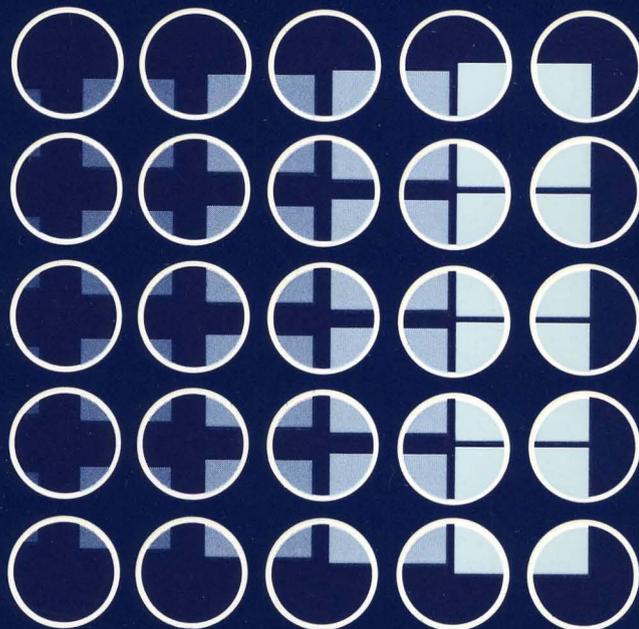


ULTRIX-32™

Programmer's Manual

ULTRIX-32™
Programmer's Manual
Sections 2, 3, 5, and 7

Order No. AA-BG54A-TE



digital
software

ULTRIX-32™
Programmer's Manual
Sections 2, 3, 5, and 7

Order No. AA-BG54A-TE

digital equipment corporation, merrimack, new hampshire

First printing, May 1984

Copyright © 1984 by Digital Equipment Corporation.

The information in this document is subject to change without notice and should not be construed as a commitment by Digital Equipment Corporation. Digital Equipment Corporation assumes no responsibility for any errors that may appear in this document.

The software described in this document is furnished under a license and may be used or copied only in accordance with the terms of such license.

No responsibility is assumed for the use or reliability of software on equipment that is not supplied by DIGITAL or its affiliated companies.

The postage-paid READER'S COMMENTS form on the last page of this document requests your critical evaluation to assist us in preparing future documentation.

The following are trademarks of Digital Equipment Corporation:

DEC	ULTRIX-32
DECUS	UNIBUS
MASSBUS	VAX
PDP	VMS
ULTRIX	VT
ULTRIX-11	digital ™

UNIX is a trademark of AT&T Bell Laboratories.

Information herein is derived from copyrighted material as permitted under a license agreement with AT&T Bell Laboratories.

This software and documentation is based in part on the Fourth Berkeley Software Distribution under license from the Regents of the University of California. We acknowledge the Electrical Engineering and Computer Sciences Departments at the Berkeley Campus of the University of California for their role in its development.

) This software and documentation is based in part on the Fourth Berkeley Software Distribution under license from The Regents of the University of California. We acknowledge the following individuals for their role in its development:

Eric Allman, Ken Arnold, Ozalp Babaoglu, Scott B. Baden, Jerry Berkman, John Breedlove, Earl T. Cohen, Robert P. Corbett, Mike Curry, Steve Feldman, Tom Ferrin, John Foderaro, Susan L. Graham, Charles Haley, Robert R. Henry, Andy Hertzfeld, Mark Horton, S.C. Johnson, William Joy, Howard Katseff, Peter Kessler, Jim Kleckner, J.E. Kulp, James Larus, Kevin Layer, Mike Lesk, Steve Levine, Jeff Levinsky, Louise Madrid, M. Kirk McKusick, Colin L. McMaster, Mikey Olson, Geoffrey Peck, Ed Pelegri-Llopart, Rob Pike, Dave Presotto, John F. Reiser, Asa Romberger, Bill Rowan, Jeff Schreibman, Eric P. Scott, Greg Shenaut, Eric Shienbrood, Kurt Shoens, Keith Sklower, Helge Skrivervik, Al Stanberger, Ken Thompson, Michael C. Toy, Richard Tuck, Bill Tuthill, Mike Urban, Edward Wang, David Wasley, Joseph Weizenbaum, Jon L. White, Glenn Wichman, Niklaus Wirth.

