNCTIONS

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
crmode()	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>
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
nocrmode()	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

<code>printw(fmt,arg1,arg2,...)</code>	<code>printf</code> on <i>stdscr</i>
<code>raw()</code>	set raw mode
<code>refresh()</code>	make current screen look like <i>stdscr</i>
<code>resetty()</code>	reset tty flags to stored value
<code>savetty()</code>	stored current tty flags
<code>scanw(fmt,arg1,arg2,...)</code>	<code>scanf</code> through <i>stdscr</i>
<code>scroll(win)</code>	scroll <i>win</i> one line
<code>scrollok(win,boolf)</code>	set scroll flag
<code>setterm(name)</code>	set term variables for name
<code>standend()</code>	end standout mode
<code>standout()</code>	start standout mode
<code>subwin(win,lines,cols,begin_y,begin_x)</code>	create a subwindow
<code>touchwin(win)</code>	"change" all of <i>win</i>
<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>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>
<code>wdelch(win,c)</code>	delete char from <i>win</i>
<code>wdeleteln(win)</code>	delete line from <i>win</i>
<code>werase(win)</code>	erase <i>win</i>
<code>wgetch(win)</code>	get a char through <i>win</i>
<code>wgetstr(win,str)</code>	get a string through <i>win</i>
<code>winch(win)</code>	get char at current (y,x) in <i>win</i>
<code>winsch(win,c)</code>	insert char into <i>win</i>
<code>winsertln(win)</code>	insert line into <i>win</i>
<code>wmove(win,y,x)</code>	set current (y,x) co-ordinates on <i>win</i>
<code>wprintw(win,fmt,arg1,arg2,...)</code>	<code>printf</code> on <i>win</i>
<code>wrefresh(win)</code>	make screen look like <i>win</i>
<code>wscanw(win,fmt,arg1,arg2,...)</code>	<code>scanf</code> through <i>win</i>
<code>wstandend(win)</code>	end standout mode on <i>win</i>
<code>wstandout(win)</code>	start standout mode on <i>win</i>

BUGS

NAME

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;
```

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))
```

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.

NAME

getdiskbyname — get disk description by its name

SYNOPSIS

```
#include <disktab.h>

struct disktab *
getdiskbyname(name)
char *name;
```

DESCRIPTION

Getdiskbyname takes a disk name (e.g. rm03) and returns a structure describing its geometry information and the standard disk partition tables. All information obtained from the *disktab(5)* file.

<*disktab.h*> has the following form:

```
/*      @(#)disktab.h 4.2 (Berkeley) 3/6/83  */

/*
 * Disk description table, see disktab(5)
 */
#define DISKTAB          "/etc/disktab"

struct disktab {
    char  *d_name;           /* drive name */
    char  *d_type;          /* drive type */
    int   d_sectsize;       /* sector size in bytes */
    int   d_ntracks;        /* # tracks/cylinder */
    int   d_nsectors;       /* # sectors/track */
    int   d_ncylinders;     /* # cylinders */
    int   d_rpm;            /* revolutions/minute */
    struct partition {
        int   p_size;       /* #sectors in partition */
        short p_bsize; /* block size in bytes */
        short p_fsize; /* frag size in bytes */
    } d_partitions[8];
};

struct disktab *getdiskbyname();
```

SEE ALSO

disktab(5)

BUGS

This information should be obtained from the system for locally available disks (in particular, the disk partition tables).

NAME

getfsent, getfsspec, getfsfile, getfstype, setfsent, endfsent — get file system descriptor file entry

SYNOPSIS

```
#include <fstab.h>
struct fstab *getfsent()
struct fstab *getfsspec(spec)
char *spec;
struct fstab *getfsfile(file)
char *file;
struct fstab *getfstype(type)
char *type;
int setfsent()
int endfsent()
```

DESCRIPTION

Getfsent, *getfsspec*, *getfstype*, and *getfsfile* each return a pointer to an object with the following structure containing the broken-out fields of a line in the file system description file, <fstab.h>.

```
struct fstab{
    char    *fs_spec;
    char    *fs_file;
    char    *fs_type;
    int     fs_freq;
    int     fs_passno;
};
```

The fields have meanings described in *fstab(5)*.

Getfsent reads the next line of the file, opening the file if necessary.

Setfsent opens and rewinds the file.

Endfsent closes the file.

Getfsspec and *getfsfile* sequentially search from the beginning of the file until a matching special file name or file system file name is found, or until EOF is encountered. *Getfstype* does likewise, matching on the file system type field.

FILES

/etc/fstab

SEE ALSO

fstab(5)

DIAGNOSTICS

Null pointer (0) returned on EOF or error.

BUGS

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

NAME

initgroups — initialize group access list

SYNOPSIS

```
initgroups(name, basegid)  
char *name;  
int basegid;
```

DESCRIPTION

Initgroups reads through the group file and sets up, using the *setgroups(2)* call, the group access list for the user specified in *name*. The *basegid* is automatically included in the groups list. Typically this value is given as the group number from the password file.

FILES

/etc/group

SEE ALSO

setgroups(2)

DIAGNOSTICS

Initgroups returns *-1* if it was not invoked by the super-user.

BUGS

Initgroups uses the routines based on *getgrent(3)*. If the invoking program uses any of these routines, the group structure will be overwritten in the call to *initgroups*.

Noone seems to keep */etc/group* up to date.

3

NAME

lib2648 — subroutines for the HP 2648 graphics terminal

SYNOPSIS

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

```
typedef char *bitmat;
FILE *trace;
```

```
cc file.c -l2648
```

DESCRIPTION

Lib2648 is a general purpose library of subroutines useful for interactive graphics on the Hewlett-Packard 2648 graphics terminal. To use it you must call the routine *tyinit()* at the beginning of execution, and *done()* at the end of execution. All terminal input and output must go through the routines *rawchar*, *readline*, *outchar*, and *oustr*.

Lib2648 does the necessary *E/F* handshaking if *getenv("TERM")* returns "hp2648", as it will if set by *tset(1)*. Any other value, including for example "2648", will disable handshaking.

Bit matrix routines are provided to model the graphics memory of the 2648. These routines are generally useful, but are specifically useful for the *update* function which efficiently changes what is on the screen to what is supposed to be on the screen. The primitive bit matrix routines are *newmat*, *mat*, and *setmat*.

The file *trace*, if non-null, is expected to be a file descriptor as returned by *fopen*. If so, *lib2648* will trace the progress of the output by writing onto this file. It is provided to make debugging output feasible for graphics programs without messing up the screen or the escape sequences being sent. Typical use of trace will include:

```
switch (argv[1][1]) {
case 'T':
    trace = fopen("trace", "w");
    break;
...
if (trace)
    fprintf(trace, "x is %d, y is %s\n", x, y);
...
dumpmat("before update", xmat);
```

ROUTINES**agoto(x, y)**

Move the alphanumeric cursor to position (x, y), measured from the upper left corner of the screen.

aoff() Turn the alphanumeric display off.

ason() Turn the alphanumeric display on.

areclear(rmin, cmin, rmax, cmax)

Clear the area on the graphics screen bordered by the four arguments. In normal mode the area is set to all black, in inverse video mode it is set to all white.

beep() Ring the bell on the terminal.

bitcopy(dest, src, rows, cols) bitmat dest,

Copy a *rows* by *cols* bit matrix from *src* to (user provided) *dest*.

cleara()

Clear the alphanumeric display.

clearg()

Clear the graphics display. Note that the 2648 will only clear the part of the screen that is visible if zoomed in.

curoff()

Turn the graphics cursor off.

curon()

Turn the graphics cursor on.

dispmsg(str, x, y, maxlen) char *str;

Display the message *str* in graphics text at position (x, y) . The maximum message length is given by *maxlen*, and is needed to for *dispmsg* to know how big an area to clear before drawing the message. The lower left corner of the first character is at (x, y) .

done() Should be called before the program exits. Restores the tty to normal, turns off graphics screen, turns on alphanumeric screen, flushes the standard output, etc.

draw(x, y)

Draw a line from the pen location to (x, y) . As with all graphics coordinates, (x, y) is measured from the bottom left corner of the screen. (x, y) coordinates represent the first quadrant of the usual Cartesian system.

drawbox(r, c, color, rows, cols)

Draw a rectangular box on the graphics screen. The lower left corner is at location (r, c) . The box is *rows* rows high and *cols* columns wide. The box is drawn if *color* is 1, erased if *color* is 0. (r, c) absolute coordinates represent row and column on the screen, with the origin at the lower left. They are equivalent to (x, y) except for being reversed in order.

dumpmat(msg, m, rows, cols) char *msg; bitmat m;

If *trace* is non-null, write a readable ASCII representation of the matrix *m* on *trace*. *Msg* is a label to identify the output.

emptyrow(m, rows, cols, r) bitmat m;

Returns 1 if row *r* of matrix *m* is all zero, else returns 0. This routine is provided because it can be implemented more efficiently with a knowledge of the internal representation than a series of calls to *mat*.

error(msg) char *msg;

Default error handler. Calls *message(msg)* and returns. This is called by certain routines in *lib2648*. It is also suitable for calling by the user program. It is probably a good idea for a fancy graphics program to supply its own error procedure which uses *setjmp(3)* to restart the program.

gdefault()

Set the terminal to the default graphics modes.

goff() Turn the graphics display off.

gon() Turn the graphics display on.

koff() Turn the keypad off.

kon() Turn the keypad on. This means that most special keys on the terminal (such as the alphanumeric arrow keys) will transmit an escape sequence instead of doing their function locally.

line(x1, y1, x2, y2)

Draw a line in the current mode from $(x1, y1)$ to $(x2, y2)$. This is equivalent to *move(x1, y1); draw(x2, y2);* except that a bug in the terminal involving repeated lines from the same point is compensated for.

lowleft()

Move the alphanumeric cursor to the lower left (home down) position.

mat(m, rows, cols, r, c) bitmat m;

Used to retrieve an element from a bit matrix. Returns 1 or 0 as the value of the $[r, c]$ element of the *rows* by *cols* matrix *m*. Bit matrices are numbered (*r, c*) from the upper left corner of the matrix, beginning at (0, 0). *R* represents the row, and *c* represents the column.

message(str) char *str;

Display the text message *str* at the bottom of the graphics screen.

minmax(g, rows, cols, rmin, cmin, rmax, cmax) bitmat g;**int *rmin, *cmin, *rmax, *cmax;**

Find the smallest rectangle that contains all the 1 (on) elements in the bit matrix *g*. The coordinates are returned in the variables pointed to by *rmin*, *cmin*, *rmax*, *cmax*.

move(x, y)

Move the pen to location (*x, y*). Such motion is internal and will not cause output until a subsequent *sync()*.

movecurs(x, y)

Move the graphics cursor to location (*x, y*).

bitmat newmat(rows, cols)

Create (with *malloc(3)*) a new bit matrix of size *rows* by *cols*. The value created (e.g. a pointer to the first location) is returned. A bit matrix can be freed directly with *free*.

outchar(c) char c;

Print the character *c* on the standard output. All output to the terminal should go through this routine or *outr*.

outr(str) char *str;

Print the string *str* on the standard output by repeated calls to *outchar*.

printg()

Print the graphics display on the printer. The printer must be configured as device 6 (the default) on the HPIB.

char rawchar()

Read one character from the terminal and return it. This routine or *readline* should be used to get all input, rather than *getchar(3)*.

rboff() Turn the rubber band line off.**rbon()** Turn the rubber band line on.**char *rdchar(c) char c;**

Return a readable representation of the character *c*. If *c* is a printing character it returns itself, if a control character it is shown in the ^X notation, if negative an apostrophe is prepended. Space returns ^, rubout returns ^?.

NOTE: A pointer to a static place is returned. For this reason, it will not work to pass *rdchar* twice to the same *sprintf/sprintf* call. You must instead save one of the values in your own buffer with *strcpy*.

readline(prompt, msg, maxlen) char *prompt, *msg;

Display *prompt* on the bottom line of the graphics display and read one line of text from the user, terminated by a newline. The line is placed in the buffer *msg*, which has size *maxlen* characters. Backspace processing is supported.

setclear()

Set the display to draw lines in erase mode. (This is reversed by inverse video mode.)

setmat(*m*, *rows*, *cols*, *r*, *c*, *val*) **bitmat** *m*;

The basic operation to store a value in an element of a bit matrix. The [*r*, *c*] element of *m* is set to *val*, which should be either 0 or 1.

setset(0)

Set the display to draw lines in normal (solid) mode. (This is reversed by inverse video mode.)

setxor(0)

Set the display to draw lines in exclusive or mode.

sync(0) Force all accumulated output to be displayed on the screen. This should be followed by `flush(stdout)`. The cursor is not affected by this function. Note that it is normally never necessary to call `sync`, since `rawchar` and `readline` call `sync()` and `flush(stdout)` automatically.

toggleid(0)

Toggle the state of video. If in normal mode, go into inverse video mode, and vice versa. The screen is reversed as well as the internal state of the library.

ttyinit(0)

Set up the terminal for processing. This routine should be called at the beginning of execution. It places the terminal in CBREAK mode, turns off echo, sets the proper modes in the terminal, and initializes the library.

update(*mold*, *mnew*, *rows*, *cols*, *baser*, *basec*) **bitmat** *mold*, *mnew*;

Make whatever changes are needed to make a window on the screen look like *mnew*. *Mold* is what the window on the screen currently looks like. The window has size *rows* by *cols*, and the lower left corner on the screen of the window is [*baser*, *basec*]. Note: `update` was not intended to be used for the entire screen. It would work but be very slow and take 64K bytes of memory just for *mold* and *mnew*. It was intended for 100 by 100 windows with objects in the center of them, and is quite fast for such windows.

vidinv(0)

Set inverse video mode.

vidnorm(0)

Set normal video mode.

zermat(*m*, *rows*, *cols*) **bitmat** *m*;

Set the bit matrix *m* to all zeros.

zoomn(*size*)

Set the hardware zoom to value *size*, which can range from 1 to 15.

zoomoff(0)

Turn zoom off. This forces the screen to zoom level 1 without affecting the current internal zoom number.

zoomon(0)

Turn zoom on. This restores the screen to the previously specified zoom size.

DIAGNOSTICS

The routine `error` is called when an error is detected. The only error currently detected is overflow of the buffer provided to `readline`.

Subscripts out of bounds to `setmat` return without setting anything.

FILES

/usr/lib/lib2648.a

SEE ALSO

fed(1)

AUTHOR

Mark Horton

BUGS

This library is not supported. It makes no attempt to use all of the features of the terminal, only those needed by fed. Contributions from users will be accepted for addition to the library.

The HP 2648 terminal is somewhat unreliable at speeds over 2400 baud, even with the ^E/^F handshaking. In an effort to improve reliability, handshaking is done every 32 characters. (The manual claims it is only necessary every 80 characters.) Nonetheless, I/O errors sometimes still occur.

There is no way to control the amount of debugging output generated on *trace* without modifying the source to the library.

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(1)* filters
- l300** GSI 300 terminal
- l300s** GSI 300S terminal
- l450** DASI 450 terminal
- l4014** Tektronix 4014 terminal

SEE ALSO

plot(5), *plot(1G)*, *graph(1G)*

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 1 is returned if the machine name is listed in the "hosts.equiv" file, or the host and remote user name are found in the ".rhosts" file; otherwise *ruserok* returns 0. If the *superuser* flag is 1, the checking of the "host.equiv" file is bypassed.

SEE ALSO

rlogin(1C), *rsh*(1C), *rexec*(3X), *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(3X)`, `rexecd(8C)`

BUGS

There is no way to specify options to the *socket* call which *rexec* makes.

NAME

`tgetent`, `tgetnum`, `tgetflag`, `tgetstr`, `tgoto`, `tputs` — terminal independent operation routines

SYNOPSIS

```
char PC;
char *BC;
char *UP;
short ospeed;

tgetent(bp, name)
char *bp, *name;

tgetnum(id)
char *id;

tgetflag(id)
char *id;

char *
tgetstr(id, area)
char *id, **area;

char *
tgoto(cm, destcol, destline)
char *cm;

tputs(cp, affcnt, outc)
register char *cp;
int affcnt;
int (*outc)();
```

DESCRIPTION

These functions extract and use capabilities from the terminal capability data base `termcap(5)`. These are low level routines; see `curses(3X)` for a higher level package.

`Tgetent` extracts the entry for terminal *name* into the buffer at *bp*. *Bp* should be a character buffer of size 1024 and must be retained through all subsequent calls to `tgetnum`, `tgetflag`, and `tgetstr`. `Tgetent` returns -1 if it cannot open the `termcap` file, 0 if the terminal name given does not have an entry, and 1 if all goes well. It will look in the environment for a TERMCAP variable. If found, and the value does not begin with a slash, and the terminal type *name* is the same as the environment string TERM, the TERMCAP string is used instead of reading the `termcap` file. If it does begin with a slash, the string is used as a path name rather than `letctermcap`. 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 `letctermcap`.

`Tgetnum` gets the numeric value of capability *id*, returning -1 if is not given for the terminal. `Tgetflag` returns 1 if the specified capability is present in the terminal's entry, 0 if it is not. `Tgetstr` gets the string value of capability *id*, placing it in the buffer at *area*, advancing the *area* pointer. It decodes the abbreviations for this field described in `termcap(5)`, except for cursor addressing and padding information.

`Tgoto` returns a cursor addressing string decoded from *cm* to go to column *destcol* in line *destline*. It uses the external variables UP (from the `up` capability) and BC (if `bc` is given rather than `bs`) if necessary to avoid placing `\n`, `^D` or `^@` in the returned string. (Programs 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 output speed of the terminal as encoded by *stty*(3). The external variable **PC** should contain a pad character to be used (from the **pc** capability) if a null (***Ⓢ**) is inappropriate.

FILES

`/usr/lib/libtermcap.a` — `ltermcap` library
`/etc/termcap` data base

SEE ALSO

`ex`(1), `curses`(3X), `termcap`(5)

AUTHOR

William Joy

3

NAME

intro — introduction to compatibility library functions

DESCRIPTION

These functions constitute the compatibility library portion of *libc*. They are automatically loaded as needed by the C compiler *cc(1)*. The link editor searches this library under the “-lc” option. Use of these routines should, for the most part, be avoided. Manual entries for the functions in this library describe the proper routine to use.

LIST OF FUNCTIONS

<i>Name</i>	<i>Appears on Page</i>	<i>Description</i>
alarm	alarm.3c	schedule signal after specified time
ftime	time.3c	get date and time
getpw	getpw.3c	get name from uid
gtty	stty.3c	set and get terminal state (defunct)
nice	nice.3c	set program priority
pause	pause.3c	stop until signal
rand	rand.3c	random number generator
signal	signal.3c	simplified software signal facilities
srand	rand.3c	random number generator
stty	stty.3c	set and get terminal state (defunct)
time	time.3c	get date and time
times	times.3c	get process times
utime	utime.3c	set file times
vlimit	vlimit.3c	control maximum system resource consumption
vtimes	vtimes.3c	get information about resource utilization

NAME

alarm — schedule signal after specified time

SYNOPSIS

alarm(seconds)
unsigned seconds;

DESCRIPTION

This interface is obsoleted by setitimer(2).

Alarm causes signal SIGALRM, see *signal(3C)*, to be sent to the invoking process in a number of seconds given by the argument. Unless caught or ignored, the signal terminates the process.

Alarm requests are not stacked; successive calls reset the alarm clock. If the argument is 0, any alarm request is canceled. Because of scheduling delays, resumption of execution of when the signal is caught may be delayed an arbitrary amount. The longest specifiable delay time is 2147483647 seconds.

The return value is the amount of time previously remaining in the alarm clock.

SEE ALSO

sigpause(2), sigvec(2), signal(3C), sleep(3)

NAME

`getpw` — get name from uid

SYNOPSIS

```
getpw(uid, buf)  
char *buf;
```

DESCRIPTION

Getpw is obsoleted by `getpwuid(3)`.

Getpw searches the password file for the (numerical) *uid*, and fills in *buf* with the corresponding line; it returns non-zero if *uid* could not be found. The line is null-terminated.

FILES

/etc/passwd

SEE ALSO

`getpwent(3)`, `passwd(5)`

DIAGNOSTICS

Non-zero return on error.

NAME

`nice` — set program priority

SYNOPSIS

`nice(incr)`

DESCRIPTION

This interface is obsoleted by `setpriority(2)`.

The scheduling priority of the process is augmented by *incr*. Positive priorities get less service than normal. Priority 10 is recommended to users who wish to execute long-running programs without flak from the administration.

Negative increments are ignored except on behalf of the super-user. The priority is limited to the range -20 (most urgent) to 20 (least).

The priority of a process is passed to a child process by *fork(2)*. For a privileged process to return to normal priority from an unknown state, *nice* should be called successively with arguments -40 (goes to priority -20 because of truncation), 20 (to get to 0), then 0 (to maintain compatibility with previous versions of this call).

SEE ALSO

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

NAME

`pause` — stop until signal

SYNOPSIS

`pause()`

DESCRIPTION

Pause never returns normally. It is used to give up control while waiting for a signal from *kill(2)* or an interval timer, see *setitimer(2)*. Upon termination of a signal handler started during a *pause*, the *pause* call will return.

RETURN VALUE

Always returns `-1`.

ERRORS

Pause always returns:

[EINTR] The call was interrupted.

SEE ALSO

kill(2), *select(2)*, *sigpause(2)*

NAME

rand, *srand* — random number generator

SYNOPSIS

srand(seed)

int seed;

rand()

DESCRIPTION

The newer *random(3)* should be used in new applications; *rand* remains for compatibility.

Rand uses a multiplicative congruential random number generator with period 2^{32} to return successive pseudo-random numbers in the range from 0 to $2^{31} - 1$.

The generator is reinitialized by calling *srand* with 1 as argument. It can be set to a random starting point by calling *srand* with whatever you like as argument.

SEE ALSO

random(3)

3

NAME

signal — simplified software signal facilities

SYNOPSIS

```
#include <signal.h>
(*signal(sig, func)) 0
void (*func)();
```

DESCRIPTION

Signal is a simplified interface to the more general *sigvec(2)* facility.

A signal is generated by some abnormal event, initiated by a user at a terminal (quit, interrupt, stop), by a program error (bus error, etc.), by request of another program (kill), or when a process is stopped because it wishes to access its control terminal while in the background (see *tty(4)*). Signals are optionally generated when a process resumes after being stopped, when the status of child processes changes, or when input is ready at the control terminal. Most signals cause termination of the receiving process if no action is taken; some signals instead cause the process receiving them to be stopped, or are simply discarded if the process has not requested otherwise. Except for the SIGKILL and SIGSTOP signals, the *signal* call allows signals either to be ignored or to cause an interrupt to a specified location. The following is a list of all signals with names as in the include file *<signal.h>*:

SIGHUP	1	hangup
SIGINT	2	interrupt
SIGQUIT	3*	quit
SIGILL	4*	illegal instruction
SIGTRAP	5*	trace trap
SIGIOT	6*	IOT instruction
SIGEMT	7*	EMT instruction
SIGFPE	8*	floating point exception
SIGKILL	9	kill (cannot be caught or ignored)
SIGBUS	10*	bus error
SIGSEGV	11*	segmentation violation
SIGSYS	12*	bad argument to system call
SIGPIPE	13	write on a pipe with no one to read it
SIGALRM	14	alarm clock
SIGTERM	15	software termination signal
SIGURG	16*	urgent condition present on socket
SIGSTOP	17†	stop (cannot be caught or ignored)
SIGTSTP	18†	stop signal generated from keyboard
SIGCONT	19*	continue after stop
SIGCHLD	20*	child status has changed
SIGTTIN	21†	background read attempted from control terminal
SIGTTOU	22†	background write attempted to control terminal
SIGIO	23*	i/o is possible on a descriptor (see <i>fcntl(2)</i>)
SIGXCPU	24	cpu time limit exceeded (see <i>setrlimit(2)</i>)
SIGXFSZ	25	file size limit exceeded (see <i>setrlimit(2)</i>)
SIGVTALRM	26	virtual time alarm (see <i>setitimer(2)</i>)
SIGPROF	27	profiling timer alarm (see <i>setitimer(2)</i>)

The starred signals in the list above cause a core image if not caught or ignored.

If *func* is SIG_DFL, the default action for signal *sig* is reinstated; this default is termination (with a core image for starred signals) except for signals marked with * or †. Signals marked with * are discarded if the action is SIG_DFL; signals marked with † cause the process to stop. If *func* is SIG_IGN the signal is subsequently ignored and pending instances of the signal are

discarded. Otherwise, when the signal occurs further occurrences of the signal are automatically blocked and *func* is called.

A return from the function unblocks the handled signal and continues the process at the point it was interrupted. **Unlike previous signal facilities, the handler *func* remains installed after a signal has been delivered.**

If a caught signal occurs during certain system calls, causing the call to terminate prematurely, the call is automatically restarted. In particular this can occur during a *read* or *write*(2) on a slow device (such as a terminal; but not a file) and during a *wait*(2).

The value of *signal* is the previous (or initial) value of *func* for the particular signal.

After a *fork*(2) or *yfork*(2) the child inherits all signals. *Execve*(2) resets all caught signals to the default action; ignored signals remain ignored.

RETURN VALUE

The previous action is returned on a successful call. Otherwise, -1 is returned and *errno* is set to indicate the error.

ERRORS

Signal will fail and no action will take place if one of the following occur:

- [EINVAL] *Sig* is not a valid signal number.
- [EINVAL] An attempt is made to ignore or supply a handler for SIGKILL or SIGSTOP.
- [EINVAL] An attempt is made to ignore SIGCONT (by default SIGCONT is ignored).

SEE ALSO

kill(1), *ptrace*(2), *kill*(2), *sigvec*(2), *sigblock*(2), *sigsetmask*(2), *sigpause*(2), *sigstack*(2), *setjmp*(3), *tty*(4)

NOTES (VAX-11)

The handler routine can be declared:

```
handler(sig, code, scp)
```

Here *sig* is the signal number, into which the hardware faults and traps are mapped as defined below. *Code* is a parameter which is either a constant as given below or, for compatibility mode faults, the code provided by the hardware. *Scp* is a pointer to the *struct sigcontext* used by the system to restore the process context from before the signal. Compatibility mode faults are distinguished from the other SIGILL traps by having PSL_CM set in the *psl*.

The following defines the mapping of hardware traps to signals and codes. All of these symbols are defined in *<signal.h>*:

Hardware condition	Signal	Code
Arithmetic traps:		
Integer overflow	SIGFPE	FPE_INTOVF_TRAP
Integer division by zero	SIGFPE	FPE_INTDIV_TRAP
Floating overflow trap	SIGFPE	FPE_FLTOVF_TRAP
Floating/decimal division by zero	SIGFPE	FPE_FLTDIV_TRAP
Floating underflow trap	SIGFPE	FPE_FLTUND_TRAP
Decimal overflow trap	SIGFPE	FPE_DECOVF_TRAP
Subscript-range	SIGFPE	FPE_SUBRNG_TRAP
Floating overflow fault	SIGFPE	FPE_FLTOVF_FAULT
Floating divide by zero fault	SIGFPE	FPE_FLTDIV_FAULT
Floating underflow fault	SIGFPE	FPE_FLTUND_FAULT
Length access control	SIGSEGV	
Protection violation	SIGBUS	

Reserved instruction	SIGILL	ILL_RESAD_FAULT
Customer-reserved instr.	SIGEMT	
Reserved operand	SIGILL	ILL_PRIVIN_FAULT
Reserved addressing	SIGILL	ILL_RESOP_FAULT
Trace pending	SIGTRAP	
Bpt instruction	SIGTRAP	
Compatibility-mode	SIGILL	hardware supplied code
Chme	SIGSEGV	
Chms	SIGSEGV	
Chmu	SIGSEGV	

NAME

stty, *gtty* — set and get terminal state (defunct)

SYNOPSIS

```
#include <sgtty.h>
```

```
stty(fd, buf)
int fd;
struct sgttyb *buf;
```

```
gtty(fd, buf)
int fd;
struct sgttyb *buf;
```

DESCRIPTION

This interface is obsoleted by *ioctl(2)*.

Stty sets the state of the terminal associated with *fd*. *Gtty* retrieves the state of the terminal associated with *fd*. To set the state of a terminal the call must have write permission.

The *stty* call is actually “*ioctl*(fd, TIOCSETP, buf)”, while the *gtty* call is “*ioctl*(fd, TIOCGETP, buf)”. See *ioctl(2)* and *tty(4)* for an explanation.

DIAGNOSTICS

If the call is successful 0 is returned, otherwise *-1* is returned and the global variable *errno* contains the reason for the failure.

SEE ALSO

ioctl(2), *tty(4)*

NAME

time, ftime — get date and time

SYNOPSIS

```
long time(0)
long time(tloc)
long *tloc;

#include <sys/types.h>
#include <sys/timeb.h>
ftime(tp)
struct timeb *tp;
```

DESCRIPTION

These interfaces are obsoleted by `gettimeofday(2)`.

Time returns the time since 00:00:00 GMT, Jan. 1, 1970, measured in seconds.

If *tloc* is nonnull, the return value is also stored in the place to which *tloc* points.