)

)

ULTRIX-32 Documentation Set

1. Organization

The ULTRIX-32 documentation set is organized into four separate binders. The documentation in each binder is so organized to better meet the needs of three separate audiences in the performance of their respective tasks. The ULTRIX-32 documentation set comprises:

Programmer's Manual Binder 1 – General Users

Section 1 – Commands

Documentation for all user-invoked programs (commands)

Section 6 – Games

Documentation for all user-invoked game programs

Programmer's Manual Binder 2 – Programmers

Section 2 – System Calls

Documentation for the system calls (entries into the kernel)

Section 3 – Subroutines

Documentation for the library subroutines

Section 5 – File Formats and Conventions

Documentation for the output and system file structures

Section 7 – Macro Packages and Language Conventions

Documentation for miscellaneous information

Programmer's Manual Binder 3A – System Managers

Installation Guide

Documentation for installing an ULTRIX-32 system

Building an ULTRIX-32 System with the Config Program

Documentation for configuring an executable kernel image

4.2BSD Line Printer Spooler

Documentation for installing the line printer spooling system

Sendmail Installation and Operations Guide

Documentation for installing and operating the sendmail system

UUCP Installation and Administration
Documentation for installing and administering the uucp system

Programmer's Manual Binder 3B – System Managers

Guidelines for System Management

Documentation for maintaining an installed ULTRIX-32 system

Section 4 – Special Files

Documentation for the special files (I/O devices and drivers)

Section 8 – Maintenance Commands

Documentation for the system maintenance programs (commands)

2. Man(1) Format

Each numbered section in Binder 1, Binder 2, and Binder 3B contains an introduction and separate entries that correspond to those created by the `man(1)` command. Each entry, following the order in their respective binders, describes a user command or game; a system call, library subroutine, file structure, or macro package; and a special file or maintenance command.

Regardless of their respective binder, all entries have a single, consistent format. First, the header provides information that quickly identifies the entry: name and section number (enclosed in parentheses). For example, `AARDVARK(6)` identifies the `aardvark(6)` game (Binder 1). Then, the listed subsections provide specific information about the entry. Although each entry lists only those subsections that are applicable, seven main subsections may be used. Finally, the footer provides paging information: section and consecutive page number. For example, the footers for section 6 (Binder 1) are 6-1 through 6-35.

The seven main subsections are:

NAME

This subsection lists the exact name and a short description of its function.

SYNTAX

This subsection lists the complete syntax. Boldface indicates literals. A minus sign (-) indicates command options. Ellipses (...) indicate that the preceding argument may be repeated. Square brackets [] indicate optional arguments.

DESCRIPTION

This subsection provides a detailed description of function and background.

FILES

This subsection lists those related files that either are part of or are used during execution.

DIAGNOSTICS

This subsection lists those diagnostic messages that may be produced. Since most self-explanatory messages are not listed, this subsection is not comprehensive.

RESTRICTIONS

This subsection lists those restrictions that are known to apply.

SEE ALSO

This subsection lists the names of the related entries and other documentation.

3. Conventions

The following conventions apply specifically to these documents:

Installation Guide (Binder 3A)
 Building an ULTRIX-32 System with the Config Program (Binder 3A)
 UUCP Installation and Administration (Binder 3A)
 Guidelines for System Management (Binder 3B)

- bold** Literals are printed in **bold type**. Literals frequently indicate a specific command option and should be entered exactly as printed.
- case** The ULTRIX-32 system differentiates between uppercase and lowercase. Therefore, enter uppercase only where specifically indicated by an example or the command syntax.
- color** Examples are printed in color. Examples represent command sequences or information that the user enters from the terminal.
- <CTRL/X>** Terminal control characters are represented by <CTRL/X>, where *X* is a single character. To generate a terminal control characters, hold down the CTRL key while entering the character.
- <DELETE>** The DELETE key or ERASE character is represented by <DELETE>.
- italics** Substitutable parameters are printed in *italics*.
- <RETURN>** The RETURN key is printed as <RETURN>. To invoke a command, enter the command sequence and depress the RETURN key.

- # The superuser prompt (normally a #) is displayed at the console when the system is in single-user mode or at a terminal when the superuser is logged in.

- >>> The console subsystem prompt is represented by three right angle brackets, >>>. For further information about console commands, read the *VAX Hardware Manual*.

NAME

intro – introduction to system calls and error numbers

SYNTAX

```
#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*.

INTRO(2)

- 8 ENOEXEC Exec format error
A request is made to execute a file which, although it has the appropriate permissions, does not start with a valid magic number, see *a.out*(5).
- 9 EBADF Bad file number
Either a file descriptor refers to no open file, or a read (resp. write) request is made to a file which is open only for writing (resp. reading).
- 10 ECHILD No children
Wait and the process has no living or unwaited-for children.
- 11 EAGAIN No more processes
In a *fork*, the system's process table is full or the user is not allowed to create any more processes.
- 12 ENOMEM Not enough core
During an *execve* or *break*, a program asks for more core or swap space than the system is able to supply. A lack of swap space is normally a temporary condition, however a lack of core is not a temporary condition; the maximum size of the text, data, and stack segments is a system parameter.
- 13 EACCES Permission denied
An attempt was made to access a file in a way forbidden by the protection system.
- 14 EFAULT Bad address
The system encountered a hardware fault in attempting to access the arguments of a system call.
- 15 ENOTBLK Block device required
A plain file was mentioned where a block device was required, e.g. in *mount*.
- 16 EBUSY Mount device busy
An attempt to mount a device that was already mounted or an attempt was made to dismount a device on which there is an active file directory. (open file, current directory, mounted-on file, active text segment).
- 17 EEXIST File exists
An existing file was mentioned in an inappropriate context, e.g. *link*.
- 18 EXDEV Cross-device link
A hard link to a file on another device was attempted.
- 19 ENODEV No such device
An attempt was made to apply an inappropriate system call to a device; e.g. read a write-only device.
- 20 ENOTDIR Not a directory
A non-directory was specified where a directory is required, for example in a path name or as an argument to *chdir*.

- 21 EISDIR Is a directory
An attempt to write on a directory.
- 22 EINVAL Invalid argument
Some invalid argument: dismounting a non-mounted device, mentioning an unknown signal in *signal*, reading or writing a file for which *seek* has generated a negative pointer. Also set by math functions, see *intro*(3).
- 23 ENFILE File table overflow
The system's table of open files is full, and temporarily no more *opens* can be accepted.
- 24 EMFILE Too many open files
Customary configuration limit is 20 per process.
- 25 ENOTTY Not a typewriter
The file mentioned in an *ioctl* is not a terminal or one of the other devices to which these calls apply.
- 26 ETXTBSY Text file busy
An attempt to execute a pure-procedure program which is currently open for writing (or reading!). Also an attempt to open for writing a pure-procedure program that is being executed.
- 27 EFBIG File too large
The size of a file exceeded the maximum (about 10^9 bytes).
- 28 ENOSPC No space left on device
During a *write* to an ordinary file, there is no free space left on the device.
- 29 ESPIPE Illegal seek
An *lseek* was issued to a pipe. This error may also be issued for other non-seekable devices.
- 30 EROFS Read-only file system
An attempt to modify a file or directory was made on a device mounted read-only.
- 31 EMLINK Too many links
An attempt to make more than 32767 hard links to a file.
- 32 EPIPE Broken pipe
A write on a pipe or socket for which there is no process to read the data. This condition normally generates a signal; the error is returned if the signal is ignored.
- 33 EDOM Math argument
The argument of a function in the math package (3M) is out of the domain of the function.
- 34 ERANGE Result too large
The value of a function in the math package (3M) is unrepresentable within machine precision.

INTRO(2)

- 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.

INTRO(2)

63 ENAMETOOLONG File name too long

A component of a path name exceeded 255 characters, or an entire path name exceeded 1023 characters.

64 ENOTEMPTY Directory not empty

A directory with entries other than “.” and “..” was supplied to a remove directory or rename call.

DEFINITIONS

Process ID

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

Parent process ID

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

Process Group ID

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

Tty Group ID

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

SYNTAX

```
#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.

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

DIAGNOSTICS

The *accept* will fail if:

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

ACCEPT(2)

SEE ALSO

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

STATUS

ACCEPT(2) currently is not supported by Digital Equipment Corporation.

NAME

`access` – determine accessibility of file

SYNTAX

```
#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.

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.

DIAGNOSTICS

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

ACCESS(2)

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)

STATUS

ACCESS(2) currently is not supported by Digital Equipment Corporation.

NAME

acct – turn accounting on or off

SYNTAX

```
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. 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. 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.