The *ftime* entry fills in a structure pointed to by its argument, as defined by `<sys/timeb.h>`:

```
/*    timeb.h    6.183/07/29*/

/*
 * Structure returned by ftime system call
 */
struct timeb
{
    time_t    time;
    unsigned short millitm;
    short    timezone;
    short    dstflag;
};
```

The structure contains the time since the epoch in seconds, up to 1000 milliseconds of more-precise interval, the local time zone (measured in minutes of time westward from Greenwich), and a flag that, if nonzero, indicates that Daylight Saving time applies locally during the appropriate part of the year.

SEE ALSO

`date(1)`, `gettimeofday(2)`, `settimeofday(2)`, `ctime(3)`

NAME

times — get process times

SYNOPSIS

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

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

```
times(buffer)
```

```
struct tms *buffer;
```

DESCRIPTION

This interface is **obsoleted** by `getrusage(2)`.

Times returns time-accounting information for the current process and for the terminated child processes of the current process. All times are in 1/HZ seconds, where HZ is 60.

This is the structure returned by *times*:

```
/* times.h 6.1 83/07/29 */
```

```
/*
```

```
 * Structure returned by times()
```

```
*/
```

```
struct tms {
```

```
    time_t tms_utime;        /* user time */
```

```
    time_t tms_stime;        /* system time */
```

```
    time_t tms_cutime;       /* user time, children */
```

```
    time_t tms_cstime;       /* system time, children */
```

```
};
```

The children times are the sum of the children's process times and their children's times.

SEE ALSO

time(1), `getrusage(2)`, `wait3(2)`, `time(3)`

NAME

utime — set file times

SYNOPSIS

```
#include <sys/types.h>
utime(file, timep)
char *file;
time_t timep[2];
```

DESCRIPTION

This interface is obsoleted by utimes(2).

The *utime* call uses the 'accessed' and 'updated' times in that order from the *timep* vector to set the corresponding recorded times for *file*.

The caller must be the owner of the file or the super-user. The 'inode-changed' time of the file is set to the current time.

SEE ALSO

utimes(2), stat(2)

NAME

`vlimit` — control maximum system resource consumption

SYNOPSIS

```
#include <sys/vlimit.h>
vlimit(resource, value)
```

DESCRIPTION

This facility is superseded by `getrlimit(2)`.

Limits the consumption by the current process and each process it creates to not individually exceed *value* on the specified *resource*. If *value* is specified as `-1`, then the current limit is returned and the limit is unchanged. The resources which are currently controllable are:

LIM_NORAISE A pseudo-limit; if set non-zero then the limits may not be raised. Only the super-user may remove the *noraise* restriction.

LIM_CPU the maximum number of cpu-seconds to be used by each process

LIM_FSIZE the largest single file which can be created

LIM_DATA the maximum growth of the data+stack region via *sbrk(2)* beyond the end of the program text

LIM_STACK the maximum size of the automatically-extended stack region

LIM_CORE the size of the largest core dump that will be created.

LIM_MAXRSS a soft limit for the amount of physical memory (in bytes) to be given to the program. If memory is tight, the system will prefer to take memory from processes which are exceeding their declared **LIM_MAXRSS**.

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

The system refuses to extend the data or stack space when the limits would be exceeded in the normal way; a *break* call fails if the data space limit is reached, or the process is killed when the stack limit is reached (since the stack cannot be extended, there is no way to send a signal!).

A file i/o operation which would create a file which is too large will cause a signal **SIGXFSZ** to be generated, this normally terminates the process, but may be caught. When the cpu time limit is exceeded, a signal **SIGXCPU** is sent to the offending process; to allow it time to process the signal it is given 5 seconds grace by raising the cpu time limit.

SEE ALSO

csh(1)

BUGS

If **LIM_NORAISE** is set, then no grace should be given when the cpu time limit is exceeded.

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

This call is peculiar to this version of UNIX. The options and specifications of this system call and even the call itself are subject to change. It may be extended or replaced by other facilities in future versions of the system.

NAME

`vtimes` — get information about resource utilization

SYNOPSIS

```
vtimes(par_vm, ch_vm)
struct vtimes *par_vm, *ch_vm;
```

DESCRIPTION

This facility is superseded by `getrusage(2)`.

`Vtimes` returns accounting information for the current process and for the terminated child processes of the current process. Either `par_vm` or `ch_vm` or both may be 0, in which case only the information for the pointers which are non-zero is returned.

After the call, each buffer contains information as defined by the contents of the include file `/usr/include/sys/vtimes.h`:

```
struct vtimes {
    int    vm_utime;           /* user time (*HZ) */
    int    vm_stime;         /* system time (*HZ) */
    /* divide next two by utime+stime to get averages */
    unsigned vm_idrssi;      /* integral of d+s rssi */
    unsigned vm_ixrssi;      /* integral of text rssi */
    int    vm_maxrssi;       /* maximum rssi */
    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_idrssi` and `vm_ixrssi` 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_idrssi` would have the value 5*60, where `vm_utime+vm_stime` would be the 60. `vm_idrssi` integrates data and stack segment usage, while `vm_ixrssi` integrates text segment usage. `vm_maxrssi` reports the maximum instantaneous sum of the text+data+stack core-resident page count.

The `vm_majflt` field gives the number of page faults which resulted in disk activity; the `vm_minflt` field gives the number of page faults incurred in simulation of reference bits; `vm_nswap` is the number of swaps which occurred. The number of file system input/output events are reported in `vm_inblk` and `vm_oublk`. These numbers account only for real i/o; data supplied by the caching mechanism is charged only to the first process to read or write the data.

SEE ALSO

`time(2)`, `wait3(2)`

BUGS

This call is peculiar to this version of UNIX. The options and specifications of this system call are subject to change. It may be extended to include additional information in future versions of the system.

NAME

intro — introduction to special files and hardware support

DESCRIPTION

This section describes the special files, related driver functions, and networking support available in the system. In this part of the manual, the SYNOPSIS section of each configurable device gives a sample specification for use in constructing a system description for the *config(8)* program. The DIAGNOSTICS section lists messages which may appear on the console and in the system error log *lusr/adm/messages* due to errors in device operation.

This section contains both devices which may be configured into the system, “4” entries, and network related information, “4N”, “4P”, and “4F” entries; The networking support is introduced in *intro(4N)*.

VAX DEVICE SUPPORT

This section describes the hardware supported on the DEC VAX-11. Software support for these devices comes in two forms. A hardware device may be supported with a character or block *device driver*, or it may be used within the networking subsystem and have a *network interface driver*. Block and character devices are accessed through files in the file system of a special type; c.f. *mknod(8)*. Network interfaces are indirectly accessed through the interprocess communication facilities provided by the system; see *socket(2)*.

A hardware device is identified to the system at configuration time and the appropriate device or network interface driver is then compiled into the system. When the resultant system is booted, the autoconfiguration facilities in the system probe for the device on either the UNIBUS or MASSBUS and, if found, enable the software support for it. If a UNIBUS device does not respond at autoconfiguration time it is not accessible at any time afterwards. To enable a UNIBUS device which did not autoconfigure, the system will have to be rebooted. If a MASSBUS device comes “on-line” after the autoconfiguration sequence it will be dynamically autoconfigured into the running system.

The autoconfiguration system is described in *autoconf(4)*. VAX specific device support is described in “4V” entries. A list of the supported devices is given below.

SEE ALSO

intro(4), intro(4N), autoconf(4), config(8)

LIST OF DEVICES

The devices listed below are supported in this incarnation of the system. Devices are indicated by their functional interface. If second vendor products provide functionally identical interfaces they should be usable with the supplied software. (Beware however that we promise the software works ONLY with the hardware indicated on the appropriate manual page.)

acc	ACC LH/DH IMP communications interface
ad	Data translation A/D interface
css	DEC IMP-11A communications interface
ct	C/A/T phototypesetter
dh	DH-11 emulators, terminal multiplexor
dmc	DEC DMC-11/DMR-11 point-to-point communications device
dmf	DEC DMF-32 terminal multiplexor
dn	DEC DN-11 autodialer interface
dz	DZ-11 terminal multiplexor
ec	3Com 10Mb/s Ethernet controller
en	Xerox 3Mb/s Ethernet controller (obsolete)
kg	KL-11/DL-11W line clock
fl	VAX-11/780 console floppy interface
hk	RK6-11/RK06 and RK07 moving head disk

hp	MASSBUS disk interface (with RP06, RM03, RM05, etc.)
ht	TM03 MASSBUS tape drive interface (with TE-16, TU-45, TU-77)
hy	DR-11B or GI-13 interface to an NSC Hyperchannel
ik	Ikonas frame buffer graphics device interface
il	Interlan 10Mb/s Ethernet controller
lp	LP-11 parallel line printer interface
mt	TM78 MASSBUS tape drive interface
pcl	DEC PCL-11 communications interface
ps	Evans and Sutherland Picture System 2 graphics interface
rx	DEC RX02 floppy interface
tm	TM-11/TE-10 tape drive interface
ts	TS-11 tape drive interface
tu	VAX-11/730 TU58 console cassette interface
uda	DEC UDA-50 disk controller
un	DR-11W interface to Ungermann-Bass
up	Emulex SC-21V UNIBUS disk controller
ut	UNIBUS TU-45 tape drive interface
uu	TU58 dual cassette drive interface (DL11)
va	Benson-Varian printer/plotter interface
vp	Versatec printer/plotter interface
vv	Proteon proNET 10Mb/s ring network interface

4

NAME

networking — introduction to networking facilities

SYNOPSIS

```
#include <sys/socket.h>
#include <net/route.h>
#include <net/if.h>
```

DESCRIPTION

This section briefly describes the networking facilities available in the system. Documentation in this part of section 4 is broken up into three areas: *protocol-families*, *protocols*, and *network interfaces*. Entries describing a protocol-family are marked “4F”, while entries describing protocol use are marked “4P”. Hardware support for network interfaces are found among the standard “4” entries.

All network protocols are associated with a specific *protocol-family*. A protocol-family provides basic services to the protocol implementation to allow it to function within a specific network environment. These services may include packet fragmentation and reassembly, routing, addressing, and basic transport. A protocol-family may support multiple methods of addressing, though the current protocol implementations do not. A protocol-family is normally comprised of a number of protocols, one per *socket(2)* type. It is not required that a protocol-family support all socket types. A protocol-family may contain multiple protocols supporting the same socket abstraction.

A protocol supports one of the socket abstractions detailed in *socket(2)*. A specific protocol may be accessed either by creating a socket of the appropriate type and protocol-family, or by requesting the protocol explicitly when creating a socket. Protocols normally accept only one type of address format, usually determined by the addressing structure inherent in the design of the protocol-family/network architecture. Certain semantics of the basic socket abstractions are protocol specific. All protocols are expected to support the basic model for their particular socket type, but may, in addition, provide non-standard facilities or extensions to a mechanism. For example, a protocol supporting the SOCK_STREAM abstraction may allow more than one byte of out-of-band data to be transmitted per out-of-band message.

A network interface is similar to a device interface. Network interfaces comprise the lowest layer of the networking subsystem, interacting with the actual transport hardware. An interface may support one or more protocol families, and/or address formats. The SYNOPSIS section of each network interface entry gives a sample specification of the related drivers for use in providing a system description to the *config(8)* program. The DIAGNOSTICS section lists messages which may appear on the console and in the system error log */usr/adm/messages* due to errors in device operation.

PROTOCOLS

The system currently supports only the DARPA Internet protocols fully. Raw socket interfaces are provided to IP protocol layer of the DARPA Internet, to the IMP link layer (1822), and to Xerox PUP-1 layer operating on top of 3Mb/s Ethernet interfaces. Consult the appropriate manual pages in this section for more information regarding the support for each protocol family.

ADDRESSING

Associated with each protocol family is an address format. The following address formats are used by the system:

```
#define AF_UNIX      1      /* local to host (pipes, portals) */
#define AF_INET      2      /* internetwork: UDP, TCP, etc. */
#define AF_IMPLINK   3      /* arpanet imp addresses */
#define AF_PUP       4      /* pup protocols: e.g. BSP */
```

ROUTING

The network facilities provided limited packet routing. A simple set of data structures comprise a "routing table" used in selecting the appropriate network interface when transmitting packets. This table contains a single entry for each route to a specific network or host. A user process, the routing daemon, maintains this data base with the aid of two socket specific *ioctl(2)* commands, *SIOCADDRT* and *SIOCDELRT*. The commands allow the addition and deletion of a single routing table entry, respectively. Routing table manipulations may only be carried out by super-user.

A routing table entry has the following form, as defined in *<netroute.h>*;

```
struct rtenry {
    u_long      rt_hash;
    struct      sockaddr rt_dst;
    struct      sockaddr rt_gateway;
    short       rt_flags;
    short       rt_refcnt;
    u_long      rt_use;
    struct      ifnet *rt_ifp;
};
```

with *rt_flags* defined from,

```
#define RTF_UP          0x1    /* route usable */
#define RTF_GATEWAY    0x2    /* destination is a gateway */
#define RTF_HOST       0x4    /* host entry (net otherwise) */
```

Routing table entries come in three flavors: for a specific host, for all hosts on a specific network, for any destination not matched by entries of the first two types (a wildcard route). When the system is booted, each network interface autoconfigured installs a routing table entry when it wishes to have packets sent through it. Normally the interface specifies the route through it is a "direct" connection to the destination host or network. If the route is direct, the transport layer of a protocol family usually requests the packet be sent to the same host specified in the packet. Otherwise, the interface may be requested to address the packet to an entity different from the eventual recipient (i.e. the packet is forwarded).

Routing table entries installed by a user process may not specify the hash, reference count, use, or interface fields; these are filled in by the routing routines. If a route is in use when it is deleted (*rt_refcnt* is non-zero), the resources associated with it will not be reclaimed until further references to it are released.

The routing code returns *EEXIST* if requested to duplicate an existing entry, *ESRCH* if requested to delete a non-existent entry, or *ENOBUFS* if insufficient resources were available to install a new route.

User processes read the routing tables through the */dev/kmem* device.

The *rt_use* field contains the number of packets sent along the route. This value is used to select among multiple routes to the same destination. When multiple routes to the same destination exist, the least used route is selected.

A wildcard routing entry is specified with a zero destination address value. Wildcard routes are used only when the system fails to find a route to the destination host and network. The combination of wildcard routes and routing redirects can provide an economical mechanism for routing traffic.

INTERFACES

Each network interface in a system corresponds to a path through which messages may be sent and received. A network interface usually has a hardware device associated with it, though certain interfaces such as the loopback interface, *lo(4)*, do not.

At boot time each interface which has underlying hardware support makes itself known to the system during the autoconfiguration process. Once the interface has acquired its address it is expected to install a routing table entry so that messages may be routed through it. Most interfaces require some part of their address specified with an *SIOCSIFADDR* *ioctl* before they will allow traffic to flow through them. On interfaces where the network-link layer address mapping is static, only the network number is taken from the *ioctl*; the remainder is found in a hardware specific manner. On interfaces which provide dynamic network-link layer address mapping facilities (e.g. 10Mb/s Ethernets), the entire address specified in the *ioctl* is used.

The following *ioctl* calls may be used to manipulate network interfaces. Unless specified otherwise, the request takes an *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

socket(2), ioctl(2), intro(4), config(8), routed(8C)

NAME

acc — ACC LH/DH IMP interface

SYNOPSIS

```
pseudo-device imp  
device acc0 at uba0 csr 167600 vector accrint accxint
```

DESCRIPTION

The *acc* device provides a Local Host/Distant Host interface to an IMP. It is normally used when participating in the DARPA Internet. The controller itself is not accessible to users, but instead provides the hardware support to the IMP interface described in *imp*(4). When configuring, the *imp* pseudo-device must also be included.

DIAGNOSTICS

acc%d: not alive. The initialization routine was entered even though the device did not autoconfigure. This indicates a system problem.

acc%d: can't initialize. Insufficient UNIBUS resources existed to initialize the device. This is likely to occur when the device is run on a buffered data path on an 11/750 and other network interfaces are also configured to use buffered data paths, or when it is configured to use buffered data paths on an 11/730 (which has none).

acc%d: imp doesn't respond, icsr=%b. The driver attempted to initialize the device, but the IMP failed to respond after 500 tries. Check the cabling.

acc%d: stray xmit interrupt, csr=%b. An interrupt occurred when no output had previously been started.

acc%d: output error, ocsr=%b, icsr=%b. The device indicated a problem sending data on output.

acc%d: input error, csr=%b. The device indicated a problem receiving data on input.

acc%d: bad length=%d. An input operation resulted in a data transfer of less than 0 or more than 1008 bytes of data into memory (according to the word count register). This should never happen as the maximum size of a host-IMP message is 1008 bytes.

NAME

ad — Data Translation A/D converter

SYNOPSIS

device ad0 at uba0 csr 0170400 vector adintr

DESCRIPTION

Ad provides the interface to the Data Translation A/D converter. This is **not** a real-time driver, but merely allows the user process to sample the board's channels one at a time. Each minor device selects a different A/D board.

The driver communicates to a user process by means of *ioctl*s. The *AD_CHAN* *ioctl* selects which channel of the board to read. For example,

```
chan = 5; ioctl(fd, AD_CHAN, &chan);
```

selects channel 5. The *AD_READ* *ioctl* actually reads the data and returns it to the user process. An example is

```
ioctl(fd, AD_READ, &data);
```

FILES

/dev/ad

DIAGNOSTICS

None.

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 on a local area network. It is used by all the 10Mb/s Ethernet interface drivers and is not directly accessible to users.

ARP caches Internet-Ethernet address mappings. When an interface requests a mapping for an address not in the cache, ARP queues the message which requires the mapping and broadcasts a message on the associated network requesting the address mapping. If a response is provided, the new mapping is cached and any pending messages are transmitted. ARP itself is not Internet or Ethernet specific; this implementation, however, is. ARP will queue at most one packet while waiting for a mapping request to be responded to; only the most recently “transmitted” packet is kept.

ARP watches passively for hosts impersonating the local host (i.e. a host which responds to an ARP mapping request for the local host's address) and will, optionally, periodically probe a network looking for impostors.

DIAGNOSTICS

duplicate IP address!! sent from ethernet address: %x %x %x %x %x %x . ARP has discovered another host on the local network which responds to mapping requests for its own Internet address.

SEE ALSO

ec(4), il(4)

NAME

autoconf — diagnostics from the autoconfiguration code

DESCRIPTION

When UNIX bootstraps it probes the innards of the machine it is running on and locates controllers, drives, and other devices, printing out what it finds on the console. This procedure is driven by a system configuration table which is processed by *config(8)* and compiled into each kernel.

Devices in NEXUS slots are normally noted, thus memory controllers, UNIBUS and MASSBUS adaptors. Devices which are not supported which are found in NEXUS slots are noted also.

MASSBUS devices are located by a very deterministic procedure since MASSBUS space is completely probe-able. If devices exist which are not configured they will be silently ignored; if devices exist of unsupported type they will be noted.

UNIBUS devices are located by probing to see if their control-status registers respond. If not, they are silently ignored. If the control status register responds but the device cannot be made to interrupt, a diagnostic warning will be printed on the console and the device will not be available to the system.

A generic system may be built which picks its root device at boot time as the “best” available device (MASSBUS disks are better than SMD UNIBUS disks are better than RK07's; the device must be drive 0 to be considered.) If such a system is booted with the RB_ASKNAME option of (see *reboot(2)*), then the name of the root device is read from the console terminal at boot time, and any available device may be used.

SEE ALSO

intro(4), config(8)

DIAGNOSTICS

cpu type %d not configured. You tried to boot UNIX on a cpu type which it doesn't (or at least this compiled version of UNIX doesn't) understand.

mba%d at tr%d. A MASSBUS adapter was found in tr%d (the NEXUS slot number). UNIX will call it mba%d.

%d mba's not configured. More MASSBUS adapters were found on the machine than were declared in the machine configuration; the excess MASSBUS adapters will not be accessible.

uba%d at tr%d. A UNIBUS adapter was found in tr%d (the NEXUS slot number). UNIX will call it uba%d.

dr32 unsupported (at tr %d). A DR32 interface was found in a NEXUS, for which UNIX does not have a driver.

mcr%d at tr%d. A memory controller was found in tr%d (the NEXUS slot number). UNIX will call it mcr%d.

5 mcr's unsupported. UNIX supports only 4 memory controllers per cpu.

mpm unsupported (at tr%d). Multi-port memory is unsupported in the sense that UNIX does not know how to poll it for ECC errors.

%s%d at mba%d drive %d. A tape formatter or a disk was found on the MASSBUS; for disks %s%d will look like “hp0”, for tape formatters like “ht1”. The drive number comes from the unit plug on the drive or in the TM formatter (not on the tape drive; see below).

%s%d at %s%d slave %d. (For MASSBUS devices). Which would look like “tu0 at ht0 slave 0”, where tu0 is the name for the tape device and ht0 is the name for the formatter. A tape slave was found on the tape formatter at the indicated drive number (on the front of the tape drive). UNIX will call the device, e.g., tu0.

%s%d at uba%d csr %o vec %o ipl %x. The device %s%d, e.g. dz0 was found on uba%d at control-status register address %o and with device vector %o. The device interrupted at priority level %x.

%s%d at uba%d csr %o zero vector. The device did not present a valid interrupt vector, rather presented 0 (a passive release condition) to the adapter.

%s%d at uba%d csr %o didn't interrupt. The device did not interrupt, likely because it is broken, hung, or not the kind of device it is advertised to be.

%s%d at %s%d slave %d. (For UNIBUS devices). Which would look like "up0 at sc0 slave 0", where up0 is the name of a disk drive and sc0 is the name of the controller. Analogous to MASSBUS case.

NAME

bk — line discipline for machine-machine communication (obsolete)

SYNOPSIS

```
pseudo-device bk
```

DESCRIPTION

This line discipline provides a replacement for the old and new tty drivers described in *ty(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 the Berkeley network. 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 disables 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 recognized 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

ty(4)

DIAGNOSTICS

None.

BUGS

The Purdue uploading line discipline, which provides 8 bits and uses timeout's to terminate uploading should be incorporated into the standard system, as it is much more suitable for microprocessor connections.

NAME

cons — VAX-11 console interface

DESCRIPTION

The console is available to the processor through the console registers. It acts like a normal terminal, except that when the local functions are not disabled, control-P puts the console in local console mode (where the prompt is ">>>"). The operation of the console in this mode varies slightly per-processor.

On an 11/780 you can return to the conversational mode using the command "set t p" (set terminal program) if the processor is still running or "continue" if it is halted. The latter command may be abbreviated "c". If you hit the break key on the console, then the console will go into ODT (console debugger mode). Hit a "P" (upper-case letter p) to get out of this mode.

On an 11/750 or an 11/730 the processor is halted whenever the console is not in conversational mode, and typing "C" returns to conversational mode. When in console mode on an 11/750 which has a remote diagnosis module, a ^D will put you in remote diagnosis mode, where the prompt will be "RDM>". The command "ret" will return from remote diagnosis mode to local console mode.

With the above proviso's the console works like any other UNIX terminal.

FILES

/dev/console

SEE ALSO

tty(4), reboot(8)
VAX Hardware Handbook

DIAGNOSTICS

None.

NAME

css — DEC IMP-11A LH/DH IMP interface

SYNOPSIS

pseudo-device *imp*
device *css0* at uba0 csr 167600 flags 10 vector *cssrint* *cssxint*

DESCRIPTION

The *css* device provides a Local Host/Distant Host interface to an IMP. It is normally used when participating in the DARPA Internet. The controller itself is not accessible to users, but instead provides the hardware support to the IMP interface described in *imp*(4). When configuring, the *imp* pseudo-device is also included.

DIAGNOSTICS

***css%d*: not alive.** The initialization routine was entered even though the device did not autoconfigure. This indicates a system problem.

***css%d*: can't initialize.** Insufficient UNIBUS resources existed to initialize the device. This is likely to occur when the device is run on a buffered data path on an 11/750 and other network interfaces are also configured to use buffered data paths, or when it is configured to use buffered data paths on an 11/730 (which has none).

***css%d*: *imp* doesn't respond, *icsr*=%b.** The driver attempted to initialize the device, but the IMP failed to respond after 500 tries. Check the cabling.

***css%d*: stray output interrupt *csr*=%b.** An interrupt occurred when no output had previously been started.

***css%d*: output error, *ocsr*=%b *icsr*=%b.** The device indicated a problem sending data on output.

***css%d*: *rcv* error, *csr*=%b.** The device indicated a problem receiving data on input.

***css%d*: bad length=%d.** An input operation resulted in a data transfer of less than 0 or more than 1008 bytes of data into memory (according to the word count register). This should never happen as the maximum size of a host-IMP message is 1008 bytes.

4

NAME

ct — phototypesetter interface

SYNOPSIS

device ct0 at uba0 csr 0167760 vector ctintr

DESCRIPTION

This provides an interface to a Graphic Systems C/A/T phototypesetter. Bytes written on the file specify font, size, and other control information as well as the characters to be flashed. The coding is not described here.

Only one process may have this file open at a time. It is write-only.

FILES

/dev/cat

SEE ALSO

troff(1)

Phototypesetter interface specification

DIAGNOSTICS

None.

NAME

dh — DH-11/DM-11 communications multiplexer

SYNOPSIS

device dh0 at uba0 csr 0160020 vector dhrint dhxint
device dm0 at uba0 csr 0170500 vector dmintr

DESCRIPTION

A dh-11 provides 16 communication lines; dm-11's may be optionally paired with dh-11's to provide modem control for the lines.

Each line attached to the DH-11 communications multiplexer 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 for a dh to say that a line is not properly connected, and that the line should be treated as hard-wired with carrier always present. Thus specifying "flags 0x0004" in the specification of dh0 would cause line ttyh2 to be treated in this way.

The dh driver normally uses input silos and polls for input at each clock tick (10 milliseconds) rather than taking an interrupt on each input character.

FILES

/dev/tty[hi][0-9a-f]
 /dev/ttyd[0-9a-f]

SEE ALSO

tty(4)

DIAGNOSTICS

dh%d: NXM. No response from UNIBUS on a dma transfer within a timeout period. This is often followed by a UNIBUS adapter error. This occurs most frequently when the UNIBUS is heavily loaded and when devices which hog the bus (such as rk07's) are present. It is not serious.

dh%d: silo overflow. The character input silo overflowed before it could be serviced. This can happen if a hard error occurs when the CPU is running with elevated priority, as the system will then print a message on the console with interrupts disabled. If the Berknet is running on a *dh* line at high speed (e.g. 9600 baud), there is only 1/15th of a second of buffering capacity in the silo, and overrun is possible. This may cause a few input characters to be lost to users and a network packet is likely to be corrupted, but the network will recover. It is not serious.

NAME

dmc — DEC DMC-11/DMR-11 point-to-point communications device

SYNOPSIS

device dmc0 at uba0 csr 167600 vector dmcrint dmcxint

DESCRIPTION

The *dmc* interface provides access to a point-to-point communications device which runs at either 1 Mb/s or 56 Kb/s. DMC-11's communicate using the DEC DDCMP link layer protocol.

The *dmc* interface driver also supports a DEC DMR-11 providing point-to-point communication running at data rates from 2.4 Kb/s to 1 Mb/s. DMR-11's are a more recent design and thus are preferred over DMC-11's.

The host's address must be specified with an SIOCSIFADDR ioctl before the interface will transmit or receive any packets.

DIAGNOSTICS

dmc%d: bad control %o. A bad parameter was passed to the *dmcload* routine.

dmc%d: unknown address type %d. An input packet was received which contained a type of address unknown to the driver.

DMC FATAL ERROR 0%o.

DMC SOFT ERROR 0%o.

dmc%d: af%d not supported. The interface was handed a message which has addresses formatted in an unsuitable address family.

SEE ALSO

intro(4N), inet(4F)

BUGS

Should allow multiple outstanding DMA requests, but due to the design of the current UNIBUS support routines this is very difficult.

NAME

`dmf` — DMF-32, terminal multiplexor

SYNOPSIS

```
device dmf0 at uba? csr 0170000
        vector dmfsrint dmfsxint dmfdaint dmfdbint dmfrint dmfxint dmflint
```

DESCRIPTION

The *dmf* device provides 8 lines of asynchronous serial line support with full modem control (the DMF-32 provides other services, but these are not supported by the driver).

Each line attached to a DMF-32 serial line port 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 for a *dmf* to say that a line is not properly connected, and that the line should be treated as hard-wired with carrier always present. Thus specifying “flags 0x0004” in the specification of *dmf0* would cause line *tyh2* to be treated in this way.

The *dmf* driver normally uses input silos and polls for input at each clock tick (10 milliseconds).

FILES

```
/dev/tty[hi][0-9a-f]
/dev/ttyd[0-9a-f]
```

SEE ALSO

ty(4)

DIAGNOSTICS

dmf%d: NXM line %d. No response from UNIBUS on a dma transfer within a timeout period. This is often followed by a UNIBUS adapter error. This occurs most frequently when the UNIBUS is heavily loaded and when devices which hog the bus (such as *rk07*'s) are present. It is not serious.

dmf%d: silo overflow. The character input silo overflowed before it could be serviced. This can happen if a hard error occurs when the CPU is running with elevated priority, as the system will then print a message on the console with interrupts disabled. If the Berknet is running on a *dh* line at high speed (e.g. 9600 baud), there is only 1/15th of a second of buffering capacity in the silo, and overrun is possible. This may cause a few input characters to be lost to users and a network packet is likely to be corrupted, but the network will recover. It is not serious.

dmfsrint.
dmfsxint.
dmfdaint.
dmfdbint.
dmflint.

One of the unsupported parts of the *dmf* interrupted; something is amiss, check your interrupt vectors for a conflict with another device.

NAME

`dn` — DN-11 autocal unit interface

SYNOPSIS

device `dn0` at `uba?` csr 0160020 vector `dnintr`

DESCRIPTION

The `dn` device provides an interface through a DEC DN-11 (or equivalent such as the Able Quadracall) to an auto-call unit (ACU). To place an outgoing call one forks a sub-process which opens the appropriate call unit file, `/dev/cua?` and writes the phone number on it. The parent process then opens the corresponding modem line `/dev/cul?`. When the connection has been established, the open on the modem line, `/dev/cul?` will return and the process will be connected. A timer is normally used to timeout the opening of the modem line.

The codes for the phone numbers are:

```
0-9 dial 0-9
* dial * ('.' is a synonym)
# dial # (';' is a synonym)
- delay 20 milliseconds
< end-of-number ('e' is a synonym)
= delay for a second dial tone ('w' is a synonym)
f force a hangup of any existing connection
```

The entire telephone number must be presented in a single *write* system call.

By convention, even numbered call units are for 300 baud modem lines, while odd numbered units are for 1200 baud lines. For example, `/dev/cua0` is associated with a 300 baud modem line, `/dev/cul0`, while `/dev/cua1` is associated with a 1200 baud modem line, `/dev/cul1`. For devices such as the Quadracall which simulate multiple DN-11 units, the minor device indicates which outgoing modem to use.

FILES

```
/dev/cua?    call units
/dev/cul?    associated modem lines
```

SEE ALSO

`tip(1C)`

DIAGNOSTICS

Two error numbers are of interest at open time.

[EBUSY] The dialer is in use.

[ENXIO] The device doesn't exist, or there's no power to it.

NAME

drum — paging device

DESCRIPTION

This file refers to the paging device in use by the system. This may actually be a subdevice of one of the disk drivers, but in a system with paging interleaved across multiple disk drives it provides an indirect driver for the multiple drives.

FILES

/dev/drum

BUGS

Reads from the drum are not allowed across the interleaving boundaries. Since these only occur every .5Mbytes or so, and since the system never allocates blocks across the boundary, this is usually not a problem.

NAME

dz — DZ-11 communications multiplexer

SYNOPSIS

device dz0 at uba0 csr 0160100 vector dzzrint dzzxint

DESCRIPTION

A dz-11 provides 8 communication lines with partial modem control, adequate for UNIX dialup use. Each line attached to the DZ-11 communications multiplexer behaves as described in *ty(4)* and may be set to run at any of 16 speeds; see *ty(4)* for the encoding.

Bit *i* of flags may be specified for a dz to say that a line is not properly connected, and that the line should be treated as hard-wired with carrier always present. Thus specifying “flags 0x04” in the specification of dz0 would cause line tty02 to be treated in this way.

The dz driver normally uses its input silos and polls for input at each clock tick (10 milliseconds) rather than taking an interrupt on each input character.

FILES

/dev/tty[0-9][0-9]
/dev/ttyd[0-9a-f] dialups

SEE ALSO

tty(4)

DIAGNOSTICS

dz%d: silo overflow. The 64 character input silo overflowed before it could be serviced. This can happen if a hard error occurs when the CPU is running with elevated priority, as the system will then print a message on the console with interrupts disabled. If the Berknet is running on a dz line at high speed (e.g. 9600 baud), there is only 1/15th of a second of buffering capacity in the silo, and overrun is possible. This may cause a few input characters to be lost to users and a network packet is likely to be corrupted, but the network will recover. It is not serious.

NAME

ec — 3Com 10 Mb/s Ethernet interface

SYNOPSIS

device ec0 at uba0 csr 161000 vector ecrint eccollide excnt

DESCRIPTION

The *ec* interface provides access to a 10 Mb/s Ethernet network through a 3com controller.

The hardware has 32 kilobytes of dual-ported memory on the UNIBUS. This memory is used for internal buffering by the board, and the interface code reads the buffer contents directly through the UNIBUS.

The host's Internet address is specified at boot time with an *SIOCSIFADDR* ioctl. The *ec* interface employs the address resolution protocol described in *arp*(4P) to dynamically map between Internet and Ethernet addresses on the local network.

The interface software implements an exponential backoff algorithm when notified of a collision on the cable. This algorithm utilizes a 16-bit mask and the VAX-11's interval timer in calculating a series of random backoff values. The algorithm is as follows:

1. Initialize the mask to be all 1's.
2. If the mask is zero, 16 retries have been made and we give up.
3. Shift the mask left one bit and formulate a backoff by masking the interval timer with the mask (this is actually the two's complement of the value).
4. Use the value calculated in step 3 to delay before retransmitting the packet. The delay is done in a software busy loop.

The interface normally tries to use a "trailer" encapsulation to minimize copying data on input and output. This may be disabled, on a per-interface basis, by setting the *IFF_NOTRAILERS* flag with an *SIOCSIFFLAGS* ioctl.

DIAGNOSTICS

ec%d: send error. After 16 retransmissions using the exponential backoff algorithm described above, the packet was dropped.

ec%d: input error (offset=%d). The hardware indicated an error in reading a packet off the cable or an illegally sized packet. The buffer offset value is printed for debugging purposes.

ec%d: can't handle af%d. The interface was handed a message with addresses formatted in an unsuitable address family; the packet was dropped.

SEE ALSO

intro(4N), *inet*(4F), *arp*(4P)

BUGS

The PUP protocol family should be added.

The hardware is not capable of talking to itself. The software implements local sending and broadcast by sending such packets to the loop interface. This is a kludge.

Backoff delays are done in a software busy loop. This can degrade the system if the network experiences frequent collisions.

NAME

`en` — Xerox 3 Mb/s Ethernet interface

SYNOPSIS

`device en0 at uba0 csr 161000 vector enrint enxint encollide`

DESCRIPTION

The `en` interface provides access to a 3 Mb/s Ethernet network. Due to limitations in the hardware, DMA transfers to and from the network must take place in the lower 64K bytes of the UNIBUS address space.

The network number is specified with a `SIOCSIFADDR` ioctl; the host's address is discovered by probing the on-board Ethernet address register. No packets will be sent or accepted until a network number is supplied.

The interface software implements an exponential backoff algorithm when notified of a collision on the cable. This algorithm utilizes a 16-bit mask and the VAX-11's interval timer in calculating a series of random backoff values. The algorithm is as follows:

1. Initialize the mask to be all 1's.
2. If the mask is zero, 16 retries have been made and we give up.
3. Shift the mask left one bit and formulate a backoff by masking the interval timer with the mask (this is actually the two's complement of the value).
4. Use the value calculated in step 3 to delay before retransmitting the packet.

The interface handles both Internet and PUP protocol families, with the interface address maintained in Internet format. PUP addresses are converted to Internet addresses by substituting PUP network and host values for Internet network and local part values.

The interface normally tries to use a "trailer" encapsulation to minimize copying data on input and output. This may be disabled, on a per-interface basis, by setting the `IFF_NOTRAILERS` flag with an `SIOCSIFFLAGS` ioctl.

DIAGNOSTICS

`en%d: output error.` The hardware indicated an error on the previous transmission.

`en%d: send error.` After 16 retransmissions using the exponential backoff algorithm described above, the packet was dropped.

`en%d: input error.` The hardware indicated an error in reading a packet off the cable.

`en%d: can't handle af%d.` The interface was handed a message with addresses formatted in an unsuitable address family; the packet was dropped.

SEE ALSO

`intro(4N)`, `inet(4F)`

BUGS

The device has insufficient buffering to handle back to back packets. This makes use in a production environment painful.

The hardware does word at a time DMA without byte swapping. To compensate, byte swapping of user data must either be done by the user or by the system. A kludge to byte swap only IP packets is provided if the `ENF_SWABIPS` flag is defined in the driver and set at boot time with an `SIOCSIFFLAGS` ioctl.

NAME

fl — console floppy interface

DESCRIPTION

This is a simple interface to the DEC RX01 floppy disk unit, which is part of the console LSI-11 subsystem for VAX-11/780's. Access is given to the entire floppy consisting of 77 tracks of 26 sectors of 128 bytes.

All i/o is raw; the seek addresses in raw transfers should be a multiple of 128 bytes and a multiple of 128 bytes should be transferred, as in other "raw" disk interfaces.

FILES

/dev/floppy

SEE ALSO

arff(8V)

DIAGNOSTICS

None.

BUGS

Multiple console floppies are not supported.

If a write is given with a count not a multiple of 128 bytes then the trailing portion of the last sector will be zeroed.

NAME

hk - RK6-11/RK06 and RK07 moving head disk

SYNOPSIS

controller hk0 at uba? csr 0177440 vector rkintr
disk rk0 at hk0 drive 0

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, etc. The standard device names begin with "hk" 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 results in exactly one I/O operation and 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 calls should specify a multiple of 512 bytes.

DISK SUPPORT

The origin and size (in sectors) of the pseudo-disks on each drive are as follows:

RK07 partitions:

Table with 4 columns: disk, start, length, cyl. Rows include hk?a, hk?b, hk?c, hk?g.

RK06 partitions

Table with 4 columns: disk, start, length, cyl. Rows include hk?a, hk?b, hk?c.

On a dual RK-07 system partition hk?a is used for the root for one drive and partition hk?g for the /usr file system. If large jobs are to be run using hk?b on both drives as swap area provides a 10Mbyte paging area. Otherwise partition hk?c on the other drive is used as a single large file system.

FILES

- /dev/hk[0-7][a-h] block files
/dev/rhk[0-7][a-h] raw files

SEE ALSO

hp(4), uda(4), up(4)

DIAGNOSTICS

rk%d%c: hard error sn%d cs2=%b ds=%b er=%b. An unrecoverable error occurred during transfer of the specified sector of the specified disk partition. The contents of the cs2, ds and er registers are printed in octal and symbolically with bits decoded. The error was either unrecoverable, or a large number of retry attempts (including offset positioning and drive recalibration) could not recover the error.

rk%d: write locked. The write protect switch was set on the drive when a write was attempted. The write operation is not recoverable.

rk%d: not ready. The drive was spun down or off line when it was accessed. The i/o operation is not recoverable.

rk%d: not ready (came back!). The drive was not ready, but after printing the message about being not ready (which takes a fraction of a second) was ready. The operation is recovered if no further errors occur.

rk%d%c: soft ecc sn%d. A recoverable ECC error occurred on the specified sector in the specified disk partition. This happens normally a few times a week. If it happens more frequently than this the sectors where the errors are occurring should be checked to see if certain cylinders on the pack, spots on the carriage of the drive or heads are indicated.

hk%d: lost interrupt. A timer watching the controller detected no interrupt for an extended period while an operation was outstanding. This indicates a hardware or software failure. There is currently a hardware/software problem with spinning down drives while they are being accessed which causes this error to occur. The error causes a UNIBUS reset, and retry of the pending operations. If the controller continues to lose interrupts, this error will recur a few seconds later.

BUGS

In raw I/O *read* and *write(2)* 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.

DEC-standard error logging should be supported.

A program to analyze the logged error information (even in its present reduced form) is needed.

The partition tables for the file systems should be read off of each pack, as they are never quite what any single installation would prefer, and this would make packs more portable.

NAME

hp — MASSBUS disk interface

SYNOPSIS

disk hp0 at mba0 drive 0

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, etc. The standard device names begin with "hp" 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 file's 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. 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* calls should specify a multiple of 512 bytes.

DISK SUPPORT

This driver handles both standard DEC controllers and Emulex SC750 and SC780 controllers. Standard DEC drive types are recognized according to the MASSBUS drive type register. For the Emulex controller the drive type register should be configured to indicate the drive is an RM02. When this is encountered, the driver checks the holding register to find out the disk geometry and, based on this information, decides what the drive type is. The following disks are supported: RM03, RM05, RP06, RM80, RP05, RP07, ML11A, ML11B, CDC 9775, CDC 9730, AMPEX Capricorn (32 sectors/track), FUJITSU Eagle (48 sectors/track), and AMPEX 9300. The origin and size (in sectors) of the pseudo-disks on each drive are as follows:

RM03 partitions

disk	start	length	cyls
hp?a	0	15884	0-99
hp?b	16000	33440	100-309
hp?c	0	131680	0-822
hp?d	49600	15884	309-408
hp?e	65440	55936	409-758
hp?f	121440	10080	759-822
hp?g	49600	82080	309-822

RM05 partitions

disk	start	length	cyls
hp?a	0	15884	0-26
hp?b	16416	33440	27-81
hp?c	0	500384	0-822
hp?d	341696	15884	562-588
hp?e	358112	55936	589-680
hp?f	414048	86176	681-822
hp?g	341696	158528	562-822
hp?h	49856	291346	82-561

RP06 partitions

disk	start	length	cyls
hp?a	0	15884	0-37
hp?b	15884	33440	38-117

hp?c	0	340670	0-814
hp?d	49324	15884	118-155
hp?e	65208	55936	156-289
hp?f	121220	219296	290-814
hp?g	49324	291192	118-814

RM80 partitions

disk	start	length	cyls
hp?a	0	15884	0-36
hp?b	16058	33440	37-114
hp?c	0	242606	0-558
hp?d	49910	15884	115-151
hp?e	68096	55936	152-280
hp?f	125888	120466	281-558
hp?g	49910	192510	115-558

RP05 partitions

disk	start	length	cyls
hp?a	0	15884	0-37
hp?b	15884	33440	38-117
hp?c	0	171798	0-410
hp?d	2242	15884	118-155
hp?e	65208	55936	156-289
hp?f	121220	50424	290-410
hp?g	2242	122320	118-410

RP07 partitions

disk	start	length	cyls
hp?a	0	15884	0-9
hp?b	16000	66880	10-51
hp?c	0	1008000	0-629
hp?d	376000	15884	235-244
hp?e	392000	307200	245-436
hp?f	699200	308600	437-629
hp?g	376000	631800	235-629
hp?h	83200	291346	52-234

CDC 9775 partitions

disk	start	length	cyls
hp?a	0	15884	0-12
hp?b	16640	66880	13-65
hp?c	0	1079040	0-842
hp?d	376320	15884	294-306
hp?e	392960	307200	307-546
hp?f	700160	378720	547-842
hp?g	376320	702560	294-842
hp?h	84480	291346	66-293

CDC 9730 partitions

disk	start	length	cyls
hp?a	0	15884	0-49
hp?b	16000	33440	50-154
hp?c	0	263360	0-822
hp?d	49600	15884	155-204
hp?e	65600	55936	205-379
hp?f	121600	141600	380-822

hp?g	49600	213600	155-822
AMPEX Capricorn partitions			
disk	start	length	cyls
hp?a	0	15884	0-31
hp?b	16384	33440	32-97
hp?c	0	524288	0-1023
hp?d	342016	15884	668-699
hp?e	358400	55936	700-809
hp?f	414720	109408	810-1023
hp?g	342016	182112	668-1023
hp?h	50176	291346	98-667
FUJITSU Eagle partitions			
disk	start	length	cyls
hp?a	0	15884	0-16
hp?b	16320	66880	17-86
hp?c	0	808320	0-841
hp?d	375360	15884	391-407
hp?e	391680	55936	408-727
hp?f	698880	109248	728-841
hp?g	375360	432768	391-841
hp?h	83520	291346	87-390
AMPEX 9300 partitions			
disk	start	length	cyl
hp?a	0	15884	0-26
hp?b	16416	33440	27-81
hp?c	0	495520	0-814
hp?d	341696	15884	562-588
hp?e	358112	55936	589-680
hp?f	414048	81312	681-814
hp?g	341696	153664	562-814
hp?h	49856	291346	82-561

It is unwise for all of these files to be present in one installation, since there is overlap in addresses and protection becomes a sticky matter. The hp?a partition is normally used for the root file system, the hp?b partition as a paging area, and the hp?c partition for pack-pack copying (it maps the entire disk). On disks larger than about 205 Megabytes, the hp?h partition is inserted prior to the hp?d or hp?g partition; the hp?g partition then maps the remainder of the pack. All disk partition tables are calculated using the *diskpart(8)* program.

FILES

/dev/hp[0-7][a-h] block files
 /dev/rhp[0-7][a-h] raw files

SEE ALSO

hk(4), uda(4), up(4)

DIAGNOSTICS

hp%d%c: hard error sn%d mbsr=%b er1=%b er2=%b. An unrecoverable error occurred during transfer of the specified sector of the specified disk partition. The MASSBUS status register is printed in hexadecimal and with the error bits decoded if any error bits other than MBEXC and DTABT are set. In any case the contents of the two error registers are also printed in octal and symbolically with bits decoded. (Note that er2 is what old rp06 manuals would call er3; the terminology is that of the rm disks). The error was either unrecoverable, or a large number of retry attempts (including offset positioning and drive recalibration) could not recover the error.

hp%d: write locked. The write protect switch was set on the drive when a write was attempted. The write operation is not recoverable.

hp%d: not ready. The drive was spun down or off line when it was accessed. The i/o operation is not recoverable.

hp%d%c: soft ecc sn%d. A recoverable ECC error occurred on the specified sector of the specified disk partition. This happens normally a few times a week. If it happens more frequently than this the sectors where the errors are occurring should be checked to see if certain cylinders on the pack, spots on the carriage of the drive or heads are indicated.

During autoconfiguration one of the following messages may appear on the console indicating the appropriate drive type was recognized. The last message indicates the drive is of a unknown type.

hp%d: 9775 (direct).

hp%d: 9730 (direct).

hp%d: 9300.

hp%d: 9762.

hp%d: capricorn.

hp%d: eagle.

hp%d: ntracks %d, nsectors %d: unknown device.

BUGS

In raw I/O *read* and *write(2)* 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.

DEC-standard error logging should be supported.

A program to analyze the logged error information (even in its present reduced form) is needed.

The partition tables for the file systems should be read off of each pack, as they are never quite what any single installation would prefer, and this would make packs more portable.

NAME

ht — TM-03/TE-16,TU-45,TU-77 MASSBUS magtape interface

SYNOPSIS

master ht0 at mba? drive ?
tape tu0 at ht0 slave 0

DESCRIPTION

The tm-03/transport combination provides a standard tape drive interface as described in *mtio(4)*. All drives provide both 800 and 1600 bpi; the TE-16 runs at 45 ips, the TU-45 at 75 ips, while the TU-77 runs at 125 ips and autoloads tapes.

SEE ALSO

mt(1), tar(1), tp(1), mtio(4), tm(4), ts(4), mt(4), ut(4)

DIAGNOSTICS

tu%d: no write ring. An attempt was made to write on the tape drive when no write ring was present; this message is written on the terminal of the user who tried to access the tape.

tu%d: not online. An attempt was made to access the tape while it was offline; this message is written on the terminal of the user who tried to access the tape.

tu%d: can't switch density in mid-tape. An attempt was made to write on a tape at a different density than is already recorded on the tape. This message is written on the terminal of the user who tried to switch the density.

tu%d: hard error bn%d mbr=%b er=%b ds=%b. A tape error occurred at block *bn*; the *ht* error register and drive status register are printed in octal with the bits symbolically decoded. Any error is fatal on non-raw tape; when possible the driver will have retried the operation which failed several times before reporting the error.

BUGS

If any non-data error is encountered on non-raw tape, it refuses to do anything more until closed.

NAME

hy — Network Systems Hyperchannel interface

SYNOPSIS

device hy0 at uba0 csr 0172410 vector hyint

DESCRIPTION

The *hy* interface provides access to a Network Systems Corporation Hyperchannel Adapter.

The network to which the interface is attached is specified at boot time with an SIOCSIFADDR ioctl. The host's address is discovered by reading the adapter status register. The interface will not transmit or receive packets until the network number is known.

DIAGNOSTICS

hy%d: unit number 0x%x port %d type %x microcode level 0x%x. Identifies the device during autoconfiguration.

hy%d: can't handle af%d. The interface was handed a message with addresses formatted in an unsuitable address family; the packet was dropped.

hy%d: can't initialize. The interface was unable to allocate UNIBUS resources. This is usually due to having too many network devices on an 11/750 where there are only 3 buffered data paths.

hy%d: NEX - Non Existent Memory. Non existent memory error returned from hardware.

hy%d: BAR overflow. Bus address register overflow error returned from hardware.

hy%d: Power Off bit set, trying to reset. Adapter has lost power, driver will reset the bit and see if power is still out in the adapter.

hy%d: Power Off Error, network shutdown. Power was really off in the adapter, network connections are dropped. Software does not shut down the network unless power has been off for a while.

hy%d: RECDV MP > MPSIZE (%d). A message proper was received that is too big. Probable a driver bug. Shouldn't happen.

hy%d: xmit error — len > hy_olen [%d > %d]. Probable driver error. Shouldn't happen.

hy%d: DRIVER BUG — INVALID STATE %d. The driver state machine reached a non-existent state. Definite driver bug.

hy%d: watchdog timer expired. A command in the adapter has taken too long to complete. Driver will abort and retry the command.

hy%d: adapter power restored. Software was able to reset the power off bit, indicating that the power has been restored.

SEE ALSO

intro(4N), inet(4F)

BUGS

If the adapter does not respond to the status command issued during autoconfigure, the adapter is assumed down. A reboot is required to recognize it.

The adapter power fail interrupt seems to occur sporadically when power has, in fact, not failed. The driver will believe that power has failed only if it can not reset the power fail latch after a "reasonable" time interval. These seem to appear about 2-4 times a day on some machines. There seems to be no correlation with adapter rev level, number of ports used etc. and whether a machine will get these "bogus powerfails". They don't seem to cause any real problems so they have been ignored.

NAME

ik — Ikonas frame buffer, graphics device interface

SYNOPSIS

device lk0 at uba? csr 0172460 vector ikintr

DESCRIPTION

Ik provides an interface to an Ikonas frame buffer graphics device. Each minor device is a different frame buffer interface board. When the device is opened, its interface registers are mapped, via virtual memory, into the user processes address space. This allows the user process very high bandwidth to the frame buffer with no system call overhead.

Bytes written or read from the device are DMA'ed from or to the interface. The frame buffer XY address, its addressing mode, etc. must be set up by the user process before calling write or read.

Other communication with the driver is via *ioctl*s. The *IK_GETADDR* *ioctl* returns the virtual address where the user process can find the interface registers. The *IK_WAITINT* *ioctl* suspends the user process until the ikonas device has interrupted (for whatever reason — the user process has to set the interrupt enables).

FILES

/dev/ik

DIAGNOSTICS

None.

BUGS

An invalid access (e.g., longword) to a mapped interface register can cause the system to crash with a machine check. A user process could possibly cause infinite interrupts hence bringing things to a crawl.

NAME

il — Interlan 10 Mb/s Ethernet interface

SYNOPSIS

device il0 at uba0 csr 161000 vector ilrint ilcint

DESCRIPTION

The *il* interface provides access to a 10 Mb/s Ethernet network through an Interlan controller.

The host's Internet address is specified at boot time with an `SIOCSIFADDR` ioctl. The *ec* interface employs the address resolution protocol described in *arp(4P)* to dynamically map between Internet and Ethernet addresses on the local network.

The interface normally tries to use a "trailer" encapsulation to minimize copying data on input and output. This may be disabled, on a per-interface basis, by setting the `IFF_NOTRAILERS` flag with an `SIOCSIFFLAGS` ioctl.

DIAGNOSTICS

il%d: input error. The hardware indicated an error in reading a packet off the cable or an illegally sized packet.

il%d: can't handle af%d. The interface was handed a message with addresses formatted in an unsuitable address family; the packet was dropped.

SEE ALSO

intro(4N), *inet(4F)*, *arp(4P)*

BUGS

The PUP protocol family should be added.

NAME

imp — 1822 network interface

SYNOPSIS

pseudo-device imp

DESCRIPTION

The *imp* interface, as described in BBN Report 1822, provides access to an intelligent message processor normally used when participating in the Department of Defense ARPA network. The network interface communicates through a device controller, usually an ACC LH/DH or DEC IMP-11A, with the IMP. The interface is “reliable” and “flow-controlled” by the host-IMP protocol.

To configure IMP support, one of *acc(4)* and *css(4)* must be included. The network number on which the interface resides is specified at boot time using the SIOCSIFADDR ioctl. The host number is discovered through receipt of NOOP messages from the IMP.

The network interface is always in one of four states: up, down, initializing, or going down. When the system is booted, the interface is marked down. If the hardware controller is successfully probed, the interface enters the initializing state and transmits three NOOP messages to the IMP. It then waits for the IMP to respond with two or more NOOP messages in reply. When it receives these messages it enters the up state. The going down state is entered only when notified by the IMP of an impending shutdown. Packets may be sent through the interface only while it is in the up state. Packets received in any other state are dropped with the error ENETDOWN returned to the caller.

DIAGNOSTICS

imp%d: leader error. The IMP reported an error in a leader (1822 message header). This causes the interface to be reset and any packets queued up for transmission to be purged.

imp%d: going down in 30 seconds.

imp%d: going down for hardware PM.

imp%d: going down for reload software.

imp%d: going down for emergency reset. The Network Control Center (NCC) is manipulating the IMP. By convention these messages are reported to all hosts on an IMP.

imp%d: reset (host %d/imp %d). The host has received a NOOP message which caused it to reset its notion of its current address. This normally occurs at boot time, though it may also occur while the system is running (for example, if the IMP-controller cable is disconnected, then reconnected).

imp%d: host dead. The IMP has noted a host, to which a prior packet was sent, is not up.

imp%d: host unreachable. The IMP has discovered a host, to which a prior packet was sent, is not accessible.

imp%d: data error. The IMP noted an error in data transmitted. The host-IMP interface is reset and the host enters the init state (awaiting NOOP messages).

imp%d: interface reset. The reset process has been completed.

imp%d: marked down. After receiving a “going down in 30 seconds” message, and waiting 30 seconds, the host has marked the IMP unavailable. Before packets may be sent to the IMP again, the IMP must notify the host, through a series of NOOP messages, that it is back up.

imp%d: can't handle af%d. The interface was handed a message with addresses formatting in an unsuitable address family; the packet was dropped.

SEE ALSO

intro(4N), inet(4F), acc(4), css(4)

NAME

`imp` – IMP raw socket interface

SYNOPSIS

```
#include <sys/socket.h>
#include <netinet/in.h>
#include <netimp/if_imp.h>

s = socket(AF_IMPLINK, SOCK_RAW, IMPLINK_IP);
```

DESCRIPTION

The raw `imp` socket provides direct access to the `imp(4)` network interface. Users send packets through the interface using the `send(2)` calls, and receive packets with the `recv(2)`, calls. All outgoing packets must have space for an 1822 96-bit leader on the front. Likewise, packets received by the user will have this leader on the front. The 1822 leader and the legal values for the various fields are defined in the include file `<netimp/if_imp.h>`.

The raw `imp` interface automatically installs the length and destination address in the 1822 leader of all outgoing packets; these need not be filled in by the user.

DIAGNOSTICS

An operation on a socket may fail with one of the following errors:

- [EISCONN] when trying to establish a connection on a socket which already has one, or when trying to send a datagram with the destination address specified and the socket is already connected;
- [ENOTCONN] when trying to send a datagram, but no destination address is specified, and the socket hasn't been connected;
- [ENOBUFS] when the system runs out of memory for an internal data structure;
- [EADDRNOTAVAIL] when an attempt is made to create a socket with a network address for which no network interface exists.

SEE ALSO

`intro(4N)`, `inet(4F)`, `imp(4)`

4

NAME

inet — Internet protocol family

SYNOPSIS

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

DESCRIPTION

The Internet protocol family is a collection of protocols layered atop the *Internet Protocol* (IP) transport layer, and utilizing the Internet address format. The Internet family provides protocol support for the SOCK_STREAM, SOCK_DGRAM, and SOCK_RAW socket types; the SOCK_RAW interface provides access to the IP protocol.

ADDRESSING

Internet addresses are four byte quantities, stored in network standard format (on the VAX these are word and byte reversed). The include file *<netinet/in.h>* defines this address as a discriminated union.

Sockets bound to the Internet protocol family utilize the following addressing structure,

```
struct sockaddr_in {
    short    sin_family;
    u_short  sin_port;
    struct   in_addr sin_addr;
    char     sin_zero[8];
};
```

Sockets may be created with the address INADDR_ANY to effect “wildcard” matching on incoming messages.

PROTOCOLS

The Internet protocol family is comprised of the IP transport protocol, Internet Control Message Protocol (ICMP), Transmission Control Protocol (TCP), and User Datagram Protocol (UDP). TCP is used to support the SOCK_STREAM abstraction while UDP is used to support the SOCK_DGRAM abstraction. A raw interface to IP is available by creating an Internet socket of type SOCK_RAW. The ICMP message protocol is not directly accessible.

SEE ALSO

tcp(4P), udp(4P), ip(4P)

CAVEAT

The Internet protocol support is subject to change as the Internet protocols develop. Users should not depend on details of the current implementation, but rather the services exported.

NAME

ip — Internet Protocol

SYNOPSIS