DIAGNOSTICS

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)

RESTRICTIONS

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

STATUS

ACCT(2) currently is not supported by Digital Equipment Corporation.

BIND(2)

NAME

`bind` – bind a name to a socket

SYNTAX

```
#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.

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.

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*.

DIAGNOSTICS

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)`

STATUS

BIND(2) currently is not supported by Digital Equipment Corporation.

NAME

brk, *sbrk* — change data segment size

SYNTAX

```
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*.)

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.

DIAGNOSTICS

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)

RESTRICTIONS

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*.

STATUS

BRK(2) currently is not supported by Digital Equipment Corporation.

CHDIR(2)

NAME

`chdir` – change current working directory

SYNTAX

`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.

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

DIAGNOSTICS

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] *Path* points outside the process's allocated address space.
- [ELOOP] Too many symbolic links were encountered in translating the pathname.

SEE ALSO

`chroot(2)`

STATUS

CHDIR(2) currently is not supported by Digital Equipment Corporation.

NAME

`chmod` – change mode of file

SYNTAX

```
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.

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

DIAGNOSTICS

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

CHMOD(2)

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.

Fchmod will fail if:

[EBADF] The descriptor is not valid.

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

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

SEE ALSO

open(2), chown(2)

STATUS

CHMOD(2) currently is not supported by Digital Equipment Corporation.

NAME

`chown` – change owner and group of a file

SYNTAX

`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`.

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*.

DIAGNOSTICS

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.

CHOWN(2)

SEE ALSO

chmod(2), flock(2)

STATUS

CHOWN(2) currently is not supported by Digital Equipment Corporation.

NAME

chroot — change root directory

SYNTAX

```
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.

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

DIAGNOSTICS

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

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

SEE ALSO

chdir(2)

STATUS

CHROOT(2) currently is not supported by Digital Equipment Corporation.

CLOSE(2)

NAME

close – delete a descriptor

SYNTAX

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.

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.

DIAGNOSTICS

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)*

STATUS

CLOSE(2) currently is not supported by Digital Equipment Corporation.

NAME

connect — initiate a connection on a socket

SYNTAX

```
#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.

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*.

DIAGNOSTICS

The call fails if:

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

CONNECT(2)

SEE ALSO

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

STATUS

CONNECT(2) currently is not supported by Digital Equipment Corporation.

NAME

`creat` – create a new file

SYNTAX

`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.

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.

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

DIAGNOSTICS

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

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

CREAT(2)

SEE ALSO

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

STATUS

CREAT(2) currently is not supported by Digital Equipment Corporation.

NAME

`dup`, `dup2` – duplicate a descriptor

SYNTAX

```
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.

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

DIAGNOSTICS

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)*

STATUS

DUP(2) currently is not supported by Digital Equipment Corporation.

EXECVE(2)

NAME

`execve` – execute a file

SYNTAX

```
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.

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.

DIAGNOSTICS

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

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

EXECVE(2)

[EFAULT] *Path*, *argv*, or *envp* point to an illegal address.

RESTRICTIONS

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)

STATUS

EXECVE(2) currently is not supported by Digital Equipment Corporation.

NAME

_exit – terminate a process

SYNTAX

```
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*.

SEE ALSO

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

STATUS

EXIT(2) currently is not supported by Digital Equipment Corporation.

FCNTL(2)

NAME

`fcntl` – file control

SYNTAX

```
#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.

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.

DIAGNOSTICS

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

[EBADF]	<i>Fildes</i> is not a valid open file descriptor.
[EMFILE]	<i>Cmd</i> is F_DUPFD and the maximum allowed number of file descriptors are currently open.
[EINVAL]	<i>Cmd</i> is F_DUPFD and <i>arg</i> is negative or greater the maximum allowable number (see <i>getdtablesize(2)</i>).

SEE ALSO

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

RESTRICTIONS

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.

STATUS

FCNTL(2) currently is not supported by Digital Equipment Corporation.

FLOCK(2)

NAME

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

SYNTAX

```
#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.

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.

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*.

DIAGNOSTICS

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)

STATUS

FLOCK(2) currently is not supported by Digital Equipment Corporation.

FORK(2)

NAME

fork – create a new process

SYNTAX

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

DESCRIPTION

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

The child process has a unique process ID.

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

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

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

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.

DIAGNOSTICS

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)

STATUS

FORK(2) currently is not supported by Digital Equipment Corporation.

NAME

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

SYNTAX

```
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.

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

DIAGNOSTICS

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)`

STATUS

FSYNC(2) currently is not supported by Digital Equipment Corporation.

GETDTABLESIZE(2)

NAME

getdtablesize – get descriptor table size

SYNTAX

```
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)

STATUS

GETDTABLESIZE(2) currently is not supported by Digital Equipment Corporation.

NAME

getgid, *getegid* – get group identity

SYNTAX

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)

STATUS

GETGID(2) currently is not supported by Digital Equipment Corporation.

GETGROUPS(2)

NAME

`getgroups` – get group access list

SYNTAX

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

DESCRIPTION

Getgroups gets the current group access list of the user process and stores it in the array *gidset*. The parameter *ngroups* indicates the number of entries which may be placed in *gidset* and is modified on return to indicate the actual number of groups returned. No more than NGRPS, as defined in *<sys/param.h>*, will ever be returned.

A value of 0 indicates that the call succeeded, and that the number of elements of *gidset* and the set itself were returned. A value of -1 indicates that an error occurred, and the error code is stored in the global variable *errno*.

DIAGNOSTICS

The possible errors for *getgroup* are:

[EFAULT] The arguments *ngroups* or *gidset* specify invalid addresses.

SEE ALSO

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

STATUS

GETGROUPS(2) currently is not supported by Digital Equipment Corporation.

NAME

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

SYNTAX

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)

STATUS

GETHOSTID(2) currently is not supported by Digital Equipment Corporation.

GETHOSTNAME(2)

NAME

gethostname, sethostname – get/set name of current host

SYNTAX

```
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.

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*.

DIAGNOSTICS

The following errors may be returned by these calls:

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

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

SEE ALSO

gethostid(2)

RESTRICTIONS

Host names are limited to 255 characters.

STATUS

GETHOSTNAME(2) currently is not supported by Digital Equipment Corporation.

NAME

getitimer, setitimer – get/set value of interval timer

SYNTAX

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

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

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

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

DESCRIPTION

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

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

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

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

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

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

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

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

GETITIMER(2)

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).

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*.

DIAGNOSTICS

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)`

STATUS

GETITIMER(2) currently is not supported by Digital Equipment Corporation.

NAME

getpagesize – get system page size

SYNTAX

```
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)

STATUS

GETPAGESIZE(2) currently is not supported by Digital Equipment Corporation.

GETPEERNAME(2)

NAME

getpeername – get name of connected peer

SYNTAX

```
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).

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

DIAGNOSTICS

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)*

RESTRICTIONS

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

STATUS

GETPEERNAME(2) currently is not supported by Digital Equipment Corporation.

NAME

getpgrp – get process group

SYNTAX

```
pgrp = getpgrp(pid)  
int prgp;  
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 *csd*(1) to create process groups in implementing job control. The *TIOCGPRP* and *TIOCSPGRP* calls described in *tty*(4) are used to get/set the process group of the control terminal.

SEE ALSO

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

STATUS

GETPGRP(2) currently is not supported by Digital Equipment Corporation.

GETPID(2)

NAME

getpid, *getppid* – get process identification

SYNTAX

```
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)

STATUS

GETPID(2) currently is not supported by Digital Equipment Corporation.

NAME

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

SYNTAX

```
#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.

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.

DIAGNOSTICS

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)`

GETPRIORITY(2)

STATUS

GETPRIORITY(2) currently is not supported by Digital Equipment Corporation.

NAME

getrlimit, setrlimit – control maximum system resource consumption

SYNTAX

```
#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 */
};
```

GETRLIMIT(2)

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

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

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

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*.

DIAGNOSTICS

The possible errors are:

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

SEE ALSO

*cs**h*(1), *quota*(2)

STATUS

GETRLIMIT(2) currently is not supported by Digital Equipment Corporation.

NAME

getrusage – get information about resource utilization

SYNTAX

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

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

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

DESCRIPTION

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

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

The fields are interpreted as follows:

ru_utime	the total amount of time spent executing in user mode.
ru_stime	the total amount of time spent in the system executing on behalf of the process(es).
ru_maxrss	the maximum resident set size utilized (in kilobytes).
ru_ixrss	an “integral” value indicating the amount of memory used which was also shared among other processes. This value is expressed in units of kilobytes

GETRUSAGE(2)

* seconds-of-execution and is calculated by summing the number of shared memory pages in use each time the internal system clock ticks and then averaging over 1 second intervals.