```
#include <sys/socket.h>
#include <netinet/in.h>
s = socket(AF_INET, SOCK_RAW, 0);
```

DESCRIPTION

IP is the transport layer protocol used by the Internet protocol family. It may be accessed through a “raw socket” when developing new protocols, or special purpose applications. IP sockets are connectionless, and are normally used with the *sendto* and *recvfrom* calls, though the *connect(2)* call may also be used to fix the destination for future packets (in which case the *read(2)* or *recv(2)* and *write(2)* or *send(2)* system calls may be used).

Outgoing packets automatically have an IP header prepended to them (based on the destination address and the protocol number the socket is created with). Likewise, incoming packets have their IP header stripped before being sent to the user.

DIAGNOSTICS

A socket operation may fail with one of the following errors returned:

- [EISCONN] when trying to establish a connection on a socket which already has one, or when trying to send a datagram with the destination address specified and the socket is already connected;
- [ENOTCONN] when trying to send a datagram, but no destination address is specified, and the socket hasn't been connected;
- [ENOBUFS] when the system runs out of memory for an internal data structure;
- [EADDRNOTAVAIL] when an attempt is made to create a socket with a network address for which no network interface exists.

SEE ALSO

send(2), *recv(2)*, *intro(4N)*, *inet(4F)*

BUGS

- One should be able to send and receive ip options.
- The protocol should be settable after socket creation.

NAME

kg — KL-11/DL-11W line clock

SYNOPSIS

device kg0 at uba0 csr 0176500 vector kglock

DESCRIPTION

A kl-11 or dl-11w can be used as an alternate real time clock source. When configured, certain system statistics and, optionally, system profiling work will be collected each time the clock interrupts. For optimum accuracy in profiling, the dl-11w should be configured to interrupt at the highest possible priority level. The *kg* device driver automatically calibrates itself to the line clock frequency.

SEE ALSO

kgmon(8), config(8)

NAME

lo — software loopback network interface

SYNOPSIS

pseudo-device loop

DESCRIPTION

The *loop* interface is a software loopback mechanism which may be used for performance analysis, software testing, and/or local communication. By default, the loopback interface is accessible at address 127.0.0.1 (non-standard); this address may be changed with the SIOCSI-FADDR ioctl.

DIAGNOSTICS

lo%d: can't handle af%d. The interface was handed a message with addresses formatted in an unsuitable address family; the packet was dropped.

SEE ALSO

intro(4N), inet(4F)

BUGS

It should handle all address and protocol families. An approved network address should be reserved for this interface.

NAME

lp — line printer

SYNOPSIS

device lp0 at uba0 csr 0177514 vector lpintr

DESCRIPTION

Lp provides the interface to any of the standard DEC line printers on an LP-11 parallel interface. When it is opened or closed, a suitable number of page ejects is generated. Bytes written are printed.

The unit number of the printer is specified by the minor device after removing the low 3 bits, which act as per-device parameters. Currently only the lowest of the low three bits is interpreted: if it is set, the device is treated as having a 64-character set, rather than a full 96-character set. In the resulting half-ASCII mode, lower case letters are turned into upper case and certain characters are escaped according to the following table:

{	{
}	}
,	,
	+
_	^

The driver correctly interprets carriage returns, backspaces, tabs, and form feeds. Lines longer than the maximum page width are truncated. The default page width is 132 columns. This may be overridden by specifying, for example, "flags 256" .

FILES

/dev/lp

SEE ALSO

lpr(1)

DIAGNOSTICS

None.



NAME

mem, *kmem* — main memory

DESCRIPTION

Mem is a special file that is an image of the main memory of the computer. It may be used, for example, to examine (and even to patch) the system.

Byte addresses in *mem* are interpreted as physical memory addresses. References to non-existent locations cause errors to be returned.

Examining and patching device registers is likely to lead to unexpected results when read-only or write-only bits are present.

The file *kmem* is the same as *mem* except that kernel virtual memory rather than physical memory is accessed.

On PDP11's, the I/O page begins at location 0160000 of *kmem* and per-process data for the current process begins at 0140000. On VAX 11/780 the I/O space begins at physical address 20000000(16); on an 11/750 I/O space addresses are of the form *fx*xxxx(16); on all VAX'en per-process data for the current process is at virtual 7ffff000(16).

FILES

/dev/*mem*
/dev/*kmem*

BUGS

On PDP11's and VAX's, memory files are accessed one byte at a time, an inappropriate method for some device registers.

NAME

mt — TM78/TU-78 MASSBUS magtape interface

SYNOPSIS

master mt0 at mba? drive ?
 tape mu0 at mt0 slave 0

DESCRIPTION

The tm78/tu-78 combination provides a standard tape drive interface as described in *mtio(4)*. Only 1600 and 6250 bpi are supported; the TU-78 runs at 125 ips and autoloads tapes.

SEE ALSO

mt(1), tar(1), tp(1), mtio(4), tm(4), ts(4), ut(4)

DIAGNOSTICS

mu%d: no write ring. An attempt was made to write on the tape drive when no write ring was present; this message is written on the terminal of the user who tried to access the tape.

mu%d: not online. An attempt was made to access the tape while it was offline; this message is written on the terminal of the user who tried to access the tape.

mu%d: can't switch density in mld-tape. An attempt was made to write on a tape at a different density than is already recorded on the tape. This message is written on the terminal of the user who tried to switch the density.

mu%d: hard error bn%d mbsr=%b er=%x ds=%b. A tape error occurred at block *bn*; the mt error register and drive status register are printed in octal with the bits symbolically decoded. Any error is fatal on non-raw tape; when possible the driver will have retried the operation which failed several times before reporting the error.

mu%d: blank tape. An attempt was made to read a blank tape (a tape without even end-of-file marks).

mu%d: offline. During an i/o operation the device was set offline. If a non-raw tape was used in the access it is closed.

BUGS

If any non-data error is encountered on non-raw tape, it refuses to do anything more until closed.

NAME

mtio — UNIX magtape interface

DESCRIPTION

The files *mt0*, ..., *mt15* refer to the UNIX magtape drives, which may be on the MASSBUS using the TM03 formatter *ht(4)*, or TM78 formatter, *mt(4)*, or on the UNIBUS using either the TM11 or TS11 formatters *tm(4)*, TU45 compatible formatters, *ut(4)*, or *ts(4)*. The following description applies to any of the transport/controller pairs. The files *mt0*, ..., *mt7* are 800bpi, *mt8*, ..., *mt15* are 1600bpi, and *mt16*, ..., *mt23* are 6250bpi. (But note that only 1600 bpi is available with the TS11.) The files *mt0*, ..., *mt3*, *mt8*, ..., *mt11*, and *mt16*, ..., *mt19* are rewound when closed; the others are not. When a file open for writing is closed, two end-of-files are written. If the tape is not to be rewound it is positioned with the head between the two tapemarks.

A standard tape consists of a series of 1024 byte records terminated by an end-of-file. To the extent possible, the system makes it possible, if inefficient, to treat the tape like any other file. Seeks have their usual meaning and it is possible to read or write a byte at a time. Writing in very small units is inadvisable, however, because it tends to create monstrous record gaps.

The *mt* files discussed above are useful when it is desired to access the tape in a way compatible with ordinary files. When foreign tapes are to be dealt with, and especially when long records are to be read or written, the 'raw' interface is appropriate. The associated files are named *rmt0*, ..., *rmt23*, but the same minor-device considerations as for the regular files still apply. A number of other ioctl operations are available on raw magnetic tape. The following definitions are from `<sys/mtio.h>`:

```
/*
 * Structures and definitions for mag tape io control commands
 */

/* structure for MTIOCTOP - mag tape op command */
struct mtop {
    short   mt_op;           /* operations defined below */
    daddr_t mt_count;       /* how many of them */
};

/* operations */
#define MTWEOF    0        /* write an end-of-file record */
#define MTFSF    1        /* forward space file */
#define MTBSF    2        /* backward space file */
#define MTFSR    3        /* forward space record */
#define MTBSR    4        /* backward space record */
#define MTREW    5        /* rewind */
#define MTOFFL   6        /* rewind and put the drive offline */
#define MTNOP    7        /* no operation, sets status only */

/* structure for MTIOCGET - mag tape get status command */

struct mtget {
    short   mt_type;        /* type of magtape device */
    /* the following two registers are grossly device dependent */
    short   mt_dsreg;       /* "drive status" register */
    short   mt_erreg;       /* "error" register */
    /* end device-dependent registers */
    short   mt_resid;       /* residual count */
};
```

```

/* the following two are not yet implemented */
    daddr_t mt_fileno; /* file number of current position */
    daddr_t mt_blkno; /* block number of current position */
/* end not yet implemented */
};

/*
 * Constants for mt_type byte
 */
#define MT_ISTS          0x01
#define MT_ISHT         0x02
#define MT_ISTM         0x03
#define MT_ISMT         0x04
#define MT_ISUT         0x05
#define MT_ISCPC        0x06
#define MT_ISAR         0x07

/* mag tape io control commands */
#define MTIOCTOP        _IOW(m, 1, struct mtop) /* do a mag tape op */
#define MTIOCGET        _IOR(m, 2, struct mtget) /* get tape status */

#ifdef KERNEL
#define DEFTAPE         "/dev/rmt12"
#endif

```

Each *read* or *write* call reads or writes the next record on the tape. In the write case the record has the same length as the buffer given. During a read, the record size is passed back as the number of bytes read, provided it is no greater than the buffer size; if the record is long, an error is indicated. In raw tape I/O seeks are ignored. A zero byte count is returned when a tape mark is read, but another read will fetch the first record of the new tape file.

FILES

/dev/mt?
/dev/rmt?

SEE ALSO

mt(1), tar(1), tp(1), ht(4), tm(4), ts(4), mt(4), ut(4)

BUGS

The status should be returned in a device independent format.

NAME

null — data sink

DESCRIPTION

Data written on a null special file is discarded.

Reads from a null special file always return 0 bytes.

FILES

/dev/null

NAME

`pcl` — DEC CSS PCL-11 B Network Interface

SYNOPSIS

`device pcl0 at uba? csr 164200 vector pclxint pclrint`

DESCRIPTION

The `pcl` device provides an IP-only interface to the DEC CSS PCL-11 time division multiplexed network bus. The controller itself is not accessible to users.

The hosts's address is specified with the `SIOCIFADDR` ioctl. The interface will not transmit or receive any data before its address is defined.

As the PCL-11 hardware is only capable of having 15 interfaces per network, a single-byte host-on-network number is used, with range [1..15] to match the TDM bus addresses of the interfaces.

The interface currently only supports the Internet protocol family and only provides "natural" (header) encapsulation.

DIAGNOSTICS

pcl%d: can't init. Insufficient UNIBUS resources existed to initialize the device. This is likely to occur when the device is run on a buffered data path on an 11/750 and other network interfaces are also configured to use buffered data paths, or when it is configured to use buffered data paths on an 11/730 (which has none).

pcl%d: can't handle af%d. The interface was handed a message with addresses formatted in an unsuitable address family; the packet was dropped.

pcl%d: stray xmit interrupt. An interrupt occurred when no output had previously been started.

pcl%d: master. The TDM bus had no station providing "bus master" timing signals, so this interface has assumed the "master" role. This message should only appear at most once per UNIBUS INIT on a single system. Unless there is a hardware failure, only one station may be master at a time.

pcl%d: send error, tcr=%b, tsr=%b. The device indicated a problem sending data on output. If a "receiver offline" error is detected, it is not normally logged unless the option `PCL_TESTING` has been selected, as this causes a lot of console chatter when sending to a down machine. However, this option is quite useful when debugging problems with the PCL interfaces.

pcl%d: rcv error, rcr=%b rsr=%b. The device indicated a problem receiving data on input.

pcl%d: bad len=%d. An input operation resulted in a data transfer of less than 0 or more than 1008 bytes of data into memory (according to the word count register). This should never happen as the maximum size of a PCL message has been agreed upon to be 1008 bytes (same as ArpaNet message).

SEE ALSO

`intro(4N)`, `inet(4F)`

NAME

ps — Evans and Sutherland Picture System 2 graphics device interface

SYNOPSIS

device ps0 at uba? csr 0172460 vector psintr

DESCRIPTION

The *ps* driver provides access to an Evans and Sutherland Picture System 2 graphics device. Each minor device is a new PS2. When the device is opened, its interface registers are mapped, via virtual memory, into a user process's address space. This allows the user process very high bandwidth to the device with no system call overhead.

DMA to and from the PS2 is not supported. All read and write system calls will fail. All data is moved to and from the PS2 via programmed I/O using the device's interface registers.

Commands are fed to and from the driver using the following ioctl's:

PSIOGETADDR

Returns the virtual address through which the user process can access the device's interface registers.

PSIOAUTOREFRESH

Start auto refreshing the screen. The argument is an address in user space where the following data resides. The first longword is a *count* of the number of static refresh buffers. The next *count* longwords are the addresses in refresh memory where the refresh buffers lie. The driver will cycle thru these refresh buffers displaying them one by one on the screen.

PSIOAUTOMAP

Start automatically passing the display file thru the matrix processor and into the refresh buffer. The argument is an address in user memory where the following data resides. The first longword is a *count* of the number of display files to operate on. The next *count* longwords are the address of these display files. The final longword is the address in refresh buffer memory where transformed coordinates are to be placed if the driver is not in double buffer mode (see below).

PSIODOUBLEBUFFER

Cause the driver to double buffer the output from the map that is going to the refresh buffer. The argument is again a user space address where the real arguments are stored. The first argument is the starting address of refresh memory where the two double buffers are located. The second argument is the length of each double buffer. The refresh mechanism displays the current double buffer, in addition to its static refresh lists, when in double buffer mode.

PSIOSINGLEREFRESH

Single step the refresh process. That is, the driver does not continually refresh the screen.

PSIOSINGLEMAP

Single step the matrix process. The driver does not automatically feed display files thru the matrix unit.

PSIOSINGLEBUFFER

Turn off double buffering.

PSIOTIMEREFRESH

The argument is a count of the number of refresh interrupts to take before turning off the screen. This is used to do time exposures.

PSIOWAITREFRESH

Suspend the user process until a refresh interrupt has occurred. If in **TIMEREFRESH**

mode, suspend until count refreshes have occurred.

PSIOSTOPREFRESH

Wait for the next refresh, stop all refreshes, and then return to user process.

PSIOWAITMAP

Wait until a map done interrupt has occurred.

PSIOSTOPMAP

Wait for a map done interrupt, do not restart the map, and then return to the user.

FILES

/dev/ps

DIAGNOSTICS

ps device intr.

ps dma intr. An interrupt was received from the device. This shouldn't happen, check your device configuration for overlapping interrupt vectors.

BUGS

An invalid access (e.g., longword) to a mapped interface register can cause the system to crash with a machine check. A user process could possibly cause infinite interrupts hence bringing things to a crawl.

NAME

pty — pseudo terminal driver

SYNOPSIS

pseudo-device pty

DESCRIPTION

The *pty* driver provides support for a device-pair termed a *pseudo terminal*. A pseudo terminal is a pair of character devices, a *master* device and a *slave* device. The slave device provides processes an interface identical to that described in *ty(4)*. However, whereas all other devices which provide the interface described in *ty(4)* have a hardware device of some sort behind them, the slave device has, instead, another process manipulating it through the master half of the pseudo terminal. That is, anything written on the master device is given to the slave device as input and anything written on the slave device is presented as input on the master device.

In configuring, if no optional “count” is given in the specification, 16 pseudo terminal pairs are configured.

The following *ioctl* calls apply only to pseudo terminals:

TIOCSTOP

Stops output to a terminal (e.g. like typing ^S). Takes no parameter.

TIOCPKT

Restarts output (stopped by TIOCSTOP or by typing ^S). Takes no parameter.

TIOCPKT

Enable/disable *packet* mode. Packet mode is enabled by specifying (by reference) a nonzero parameter and disabled by specifying (by reference) a zero parameter. When applied to the master side of a pseudo terminal, each subsequent *read* from the terminal will return data written on the slave part of the pseudo terminal preceded by a zero byte (symbolically defined as TIOCPKT_DATA), or a single byte reflecting control status information. In the latter case, the byte is an inclusive-or of zero or more of the bits:

TIOCPKT_FLUSHREAD

whenever the read queue for the terminal is flushed.

TIOCPKT_FLUSHWRITE

whenever the write queue for the terminal is flushed.

TIOCPKT_STOP

whenever output to the terminal is stopped by a ^S.

TIOCPKT_START

whenever output to the terminal is restarted.

TIOCPKT_DOSTOP

whenever *t_stop* is ^S and *t_start* is ^Q.

TIOCPKT_NOSTOP

whenever the start and stop characters are not ^S/^Q.

This mode is used by *rlogin(1C)* and *rlogind(8C)* to implement a remote-echoed, locally ^S/^Q flow-controlled remote login with proper back-flushing of output; it can be used by other similar programs.

TIOCREMOTE

A mode for the master half of a pseudo terminal, independent of TIOCPKT. This mode causes input to the pseudo terminal to be flow controlled and not input edited (regardless of the terminal mode). Each write to the control terminal produces a record boundary for the process reading the terminal. In normal usage, a write of data is like the data typed as a line on the terminal; a write of 0 bytes is like typing an end-of-file

character. TIOCREMOTE can be used when doing remote line editing in a window manager, or whenever flow controlled input is required.

FILES

/dev/pty[p-r][0-9a-f] master pseudo terminals
/dev/tty[p-r][0-9a-f] slave pseudo terminals

DIAGNOSTICS

None.

BUGS

It is not possible to send an EOT.

NAME

pup — Xerox PUP-I protocol family

SYNOPSIS

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

DESCRIPTION

The PUP-I protocol family is a collection of protocols layered atop the PUP Level-0 packet format, and utilizing the PUP Internet address format. The PUP family is currently supported only by a raw interface.

ADDRESSING

PUP addresses are composed of network, host, and port portions. The include file `<netpup/pup.h>` defines this address as,

```
struct pupport {
    u_char    pup_net;
    u_char    pup_host;
    u_char    pup_socket[4];
};
```

Sockets bound to the PUP protocol family utilize the following addressing structure,

```
struct sockaddr_pup {
    short     spup_family;
    short     spup_zero1;
    u_char    spup_net;
    u_char    spup_host;
    u_char    spup_sock[4];
    char      spup_zero2[4];
};
```

HEADERS

The current PUP support provides only raw access to the 3Mb/s Ethernet. Packets sent through this interface must have space for the following packet header present at the front of the message,

```
struct pup_header {
    u_short   pup_length;
    u_char    pup_tcontrol;    /* transport control */
    u_char    pup_type;        /* protocol type */
    u_long    pup_id;          /* used by protocols */
    u_char    pup_dnet;        /* destination */
    u_char    pup_dhost;
    u_char    pup_dsock[4];
    u_char    pup_snet;        /* source */
    u_char    pup_shost;
    u_char    pup_ssock[4];
};
```

The sender should fill in the `pup_tcontrol`, `pup_type`, and `pup_id` fields. The remaining fields are filled in by the system. The system checks the message to insure its size is valid and, calculates a checksum for the message. If no checksum should be calculated, the checksum field (the last 16-bit word in the message) should be set to PUP_NOCKSUM.

The *pup_tcontrol* field is restricted to be 0 or PUP_TRACE; PUP_TRACE indicates packet tracing should be performed. The *pup_type* field may not be 0.

On input, the entire packet, including header, is provided the user. No checksum validation is performed.

SEE ALSO

intro(4N), pup(4P), en(4)

BUGS

The only interface which currently supports use of pup's is the Xerox 3 Mb/s *en(4)* interface.

With the release of the second generation, PUP-II, protocols it is not clear what future PUP-I has. Consequently, there has been little motivation to provide extensive kernel support.

NAME

pup — raw PUP socket interface

SYNOPSIS

```
#include <sys/socket.h>
#include <netpup/pup.h>
socket(AF_PUP, SOCK_RAW, PUPPROTO_BSP);
```

DESCRIPTION

A raw pup socket provides PUP-I access to an Ethernet network. Users send packets using the *sendto* call, and receive packets with the *recvfrom* call. All outgoing packets must have space present at the front of the packet to allow the PUP header to be filled in. The header format is described in *pup(4F)*. Likewise, packets received by the user will have the PUP header on the front. The PUP header and legal values for the various fields are defined in the include file *<netpup/pup.h>*.

The raw pup interface automatically installs the length and source and destination addresses in the PUP header of all outgoing packets; these need not be filled in by the user. The only control bit that may be set in the *tcontrol* field of outgoing packets is the "trace" bit. A checksum is calculated unless the sender sets the checksum field to PUP_NOCKSUM.

DIAGNOSTICS

A socket operation may fail and one of the following will be returned:

- [EISCONN] when trying to establish a connection on a socket which already has one, or when trying to send a datagram with the destination address specified and the socket is already connected;
- [ENOTCONN] when trying to send a datagram, but no destination address is specified, and the socket hasn't been connected;
- [ENOBUFS] when the system runs out of memory for an internal data structure;
- [EADDRNOTAVAIL] when an attempt is made to create a socket with a network address for which no network interface exists.

A *sendto* operation may fail if one of the following is true:

- [EINVAL] Insufficient space was left by the user for the PUP header.
- [EINVAL] The *pup_type* field was 0 or the *pup_tcontrol* field had a bit other than PUP_TRACE set.
- [EMSGSIZE] The message was not an even number of bytes, smaller than MINPUPSIZ, or large than MAXPUPSIZ.
- [ENETUNREACH] The destination address was on a network which was not directly reachable (the raw interface provides no routing support).

SEE ALSO

send(2), *recv(2)*, *intro(4N)*, *pup(4F)*

BUGS

The interface is untested against other PUP implementations.

NAME

rx — DEC RX02 floppy disk interface

SYNOPSIS

```
controller fx0 at uba0 csr 0177170 vector rxintr
disk rx0 at fx0 drive 0
disk rx1 at fx0 drive 1
```

DESCRIPTION

The *rx* device provides access to a DEC RX02 floppy disk unit with M8256 interface module (RX211 configuration). The RX02 uses 8-inch, single-sided, soft-sectored floppy disks (with pre-formatted industry-standard headers) in either single or double density.

Floppy disks handled by the RX02 contain 77 tracks, each with 26 sectors (for a total of 2,002 sectors). The sector size is 128 bytes for single density, 256 bytes for double density. Single density disks are compatible with the RX01 floppy disk unit and with IBM 3740 Series Diskette 1 systems.

In addition to normal ('block' and 'raw') i/o, the driver supports formatting of disks for either density and the ability to invoke a 2 for 1 interleaved sector mapping compatible with the DEC operating system RT-11.

The minor device number is interpreted as follows:

Bit	Description
0	Sector interleaving (1 disables interleaving)
1	Logical sector 1 is on track 1 (0 no, 1 yes)
2	Not used, reserved
Other	Drive number

The two drives in a single RX02 unit are treated as two disks attached to a single controller. Thus, if there are two RX02's on a system, the drives on the first RX02 are "rx0" and "rx1", while the drives on the second are "rx2" and "rx3".

When the device is opened, the density of the disk currently in the drive is automatically determined. If there is no floppy in the device, open will fail.

The interleaving parameters are represented in raw device names by the letters 'a' through 'd'. Thus, unit 0, drive 0 is called by one of the following names:

Mapping	Device name	Starting track
interleaved	/dev/rxx0a	0
direct	/dev/rxx0b	0
interleaved	/dev/rxx0c	1
direct	/dev/rxx0d	1

The mapping used on the 'c' device is compatible with the DEC operating system RT-11. The 'b' device accesses the sectors of the disk in strictly sequential order. The 'a' device is the most efficient for disk-to-disk copying. This mapping is always used by the block device.

I/O requests must start on a sector boundary, involve an integral number of complete sectors, and not go off the end of the disk.

NOTES

Even though the storage capacity on a floppy disk is quite small, it is possible to make filesystems on double density disks. For example, the command

```
% mkfs /dev/rx0 1001 13 1 4096 512 32 0 4
```

makes a file system on the double density disk in rx0 with 436 kbytes available for file storage. Using *tar*(1) gives a more efficient utilization of the available space for file storage. Single density diskettes do not provide sufficient storage capacity to hold file systems.

A number of *ioctl(2)* calls apply to the rx devices, and have the form

```
#include <vaxuba/rxreg.h>
ioctl(fildes, code, arg)
int *arg;
```

The applicable codes are:

RXIOC_FORMAT Format the diskette. The density to use is specified by the *arg* argument, zero gives single density while non-zero gives double density.

RXIOC_GETDENS

Return the density of the diskette (zero or non-zero as above).

RXIOC_WDDMK On the next write, include a *deleted data address mark* in the header of the first sector.

RXIOC_RDDMK Return non-zero if the last sector read contained a *deleted data address mark* in its header, otherwise return 0.

ERRORS

The following errors may be returned by the driver:

- [ENODEV] Drive not ready; usually because no disk is in the drive or the drive door is open.
- [ENXIO] Nonexistent drive (on open); offset is too large or not on a sector boundary or byte count is not a multiple of the sector size (on read or write); or bad (undefined) *ioctl* code.
- [EIO] A physical error other than "not ready", probably bad media or unknown format.
- [EBUSY] Drive has been opened for exclusive access.
- [EBADF] No write access (on format), or wrong density; the latter can only happen if the disk is changed without *closing* the device (i.e., calling *close(2)*).

FILES

```
/dev/rx?
/dev/rxx?[a-d]
```

SEE ALSO

rxformat(8V), *newfs(8)*, *mkfs(8)*, *tar(1)*, *arff(8V)*

DIAGNOSTICS

rx%d: hard error, trk %d psec %d cs=%b, db=%b, err=%x, %x, %x, %x. An unrecoverable error was encountered. The track and physical sector numbers, the device registers and the extended error status are displayed.

rx%d: state %d (reset). The driver entered a bogus state. This should not happen.

BUGS

A floppy may not be formatted if the header info on sector 1, track 0 has been damaged. Hence, it is not possible to format completely degaussed disks or disks with other formats than the two known by the hardware.

If the drive subsystem is powered down when the machine is booted, the controller won't interrupt.

NAME

tcp — Internet Transmission Control Protocol

SYNOPSIS

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

s = socket(AF_INET, SOCK_STREAM, 0);
```

DESCRIPTION

The TCP protocol provides reliable, flow-controlled, two-way transmission of data. It is a byte-stream protocol used to support the SOCK_STREAM abstraction. TCP uses the standard Internet address format and, in addition, provides a per-host collection of “port addresses”. Thus, each address is composed of an Internet address specifying the host and network, with a specific TCP port on the host identifying the peer entity.

Sockets utilizing the tcp protocol are either “active” or “passive”. Active sockets initiate connections to passive sockets. By default TCP sockets are created active; to create a passive socket the *listen(2)* system call must be used after binding the socket 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.

DIAGNOSTICS

A socket operation may fail with one of the following errors returned:

[EISCONN]	when trying to establish a connection on a socket which already has one;
[ENOBUFS]	when the system runs out of memory for an internal data structure;
[ETIMEDOUT]	when a connection was dropped due to excessive retransmissions;
[ECONNRESET]	when the remote peer forces the connection to be closed;
[ECONNREFUSED]	when the remote peer actively refuses connection establishment (usually because no process is listening to the port);
[EADDRINUSE]	when an attempt is made to create a socket with a port which has already been allocated;
[EADDRNOTAVAIL]	when an attempt is made to create a socket with a network address for which no network interface exists.

SEE ALSO

intro(4N), inet(4F)

BUGS

It should be possible to send and receive TCP options. The system always tries to negotiate the maximum TCP segment size to be 1024 bytes. This can result in poor performance if an intervening network performs excessive fragmentation.

NAME

tm — TM-11/TE-10 magtape interface

SYNOPSIS

controller **tm0** at uba? csr 0172520 vector **tmintr**
tape **te0** at **tm0** drive 0

DESCRIPTION

The tm-11/te-10 combination provides a standard tape drive interface as described in *mtio(4)*. Hardware implementing this on the VAX is typified by the Emulex TC-11 controller operating with a Kennedy model 9300 tape transport, providing 800 and 1600 bpi operation at 125 ips.

SEE ALSO

mt(1), tar(1), tp(1), mtio(4), ht(4), ts(4), mt(4), ut(4)

DIAGNOSTICS

te%d: no write ring. An attempt was made to write on the tape drive when no write ring was present; this message is written on the terminal of the user who tried to access the tape.

te%d: not online. An attempt was made to access the tape while it was offline; this message is written on the terminal of the user who tried to access the tape.

te%d: can't switch density in mid-tape. An attempt was made to write on a tape at a different density than is already recorded on the tape. This message is written on the terminal of the user who tried to switch the density.

te%d: hard error bn%d er=%b. A tape error occurred at block *bn*, the tm error register is printed in octal with the bits symbolically decoded. Any error is fatal on non-raw tape; when possible the driver will have retried the operation which failed several times before reporting the error.

te%d: lost interrupt. A tape operation did not complete within a reasonable time, most likely because the tape was taken off-line during rewind or lost vacuum. The controller should, but does not, give an interrupt in these cases. The device will be made available again after this message, but any current open reference to the device will return an error as the operation in progress aborts.

BUGS

If any non-data error is encountered on non-raw tape, it refuses to do anything more until closed.

NAME

ts — TS-11 magtape interface

SYNOPSIS

controller zs0 at uba? csr 0172520 vector tsintr
tape ts0 at zs0 drive 0

DESCRIPTION

The ts-11 combination provides a standard tape drive interface as described in *mtio(4)*. The ts-11 operates only at 1600 bpi, and only one transport is possible per controller.

SEE ALSO

mt(1), *tar(1)*, *tp(1)*, *mtio(4)*, *ht(4)*, *tm(4)*, *mt(4)*, *ut(4)*

DIAGNOSTICS

ts%d: no write ring. An attempt was made to write on the tape drive when no write ring was present; this message is written on the terminal of the user who tried to access the tape.

ts%d: not online. An attempt was made to access the tape while it was offline; this message is written on the terminal of the user who tried to access the tape.

ts%d: hard error bn%d xs0=%b. A hard error occurred on the tape at block *bn*; status register 0 is printed in octal and symbolically decoded as bits.

BUGS

If any non-data error is encountered on non-raw tape, it refuses to do anything more until closed.

The device lives at the same address as a tm-11 *tm(4)*; as it is very difficult to get this device to interrupt, a generic system assumes that a ts is present whenever no tm-11 exists but the csr responds and a ts-11 is configured. This does no harm as long as a non-existent ts-11 is not accessed.

NAME

tty — general terminal interface

SYNOPSIS

```
#include <sgtty.h>
```

DESCRIPTION

This section describes both a particular special file `/dev/tty` and the terminal drivers used for conversational computing.

Line disciplines.

The system provides different *line disciplines* for controlling communications lines. In this version of the system there are three disciplines available:

- old The old (standard) terminal driver. This is used when using the standard shell *sh*(1) and for compatibility with other standard version 7 UNIX systems.
- new A newer terminal driver, with features for job control; this must be used when using *cs**h*(1).
- net A line discipline used for networking and loading data into the system over communications lines. It allows high speed input at very low overhead, and is described in *bk*(4).

Line discipline switching is accomplished with the TIOCSETD *ioctl*:

```
int ldisc = LDISC; ioctl(filedes, TIOCSETD, &ldisc);
```

where LDISC is OTTYDISC for the standard tty driver, NTTYDISC for the new driver and NETLDISC for the networking discipline. The standard (currently old) tty driver is discipline 0 by convention. The current line discipline can be obtained with the TIOCGETD *ioctl*. Pending input is discarded when the line discipline is changed.

All of the low-speed asynchronous communications ports can use any of the available line disciplines, no matter what hardware is involved. The remainder of this section discusses the “old” and “new” disciplines.

The control terminal.

When a terminal file is opened, it causes the process to wait until a connection is established. In practice, user programs seldom open these files; they are opened by *init*(8) and become a user's standard input and output file.

If a process which has no control terminal opens a terminal file, then that terminal file becomes the control terminal for that process. The control terminal is thereafter inherited by a child process during a *fork*(2), even if the control terminal is closed.

The file `/dev/tty` is, in each process, a synonym for a *control terminal* associated with that process. It is useful for programs that wish to be sure of writing messages on the terminal no matter how output has been redirected. It can also be used for programs that demand a file name for output, when typed output is desired and it is tiresome to find out which terminal is currently in use.

Process groups.

Command processors such as *cs**h*(1) can arbitrate the terminal between different *jobs* by placing related jobs in a single process group and associating this process group with the terminal. A terminals associated process group may be set using the TIOCSPGRP *ioctl*(2):

```
ioctl(fildes, TIOCSPGRP, &pgrp)
```

or examined using TIOCGPGRP rather than TIOCSPGRP, returning the current process group in *pgrp*. The new terminal driver aids in this arbitration by restricting access to the terminal by processes which are not in the current process group; see **Job access control** below.

Modes.

The terminal drivers have three major modes, characterized by the amount of processing on the input and output characters:

- cooked The normal mode. In this mode lines of input are collected and input editing is done. The edited line is made available when it is completed by a newline or when an EOT (control-D, hereafter `^D`) is entered. A carriage return is usually made synonymous with newline in this mode, and replaced with a newline whenever it is typed. All driver functions (input editing, interrupt generation, output processing such as delay generation and tab expansion, etc.) are available in this mode.
- CBREAK This mode eliminates the character, word, and line editing input facilities, making the input character available to the user program as it is typed. Flow control, literal-next and interrupt processing are still done in this mode. Output processing is done.
- RAW This mode eliminates all input processing and makes all input characters available as they are typed; no output processing is done either.

The style of input processing can also be very different when the terminal is put in non-blocking i/o mode; see *fcntl(2)*. In this case a *read(2)* from the control terminal will never block, but rather return an error indication (EWOULDBLOCK) if there is no input available.

A process may also request a SIGIO signal be sent it whenever input is present. To enable this mode the FASYNC flag should be set using *fcntl(2)*.

Input editing.

A UNIX terminal ordinarily operates in full-duplex mode. Characters may be typed at any time, even while output is occurring, and are only lost when the system's character input buffers become completely choked, which is rare, or when the user has accumulated the maximum allowed number of input characters that have not yet been read by some program. Currently this limit is 256 characters. In the old terminal driver all the saved characters are thrown away when the limit is reached, without notice; the new driver simply refuses to accept any further input, and rings the terminal bell.

Input characters are normally accepted in either even or odd parity with the parity bit being stripped off before the character is given to the program. By clearing either the EVEN or ODD bit in the flags word it is possible to have input characters with that parity discarded (see the **Summary** below.)

In all of the line disciplines, it is possible to simulate terminal input using the TIOCSTI ioctl, which takes, as its third argument, the address of a character. The system pretends that this character was typed on the argument terminal, which must be the control terminal except for the super-user (this call is not in standard version 7 UNIX).

Input characters are normally echoed by putting them in an output queue as they arrive. This may be disabled by clearing the ECHO bit in the flags word using the *stty(3)* call or the TIOCSETN or TIOCSETP ioctls (see the **Summary** below).

In cooked mode, terminal input is processed in units of lines. A program attempting to read will normally be suspended until an entire line has been received (but see the description of SIGTTIN in **Modes** above and FIONREAD in **Summary** below.) No matter how many characters are requested in the read call, at most one line will be returned. It is not, however, necessary to read a whole line at once; any number of characters may be requested in a read, even one, without losing information.

During input, line editing is normally done, with the character '#' logically erasing the last character typed and the character '@' logically erasing the entire current input line. These are often reset on crt's, with ^H replacing #, and ^U replacing @. These characters never erase beyond the beginning of the current input line or an ^D. These characters may be entered literally by preceding them with '\'; in the old teletype driver both the '\' and the character entered literally will appear on the screen; in the new driver the '\' will normally disappear.

The drivers normally treat either a carriage return or a newline character as terminating an input line, replacing the return with a newline and echoing a return and a line feed. If the CRMOD bit is cleared in the local mode word then the processing for carriage return is disabled, and it is simply echoed as a return, and does not terminate cooked mode input.

In the new driver there is a literal-next character ^V which can be typed in both cooked and CBREAK mode preceding any character to prevent its special meaning. This is to be preferred to the use of '\' escaping erase and kill characters, but '\' is (at least temporarily) retained with its old function in the new driver for historical reasons.

The new terminal driver also provides two other editing characters in normal mode. The word-erase character, normally ^W, erases the preceding word, but not any spaces before it. For the purposes of ^W, a word is defined as a sequence of non-blank characters, with tabs counted as blanks. Finally, the reprint character, normally ^R, retypes the pending input beginning on a new line. Retyping occurs automatically in cooked mode if characters which would normally be erased from the screen are fouled by program output.

Input echoing and redisplay

In the old terminal driver, nothing special occurs when an erase character is typed; the erase character is simply echoed. When a kill character is typed it is echoed followed by a new-line (even if the character is not killing the line, because it was preceded by a '\'.!)

The new terminal driver has several modes for handling the echoing of terminal input, controlled by bits in a local mode word.

Hardcopy terminals. When a hardcopy terminal is in use, the LPRTERA bit is normally set in the local mode word. Characters which are logically erased are then printed out backwards preceded by '\' and followed by '/' in this mode.

Crt terminals. When a crt terminal is in use, the LCRTBS bit is normally set in the local mode word. The terminal driver then echoes the proper number of erase characters when input is erased; in the normal case where the erase character is a ^H this causes the cursor of the terminal to back up to where it was before the logically erased character was typed. If the input has become fouled due to interspersed asynchronous output, the input is automatically retyped.

Erasing characters from a crt. When a crt terminal is in use, the LCRTERA bit may be set to cause input to be erased from the screen with a "backspace-space-backspace" sequence when character or word deleting sequences are used. A LCRTKIL bit may be set as well, causing the input to be erased in this manner on line kill sequences as well.

Echoing of control characters. If the LCTLECH bit is set in the local state word, then non-printing (control) characters are normally echoed as ^X (for some X) rather than being echoed unmodified; delete is echoed as ^?.

The normal modes for using the new terminal driver on crt terminals are speed dependent. At speeds less than 1200 baud, the LCRTERA and LCRTKILL processing is painfully slow, so *stty*(1) normally just sets LCRTBS and LCTLECH; at speeds of 1200 baud or greater all of these bits are normally set. *Stty*(1) summarizes these option settings and the use of the new terminal driver as "newcrt."

Output processing.

When one or more characters are written, they are actually transmitted to the terminal as soon as previously-written characters have finished typing. (As noted above, input characters are normally echoed by putting them in the output queue as they arrive.) When a process produces characters more rapidly than they can be typed, it will be suspended when its output queue exceeds some limit. When the queue has drained down to some threshold the program is resumed. Even parity is normally generated on output. The EOT character is not transmitted in cooked mode to prevent terminals that respond to it from hanging up; programs using raw or cbreak mode should be careful.

The terminal drivers provide necessary processing for cooked and CBREAK mode output including delay generation for certain special characters and parity generation. Delays are available after backspaces `^H`, form feeds `^L`, carriage returns `^M`, tabs `^I` and newlines `^J`. The driver will also optionally expand tabs into spaces, where the tab stops are assumed to be set every eight columns. These functions are controlled by bits in the `ty` flags word; see **Summary** below.

The terminal drivers provide for mapping between upper and lower case on terminals lacking lower case, and for other special processing on deficient terminals.

Finally, in the new terminal driver, there is a output flush character, normally `^O`, which sets the LFLUSHO bit in the local mode word, causing subsequent output to be flushed until it is cleared by a program or more input is typed. This character has effect in both cooked and CBREAK modes and causes pending input to be retyped if there is any pending input. An `ioctl` to flush the characters in the input and output queues TIOCFLUSH, is also available.

Upper case terminals and Hazeltines

If the LCASE bit is set in the `ty` flags, then all upper-case letters are mapped into the corresponding lower-case letter. The upper-case letter may be generated by preceding it by `^`. If the new terminal driver is being used, then upper case letters are preceded by a `^` when output. In addition, the following escape sequences can be generated on output and accepted on input:

for	<code>^</code>	<code>^!</code>	<code>^~</code>	<code>^{</code>	<code>^}</code>
use	<code>^</code>	<code>^!</code>	<code>^~</code>	<code>^(</code>	<code>^)</code>

To deal with Hazeltine terminals, which do not understand that `^` has been made into an ASCII character, the LTIILDE bit may be set in the local mode word when using the new terminal driver; in this case the character `^` will be replaced with the character `^` on output.

Flow control.

There are two characters (the stop character, normally `^S`, and the start character, normally `^Q`) which cause output to be suspended and resumed respectively. Extra stop characters typed when output is already stopped have no effect, unless the start and stop characters are made the same, in which case output resumes.

A bit in the flags word may be set to put the terminal into TANDEM mode. In this mode the system produces a stop character (default `^S`) when the input queue is in danger of overflowing, and a start character (default `^Q`) when the input has drained sufficiently. This mode is useful when the terminal is actually another machine that obeys the conventions.

Line control and breaks.

There are several `ioctl` calls available to control the state of the terminal line. The TIOCSBRK `ioctl` will set the break bit in the hardware interface causing a break condition to exist; this can be cleared (usually after a delay with `sleep(3)`) by TIOCCBRK. Break conditions in the input are reflected as a null character in RAW mode or as the interrupt character in cooked or CBREAK mode. The TIOCCDTR `ioctl` will clear the data terminal ready condition; it can be

set again by TIOCSDTR.

When the carrier signal from the dataset drops (usually because the user has hung up his terminal) a SIGHUP hangup signal is sent to the processes in the distinguished process group of the terminal; this usually causes them to terminate (the SIGHUP can be suppressed by setting the LNOHANG bit in the local state word of the driver.) Access to the terminal by other processes is then normally revoked, so any further reads will fail, and programs that read a terminal and test for end-of-file on their input will terminate appropriately.

When using an ACU it is possible to ask that the phone line be hung up on the last close with the TIOCHPCL ioctl; this is normally done on the outgoing line.

Interrupt characters.

There are several characters that generate interrupts in cooked and CBREAK mode; all are sent the processes in the control group of the terminal, as if a TIOCGPRP 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 (*'i'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, although this is not often done.

- `^?` **t_intrc** (Delete) generates a SIGINT signal. This is the normal way to stop a process which is no longer interesting, or to regain control in an interactive program.
- `^\` **t_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 driver, if a process which is not in the distinguished process group of its control terminal attempts to read from that terminal its process group is sent a SIGTTIN signal. This signal normally causes the members of that process group to stop. If, however, the process is ignoring SIGTTIN, has SIGTTIN blocked, is an *orphan process*, or is in the middle of process creation using *vfork(2)*, it is instead returned an end-of-file. (An *orphan process* is a process whose parent has exited and has been inherited by the *init(8)* process.) Under older UNIX systems these processes would typically have had their input files reset to */dev/null*, so this is a compatible change.

When using the new terminal driver with the LTOSTOP bit set in the local modes, a process is prohibited from writing on its control terminal if it is not in the distinguished process group for that terminal. Processes which are holding or ignoring SIGTTOU signals, which are orphans, or which are in the middle of a *vfork(2)* are excepted and allowed to produce output.

Summary of modes.

Unfortunately, due to the evolution of the terminal driver, there are 4 different structures which contain various portions of the driver data. The first of these (*sgttyb*) contains that part of the information largely common between version 6 and version 7 UNIX systems. The second contains additional control characters added in version 7. The third is a word of local state peculiar to the new terminal driver, and the fourth is another structure of special characters added for the new driver. In the future a single structure may be made available to programs which need to access all this information; most programs need not concern themselves with all this state.

Basic modes: `sgtty`.

The basic *ioctls* use the structure defined in `<sgtty.h>`:

```
struct 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 `<sgtty.h>`.

B0	0	(hang up dataphone)
B50	1	50 baud
B75	2	75 baud
B110	3	110 baud
B134	4	134.5 baud
B150	5	150 baud
B200	6	200 baud
B300	7	300 baud
B600	8	600 baud
B1200	9	1200 baud
B1800	10	1800 baud
B2400	11	2400 baud
B4800	12	4800 baud
B9600	13	9600 baud
EXTA	14	External A
EXTB	15	External B

In the current configuration, only 110, 150, 300 and 1200 baud are really supported on dial-up lines. Code conversion and line control required for IBM 2741's (134.5 baud) must be implemented by the user's program. The half-duplex line discipline required for the 202 dataset (1200 baud) is not supplied; full-duplex 212 datasets work fine.

The *sg_erase* and *sg_kill* fields of the argument structure specify the erase and kill characters respectively. (Defaults are `#` and `@`.)

The *sg_flags* field of the argument structure contains several bits that determine the system's treatment of the terminal:

ALLDELAY	0177400	Delay algorithm selection
BSDELAY	0100000	Select backspace delays (not implemented):
BS0	0	
BS1	0100000	
VTDELAY	0040000	Select form-feed and vertical-tab delays:
VF0	0	
VF1	0100000	
CRDELAY	0030000	Select carriage-return delays:
CR0	0	
CR1	0010000	
CR2	0020000	
CR3	0030000	

TBDELAY	0006000	Select tab delays:
TAB0	0	
TAB1	0001000	
TAB2	0004000	
XTABS	0006000	
NLDELAY	0001400	Select new-line delays:
NL0	0	
NL1	0000400	
NL2	0001000	
NL3	0001400	
EVENP	0000200	Even parity allowed on input (most terminals)
ODDP	0000100	Odd parity allowed on input
RAW	0000040	Raw mode: wake up on all characters, 8-bit interface
CRMOD	0000020	Map CR into LF; echo LF or CR as CR-LF
ECHO	0000010	Echo (full duplex)
LCASE	0000004	Map upper case to lower on input
CBREAK	0000002	Return each character as soon as typed
TANDEM	0000001	Automatic flow control

The delay bits specify how long transmission stops to allow for mechanical or other movement when certain characters are sent to the terminal. In all cases a value of 0 indicates no delay.

Backspace delays are currently ignored but might be used for Terminet 300's.

If a form-feed/vertical tab delay is specified, it lasts for about 2 seconds.

Carriage-return delay type 1 lasts about .08 seconds and is suitable for the Terminet 300. Delay type 2 lasts about .16 seconds and is suitable for the VT05 and the TI 700. Delay type 3 is suitable for the concept-100 and pads lines to be at least 9 characters at 9600 baud.

New-line delay type 1 is dependent on the current column and is tuned for Teletype model 37's. Type 2 is useful for the VT05 and is about .10 seconds. Type 3 is unimplemented and is 0.

Tab delay type 1 is dependent on the amount of movement and is tuned to the Teletype model 37. Type 3, called XTABS, is not a delay at all but causes tabs to be replaced by the appropriate number of spaces on output.

Input characters with the wrong parity, as determined by bits 200 and 100, are ignored in cooked and CBREAK mode.

RAW disables all processing save output flushing with LFLUSHO; full 8 bits of input are given as soon as it is available; all 8 bits are passed on output. A break condition in the input is reported as a null character. If the input queue overflows in raw mode it is discarded; this applies to both new and old drivers.

CRMOD causes input carriage returns to be turned into new-lines; input of either CR or LF causes LF-CR both to be echoed (for terminals with a new-line function).

CBREAK is a sort of half-cooked (rare?) mode. Programs can read each character as soon as typed, instead of waiting for a full line; all processing is done except the input editing: character and word erase and line kill, input reprint, and the special treatment of \ or EOT are disabled.

TANDEM mode causes the system to produce a stop character (default ^S) whenever the input queue is in danger of overflowing, and a start character (default ^Q) when the input queue has drained sufficiently. It is useful for flow control when the 'terminal' is really another computer which understands the conventions.

Basic ioctls

In addition to the TIOCSETD and TIOCGETD disciplines discussed in **Line disciplines** above, a large number of other *ioctl(2)* calls apply to terminals, and have the general form:

```
#include <sgtty.h>
ioctl(fildes, code, arg)
struct sgttyb *arg;
```

The applicable codes are:

- 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 the *arg* is ignored.

- TIOCEXCL Set "exclusive-use" mode: no further opens are permitted until the file has been closed.
- TIOCNXCL Turn off "exclusive-use" mode.
- TIOCHPCL When the file is closed for the last time, hang up the terminal. This is useful when the line is associated with an ACU used to place outgoing calls.
- TIOCFDUSH All characters waiting in input or output queues are flushed.

The remaining calls are not available in vanilla version 7 UNIX. In cases where arguments are required, they are described; *arg* should otherwise be given as 0.

- TIOCSTI the argument is the address of a character which the system pretends was typed on the terminal.
- TIOCSBRK the break bit is set in the terminal.
- TIOCCBRK the break bit is cleared.
- TIOCSDTR data terminal ready is set.
- TIOCCDTR data terminal ready is cleared.
- TIOCGPGRP *arg* is the address of a word into which is placed the process group number of the control terminal.
- TIOCSGRP *arg* is a word (typically a process id) which becomes the process group for the control terminal.
- FIONREAD returns in the long integer whose address is *arg* the number of immediately readable characters from the argument unit. This works for files, pipes, and terminals, but not (yet) for multiplexed channels.

Tchars

The second structure associated with each terminal specifies characters that are special in both the old and new terminal interfaces: The following structure is defined in *<sys/ioctl.h>*, which is automatically included in *<sgtty.h>*:

```
struct tchars {
    char t_intrc;     /* interrupt */
    char t_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 `^?`, `^\`, `^Q`, `^S`, `^D`, and `-1`. A character value of `-1` eliminates the effect of that character. The `t_brkc` character, by default `-1`, acts like a new-line in that it terminates a 'line,' is echoed, and is passed to the program. The 'stop' and 'start' characters may be the same, to produce a toggle effect. It is probably counterproductive to make other special characters (including erase and kill) identical. The applicable `ioctl` calls are:

TIOCGETC Get the special characters and put them in the specified structure.

TIOCSETC Set the special characters to those given in the structure.

Local mode

The third structure associated with each terminal is a local mode word; except for the **LNOHANG** bit, this word is interpreted only when the new driver is in use. The bits of the local mode word are:

LCRTBS	000001	Backspace on erase rather than echoing erase
LPRTERA	000002	Printing terminal erase mode
LCRTERA	000004	Erase character echoes as backspace-space-backspace
LTILDE	000010	Convert <code>~</code> to <code>`</code> on output (for Hazeltine terminals)
LMDMBUF	000020	Stop/start output when carrier drops
LLITOUT	000040	Suppress output translations
LTOSTOP	000100	Send SIGTTOU for background output
LFLUSHO	000200	Output is being flushed
LNOHANG	000400	Don't send hangup when carrier drops
LETXACK	001000	Diablo style buffer hacking (unimplemented)
LCRCKIL	002000	BS-space-BS erase entire line on line kill
LINTRUP	004000	Generate interrupt SIGTINT when input ready to read
LCTLECH	010000	Echo input control chars as <code>^X</code> , delete as <code>^?</code>
LPENDIN	020000	Retype pending input at next read or input character
LDECCTQ	040000	Only <code>^Q</code> restarts output after <code>^S</code> , like DEC systems

The applicable `ioctl` functions are:

TIOCLBIS `arg` is the address of a mask which is the bits to be set in the local mode word.

TIOCLBIC `arg` is the address of a mask of bits to be cleared in the local mode word.

TIOCLSET `arg` is the address of a mask to be placed in the local mode word.

TIOCLGET `arg` is the address of a word into which the current mask is placed.

Local special chars

The final structure associated with each terminal is the `lchars` structure which defines interrupt characters for the new terminal driver. Its structure is:

```

struct lchars {
char  t_suspc;     /* stop process signal */
char  t_dsuspc;    /* delayed stop process signal */
char  t_rprntc;    /* reprint line */
char  t_flushc;    /* flush output (toggles) */
};

```

```

    char  t_werasc;    /* word erase */
    char  t_inextc;   /* literal next character */
};

```

The default values for these characters are ^Z, ^Y, ^R, ^O, ^W, and ^V. A value of -1 disables the character.

The applicable *ioctl* functions are:

TIOCSLTC args is the address of a *ltchars* structure which defines the new local special characters.

TIOCGLTC args is the address of a *ltchars* structure into which is placed the current set of local special characters.

FILES

```

/dev/tty
/dev/tty*
/dev/console

```

SEE ALSO

csh(1), stty(1), ioctl(2), sigvec(2), stty(3C), getty(8), init(8)

BUGS

Half-duplex terminals are not supported.

NAME

tu — VAX-11/730 and VAX-11/750 TU58 console cassette interface

SYNOPSIS

options MRSP (for VAX-11/750's with an MRSP prom)

DESCRIPTION

The *tu* interface provides access to the VAX 11/730 and 11/750 TU58 console cassette drive(s).

The interface supports only block i/o to the TU58 cassettes. The devices are normally manipulated with the *arff*(8V) program using the "f" and "m" options.

The device driver is automatically included when a system is configured to run on an 11/730 or 11/750.

The TU58 on an 11/750 uses the Radial Serial Protocol (RSP) to communicate with the cpu over a serial line. This protocol is inherently unreliable as it has no flow control measures built in. On an 11/730 the Modified Radial Serial Protocol is used. This protocol incorporates flow control measures which insure reliable data transfer between the cpu and the device. Certain 11/750's have been modified to use the MRSP prom used in the 11/730. To reliably use the console TU58 on an 11/750 under UNIX, the MRSP prom is required. For those 11/750's without an MRSP prom, an unreliable but often useable interface has been developed. This interface uses an assembly language "pseudo-dma" routine to minimize the receiver interrupt service latency. To include this code in the system, the configuration must **not** specify the system will run on an 11/730 or use an MRSP prom. This unfortunately makes it impossible to configure a single system which will properly handle TU58's on both an 11/750 and an 11/730 (unless both machines have MRSP prompts).

FILES

/dev/tu0
/dev/tu1 (only on a VAX-11/730)

SEE ALSO

arff(8V)

DIAGNOSTICS

tu%d: no bp, active %d. A transmission complete interrupt was received with no outstanding i/o request. This indicates a hardware problem.

tu%d protocol error, state=%s, op=%x, cnt=%d, block=%d. The driver entered an illegal state. The information printed indicates the illegal state, operation currently being executed, the i/o count, and the block number on the cassette.

tu%d receive state error, state=%s, byte=%x. The driver entered an illegal state in the receiver finite state machine. The state is shown along with the control byte of the received packet.

tu%d: read stalled. A timer watching the controller detected no interrupt for an extended period while an operation was outstanding. This usually indicates that one or more receiver interrupts were lost and the transfer is restarted (11/750 only).

tu%d: hard error bn%d, pk_mod %o. The device returned a status code indicating a hard error. The actual error code is shown in octal. No retries are attempted by the driver.

BUGS

The VAX-11/750 console interface without MRSP prom is unuseable while the system is multi-user; it should be used only with the system running single-user and, even then, with caution.

NAME

uda — UDA-50 disk controller interface

SYNOPSIS

controller **uda0** at **uba0** csr **0172150** vector **udintr**
 disk **ra0** at **uda0** drive **0**

DESCRIPTION

This is a driver for the DEC UDA-50 disk controller. The UDA-50 communicates with the host through a packet oriented protocol termed the Mass Storage Control Protocol (MSCP). Consult the file `<vx/mscp.h>` for a detailed description of this protocol.

Files with minor device numbers 0 through 7 refer to various portions of drive 0; minor devices 8 through 15 refer to drive 1, etc. The standard device names begin with "ra" 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 results in exactly one I/O operation and 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* calls should specify a multiple of 512 bytes.

DISK SUPPORT

This driver handles all drives which may be connected to the UDA-50. Drive types per se are not recognized, but rather the variable length partitions are defined as having an "infinite" length and the controller is relied on to return an error when an inaccessible block is requested. For constructing file systems, however the partitions sizes are required. The origin and size (in sectors) of the pseudo-disks on each drive are shown below. Partitions are not rounded to cylinder boundaries, as on other drives, because the type of drive attached to the controller is discovered too late in the autoconfiguration process to maintain separate partition tables for each drive. (The lack of proper drive type recognition would be more easily dealt with if the partition tables were read off the drive.)

RA60 partitions

disk	start	length
ra?a	0	15884
ra?b	15884	33440
ra?c	0	400176
ra?g	49324	82080
ra?h	131404	268772

RA80 partitions

disk	start	length
ra?a	0	15884
ra?b	15884	33440
ra?c	0	242606
ra?g	49324	82080
ra?h	131404	111202

RA81 partitions

disk	start	length
ra?a	0	15884
ra?b	15884	33440

ra?c	0	891072
ra?d	340670	15884
ra?e	356554	55936
ra?f	412490	478582
ra?g	49324	82080
ra?h	131404	759668

The ra?a partition is normally used for the root file system, the ra?b partition as a paging area, and the ra?c partition for pack-pack copying (it maps the entire disk).

FILES

/dev/ra[0-9][a-f]
/dev/rra[0-9][a-f]

DIAGNOSTICS

uda: ubinfo %x. (VAX 11/750 only.) When allocating UNIBUS resources, the driver found it already had resources previously allocated. This indicates a bug in the driver.

udasa %o, state %d. (Additional status information given after a hard i/o error.) The values of the UDA-50 status register and the internal driver state are printed.

uda%d: random interrupt ignored. An unexpected interrupt was received (e.g. when no i/o was pending). The interrupt is ignored.

uda%d: interrupt in unknown state %d ignored. An interrupt was received when the driver was in an unknown internal state. Indicates a hardware problem or a driver bug.

uda%d: fatal error (%o). The UDA-50 indicated a "fatal error" in the status returned to the host. The contents of the status register are displayed.

OFFLINE. (Additional status information given after a hard i/o error.) A hard i/o error occurred because the drive was not on-line.

status %o. (Additional status information given after a hard i/o error.) The status information returned from the UDA-50 is tacked onto the end of the hard error message printed on the console.

uda: unknown packet. An MSCP packet of unknown type was received from the UDA-50. Check the cabling to the controller.

The following errors are interpretations of MSCP error messages returned by the UDA-50 to the host.

uda%d: %s error, controller error, event 0%o.

uda%d: %s error, host memory access error, event 0%o, addr 0%o.

uda%d: %s error, disk transfer error, unit %d.

uda%d: %s error, SDI error, unit %d, event 0%o.

uda%d: %s error, small disk error, unit %d, event 0%o, cyl %d.

uda%d: %s error, unknown error, unit %d, format 0%o, event 0%o.

BUGS

The partition tables are so poorly laid out that they almost certainly force each site to tailor them to their individual needs. The problem is even worse when a site has a mixed collection of drives. The best solution would be to read the partition tables off the drive.