<code>ru_idrss</code>	an integral value of the amount of unshared memory residing in the data segment of a process (expressed in units of kilobytes * seconds-of-execution).
<code>ru_isrss</code>	an integral value of the amount of unshared memory residing in the stack segment of a process (expressed in units of kilobytes * seconds-of-execution).
<code>ru_minflt</code>	the number of page faults serviced without any i/o activity; here i/o activity is avoided by "reclaiming" a page frame from the list of pages awaiting reallocation.
<code>ru_majflt</code>	the number of page faults serviced which required i/o activity.
<code>ru_nswap</code>	the number of times a process was "swapped" out of main memory.
<code>ru_inblock</code>	the number of times the file system had to perform input.
<code>ru_outblock</code>	the number of times the file system had to perform output.
<code>ru_msgsnd</code>	the number of ipc messages sent.
<code>ru_msgrcv</code>	the number of ipc messages received.
<code>ru_signals</code>	the number of signals delivered.
<code>ru_nvcsw</code>	the number of times a context switch resulted due to a process voluntarily giving up the processor before its time slice was completed (usually to await availability of a resource).
<code>ru_nivcsw</code>	the number of times a context switch resulted due to a higher priority process becoming runnable or because the current process exceeded its time slice.

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

SEE ALSO

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

RESTRICTIONS

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

STATUS

GETRUSAGE(2) currently is not supported by Digital Equipment Corporation.

NAME

getsockname – get socket name

SYNTAX

```
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).

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

DIAGNOSTICS

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)

RESTRICTIONS

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

STATUS

GETSOCKNAME(2) currently is not supported by Digital Equipment Corporation.

GETSOCKOPT(2)

NAME

getsockopt, setsockopt – get and set options on sockets

SYNTAX

```
#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).

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

DIAGNOSTICS

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)

STATUS

GETSOCKOPT(2) currently is not supported by Digital Equipment Corporation.

GETTIMEOFDAY(2)

NAME

gettimeofday, settimeofday – get/set date and time

SYNTAX

```
#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.

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*.

DIAGNOSTICS

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)

STATUS

GETTIMEOFDAY(2) currently is not supported by Digital Equipment Corporation.

NAME

getuid, geteuid – get user identity

SYNTAX

```
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)

STATUS

GETUID(2) currently is not supported by Digital Equipment Corporation.

IOCTL (2)

NAME

`ioctl` – control device

SYNTAX

```
#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>`.

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

DIAGNOSTICS

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)`

STATUS

IOCTL(2) currently is not supported by Digital Equipment Corporation.

NAME

kill – send signal to a process

SYNTAX

```
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.

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

DIAGNOSTICS

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

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

SEE ALSO

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

STATUS

KILL(2) currently is not supported by Digital Equipment Corporation.

KILLPG(2)

NAME

killpg – send signal to a process group

SYNTAX

```
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.

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.

DIAGNOSTICS

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)

STATUS

KILLPG(2) currently is not supported by Digital Equipment Corporation.

NAME

link – make a hard link to a file

SYNTAX

```
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.

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

DIAGNOSTICS

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)

LINK(2)

STATUS

LINK(2) currently is not supported by Digital Equipment Corporation.

NAME

listen – listen for connections on a socket

SYNTAX

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.

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

DIAGNOSTICS

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)*

RESTRICTIONS

The *backlog* is currently limited to 5.

STATUS

LISTEN(2) currently is not supported by Digital Equipment Corporation.

LSEEK(2)

NAME

`lseek` – move read/write pointer

SYNTAX

```
#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.

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

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.

DIAGNOSTICS

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)`

STATUS

LSEEK(2) currently is not supported by Digital Equipment Corporation.

NAME

`mkdir` – make a directory file

SYNTAX

```
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)*.

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

DIAGNOSTICS

Mkdir will fail and no directory will be created if:

- [EPERM] The process's effective user ID is not super-user.
- [EPERM] The *path* 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] *Path* 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)*

STATUS

MKDIR(2) currently is not supported by Digital Equipment Corporation.