NAME

udp — Internet User Datagram Protocol

SYNOPSIS

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

s = socket(AF_INET, SOCK_DGRAM, 0);
```

DESCRIPTION

UDP is a simple, unreliable datagram protocol which is used to support the SOCK_DGRAM abstraction for the Internet protocol family. UDP sockets are connectionless, and are normally used with the *sendto* and *recvfrom* calls, though the *connect(2)* call may also be used to fix the destination for future packets (in which case the *recv(2)* or *read(2)* and *send(2)* or *write(2)* system calls may be used).

UDP address formats are identical to those used by TCP. In particular UDP provides a port identifier in addition to the normal Internet address format. Note that the UDP port space is separate from the TCP port space (i.e. a UDP port may not be “connected” to a TCP port). In addition broadcast packets may be sent (assuming the underlying network supports this) by using a reserved “broadcast address”; this address is network interface dependent.

DIAGNOSTICS

A socket operation may fail with one of the following errors returned:

- [EISCONN] when trying to establish a connection on a socket which already has one, or when trying to send a datagram with the destination address specified and the socket is already connected;
- [ENOTCONN] when trying to send a datagram, but no destination address is specified, and the socket hasn't been connected;
- [ENOBUFS] when the system runs out of memory for an internal data structure;
- [EADDRINUSE] when an attempt is made to create a socket with a port which has already been allocated;
- [EADDRNOTAVAIL] when an attempt is made to create a socket with a network address for which no network interface exists.

SEE ALSO

send(2), *recv(2)*, *intro(4N)*, *inet(4F)*

NAME

`un` — Ungermann-Bass interface

SYNOPSIS

`device un0 at uba0 csr 0160210 vector unintr`

DESCRIPTION

The `un` interface provides access to a 4 Mb/s baseband network. The hardware uses a standard DEC DR11-W DMA interface in communicating with the host. The Ungermann-Bass hardware incorporates substantial protocol software in the network device in an attempt to offload protocol processing from the host.

The network number on which the interface resides must be specified at boot time with an `SIOCSIFADDR` ioctl. The host's address is discovered by communicating with the interface. The interface will not transmit or receive any packets before the network number has been defined.

DIAGNOSTICS

un%d: can't initialize. Insufficient UNIBUS resources existed for the device to complete initialization. Usually caused by having multiple network interfaces configured using buffered data paths on a data path poor machine such as the 11/750.

un%d: unexpected reset. The controller indicated a reset when none had been requested. Check the hardware (but see the bugs section below).

un%d: stray interrupt. An unexpected interrupt was received. The interrupt was ignored.

un%d: input error csr=%b. The controller indicated an error on moving data from the device to host memory.

un%d: bad packet type %d. A packet was received with an unknown packet type. The packet is discarded.

un%d: output error csr=%b. The device indicated an error on moving data from the host to device memory.

un%d: invalid state %d csr=%b. The driver found itself in an invalid internal state. The state is reset to a base state.

un%d: can't handle af%d. A request was made to send a message with an address format which the driver does not understand. The message is discarded and an error is returned to the user.

un%d: error limit exceeded. Too many errors were encountered in normal operation. The driver will attempt to reset the device, desist from attempting any i/o for approximately 60 seconds, then reset itself to a base state in hopes of resyncing itself up with the hardware.

un%d: restarting. After exceeding its error limit and resetting the device, the driver is restarting operation.

SEE ALSO

`intro(4N)`, `inet(4F)`

BUGS

The device does not reset itself properly resulting in the interface getting hung up in a state from which the only recourse is to reboot the system.

NAME

up — unibus storage module controller/drives

SYNOPSIS

controller sc0 at uba? csr 0176700 vector upintr
disk up0 at sc0 drive 0

DESCRIPTION

This is a generic UNIBUS storage module disk driver. It is specifically designed to work with the Emulex SC-21 controller. It can be easily adapted to other controllers (although bootstrapping will not necessarily be directly possible.)

Files with minor device numbers 0 through 7 refer to various portions of drive 0; minor devices 8 through 15 refer to drive 1, etc. The standard device names begin with "up" 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 results in exactly one I/O operation and 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* calls should specify a multiple of 512 bytes.

DISK SUPPORT

The driver interrogates the controller's holding register to determine the type of drive attached. The driver recognizes four different drives: AMPEX 9300, CDC 9766, AMPEX Capricorn, and FUJITSU 160. The origin and size of the pseudo-disks on each drive are as follows:

CDC 9766 300M drive partitions:

disk	start	length	cyl
up?a	0	15884	0-26
up?b	16416	33440	27-81
up?c	0	500384	0-822
up?d	341696	15884	562-588
up?e	358112	55936	589-680
up?f	414048	861760	681-822
up?g	341696	158528	562-822
up?h	49856	291346	82-561

AMPEX 9300 300M drive partitions:

disk	start	length	cyl
up?a	0	15884	0-26
up?b	16416	33440	27-81
up?c	0	495520	0-814
up?d	341696	15884	562-588
up?e	358112	55936	589-680
up?f	414048	81312	681-814
up?g	341696	153664	562-814
up?h	49856	291346	82-561

AMPEX Capricorn 330M drive partitions:

disk	start	length	cyl
hp?a	0	15884	0-31
hp?b	16384	33440	32-97

hp?c	0	524288	0-1023
hp?d	342016	15884	668-699
hp?e	358400	55936	700-809
hp?f	414720	109408	810-1023
hp?g	342016	182112	668-1023
hp?h	50176	291346	98-667

FUJITSU 160M drive partitions:

disk	start	length	cyl
up?a	0	15884	0-49
up?b	16000	33440	50-154
up?c	0	263360	0-822
up?d	49600	15884	155-204
up?e	65600	55936	205-379
up?f	121600	141600	380-822
up?g	49600	213600	155-822

It is unwise for all of these files to be present in one installation, since there is overlap in addresses and protection becomes a sticky matter. The up?a partition is normally used for the root file system, the up?b partition as a paging area, and the up?c partition for pack-pack copying (it maps the entire disk). On 160M drives the up?g partition maps the rest of the pack. On other drives both up?g and up?h are used to map the remaining cylinders.

FILES

/dev/up[0-7][a-h] block files
/dev/rup[0-7][a-h] raw files

SEE ALSO

hk(4), hp(4), uda(4)

DIAGNOSTICS

up%d%c: hard error sn%d cs2=%b er1=%b er2=%b. An unrecoverable error occurred during transfer of the specified sector in the specified disk partition. The contents of the cs2, er1 and er2 registers are printed in octal and symbolically with bits decoded. The error was either unrecoverable, or a large number of retry attempts (including offset positioning and drive recalibration) could not recover the error.

up%d: write locked. The write protect switch was set on the drive when a write was attempted. The write operation is not recoverable.

up%d: not ready. The drive was spun down or off line when it was accessed. The i/o operation is not recoverable.

up%d: not ready (flakey). The drive was not ready, but after printing the message about being not ready (which takes a fraction of a second) was ready. The operation is recovered if no further errors occur.

up%d%c: soft ecc sn%d. A recoverable ECC error occurred on the specified sector of the specified disk partition. This happens normally a few times a week. If it happens more frequently than this the sectors where the errors are occurring should be checked to see if certain cylinders on the pack, spots on the carriage of the drive or heads are indicated.

sc%d: lost interrupt. A timer watching the controller detecting no interrupt for an extended period while an operation was outstanding. This indicates a hardware or software failure. There is currently a hardware/software problem with spinning down drives while they are being accessed which causes this error to occur. The error causes a UNIBUS reset, and retry of the pending operations. If the controller continues to lose interrupts, this error will recur a few seconds later.

BUGS

In raw I/O *read* and *write(2)* 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.

DEC-standard error logging should be supported.

A program to analyze the logged error information (even in its present reduced form) is needed.

The partition tables for the file systems should be read off of each pack, as they are never quite what any single installation would prefer, and this would make packs more portable.

NAME

ut — UNIBUS TU45 tri-density tape drive interface

SYNOPSIS

controller ut0 at uba0 csr 0172440 vector utintr
tape tj0 at ut0 drive 0

DESCRIPTION

The *ut* interface provides access to a standard tape drive interface as describe in *mtio(4)*. Hardware implementing this on the VAX is typified by the System Industries SI 9700 tape subsystem. Tapes may be read or written at 800, 1600, and 6250 bpi.

SEE ALSO

mt(1), mtio(4)

DIAGNOSTICS

tj%d: no write ring. An attempt was made to write on the tape drive when no write ring was present; this message is written on the terminal of the user who tried to access the tape.

tj%d: not online. An attempt was made to access the tape while it was offline; this message is written on the terminal of the user who tried to access the tape.

tj%d: can't change density in mid-tape. An attempt was made to write on a tape at a different density than is already recorded on the tape. This message is written on the terminal of the user who tried to switch the density.

ut%d: soft error bn%d cs1=%b er=%b cs2=%b ds=%b. The formatter indicated a corrected error at a density other than 800bpi. The data transferred is assumed to be correct.

ut%d: hard error bn%d cs1=%b er=%b cs2=%b ds=%b. A tape error occurred at block *bn*. Any error is fatal on non-raw tape; when possible the driver will have retried the operation which failed several times before reporting the error.

tj%d: lost interrupt. A tape operation did not complete within a reasonable time, most likely because the tape was taken off-line during rewind or lost vacuum. The controller should, but does not, give an interrupt in these cases. The device will be made available again after this message, but any current open reference to the device will return an error as the operation in progress aborts.

BUGS

If any non-data error is encountered on non-raw tape, it refuses to do anything more until closed.

NAME

`uu` — TU58/DECtape II UNIBUS cassette interface

SYNOPSIS

options UUDMA

device `uu0` at `uba0` **csr** 0176500 **vector** `uurintr` `uuxintr`

DESCRIPTION

The `uu` device provides access to dual DEC TU58 tape cartridge drives connected to the UNIBUS via a DL11-W interface module.

The interface supports only block i/o to the TU58 cassettes. The drives are normally manipulated with the `arff(8V)` program using the “m” and “f” options.

The driver provides for an optional write and verify (read after write) mode that is activated by specifying the “a” device.

The TU58 is treated as a single device by the system even though it has two separate drives, “uu0” and “uu1”. If there is more than one TU58 unit on a system, the extra drives are named “uu2”, “uu3” etc.

NOTES

Assembly language code to assist the driver in handling the receipt of data (using a pseudo-dma approach) should be included when using this driver; specify “options UUDMA” in the configuration file.

ERRORS

The following errors may be returned:

[ENXIO] Nonexistent drive (on open); offset is too large or bad (undefined) ioctl code.

[EIO] Open failed, the device could not be reset.

[EBUSY] Drive in use.

FILES

`/dev/uu?`

`/dev/uu?a`

SEE ALSO

`tu(4)`, `arff(8V)`

DIAGNOSTICS

uu%d: no bp, active %d. A transmission complete interrupt was received with no outstanding i/o request. This indicates a hardware problem.

uu%d protocol error, state=%s, op=%x, cnt=%d, block=%d. The driver entered an illegal state. The information printed indicates the illegal state, the operation currently being executed, the i/o count, and the block number on the cassette.

uu%d: break received, transfer restarted. The TU58 was sending a continuous break signal and had to be reset. This may indicate a hardware problem, but the driver will attempt to recover from the error.

uu%d receive state error, state=%s, byte=%x. The driver entered an illegal state in the receiver finite state machine. The state is shown along with the control byte of the received packet.

uu%d: read stalled. A timer watching the controller detected no interrupt for an extended period while an operation was outstanding. This usually indicates that one or more receiver interrupts were lost and the transfer is restarted.

uu%d: hard error bn%d, pk_mod %o. The device returned a status code indicating a hard error. The actual error code is shown in octal. No retries are attempted by the driver.

4

NAME

va — Benson-Varian interface

SYNOPSIS

controller va0 at uba0 csr 0164000 vector vaintr
disk vz0 at va0 drive 0

DESCRIPTION

(NOTE: the configuration description, while counter-intuitive, is actually as shown above.)

The Benson-Varian printer/plotter is normally used with the programs *vpr(1)*, *vprint(1)* or *vtroff(1)*. This description is designed for those who wish to drive the Benson-Varian directly.

In print mode, the Benson-Varian uses a modified ASCII character set. Most control characters print various non-ASCII graphics such as daggers, sigmas, copyright symbols, etc. Only LF and FF are used as format effectors. LF acts as a newline, advancing to the beginning of the next line, and FF advances to the top of the next page.

In plot mode, the Benson-Varian prints one raster line at a time. An entire raster line of bits (2112 bits = 264 bytes) is sent, and then the Benson-Varian advances to the next raster line.

Note: The Benson-Varian must be sent an even number of bytes. If an odd number is sent, the last byte will be lost. Nulls can be used in print mode to pad to an even number of bytes.

To use the Benson-Varian yourself, you must realize that you cannot open the device, */dev/va0* if there is a daemon active. You can see if there is an active daemon by doing a *lpq(1)* and seeing if there are any files being printed.

To set the Benson-Varian into plot mode include the file *<sys/vcmd.h>* and use the following *ioctl(2)* call

```
ioctl(fileno(va), VSETSTATE, plotmd);
```

where *plotmd* is defined to be

```
int plotmd[] = { VPLOT, 0, 0 };
```

and *va* is the result of a call to *open* on *stdio*. When you finish using the Benson-Varian in plot mode you should advance to a new page by sending it a FF after putting it back into print mode, i.e. by

```
int prtmd[] = { VPRINT, 0, 0 };
...
flush(va);
ioctl(fileno(va), VSETSTATE, prtmd);
write(fileno(va), "\f\0", 2);
```

N.B.: If you use the standard I/O library with the Benson-Varian you **must** do

```
setbuf(vp, vpbuf);
```

where *vpbuf* is declared

```
char vpbuf[BUFSIZ];
```

otherwise the standard I/O library, thinking that the Benson-Varian is a terminal (since it is a character special file) will not adequately buffer the data you are sending to the Benson-Varian. This will cause it to run **extremely** slowly and tend to grind the system to a halt.

FILES

/dev/va0

SEE ALSO

vfont(5), *lpr(1)*, *lpd(8)*, *vtroff(1)*, *vp(4)*

DIAGNOSTICS

The following error numbers are significant at the time the device is opened.

[ENXIO] The device is already in use.

[EIO] The device is offline.

The following message may be printed on the console.

va%d: npr timeout. The device was not able to get data from the UNIBUS within the timeout period, most likely because some other device was hogging the bus. (But see BUGS below).

BUGS

The l's (one's) and l's (lower-case el's) in the Benson-Varian's standard character set look very similar; caution is advised.

The interface hardware is rumored to have problems which can play havoc with the UNIBUS. We have intermittent minor problems on the UNIBUS where our va lives, but haven't ever been able to pin them down completely.

NAME

vp — Versatec interface

SYNOPSIS

```
device vp0 at uba0 csr 0177510 vector vpintr vpintr
```

DESCRIPTION

The Versatec printer/plotter is normally used with the programs *vpr*(1), *vprint*(1) or *vtroff*(1). This description is designed for those who wish to drive the Versatec directly.

To use the Versatec yourself, you must realize that you cannot open the device, */dev/vp0* if there is a daemon active. You can see if there is a daemon active by doing a *lpq*(1), and seeing if there are any files being sent.

To set the Versatec into plot mode you should include `<sys/vcmd.h>` and use the *ioctl*(2) call

```
ioctl(fileno(vp), VSETSTATE, plotmd);
```

where *plotmd* is defined to be

```
int plotmd[] = { VPLOT, 0, 0 };
```

and *vp* is the result of a call to *fopen* on *stdio*. When you finish using the Versatec in plot mode you should eject paper by sending it a EOT after putting it back into print mode, i.e. by

```
int prtmd[] = { VPRINT, 0, 0 };
```

```
...
```

```
fflush(vp);
```

```
ioctl(fileno(vp), VSETSTATE, prtmd);
```

```
write(fileno(vp), "\04", 1);
```

N.B.: If you use the standard I/O library with the Versatec you **must** do

```
setbuf(vp, vpbuf);
```

where *vpbuf* is declared

```
char vpbuf[BUFSIZ];
```

otherwise the standard I/O library, thinking that the Versatec is a terminal (since it is a character special file) will not adequately buffer the data you are sending to the Versatec. This will cause it to run **extremely** slowly and tends to grind the system to a halt.

FILES

/dev/vp0

SEE ALSO

vfont(5), *lpr*(1), *lpd*(8), *vtroff*(1), *va*(4)

DIAGNOSTICS

The following error numbers are significant at the time the device is opened.

[ENXIO] The device is already in use.

[EIO] The device is offline.

BUGS

The configuration part of the driver assumes that the device is set up to vector print mode through 0174 and plot mode through 0200. As the configuration program can't be sure which vector interrupted at boot time, we specify that it has two interrupt vectors, and if an interrupt comes through 0200 it is reset to 0174. This is safe for devices with one or two vectors at these two addresses. Other configurations with 2 vectors may require changes in the driver.

NAME

vv — Proteon proNET 10 Megabit ring

SYNOPSIS

device vv0 at uba0 csr 161000 vector vvrint vvxint

DESCRIPTION

The **vv** interface provides access to a 10 Mb/s Proteon proNET ring network.

The network number to which the interface is attached must be specified with an **SIOCSI-FADDR** ioctl before data can be transmitted or received. The host's address is discovered by putting the interface in digital loopback mode (not joining the ring) and sending a broadcast packet from which the source address is extracted. the Internet address of the interface would be 128.3.0.24.

The interface software implements error-rate limiting on the input side. This provides a defense against situations where other hosts or interface hardware failures cause a machine to be inundated with garbage packets. The scheme involves an exponential backoff where the input side of the interface is disabled for longer and longer periods. In the limiting case, the interface is turned on every two minutes or so to see if operation can resume.

If the installation is running CTL boards which use the old broadcast address of 0 instead of the new address of 0xff, the define **OLD_BROADCAST** should be specified in the driver.

If the installation has a Wirecenter, the define **WIRECENTER** should be specified in the driver. **N.B.:** Incorrect definition of **WIRECENTER** can cause hardware damage.

The interface normally tries to use a "trailer" encapsulation to minimize copying data on input and output. This may be disabled, on a per-interface basis, by setting the **IFF_NOTRAILERS** flag with an **SIOCSIFFLAGS** ioctl.

DIAGNOSTICS

vv%d: host %d. The software announces the host address discovered during autoconfiguration.

vv%d: can't initialize. The software was unable to discover the address of this interface, so it deemed "dead" will not be enabled.

vv%d: error vvocsr=%b. The hardware indicated an error on the previous transmission.

vv%d: output timeout. The token timer has fired and the token will be recreated.

vv%d: error vvicsr=%b. The hardware indicated an error in reading a packet off the ring.

en%d: can't handle af%d. The interface was handed a message with addresses formatted in an unsuitable address family; the packet was dropped.

vv%d: vs_olen=%d. The ring output routine has been handed a message with a preposterous length. This results in an immediate *panic: vs_olen*.

SEE ALSO

intro(4N), inet(4F)

NAME

a.out — assembler and link editor output

SYNOPSIS

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

DESCRIPTION

a.out is the output file of the assembler *as*(1) and the link editor *ld*(1). Both programs make *a.out* executable if there were no errors and no unresolved external references. Layout information as given in the include file for the VAX-11 is:

```
/*
 * Header prepended to each a.out file.
 */
struct exec {
    long    a_magic; /* magic number */
    unsigned a_text; /* size of text segment */
    unsigned a_data; /* size of initialized data */
    unsigned a_bss; /* size of uninitialized data */
    unsigned a_syms; /* size of symbol table */
    unsigned a_entry; /* entry point */
    unsigned a_trsize; /* size of text relocation */
    unsigned a_drsize; /* size of data relocation */
};

#define OMAGIC 0407 /* old impure format */
#define NMAGIC 0410 /* read-only text */
#define ZMAGIC 0413 /* demand load format */

/*
 * Macros which take exec structures as arguments and tell whether
 * the file has a reasonable magic number or offsets to text|symbols|strings.
 */
#define N_BADMAG(x) \
    (((x).a_magic)!=OMAGIC && ((x).a_magic)!=NMAGIC && ((x).a_magic)!=ZMAGIC)

#define N_TXTOFF(x) \
    ((x).a_magic==ZMAGIC ? 1024 : sizeof (struct exec))
#define N_SYMOFF(x) \
    (N_TXTOFF(x) + (x).a_text + (x).a_data + (x).a_trsize + (x).a_drsize)
#define N_STROFF(x) \
    (N_SYMOFF(x) + (x).a_syms)
```

The file has five sections: a header, the program text and data, relocation information, a symbol table and a string table (in that order). The last three may be omitted if the program was loaded with the '-s' option of *ld* or if the symbols and relocation have been removed by *strip*(1).

In the header the sizes of each section are given in bytes. The size of the header is not included in any of the other sizes.

When an *a.out* file is executed, three logical segments are set up: the text segment, the data segment (with uninitialized data, which starts off as all 0, following initialized), and a stack. The text segment begins at 0 in the core image; the header is not loaded. If the magic number in the header is OMAGIC (0407), it indicates that the text segment is not to be write-protected and shared, so the data segment is immediately contiguous with the text segment. This is the

oldest kind of executable program and is rarely used. If the magic number is NMAGIC (0410) or ZMAGIC (0413), the data segment begins at the first 0 mod 1024 byte boundary following the text segment, and the text segment is not writable by the program; if other processes are executing the same file, they will share the text segment. For ZMAGIC format, the text segment begins at a 0 mod 1024 byte boundary in the *a.out* file, the remaining bytes after the header in the first block are reserved and should be zero. In this case the text and data sizes must both be multiples of 1024 bytes, and the pages of the file will be brought into the running image as needed, and not pre-loaded as with the other formats. This is especially suitable for very large programs and is the default format produced by *ld(1)*.

The stack will occupy the highest possible locations in the core image: growing downwards from 0x7fff000. The stack is automatically extended as required. The data segment is only extended as requested by *brk(2)*.

After the header in the file follow the text, data, text relocation data relocation, symbol table and string table in that order. The text begins at the byte 1024 in the file for ZMAGIC format or just after the header for the other formats. The `N_TXTOFF` macro returns this absolute file position when given the name of an exec structure as argument. The data segment is contiguous with the text and immediately followed by the text relocation and then the data relocation information. The symbol table follows all this; its position is computed by the `N_SYMOFF` macro. Finally, the string table immediately follows the symbol table at a position which can be gotten easily using `N_STROFF`. The first 4 bytes of the string table are not used for string storage, but rather contain the size of the string table; this size INCLUDES the 4 bytes, the minimum string table size is thus 4.

The layout of a symbol table entry and the principal flag values that distinguish symbol types are given in the include file as follows:

```
/*
 * Format of a symbol table entry.
 */
struct nlist {
    union {
        char    *n_name; /* for use when in-core */
        long    n_strx;  /* index into file string table */
    } n_un;
    unsigned char n_type; /* type flag, i.e. N_TEXT etc; see below */
    char          n_other;
    short        n_desc; /* see <stab.h> */
    unsigned     n_value; /* value of this symbol (or offset) */
};
#define n_hash    n_desc /* used internally by ld */

/*
 * Simple values for n_type.
 */
#define N_UNDF    0x0    /* undefined */
#define N_ABS    0x2    /* absolute */
#define N_TEXT    0x4    /* text */
#define N_DATA    0x6    /* data */
#define N_BSS    0x8    /* bss */
#define N_COMM    0x12   /* common (internal to ld) */
#define N_FN      0x1f   /* file name symbol */

#define N_EXT     01     /* external bit, or'ed in */
```

```

#define N_TYPE      0x1e      /* mask for all the type bits */

/*
 * Other permanent symbol table entries have some of the N_STAB bits set.
 * These are given in <stab.h>
 */
#define N_STAB      0xe0      /* if any of these bits set, don't discard */

/*
 * Format for namelist values.
 */
#define N_FORMAT    "%08x"

```

In the *a.out* file a symbol's `n_un.n_strx` field gives an index into the string table. A `n_strx` value of 0 indicates that no name is associated with a particular symbol table entry. The field `n_un.n_name` can be used to refer to the symbol name only if the program sets this up using `n_strx` and appropriate data from the string table.

If a symbol's type is undefined external, and the value field is non-zero, the symbol is interpreted by the loader *ld* as the name of a common region whose size is indicated by the value of the symbol.

The value of a byte in the text or data which is not a portion of a reference to an undefined external symbol is exactly that value which will appear in memory when the file is executed. If a byte in the text or data involves a reference to an undefined external symbol, as indicated by the relocation information, then the value stored in the file is an offset from the associated external symbol. When the file is processed by the link editor and the external symbol becomes defined, the value of the symbol will be added to the bytes in the file.

If relocation information is present, it amounts to eight bytes per relocatable datum as in the following structure:

```

/*
 * Format of a relocation datum.
 */
struct relocation_info {
    int      r_address;      /* address which is relocated */
    unsigned r_symbolnum:24, /* local symbol ordinal */
           r_pcrel:1,      /* was relocated pc relative already */
           r_length:2,     /* 0=byte, 1=word, 2=long */
           r_extern:1,     /* does not include value of sym referenced */
           :4;             /* nothing, yet */
};

```

There is no relocation information if `a_trsize+a_drsize==0`. If `r_extern` is 0, then `r_symbolnum` is actually a `n_type` for the relocation (i.e. `N_TEXT` meaning relative to segment text origin.)

SEE ALSO

`adb(1)`, `as(1)`, `ld(1)`, `nm(1)`, `dbx(1)`, `stab(5)`, `strip(1)`

BUGS

Not having the size of the string table in the header is a loss, but expanding the header size would have meant stripped executable file incompatibility, and we couldn't hack this just now.

NAME

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       4.5           82/10/10*/

/*
 * 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_untime;   /* Accounting user time */
    comp_t    ac_stime;    /* 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 */

#define ACCTLO    30 /* acctg off when space < this */
#define ACCTHI    100 /* acctg resumes at this level */

#ifdef KERNEL
struct acct      acctbuf;
struct inode     *acctp;
#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 — aliases file for sendmail

SYNOPSIS

/usr/lib/aliases

DESCRIPTION

This file describes user id aliases used by *usr/lib/sendmail*. It is formatted as a series of lines of the form

name: name_1, name2, name_3, . . .

The *name* is the name to alias, and the *name_n* are the aliases for that name. Lines beginning with white space are continuation lines. Lines beginning with '#' are comments.

Aliasing occurs only on local names. Loops can not occur, since no message will be sent to any person more than once.

After aliasing has been done, local and valid recipients who have a ".forward" file in their home directory have messages forwarded to the list of users defined in that file.

This is only the raw data file; the actual aliasing information is placed into a binary format in the files *usr/lib/aliases.dir* and *usr/lib/aliases.pag* using the program *newaliases(1)*. A *newaliases* command should be executed each time the aliases file is changed for the change to take effect.

SEE ALSO

newaliases(1), *dbm(3X)*, *sendmail(8)*
SENDMAIL Installation and Operation Guide.
SENDMAIL An Internetwork Mail Router.

BUGS

Because of restrictions in *dbm(3X)* a single alias cannot contain more than about 1000 bytes of information. You can get longer aliases by "chaining"; that is, make the last name in the alias be a dummy name which is a continuation alias.

NAME

ar — archive (library) file format

SYNOPSIS

```
#include <ar.h>
```

DESCRIPTION

The archive command *ar* combines several files into one. Archives are used mainly as libraries to be searched by the link-editor *ld*.

A file produced by *ar* has a magic string at the start, followed by the constituent files, each preceded by a file header. The magic number and header layout as described in the include file are:

```
#define ARMAG "!<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 <machine/param.h>
```

DESCRIPTION

The UNIX System writes out a memory image of a terminated process when any of various errors occur. See *sigvec(2)* for the list of reasons; the most common are memory violations, illegal instructions, bus errors, and user-generated quit signals. The memory image is called 'core' and is written in the process's working directory (provided it can be; normal access controls apply).

The maximum size of a *core* file is limited by *setrlimit(2)*. Files which would be larger than the limit are not created.

The core file consists of the *u*. area, whose size (in pages) is defined by the UPAGES manifest in the *<machine/param.h>* file. The *u*. area starts with a *user* structure as given in *<sys/user.h>*. The remainder of the core file consists first of the data pages and then the stack pages of the process image. The amount of data space image in the core file is given (in pages) by the variable *u_dsize* in the *u*. area. The amount of stack image in the core file is given (in pages) by the variable *u_ssize* in the *u*. area.

In general the debugger *adb(1)* is sufficient to deal with core images.

SEE ALSO

adb(1), *dbx(1)*, *sigvec(2)*, *setrlimit(2)*

NAME

dir — format of directories

SYNOPSIS

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

DESCRIPTION

A directory behaves exactly like an ordinary file, save that no user may write into a directory. The fact that a file is a directory is indicated by a bit in the flag word of its i-node entry; see *fs(5)*. The structure of a directory entry as given in the include file is:

```
/*
 * A directory consists of some number of blocks of DIRBLKSIZ
 * bytes, where DIRBLKSIZ is chosen such that it can be transferred
 * to disk in a single atomic operation (e.g. 512 bytes on most machines).
 *
 * Each DIRBLKSIZ byte block contains some number of directory entry
 * structures, which are of variable length. Each directory entry has
 * a struct direct at the front of it, containing its inode number,
 * the length of the entry, and the length of the name contained in
 * the entry. These are followed by the name padded to a 4 byte boundary
 * with null bytes. All names are guaranteed null terminated.
 * The maximum length of a name in a directory is MAXNAMLEN.
 *
 * The macro DIRSIZ(dp) gives the amount of space required to represent
 * a directory entry. Free space in a directory is represented by
 * entries which have dp->d_reclen > DIRSIZ(dp). All DIRBLKSIZ bytes
 * in a directory block are claimed by the directory entries. This
 * usually results in the last entry in a directory having a large
 * dp->d_reclen. When entries are deleted from a directory, the
 * space is returned to the previous entry in the same directory
 * block by increasing its dp->d_reclen. If the first entry of
 * a directory block is free, then its dp->d_ino is set to 0.
 * Entries other than the first in a directory do not normally have
 * dp->d_ino set to 0.
 */
#ifdef KERNEL
#define DIRBLKSIZ DEV_BSIZE
#else
#define DIRBLKSIZ 512
#endif

#define MAXNAMLEN 255

/*
 * The DIRSIZ macro gives the minimum record length which will hold
 * the directory entry. This requires the amount of space in struct direct
 * without the d_name field, plus enough space for the name with a terminating
 * null byte (dp->d_namlen + 1), rounded up to a 4 byte boundary.
 */
#undef DIRSIZ
#define DIRSIZ(dp) \
    ((sizeof (struct direct) - (MAXNAMLEN+1)) + (((dp)->d_namlen+1 + 3) &~ 3))
```

```
struct  direct {
    u_long  d_ino;
    short   d_reclen;
    short   d_namlen;
    char    d_name[MAXNAMLEN + 1];
    /* typically shorter */
};

struct _dirdesc {
    int      dd_fd;
    long     dd_loc;
    long     dd_size;
    char     dd_buf[DIRBLKSIZ];
};
```

By convention, the first two entries in each directory are for '.' and '..'. The first is an entry for the directory itself. The second is for the parent directory. The meaning of '..' is modified for the root directory of the master file system ("/"), where '..' has the same meaning as '.'.

SEE ALSO

fs(5)

NAME

disktab — disk description file

SYNOPSIS

```
#include <disktab.h>
```

DESCRIPTION

Disktab is a simple data base which describes disk geometries and disk partition characteristics. The format is patterned after the *termcap(5)* terminal data base. Entries in *disktab* consist of a number of ':' separated fields. The first entry for each disk gives the names which are known for the disk, separated by '|' characters. The last name given should be a long name fully identifying the disk.

The following list indicates the normal values stored for each disk entry.

Name Type Description

ns	num	Number of sectors per track
nt	num	Number of tracks per cylinder
nc	num	Total number of cylinders on the disk
ba	num	Block size for partition 'a' (bytes)
bd	num	Block size for partition 'd' (bytes)
be	num	Block size for partition 'e' (bytes)
bf	num	Block size for partition 'f' (bytes)
bg	num	Block size for partition 'g' (bytes)
bh	num	Block size for partition 'h' (bytes)
fa	num	Fragment size for partition 'a' (bytes)
fd	num	Fragment size for partition 'd' (bytes)
fe	num	Fragment size for partition 'e' (bytes)
ff	num	Fragment size for partition 'f' (bytes)
fg	num	Fragment size for partition 'g' (bytes)
fh	num	Fragment size for partition 'h' (bytes)
pa	num	Size of partition 'a' in sectors
pb	num	Size of partition 'b' in sectors
pc	num	Size of partition 'c' in sectors
pd	num	Size of partition 'd' in sectors
pe	num	Size of partition 'e' in sectors
pf	num	Size of partition 'f' in sectors
pg	num	Size of partition 'g' in sectors
ph	num	Size of partition 'h' in sectors
se	num	Sector size in bytes
ty	str	Type of disk (e.g. removable, winchester)

Disktab entries may be automatically generated with the *diskpart* program.

FILES

/etc/disktab

SEE ALSO

newfs(8), diskpart(8)

BUGS

This file shouldn't exist, the information should be stored on each disk pack.

NAME

dump, dumpdates — incremental dump format

SYNOPSIS

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

DESCRIPTION

Tapes used by *dump* and *restore(8)* contain:

- a header record
- two groups of bit map records
- a group of records describing directories
- a group of records describing files

The format of the header record and of the first record of each description as given in the include file *<dumprestor.h>* is:

```
#define NTREC      10
#define MLEN       16
#define MSIZ       4096

#define TS_TAPE    1
#define TS_INODE   2
#define TS_BITS    3
#define TS_ADDR    4
#define TS_END     5
#define TS_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 number of 1024 byte records in a physical tape block. MLEN is the number of bits in a bit map word. MSIZ is the number of bit map words.