MKNOD(2)

NAME

`mknod` – make a special file

SYNTAX

```
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.

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

DIAGNOSTICS

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] *Path* 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)`

STATUS

MKNOD(2) currently is not supported by Digital Equipment Corporation.

NAME

mount, umount – mount or remove file system

SYNTAX

mount(special, name, rwflag)

char *special, *name;

int rwflag;

umount(special)

char *special;

DESCRIPTION

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

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

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

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

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.

DIAGNOSTICS

Mount will fail when one of the following occurs:

[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).

[EPERM] The pathname contains a character with the high-order bit set.

[ENOTDIR] A component of the path prefix in *name* is not a directory.

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

[EBUSY] *Name* is not a directory, or another process currently holds a reference to it.

MOUNT(2)

- [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)

STATUS

MOUNT(2) currently is not supported by Digital Equipment Corporation.

NAME

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

SYNTAX

```
#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.

DIAGNOSTICS

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

- [E`PERM`] The pathname contains a character with the high-order bit set.
- [E`NOTDIR`] A component of the path prefix is not a directory.

OPEN(2)

- [ENOENT] O_CREAT is not set and the named file does not exist.
- [EACCES] A component of the path prefix denies search permission.
- [EACCES] The required permissions (for reading and/or writing) are denied for the named flag.
- [EISDIR] The named file is a directory, and the arguments specify it is to be opened for writing.
- [EROFS] The named file resides on a read-only file system, and the file is to be modified.
- [EMFILE] {OPEN_MAX} file descriptors are currently open.
- [ENXIO] The named file is a character special or block special file, and the device associated with this special file does not exist.
- [ETXTBSY] The file is a pure procedure (shared text) file that is being executed and the *open* call requests write access.
- [EFAULT] *Path* points outside the process's allocated address space.
- [ELOOP] Too many symbolic links were encountered in translating the pathname.
- [EEXIST] O_EXCL was specified and the file exists.
- [ENXIO] The O_NDELAY flag is given, and the file is a communications device on which there is no carrier present.
- [EOPNOTSUPP] An attempt was made to open a socket (not currently implemented).

SEE ALSO

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

STATUS

OPEN(2) currently is not supported by Digital Equipment Corporation.

NAME

pipe – create an interprocess communication channel

SYNTAX

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

DESCRIPTION

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

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

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

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

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

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

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

DIAGNOSTICS

The *pipe* call will fail if:

- [EMFILE] Too many descriptors are active.
- [EFAULT] The *fildes* buffer is in an invalid area of the process's address space.

SEE ALSO

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

RESTRICTIONS

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

STATUS

PIPE(2) currently is not supported by Digital Equipment Corporation.

PROFIL(2)

NAME

profil – execution time profile

SYNTAX

```
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.

A 0, indicating success, is always returned.

SEE ALSO

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

STATUS

PROFIL(2) currently is not supported by Digital Equipment Corporation.

NAME

`ptrace` – process trace

SYNTAX

```
#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

PTRACE(2)

(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.

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.

DIAGNOSTICS

- | | |
|----------|---|
| [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)

STATUS

PTRACE(2) currently is not supported by Digital Equipment Corporation.

NAME

quota – manipulate disk quotas

SYNTAX

```
#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

QUOTA(2)

checked. Otherwise, *arg* indicates a specific major-minor device to be checked. This call is restricted to the super-user.

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.

DIAGNOSTICS

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 *arg* parameter is being interpreted as a major-minor device and it indicates an unmounted file system.
- [EFAULT] An invalid *addr* 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)

STATUS

QUOTA(2) currently is not supported by Digital Equipment Corporation.

NAME

read, readv – read input

SYNTAX

```
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.

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.

READ(2)

DIAGNOSTICS

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

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

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

- [EINVAL] *Iovcnt* was less than or equal to 0, or greater than 16.
- [EINVAL] One of the *iov len* values in the *iov* array was negative.
- [EINVAL] The sum of the *iov len* values in the *iov* array overflowed a 32-bit integer.

SEE ALSO

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

STATUS

READ(2) currently is not supported by Digital Equipment Corporation.

NAME

readlink – read value of a symbolic link

SYNTAX

```
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.

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*.

DIAGNOSTICS

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

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

SEE ALSO

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

STATUS

READLINK(2) currently is not supported by Digital Equipment Corporation.

REBOOT(2)

NAME

reboot – reboot system or halt processor

SYNTAX

```
#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.