The TS_ entries are used in the *c_type* field to indicate what sort of header this is. The types and their meanings are as follows:

TS_TAPE Tape volume label
 TS_INODE A file or directory follows. The *c_dinode* 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 *c_addr* 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 *c_magic*.
 CHECKSUM Header records checksum to this value.

The fields of the header structure are as follows:

c_type The type of the header.
c_date The date the dump was taken.
c_ddate The date the file system was dumped from.
c_volume The current volume number of the dump.
c_tapea The current number of this (1024-byte) record.
c_inumber The number of the inode being dumped if this is of type TS_INODE.
c_magic This contains the value MAGIC above, truncated as needed.
c_checksum This contains whatever value is needed to make the record sum to CHECKSUM.
c_dinode This is a copy of the inode as it appears on the file system; see *fs(5)*.
c_count The count of characters in *c_addr*.
c_addr An array of characters describing the blocks of the dumped file. A character is zero if the block associated with that character was not present on the file system, otherwise the character is non-zero. If the block was not present on the file system, no block was dumped; the block will be restored as a hole in the file. If there is not sufficient space in this record to describe all of the blocks in a file, TS_ADDR records will be scattered through the file, each one picking up where the last left off.

Each volume except the last ends with a tapemark (read as an end of file). The last volume ends with a TS_END record and then the tapemark.

The structure *idates* describes an entry in the file */etc/dumpdates* where dump history is kept. The fields of the structure are:

id_name The dumped filesystem is *'/dev/id_nam'*.
id_incno The level number of the dump tape; see *dump(8)*.
id_ddate The date of the incremental dump in system format see *types(5)*.

FILES

/etc/dumpdates

SEE ALSO

dump(8), *restore(8)*, *fs(5)*, *types(5)*

NAME

fs, inode — format of file system volume

SYNOPSIS

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

DESCRIPTION

Every file system storage volume (disk, nine-track tape, for instance) has a common format for certain vital information. Every such volume is divided into a certain number of blocks. The block size is a parameter of the file system. Sectors 0 to 15 on a file system are used to contain primary and secondary bootstrapping programs.

The actual file system begins at sector 16 with the *super block*. The layout of the super block as defined by the include file <sys/fs.h> is:

```
#define FS_MAGIC    0x011954
struct fs {
    struct fs *fs_link;          /* linked list of file systems */
    struct fs *fs_rlink;        /* used for incore super blocks */
    daddr_t fs_sbknno;          /* addr of super-block in filesys */
    daddr_t fs_cblkno;          /* offset of cyl-block in filesys */
    daddr_t fs_ibknno;          /* 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;             /* "lbnknno" calc of logical blkno */
    long fs_fshift;             /* "numfrags" calc number of frags */

    /* these are configuration parameters */
    long fs_maxcontig;          /* max number of contiguous blks */
    long fs_maxbpg;             /* max number of blks per cyl group */

    /* these fields can be computed from the others */
    long fs_fragshift;          /* block to frag shift */
    long fs_fsbtodb;           /* fsbtodb and dbtofsb shift constant */
    long fs_sbsize;             /* actual size of super block */
    long fs_csumblock;          /* 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_sparecon[6];      /* reserved for future constants */
/* sizes determined by number of cylinder groups and their sizes */
    daddr_t fs_csaddr;         /* blk addr of cyl grp summary area */
    long   fs_cssize;          /* size of cyl grp summary area */
    long   fs_cgsize;          /* cylinder group size */
/* these fields should be derived from the hardware */
    long   fs_ntrak;           /* tracks per cylinder */
    long   fs_nsect;           /* sectors per track */
    long   fs_spc;             /* sectors per cylinder */
/* this comes from the disk driver partitioning */
    long   fs_ncyl;            /* cylinders in file system */
/* these fields can be computed from the others */
    long   fs_cpg;             /* cylinders per group */
    long   fs_ipg;             /* inodes per group */
    long   fs_fpg;             /* blocks per group * fs_frag */
/* this data must be re-computed after crashes */
    struct csum fs_cstotal; /* cylinder summary information */
/* these fields are cleared at mount time */
    char   fs_fmod;            /* super block modified flag */
    char   fs_clean;           /* file system is clean flag */
    char   fs_roonly;          /* mounted read-only flag */
    char   fs_flags;           /* currently unused flag */
    char   fs_fsmnt[MAXMNTLEN]; /* name mounted on */
/* these fields retain the current block allocation info */
    long   fs_cgrotor;         /* last cg searched */
    struct csum fs_csp[MAXCSBUFS]; /* list of fs_cs info buffers */
    long   fs_cpc;             /* cyl per cycle in postbl */
    short  fs_postbl[MAXCPG][NRPOS]; /* head of blocks for each rotation */
    long   fs_magic;           /* magic number */
    u_char fs_rotbl[1];        /* list of blocks for each rotation */
/* actually longer */
};

```

Each disk drive contains some number of file systems. A file system consists of a number of cylinder groups. Each cylinder group has inodes and data.

A file system is described by its super-block, which in turn describes the cylinder groups. The super-block is critical data and is replicated in each cylinder group to protect against catastrophic loss. This is done at file system creation time and the critical super-block data does not change, so the copies need not be referenced further unless disaster strikes.

Addresses stored in inodes are capable of addressing fragments of 'blocks'. File system blocks of at most size MAXBSIZE can be optionally broken into 2, 4, or 8 pieces, each of which is addressable; these pieces may be DEV_BSIZE, or some multiple of a DEV_BSIZE unit.

Large files consist of exclusively large data blocks. To avoid undue wasted disk space, the last data block of a small file is allocated as only as many fragments of a large block as are necessary. The file system format retains only a single pointer to such a fragment, which is a piece of a single large block that has been divided. The size of such a fragment is determinable from information in the inode, using the "blksize(fs, ip, lbn)" macro.

The file system records space availability at the fragment level; to determine block availability, aligned fragments are examined.

The root inode is the root of the file system. Inode 0 can't be used for normal purposes and historically bad blocks were linked to inode 1, thus the root inode is 2 (inode 1 is no longer used for this purpose, however numerous dump tapes make this assumption, so we are stuck with it). The *lost+found* directory is given the next available inode when it is initially created by *mkfs*.

fs_minfree gives the minimum acceptable percentage of file system blocks which may be free. If the freelist drops below this level only the super-user may continue to allocate blocks. This may be set to 0 if no reserve of free blocks is deemed necessary, however severe performance degradations will be observed if the file system is run at greater than 90% full; thus the default value of *fs_minfree* is 10%.

Empirically the best trade-off between block fragmentation and overall disk utilization at a loading of 90% comes with a fragmentation of 4, thus the default fragment size is a fourth of the block size.

Cylinder group related limits: Each cylinder keeps track of the availability of blocks at different rotational positions, so that sequential blocks can be laid out with minimum rotational latency. NRPOS is the number of rotational positions which are distinguished. With NRPOS 8 the resolution of the summary information is 2ms for a typical 3600 rpm drive.

fs_rotdelay gives the minimum number of milliseconds to initiate another disk transfer on the same cylinder. It is used in determining the rotationally optimal layout for disk blocks within a file; the default value for *fs_rotdelay* is 2ms.

Each file system has a statically allocated number of inodes. An inode is allocated for each NBPI bytes of disk space. The inode allocation strategy is extremely conservative.

MAXIPG bounds the number of inodes per cylinder group, and is needed only to keep the structure simpler by having the only a single variable size element (the free bit map).

N.B.: MAXIPG must be a multiple of INOPB(fs).

MINBSIZE is the smallest allowable block size. With a MINBSIZE of 4096 it is possible to create files of size 2^{32} with only two levels of indirection. MINBSIZE must be big enough to hold a cylinder group block, thus changes to (struct cg) must keep its size within MINBSIZE. MAXCPG is limited only to dimension an array in (struct cg); it can be made larger as long as that structure's size remains within the bounds dictated by MINBSIZE. Note that super blocks are never more than size SBSIZE.

The path name on which the file system is mounted is maintained in *fs_fsmnt*. MAXMNTLEN defines the amount of space allocated in the super block for this name. The limit on the amount of summary information per file system is defined by MAXCSBUFS. It is currently parameterized for a maximum of two million cylinders.

Per cylinder group information is summarized in blocks allocated from the first cylinder group's data blocks. These blocks are read in from *fs_csaddr* (size *fs_cssize*) in addition to the super block.

N.B.: sizeof (struct csum) must be a power of two in order for the "fs_cs" macro to work.

Super block for a file system: MAXBPC bounds the size of the rotational layout tables and is limited by the fact that the super block is of size SBSIZE. The size of these tables is *inversely* proportional to the block size of the file system. The size of the tables is increased when sector sizes are not powers of two, as this increases the number of cylinders included before the rotational pattern repeats (*fs_cpc*). The size of the rotational layout tables is derived from the number of bytes remaining in (struct fs).

MAXBPG bounds the number of blocks of data per cylinder group, and is limited by the fact that cylinder groups are at most one block. The size of the free block table is derived from the size of blocks and the number of remaining bytes in the cylinder group structure (struct cg).

Inode: The inode is the focus of all file activity in the UNIX file system. There is a unique inode allocated for each active file, each current directory, each mounted-on file, text file, and the root. An inode is 'named' by its device/i-number pair. For further information, see the include file `<sys/inode.h>`.

NAME

`fstab` — static information about the filesystems

SYNOPSIS

```
#include <fstab.h>
```

DESCRIPTION

The file */etc/fstab* contains descriptive information about the various file systems. */etc/fstab* is only read by programs, and not written; it is the duty of the system administrator to properly create and maintain this file. The order of records in */etc/fstab* is important because *fsck*, *mount*, and *umount* sequentially iterate through */etc/fstab* doing their thing.

The special file name is the **block** special file name, and not the character special file name. If a program needs the character special file name, the program must create it by appending a “r” after the last “/” in the special file name.

If *fs_type* is “rw” or “ro” then the file system whose name is given in the *fs_file* field is normally mounted read-write or read-only on the specified special file. If *fs_type* is “rq”, then the file system is normally mounted read-write with disk quotas enabled. The *fs_freq* field is used for these file systems by the *dump*(8) command to determine which file systems need to be dumped. The *fs_passno* field is used by the *fsck*(8) program to determine the order in which file system checks are done at reboot time. The root file system should be specified with a *fs_passno* of 1, and other file systems should have larger numbers. File systems within a drive should have distinct numbers, but file systems on different drives can be checked on the same pass to utilize parallelism available in the hardware.

If *fs_type* is “sw” then the special file is made available as a piece of swap space by the *swapon*(8) command at the end of the system reboot procedure. The fields other than *fs_spec* and *fs_type* are not used in this case.

If *fs_type* is “rq” then at boot time the file system is automatically processed by the *quota-check*(8) command and disk quotas are then enabled with *quotaon*(8). File system quotas are maintained in a file “quotas”, which is located at the root of the associated file system.

If *fs_type* is specified as “xx” the entry is ignored. This is useful to show disk partitions which are currently not used.

```
#define FSTAB_RW    "rw"    /* read-write device */
#define FSTAB_RO    "ro"    /* read-only device */
#define FSTAB_RQ    "rq"    /* read-write with quotas */
#define FSTAB_SW    "sw"    /* swap device */
#define FSTAB_XX    "xx"    /* ignore totally */

struct fstab {
    char *fs_spec; /* block special device name */
    char *fs_file; /* file system path prefix */
    char *fs_type; /* rw,ro,sw or xx */
    int fs_freq; /* dump frequency, in days */
    int fs_passno; /* pass number on parallel dump */
};
```

The proper way to read records from */etc/fstab* is to use the routines *getfsent*(), *getfsspec*(), *getfstype*(), and *getfsfile*() .

FILES

/etc/fstab

SEE ALSO

getfsent(3X)

5

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 environment
f0	num	unused	tty mode flags to write messages
f1	num	unused	tty mode flags to read login name
f2	num	unused	tty mode flags to leave terminal as
fd	num	0	form-feed (vertical motion) delay
fl	str	^O	output flush character
hc	bool	false	do NOT hangup line on last close
he	str	NULL	hostname editing string
hn	str	hostname	hostname
ht	bool	false	terminal has real tabs
ig	bool	false	ignore garbage characters in login name
im	str	NULL	initial (banner) message
in	str	^C	interrupt character
is	num	unused	input speed
kl	str	^U	kill character
lc	bool	false	terminal has lower case
lm	str	login:	login prompt
ln	str	^V	"literal next" character
lo	str	/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
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	unused	line speed (input and output)
su	str	^Z	suspend character
tc	str	none	table continuation
to	num	0	timeout (seconds)
tt	str	NULL	terminal type (for environment)
ub	bool	false	do unbuffered output (of prompts etc)
uc	bool	false	terminal is known upper case only
we	str	^W	word erase character
xc	bool	false	do NOT echo control chars as ^X
xf	str	^S	XOFF (stop output) character
xn	str	^Q	XON (start output) character

If no line speed is specified, speed will not be altered from that which prevails when `getty` is entered. Specifying an input or output speed will override 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, `lm` 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.

SEE ALSO

termcap(5), *getty*(8).

BUGS

Some ignorant peasants insist on changing the default special characters, so it is wise to always specify (at least) the erase, kill, and interrupt characters in the *default* table. In all cases, '#' or '^H' typed in a login name will be treated as an erase character, and '@' will be treated as a kill character.

The delay stuff is a real crock. Apart from its general lack of flexibility, some of the delay algorithms are not implemented. The terminal driver should support sane delay settings.

Currently *login*(1) stomps on the environment, so there is no point setting it in *gettytab*.

The *he* capability is stupid.

Termcap format is horrid, something more rational should have been chosen.

NAME

group — group file

DESCRIPTION

Group contains for each group the following information:

group name
encrypted password
numerical group ID
a comma separated list of all users allowed in the group

This is an ASCII file. The fields are separated by colons; Each group is separated from the next by a new-line. If the password field is null, no password is demanded.

This file resides in directory */etc*. Because of the encrypted passwords, it can and does have general read permission and can be used, for example, to map numerical group ID's to names.

FILES

/etc/group

SEE ALSO

setgroups(2), *initgroups(3X)*, *crypt(3)*, *passwd(1)*, *passwd(5)*

BUGS

The *passwd(1)* command won't change the passwords.

NAME

hosts — host name data base

DESCRIPTION

The *hosts* file contains information regarding the known hosts on the DARPA Internet. For each host a single line should be present with the following information:

official host name

Internet address

aliases

Items are separated by any number of blanks and/or tab characters. A “#” indicates the beginning of a comment; characters up to the end of the line are not interpreted by routines which search the file. This file is normally created from the official host data base maintained at the Network Information Control Center (NIC), though local changes may be required to bring it up to date regarding unofficial aliases and/or unknown hosts.

Network addresses are specified in the conventional “.” notation using the *inet_addr()* routine from the Internet address manipulation library, *inet(3N)*. Host names may contain any printable character other than a field delimiter, newline, or comment character.

FILES

/etc/hosts

SEE ALSO

gethostent(3N)

BUGS

A name server should be used instead of a static file. A binary indexed file format should be available for fast access.

NAME

mtab — mounted file system table

SYNOPSIS

```
#include <fstab.h>
#include <mtab.h>
```

DESCRIPTION

Mtab resides in directory */etc* and contains a table of devices mounted by the *mount* command. *Umount* removes entries.

The table is a series of *mtab* structures, as defined in *<mtab.h>*. Each entry contains the null-padded name of the place where the special file is mounted, the null-padded name of the special file, and a type field, one of those defined in *<fstab.h>*. The special file has all its directories stripped away; that is, everything through the last *'/'* is thrown away. The type field indicates if the file system is mounted read-only, read-write, or read-write with disk quotas enabled.

This table is present only so people can look at it. It does not matter to *mount* if there are duplicated entries nor to *umount* if a name cannot be found.

FILES

/etc/mtab

SEE ALSO

mount(8)

5

NAME

networks — network name data base

DESCRIPTION

The *networks* file contains information regarding the known networks which comprise the DARPA Internet. For each network a single line should be present with the following information:

official network name
network number
aliases

Items are separated by any number of blanks and/or tab characters. A “#” indicates the beginning of a comment; characters up to the end of the line are not interpreted by routines which search the file. This file is normally created from the official network data base maintained at the Network Information Control Center (NIC), though local changes may be required to bring it up to date regarding unofficial aliases and/or unknown networks.

Network number may be specified in the conventional “.” notation using the *inet_network()* routine from the Internet address manipulation library, *inet(3N)*. Network names may contain any printable character other than a field delimiter, newline, or comment character.

FILES

/etc/networks

SEE ALSO

getnetent(3N)

BUGS

A name server should be used instead of a static file. A binary indexed file format should be available for fast access.

NAME

passwd — password file

DESCRIPTION

Passwd contains for each user the following information:

name (login name, contains no upper case)
 encrypted password
 numerical user ID
 numerical group ID
 user's real name, office, extension, home phone.
 initial working directory
 program to use as Shell

The name may contain '&', meaning insert the login name. This information is set by the *chfn*(1) command and used by the *finger*(1) command.

This is an ASCII file. Each field within each user's entry is separated from the next by a colon. Each user is separated from the next by a new-line. If the password field is null, no password is demanded; if the Shell field is null, then */bin/sh* is used.

This file resides in directory */etc*. Because of the encrypted passwords, it can and does have general read permission and can be used, for example, to map numerical user ID's to names.

Appropriate precautions must be taken to lock the file against changes if it is to be edited with a text editor; *vipw*(8) does the necessary locking.

FILES

/etc/passwd

SEE ALSO

getpwent(3), *login*(1), *crypt*(3), *passwd*(1), *group*(5), *chfn*(1), *finger*(1), *vipw*(8), *adduser*(8)

BUGS

A binary indexed file format should be available for fast access.

User information (name, office, etc.) should be stored elsewhere.

NAME

phones — remote host phone number data base

DESCRIPTION

The file `/etc/phones` contains the system-wide private phone numbers for the `tip(1C)` program. This file is normally unreadable, and so may contain privileged information. The format of the file 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.

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 'dottedashed.' 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 used to describe line 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 is used to describe 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 an option, **-Pprinter**, to allow explicit naming of a destination printer.

Refer to the *4.2BSD Line Printer Spooler Manual* for a complete discussion on how setup the database for a given printer.

CAPABILITIES

Refer to *termcap* for a description of the file layout.

Name	Type	Default	Description
af	str	NULL	name of accounting file
br	num	none	if <i>lp</i> is a tty, set the baud rate (ioctl call)
cf	str	NULL	cifplot data filter
df	str	NULL	tex data filter (DVI format)
fc	num	0	if <i>lp</i> 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 (plot (3X) format)
ic	bool	false	driver supports (non standard) ioctl to indent printout
if	str	NULL	name of text filter which does accounting
lf	str	"/dev/console"	error logging file name
lo	str	"lock"	name of lock file
lp	str	"/dev/lp"	device name to open for output
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 the printer device for reading and writing
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.

SEE ALSO

termcap(5), lpc(8), lpd(8), pac(8), lpr(1), lpq(1), lprm(1)
4.2BSD Line Printer Spooler Manual

NAME

protocols — protocol name data base

DESCRIPTION

The *protocols* file contains information regarding the known protocols used in the DARPA Internet. For each protocol a single line should be present with the following information:

official protocol name
protocol number
aliases

Items are separated by any number of blanks and/or tab characters. A “#” indicates the beginning of a comment; characters up to the end of the line are not interpreted by routines which search the file.

Protocol names may contain any printable character other than a field delimiter, newline, or comment character.

FILES

/etc/protocols

SEE ALSO

getprotoent(3N)

BUGS

A name server should be used instead of a static file. A binary indexed file format should be available for fast access.

NAME

remote — remote host description file

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*; e.g. "dv=/dev/harris". A numeric capability is specified by *capability#value*; e.g. "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.
- dl** (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. '^' escapes are only recognized by *tip* 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 even parity.
- pn** (str) Telephone number(s) for this host. If the telephone number field contains an @ sign, *tip* searches the file */etc/phones* file for a list of telephone numbers; c.f. *phones*(5).
- 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/cau0:el=^D^U^C^S^Q^O@:du:at=ventel:ie=#$:oe=^D:br#1200:
arpavax:\
:pn=7654321%:tc=UNIX-1200
```

FILES

/etc/remote

SEE ALSO

tip(1C), phones(5)

NAME

services — service name data base

DESCRIPTION

The *services* file contains information regarding the known services available in the DARPA Internet. For each service a single line should be present with the following information:

official service name
port number
protocol name
aliases

Items are separated by any number of blanks and/or tab characters. The port number and protocol name are considered a single *item*; a “/” is used to separate the port and protocol (e.g. “512/tcp”). A “#” indicates the beginning of a comment; characters up to the end of the line are not interpreted by routines which search the file.

Service names may contain any printable character other than a field delimiter, newline, or comment character.

FILES

/etc/services

SEE ALSO

getservent(3N)

BUGS

A name server should be used instead of a static file. A binary indexed file format should be available for fast access.

NAME

stab — symbol table types

SYNOPSIS

#include <stab.h>

DESCRIPTION

Stab.h defines some values of the *n_type* field of the symbol table of *a.out* files. These are the types for permanent symbols (i.e. not local labels, etc.) used by the old debugger *sdb* and the Berkeley Pascal compiler *pc(1)*. Symbol table entries can be produced by the *.stabs* assembler directive. This allows one to specify a double-quote delimited name, a symbol type, one char and one short of information about the symbol, and an unsigned long (usually an address). To avoid having to produce an explicit label for the address field, the *.stabd* directive can be used to implicitly address the current location. If no name is needed, symbol table entries can be generated using the *.stabn* directive. The loader promises to preserve the order of symbol table entries produced by *.stab* directives. As described in *a.out(5)*, an element of the symbol table consists of the following structure:

```
/*
 * Format of a symbol table entry.
 */
struct nlist {
    union {
        char *n_name; /* for use when in-core */
        long n_strx; /* index into file string table */
    } n_un;
    unsigned char n_type; /* type flag */
    char n_other; /* unused */
    short n_desc; /* see struct desc, below */
    unsigned n_value; /* address or offset or line */
};
```

The low bits of the *n_type* field are used to place a symbol into at most one segment, according to the following masks, defined in *<a.out.h>*. A symbol can be in none of these segments by having none of these segment bits set.

```
/*
 * Simple values for n_type.
 */
#define N_UNDF 0x0 /* undefined */
#define N_ABS 0x2 /* absolute */
#define N_TEXT 0x4 /* text */
#define N_DATA 0x6 /* data */
#define N_BSS 0x8 /* bss */

#define N_EXT 01 /* external bit, or'ed in */
```

The *n_value* field of a symbol is relocated by the linker, *ld(1)* as an address within the appropriate segment. *N_value* fields of symbols not in any segment are unchanged by the linker. In addition, the linker will discard certain symbols, according to rules of its own, unless the *n_type* field has one of the following bits set:

```
/*
 * Other permanent symbol table entries have some of the N_STAB bits set.
 * These are given in <stab.h>
 */
#define N_STAB 0xe0 /* if any of these bits set, don't discard */
```

This allows up to 112 (7 * 16) symbol types, split between the various segments. Some of these have already been claimed. The old symbolic debugger, *sdb*, uses the following *n_type* values:

```
#define N_GSYM 0x20 /* global symbol: name,,0,type,0 */
#define N_FNAME 0x22 /* procedure name (f77 kludge): name,,0 */
#define N_FUN 0x24 /* procedure: name,,0,linenumber,address */
#define N_STSYM 0x26 /* static symbol: name,,0,type,address */
#define N_LCSYM 0x28 /* .lcomm symbol: name,,0,type,address */
#define N_RSYM 0x40 /* register sym: name,,0,type,register */
#define N_SLINE 0x44 /* src line: 0,,0,linenumber,address */
#define N_SSYM 0x60 /* structure elt: name,,0,type,struct_offset */
#define N_SO 0x64 /* source file name: name,,0,0,address */
#define N_LSYM 0x80 /* local sym: name,,0,type,offset */
#define N_SOL 0x84 /* #included file name: name,,0,0,address */
#define N_PSYM 0xa0 /* parameter: name,,0,type,offset */
#define N_ENTRY 0xa4 /* alternate entry: name,linenumber,address */
#define N_LBRAC 0xc0 /* left bracket: 0,,0,nesting level,address */
#define N_RBRAC 0xe0 /* right bracket: 0,,0,nesting level,address */
#define N_BCOMM 0xe2 /* begin common: name,, */
#define N_ECOMM 0xe4 /* end common: name,, */
#define N_ECOML 0xe8 /* end common (local name): ,,address */
#define N_LENG 0xfe /* second stab entry with length information */
```

where the comments give *sdb* conventional use for *.stabs* and the *n_name*, *n_other*, *n_desc*, and *n_value* fields of the given *n_type*. *Sdb* uses the *n_desc* field to hold a type specifier in the form used by the Portable C Compiler, *cc(1)*, in which a base type is qualified in the following structure:

```
struct desc {
    short q6:2,
          q5:2,
          q4:2,
          q3:2,
          q2:2,
          q1:2,
          basic:4;
};
```

There are four qualifications, with *q1* the most significant and *q6* the least significant:

```
0 none
1 pointer
2 function
3 array
```

The sixteen basic types are assigned as follows:

```
0 undefined
1 function argument
2 character
3 short
4 int
5 long
6 float
7 double
8 structure
9 union
```


- 10 enumeration
- 11 member of enumeration
- 12 unsigned character
- 13 unsigned short
- 14 unsigned int
- 15 unsigned long

The Berkeley Pascal compiler, *pc*(1), uses the following *n_type* value:

```
#define N_PC 0x30 /* global pascal symbol: name,,0,subtype,line */
```

and uses the following subtypes to do type checking across separately compiled files:

- 1 source file name
- 2 included file name
- 3 global label
- 4 global constant
- 5 global type
- 6 global variable
- 7 global function
- 8 global procedure
- 9 external function
- 10 external procedure
- 11 library variable
- 12 library routine

SEE ALSO

as(1), *ld*(1), *dbx*(1), *a.out*(5)

BUGS

Sdb assumes that a symbol of type *N_GSYM* with name *name* is located at address *_name*.

More basic types are needed.

NAME

tar — tape archive file format

DESCRIPTION

Tar, (the tape archive command) dumps several files into one, in a medium suitable for transportation.

A “tar tape” or file is a series of blocks. Each block is of size **TBLOCK**. A file on the tape is represented by a header block which describes the file, followed by zero or more blocks which give the contents of the file. At the end of the tape are two blocks filled with binary zeros, as an end-of-file indicator.

The blocks are grouped for physical I/O operations. Each group of *n* blocks (where *n* is set by the **b** keyletter on the *tar*(1) command line — default is 20 blocks) is written with a single system call; on nine-track tapes, the result of this write is a single tape record. The last group is always written at the full size, so blocks after the two zero blocks contain random data. On reading, the specified or default group size is used for the first read, but if that read returns less than a full tape block, the reduced block size is used for further reads.

The header block looks like:

```
#define TBLOCK      512
#define NAMSIZ      100

union hblock {
    char dummy[TBLOCK];
    struct header {
        char name[NAMSIZ];
        char mode[8];
        char uid[8];
        char gid[8];
        char size[12];
        char mtime[12];
        char chksum[8];
        char linkflag;
        char linkname[NAMSIZ];
    } dbuf;
};
```

Name is a null-terminated string. The other fields are zero-filled octal numbers in ASCII. Each field (of width *w*) contains *w*-2 digits, a space, and a null, except *size* and *mtime*, which do not contain the trailing null. *Name* is the name of the file, as specified on the *tar* command line. Files dumped because they were in a directory which was named in the command line have the directory name as prefix and *filename* as suffix. *Mode* is the file mode, with the top bit masked off. *Uid* and *gid* are the user and group numbers which own the file. *Size* is the size of the file in bytes. Links and symbolic links are dumped with this field specified as zero. *Mtime* is the modification time of the file at the time it was dumped. *Chksum* is a decimal ASCII value which represents the sum of all the bytes in the header block. When calculating the checksum, the *chksum* field is treated as if it were all blanks. *Linkflag* is ASCII ‘0’ if the file is “normal” or a special file, ASCII ‘1’ if it is an hard link, and ASCII ‘2’ if it is a symbolic link. The name linked-to, if any, is in *linkname*, with a trailing null. Unused fields of the header are binary zeros (and are included in the checksum).

The first time a given i-node number is dumped, it is dumped as a regular file. The second and subsequent times, it is dumped as a link instead. Upon retrieval, if a link entry is retrieved, but not the file it was linked to, an error message is printed and the tape must be manually re-scanned to retrieve the linked-to file.

The encoding of the header is designed to be portable across machines.

SEE ALSO

tar(1)

BUGS

Names or linknames longer than NAMSIZ produce error reports and cannot be dumped.

NAME

termcap — terminal capability data base

SYNOPSIS

/etc/termcap

DESCRIPTION

Termcap is a data base describing terminals, used, *e.g.*, by *vi(1)* and *curses(3X)*. Terminals are described in *termcap* by giving a set of capabilities which they have, and by describing how operations are performed. Padding requirements and initialization sequences are included in *termcap*.

Entries in *termcap* consist of a number of ':' separated fields. The first entry for each terminal gives the names which are known for the terminal, separated by '|' characters. The first name is always 2 characters long and is used by older version 6 systems which store the terminal type in a 16 bit word in a systemwide data base. The second name given is the most common abbreviation for the terminal, and the last name given should be a long name fully identifying the terminal. The second name should contain no blanks; the last name may well contain blanks for readability.

CAPABILITIES

(P) indicates padding may be specified

(P*) indicates that padding may be based on no. lines affected

Name Type Pad? Description

ae	str	(P)	End alternate character set
al	str	(P*)	Add new blank line
am	bool		Terminal has automatic margins
as	str	(P)	Start alternate character set
bc	str		Backspace if not ^H
bs	bool		Terminal can backspace with ^H
bt	str	(P)	Back tab
bw	bool		Backspace wraps from column 0 to last column
CC	str		Command character in prototype if terminal settable
cd	str	(P*)	Clear to end of display
ce	str	(P)	Clear to end of line
ch	str	(P)	Like cm but horizontal motion only, line stays same
cl	str	(P*)	Clear screen
cm	str	(P)	Cursor motion
co	num		Number of columns in a line
cr	str	(P*)	Carriage return, (default ^M)
cs	str	(P)	Change scrolling region (vt100), like cm
cv	str	(P)	Like ch but vertical only.
da	bool		Display may be retained above
dB	num		Number of 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
li	num	Number of lines on screen or page
ll	str	Last line, first column (if no cm)
ma	str	Arrow key map, used by vi version 2 only
mi	bool	Safe to move while in insert mode
ml	str	Memory lock on above cursor.
ms	bool	Safe to move while in standout and underline mode
mu	str	Memory unlock (turn off memory lock).
nc	bool	No correctly working carriage return (DM2500,H2000)
nd	str	Non-destructive space (cursor right)
nl	str	(P*) Newline character (default \n)
ns	bool	Terminal is a CRT but doesn't scroll.
os	bool	Terminal overstrikes
pc	str	Pad character (rather than null)
pt	bool	Has hardware tabs (may need to be set with is)
se	str	End stand out mode
sf	str	(P) Scroll forwards
sg	num	Number of blank chars left by so or se
so	str	Begin stand out mode
sr	str	(P) Scroll reverse (backwards)
ta	str	(P) Tab (other than ^I or with padding)
tc	str	Entry of similar terminal - must be last
te	str	String to end programs that use cm
ti	str	String to begin programs that use cm
uc	str	Underscore one char and move past it
ue	str	End underscore mode
ug	num	Number of blank chars left by us or ue

ul	bool	Terminal underlines even though it doesn't overstrike
up	str	Upline (cursor up)
us	str	Start underscore mode
vb	str	Visible bell (may not move cursor)
ve	str	Sequence to end open/visual mode
vs	str	Sequence to start open/visual mode
xb	bool	Beehive (f1=escape, f2=ctrl C)
xn	bool	A newline is ignored after a wrap (Concept)
xr	bool	Return acts like <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 (Telera 1061)

A Sample Entry

The following entry, which describes the Concept-100, is among the more complex entries in the *termcap* file as of this writing. (This particular concept entry is outdated, and is used as an example only.)

```
c|c100|concept100:is=\EU\Ef\E7\E5\E8\EI\ENH\EK\E\200\Eo&\200:\
:al=3*\E^R:am:bs:cd=16*\E^C:ce=16\E^S:cl=2*\L:cm=\Ea%+ %+ :co#80:\
:dc=16\E^A:dl=3*\E^B:ei=\E\200:eo:im=\E^P:in:ip=16*:li#24:mi:nd=\E=:\
:se=\Ed\Ee:so=\ED\EE:ta=8\t:ul:up=\E;:vb=\Ek\EK:xn:
```

Entries may continue onto multiple lines by giving a `\` as the last character of a line, and that empty fields may be included for readability (here between the last field on a line and the first field on the next). Capabilities in *termcap* are of three types: Boolean capabilities which indicate that the terminal has some particular feature, numeric capabilities giving the size of the terminal or the size of particular delays, and string capabilities, which give a sequence which can be used to perform particular terminal operations.

Types of Capabilities

All capabilities have two letter codes. For instance, the fact that the Concept has "automatic margins" (i.e. an automatic return and linefeed when the end of a line is reached) is indicated by the capability **am**. Hence the description of the Concept includes **am**. Numeric capabilities are followed by the character '#' and then the value. Thus **co** which indicates the number of columns the terminal has gives the value '80' for the Concept.

Finally, string valued capabilities, such as **ce** (clear to end of line sequence) are given by the two character code, an '=', and then a string ending at the next following '.'. A delay in milliseconds may appear after the '=' in such a capability, and padding characters are supplied by the editor after the remainder of the string is sent to provide this delay. The delay can be either a integer, e.g. '20', or an integer followed by an '*', i.e. '3*'. A '*' indicates that the padding required is proportional to the number of lines affected by the operation, and the amount given is the per-affected-unit padding required. When a '*' is specified, it is sometimes useful to give a delay of the form '3.5' specify a delay per unit to tenths of milliseconds.

A number of escape sequences are provided in the string valued capabilities for easy encoding of characters there. A `\E` maps to an ESCAPE character, `^x` maps to a control-x for any appropriate x, and the sequences `\n \r \t \b \f` give a newline, return, tab, backspace and formfeed. Finally, characters may be given as three octal digits after a `\`, and the characters `^` and `\` may be given as `\^` and `\\`. If it is necessary to place a `:` in a capability it must be escaped in octal as `\072`. If it is necessary to place a null character in a string capability it must be encoded as `\200`. The routines which deal with *termcap* use C strings, and strip the high bits of the output very late so that a `\200` comes out as a `\000` would.

Preparing Descriptions

We now outline how to prepare descriptions of terminals. The most effective way to prepare a terminal description is by imitating the description of a similar terminal in *termcap* and to build up a description gradually, using partial descriptions with *ex* to check that they are correct. Be aware that a very unusual terminal may expose deficiencies in the ability of the *termcap* file to describe it or bugs in *ex*. To easily test a new terminal description you can set the environment variable *TERMCAP* to a pathname of a file containing the description you are working on and the editor will look there rather than in *letc/termcap*. *TERMCAP* can also be set to the *termcap* entry itself to avoid reading the file when starting up the editor. (This only works on version 7 systems.)

Basic capabilities

The number of columns on each line for the terminal is given by the *co* numeric capability. If the terminal is a CRT, then the number of lines on the screen is given by the *li* capability. If the terminal wraps around to the beginning of the next line when it reaches the right margin, then it should have the *am* capability. If the terminal can clear its screen, then this is given by the *cl* string capability. If the terminal can backspace, then it should have the *bs* capability, unless a backspace is accomplished by a character other than ^H (ugh) in which case you should give this character as the *bc* string capability. If it overstrikes (rather than clearing a position when a character is struck over) then it should have the *os* capability.

A very important point here is that the local cursor motions encoded in *termcap* are undefined at the left and top edges of a CRT terminal. The editor will never attempt to backspace around the left edge, nor will it attempt to go up locally off the top. The editor assumes that feeding off the bottom of the screen will cause the screen to scroll up, and the *am* capability tells whether the cursor sticks at the right edge of the screen. If the terminal has switch selectable automatic margins, the *termcap* file usually assumes that this is on, i.e. *am*.

These capabilities suffice to describe hardcopy and "glass-ty" 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|si adm3:am:bs:cl=^Z:li#24:co#80
```

Cursor addressing

Cursor addressing in the terminal is described by a *cm* string capability, with *printf(3S)* like escapes *%x* in it. These substitute to encodings of the current line or column position, while other characters are passed through unchanged. If the *cm* string is thought of as being a function, then its arguments are the line and then the column to which motion is desired, and the *%* encodings have the following meanings:

```
%d    as in printf, 0 origin
%2    like %2d
%3    like %3d
%     like %c
%+x   adds x to value, then %
%>x  if value > x adds y, no output.
%r    reverses order of line and column, no output
%i    increments line/column (for 1 origin)
%%    gives a single %
%n    exclusive or row and column with 0140 (DM2500)
%B    BCD (16*(x/10) + (x%10)), no output.
%D    Reverse coding (x-2*(x%16)), no output. (Delta Data).
```

Consider the HP2645, which, to get to row 3 and column 12, needs to be sent `\E&a12c03Y` padded for 6 milliseconds. Note that the order of the rows and columns is inverted here, and that the row and column are printed as two digits. Thus its `cm` capability is `"cm=6\E&%r%2c%2Y"`. The Microterm ACT-IV needs the current row and column sent preceded by a `^T`, with the row and column simply encoded in binary, `"cm=^T%.%"`. Terminals which use `"%."` need to be able to backspace the cursor (`bs` or `bc`), and to move the cursor up one line on the screen (`up` introduced below). This is necessary because it is not always safe to transmit `\t`, `\n ^D` and `\r`, as the system may change or discard them.

A final example is the LSI ADM-3a, which uses row and column offset by a blank character, thus `"cm=\E=% + % + "`.

Cursor motions

If the terminal can move the cursor one position to the right, leaving the character at the current position unchanged, then this sequence should be given as `nd` (non-destructive space). If it can move the cursor up a line on the screen in the same column, this should be given as `up`. If the terminal has no cursor addressing capability, but can home the cursor (to very upper left corner of screen) then this can be given as `ho`; similarly a fast way of getting to the lower left hand corner can be given as `ll`; this may involve going up with `up` from the home position, but the editor will never do this itself (unless `ll` does) because it makes no assumption about the effect of moving up from the home position.

Area clears

If the terminal can clear from the current position to the end of the line, leaving the cursor where it is, this should be given as `ce`. If the terminal can clear from the current position to the end of the display, then this should be given as `ed`. The editor only uses `ed` from the first column of a line.

Insert/delete line

If the terminal can open a new blank line before the line where the cursor is, this should be given as `al`; this is done only from the first position of a line. The cursor must then appear on the newly blank line. If the terminal can delete the line which the cursor is on, then this should be given as `dl`; this is done only from the first position on the line to be deleted. If the terminal can scroll the screen backwards, then this can be given as `sb`, but just `al` suffices. If the terminal can retain display memory above then the `da` capability should be given; if display memory can be retained below then `db` should be given. These let the editor understand that deleting a line on the screen may bring non-blank lines up from below or that scrolling back with `sb` may bring down non-blank lines.

Insert/delete character

There are two basic kinds of intelligent terminals with respect to insert/delete character which can be described using *termcap*. The most common insert/delete character operations affect only the characters on the current line and shift characters off the end of the line rigidly. Other terminals, such as the Concept 100 and the Perkin Elmer Owl, make a distinction between typed and untyped blanks on the screen, shifting upon an insert or delete only to an untyped blank on the screen which is either eliminated, or expanded to two untyped blanks. You can find out which kind of terminal you have by clearing the screen and then typing text separated by cursor motions. Type `"abc def"` using local cursor motions (not spaces) between the `"abc"` and the `"def"`. Then position the cursor before the `"abc"` and put the terminal in insert mode. If typing characters causes the rest of the line to shift rigidly and characters to fall off the end, then your terminal does not distinguish between blanks and untyped positions. If the `"abc"` shifts over to the `"def"` which then move together around the end of the current line and onto the next as you insert, you have the second type of terminal, and should give the capability `in`, which stands for "insert null". If your terminal does something different and unusual then you

may have to modify the editor to get it to use the insert mode your terminal defines. We have seen no terminals which have an insert mode not falling into one of these two classes.

The editor can handle both terminals which have an insert mode, and terminals which send a simple sequence to open a blank position on the current line. Give as **im** the sequence to get into insert mode, or give it an empty value if your terminal uses a sequence to insert a blank position. Give as **ei** the sequence to leave insert mode (give this, with an empty value also if you gave **im** so). Now give as **ic** any sequence needed to be sent just before sending the character to be inserted. Most terminals with a true insert mode will not give **ic**, terminals which send a sequence to open a screen position should give it here. (Insert mode is preferable to the sequence to open a position on the screen if your terminal has both.) If post insert padding is needed, give this as a number of milliseconds in **ip** (a string option). Any other sequence which may need to be sent after an insert of a single character may also be given in **ip**.

It is occasionally necessary to move around while in insert mode to delete characters on the same line (e.g. if there is a tab after the insertion position). If your terminal allows motion while in insert mode you can give the capability **mi** to speed up inserting in this case. Omitting **mi** will affect only speed. Some terminals (notably Datamedia's) must not have **mi** because of the way their insert mode works.

Finally, you can specify delete mode by giving **dm** and **ed** to enter and exit delete mode, and **dc** to delete a single character while in delete mode.

Highlighting, underlining, and visible bells

If your terminal has sequences to enter and exit standout mode these can be given as **so** and **se** respectively. If there are several flavors of standout mode (such as inverse video, blinking, or underlining — half bright is not usually an acceptable “standout” mode unless the terminal is in inverse video mode constantly) the preferred mode is inverse video by itself. If the code to change into or out of standout mode leaves one or even two blank spaces on the screen, as the TVI 912 and Teleray 1061 do, then **ug** should be given to tell how many spaces are left.

Codes to begin underlining and end underlining can be given as **us** and **ue** respectively. If the terminal has a code to underline the current character and move the cursor one space to the right, such as the Microterm Mime, this can be given as **uc**. (If the underline code does not move the cursor to the right, give the code followed by a nondestructive space.)

Many terminals, such as the HP 2621, automatically leave standout mode when they move to a new line or the cursor is addressed. Programs using standout mode should exit standout mode before moving the cursor or sending a newline.

If the terminal has a way of flashing the screen to indicate an error quietly (a bell replacement) then this can be given as **vb**; it must not move the cursor. If the terminal should be placed in a different mode during open and visual modes of *ex*, this can be given as **vs** and **ve**, sent at the start and end of these modes respectively. These can be used to change, e.g., from a underline to a block cursor and back.

If the terminal needs to be in a special mode when running a program that addresses the cursor, the codes to enter and exit this mode can be given as **ti** and **te**. This arises, for example, from terminals like the Concept with more than one page of memory. If the terminal has only memory relative cursor addressing and not screen relative cursor addressing, a one screen-sized window must be fixed into the terminal for cursor addressing to work properly.

If your terminal correctly generates underlined characters (with no special codes needed) even though it does not overstrike, then you should give the capability **ul**. If overstrikes are erasable with a blank, then this should be indicated by giving **eo**.

Keypad

If the terminal has a keypad that transmits codes when the keys are pressed, this information can be given. Note that it is not possible to handle terminals where the keypad only works in local (this applies, for example, to the unshifted HP 2621 keys). If the keypad can be set to transmit or not transmit, give these codes as **ks** and **ke**. Otherwise the keypad is assumed to always transmit. The codes sent by the left arrow, right arrow, up arrow, down arrow, and home keys can be given as **kl**, **kr**, **ku**, **kd**, and **kh** respectively. If there are function keys such as **f0**, **f1**, ..., **f9**, the codes they send can be given as **k0**, **k1**, ..., **k9**. If these keys have labels other than the default **f0** through **f9**, the labels can be given as **l0**, **l1**, ..., **l9**. If there are other keys that transmit the same code as the terminal expects for the corresponding function, such as clear screen, the *termcap* 2 letter codes can be given in the **ko** capability, for example, “:ko=cl,ll,sf,sb:”, which says that the terminal has clear, home down, scroll down, and scroll up keys that transmit the same thing as the **cl**, **ll**, **sf**, and **sb** entries.

The **ma** entry is also used to indicate arrow keys on terminals which have single character arrow keys. It is obsolete but still in use in version 2 of *vi*, which must be run on some minicomputers due to memory limitations. This field is redundant with **kl**, **kr**, **ku**, **kd**, and **kh**. It consists of groups of two characters. In each group, the first character is what an arrow key sends, the second character is the corresponding *vi* command. These commands are **h** for **kl**, **j** for **kd**, **k** for **ku**, **l** for **kr**, and **H** for **kh**. For example, the mime would be **:ma=^Kj^Zk^Xl**: indicating arrow keys left (^H), down (^K), up (^Z), and right (^X). (There is no home key on the mime.)

Miscellaneous

If the terminal requires other than a null (zero) character as a pad, then this can be given as **pc**.

If tabs on the terminal require padding, or if the terminal uses a character other than **^I** to tab, then this can be given as **ta**.

Hazeltine terminals, which don't allow “” characters to be printed should indicate **hz**. Datamedia terminals, which echo carriage-return linefeed for carriage return and then ignore a following linefeed should indicate **nc**. Early Concept terminals, which ignore a linefeed immediately after an **am** wrap, should indicate **xn**. If an erase-eol is required to get rid of stand-out (instead of merely writing on top of it), **xs** should be given. Teleray terminals, where tabs turn all characters moved over to blanks, should indicate **xt**. Other specific terminal problems may be corrected by adding more capabilities of the form **xx**.

Other capabilities include **is**, an initialization string for the terminal, and **if**, the name of a file containing long initialization strings. These strings are expected to properly clear and then set the tabs on the terminal, if the terminal has settable tabs. If both are given, **is** will be printed before **if**. This is useful where **if** is *lusr/lib/tabset/std* but **is** clears the tabs first.

Similar Terminals

If there are two very similar terminals, one can be defined as being just like the other with certain exceptions. The string capability **tc** can be given with the name of the similar terminal. This capability must be *last* and the combined length of the two entries must not exceed 1024. Since *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)

AUTHOR

William Joy

Mark Horton added underlining and keypad support

BUGS

Ex allows only 256 characters for string capabilities, and the routines in *termcap*(3X) do not check for overflow of this buffer. The total length of a single entry (excluding only escaped newlines) may not exceed 1024.

The **ma**, **vs**, and **ve** entries are specific to the *vi* program.

Not all programs support all entries. There are entries that are not supported by any program.

NAME

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 $(\text{size} + 511) / 512$ 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), *tp(1)*

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.

The first character of a line in the *ttys* file is either '0' or '1'. If the first character on the line is a '0', the *init* program ignores that line. If the first character on the line is a '1', the *init* program creates a login process for that line. The second character on each line is used as an argument to *getty*(8), which performs such tasks as baud-rate recognition, reading the login name, and calling *login*. For normal lines, the character is '0'; other characters can be used, for example, with hard-wired terminals where speed recognition is unnecessary or which have special characteristics. (*Getty* will have to be fixed in such cases.) The remainder of the line is the terminal's entry in the device directory, */dev*.

FILES

/etc/ttys

SEE ALSO

gettytab(5), *init*(8), *getty*(8), *login*(1)

NAME

ttytype — data base of terminal types by port

SYNOPSIS

/etc/ttytype

DESCRIPTION

Ttytype is a database containing, for each tty port on the system, the kind of terminal that is attached to it. There is one line per port, containing the terminal kind (as a name listed in *termcap* (5)), a space, and the name of the tty, minus */dev/*.

This information is read by *tset*(1) and by *login*(1) to initialize the TERM variable at login time.

SEE ALSO

tset(1), *login*(1)

BUGS

Some lines are merely known as “dialup” or “plugboard”.

NAME

types — primitive system data types

SYNOPSIS

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

DESCRIPTION

The data types defined in the include file are used in UNIX system code; some data of these types are accessible to user code:

```
/*      types.h      6.1      83/07/29*/

/*
 * Basic system types and major/minor device constructing/busting macros.
 */

/* major part of a device */
#define major(x) ((int)(((unsigned)(x)>>8)&0377))

/* minor part of a device */
#define minor(x) ((int)((x)&0377))

/* make a device number */
#define makedev(x,y) ((dev_t)(((x)<<8)|(y)))

typedef unsigned char    u_char;
typedef unsigned short   u_short;
typedef unsigned int     u_int;
typedef unsigned long    u_long;
typedef unsigned short   ushort; /* sys III compat */

#ifdef vax
typedef struct          _physadr { int r[1]; } *physadr;
typedef struct          label_t {
    int                 val[14];
} label_t;
#endif
typedef struct          _quad { long val[2]; } quad;
typedef long            daddr_t;
typedef char *          caddr_t;
typedef u_long          ino_t;
typedef long            swblk_t;
typedef int             size_t;
typedef int             time_t;
typedef short          dev_t;
typedef int             off_t;

typedef struct          fd_set { int fds_bits[1]; } fd_set;
```

The form *daddr_t* is used for disk addresses except in an i-node on disk, see *fs(5)*. Times are encoded in seconds since 00:00:00 GMT, January 1, 1970. The major and minor parts of a device code specify kind and unit number of a device and are installation-dependent. Offsets are measured in bytes from the beginning of a file. The *label_t* variables are used to save the processor state while another process is running.

SEE ALSO

fs(5), time(3), lseek(2), adb(1)

NAME

utmp, wtmp — login records

SYNOPSIS

```
#include <utmp.h>
```

DESCRIPTION

The *utmp* file records information about who is currently using the system. The file is a sequence of entries with the following structure declared in the include file:

```
/*    utmp.h 4.2    83/05/22    */

/*
 * Structure of utmp and wtmp files.
 *
 * Assuming the number 8 is unwise.
 */
struct utmp {
    char    ut_line[8];        /* tty name */
    char    ut_name[8];       /* user id */
    char    ut_host[16];      /* host name, if remote */
    long    ut_time;         /* time on */
};
```

This structure gives the name of the special file associated with the user's terminal, the user's login name, and the time of the login in the form of *time(3C)*.

The *wtmp* file records all logins and logouts. A null user name indicates a logout on the associated terminal. Furthermore, the terminal name '' indicates that the system was rebooted at the indicated time; the adjacent pair of entries with terminal names '[' and ']' indicate the system-maintained time just before and just after a *date* command has changed the system's idea of the time.

Wtmp is maintained by *login(1)* and *init(8)*. Neither of these programs creates the file, so if it is removed record-keeping is turned off. It is summarized by *ac(8)*.

FILES

```
/etc/utmp
/usr/adm/wtmp
```

SEE ALSO

login(1), *init(8)*, *who(1)*, *ac(8)*

NAME

uencode — format of an encoded uencode file

DESCRIPTION

Files output by *uencode(1C)* consist of a header line, followed by a number of body lines, and a trailer line. *Udecode(1C)* will ignore any lines preceding the header or following the trailer. Lines preceding a header must not, of course, look like a header.

The header line is distinguished by having the first 6 characters “begin”. The word *begin* is followed by a mode (in octal), and a string which names the remote file. A space separates 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

uencode(1C), udecode(1C), uuse(1C), uucp(1C), mail(1)

NAME

vfont — font formats for the Benson-Varian or Versatec

SYNOPSIS

`/usr/lib/vfont/*`

DESCRIPTION

The fonts for the printer/plotters have the following format. Each file contains a header, an array of 256 character description structures, and then the bit maps for the characters themselves. The header has the following format:

```

struct header {
    short      magic;
    unsigned short size;
    short      maxx;
    short      maxy;
    short      xtnd;
} header;

```

The *magic* number is 0436 (octal). The *maxx*, *maxy*, and *xtnd* fields are not used at the current time. *Maxx* and *maxy* are intended to be the maximum horizontal and vertical size of any glyph in the font, in raster lines. The *size* is the size of the bit maps for the characters in bytes. Before the maps for the characters is an array of 256 structures for each of the possible characters in the font. Each element of the array has the form:

```

struct dispatch {
    unsigned short addr;
    short      nbytes;
    char      up;
    char      down;
    char      left;
    char      right;
    short     width;
};

```

The *nbytes* field is nonzero for characters which actually exist. For such characters, the *addr* field is an offset into the rest of the file where the data for that character begins. There are *up+down* rows of data for each character, each of which has *left+right* bits, rounded up to a number of bytes. The *width* field is not used by vcat, although it is to make width tables for troff. It represents the logical width of the glyph, in raster lines, and shows where the base point of the next glyph would be.

FILES

`/usr/lib/vfont/*`

SEE ALSO

troff(1), pti(1), vpr(1), vtroff(1), vfontinfo(1)

NAME

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)*.

FIELDS

The following table names and describes each field.

Name Type Description

pb	str	regular expression for start of a procedure
bb	str	regular expression for start of a lexical 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
sb	str	regular expression for the start of a string
se	str	regular expression for the end of a string
lb	str	regular expression for the start of a character constant
le	str	regular expression for the end of a character constant
tl	bool	present means procedures are only defined at the top lexical level
oc	bool	present means upper and lower case are equivalent
kw	str	a list of keywords separated by spaces

Example

The following entry, which describes the C language, is typical of a language entry.

```
Cc:  :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".

Entries may continue onto multiple lines by giving a \ as the last character of a line. 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.

REGULAR EXPRESSIONS

Vgrindefs uses regular expression which are very similar to those of *ex(1)* and *lex(1)*. The characters '^', '\$', '.', and '\' are reserved characters and must be "quoted" with a preceding \ if they are to be included as normal characters. The metasympols and their meanings are:

\$	the end of a line
^	the beginning of a line
\\d	a delimiter (space, tab, newline, start of line)
\\a	matches any string of symbols (like .* in lex)
\\p	matches any alphanumeric name. In a procedure definition (pb) 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) which can include the string delimiter in a string b 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".

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.

FILES

/usr/lib/vgrindefs file containing terminal descriptions

SEE ALSO

vgrind(1), troff(1)

AUTHOR

Dave Presotto

BUGS