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

DIAGNOSTICS

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

SEE ALSO

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

STATUS

`REBOOT(2)` currently is not supported by Digital Equipment Corporation.

NAME

`recv`, `recvfrom`, `recvmsg` – receive a message from a socket

SYNTAX

```
#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>`:

RECV(2)

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

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

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

DIAGNOSTICS

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)*

STATUS

RECV(2) currently is not supported by Digital Equipment Corporation.

NAME

rename – change the name of a file

SYNTAX

```
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.

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

DIAGNOSTICS

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)

RESTRICTIONS

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

RENAME(2)

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.

STATUS

RENAME(2) currently is not supported by Digital Equipment Corporation.

NAME

`rmdir` – remove a directory file

SYNTAX

```
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 “..”.

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

DIAGNOSTICS

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)`

STATUS

RMDIR(2) currently is not supported by Digital Equipment Corporation.

SELECT(2)

NAME

select – synchronous i/o multiplexing

SYNTAX

```
#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.

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.

DIAGNOSTICS

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)*

RESTRICTIONS

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

STATUS

SELECT(2) currently is not supported by Digital Equipment Corporation.

NAME

send, *sendto*, *sendmsg* – send a message from a socket

SYNTAX

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

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

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

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

DESCRIPTION

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

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

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

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

The *flags* parameter may be set to SOF_OOB to send “out-of-band” data on sockets which support this notion (e.g. SOCK_STREAM).

See *recv(2)* for a description of the *msghdr* structure.

The call returns the number of characters sent, or -1 if an error occurred.

DIAGNOSTICS

[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.

SEND(2)

[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)

STATUS

SEND(2) currently is not supported by Digital Equipment Corporation.

NAME

setgroups – set group access list

SYNTAX

```
#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.

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

DIAGNOSTICS

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)

STATUS

SETGROUPS(2) currently is not supported by Digital Equipment Corporation.

SETPGRP(2)

NAME

setpgrp – set process group

SYNTAX

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.

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

DIAGNOSTICS

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)

STATUS

SETPGRP(2) currently is not supported by Digital Equipment Corporation.

NAME

setquota – enable/disable quotas on a file system

SYNTAX

```
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.

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.

DIAGNOSTICS

Setquota will fail when one of the following occurs:

- [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).
- [EPERM] The pathname contains a character with the high-order bit set.
- [ENOTDIR] A component of the path prefix in *file* is not a directory.
- [EROFS] *File* resides on a read-only file system.
- [EACCES] *File* resides on a file system different from *special*.
- [EACCES] *File* is not a plain file.

SEE ALSO

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

STATUS

SETQUOTA(2) currently is not supported by Digital Equipment Corporation.

SETREGID(2)

NAME

setregid – set real and effective group ID

SYNTAX

```
setregid(rgid, egid)
int rgid, egid;
```

DESCRIPTION

The real and effective group ID's of the current process are set to the arguments. Only the super-user may change the real group ID of a process. Unprivileged users may change the effective group ID to the real group ID, but to no other.

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

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

DIAGNOSTICS

[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)

STATUS

SETREGID(2) currently is not supported by Digital Equipment Corporation.

NAME

setreuid – set real and effective user ID's

SYNTAX

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

DESCRIPTION

The real and effective user ID's of the current process are set according to the arguments. If *ruid* or *euid* is -1 , the current uid is filled in by the system. Only the super-user may modify the real uid of a process. Users other than the super-user may change the effective uid of a process only to the real uid.

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

DIAGNOSTICS

[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)

STATUS

SETREUID(2) currently is not supported by Digital Equipment Corporation.

SHUTDOWN(2)

NAME

shutdown – shut down part of a full-duplex connection

SYNTAX

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.

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

DIAGNOSTICS

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)

STATUS

SHUTDOWN(2) currently is not supported by Digital Equipment Corporation.

NAME

sigblock – block signals

SYNTAX

```
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.

The previous set of masked signals is returned.

SEE ALSO

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

STATUS

SIGBLOCK(2) currently is not supported by Digital Equipment Corporation.

SIGPAUSE(2)

NAME

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

SYNTAX

```
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)`

STATUS

SIGPAUSE(2) currently is not supported by Digital Equipment Corporation.

NAME

sigsetmask – set current signal mask

SYNTAX

```
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.

The previous set of masked signals is returned.

SEE ALSO

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

STATUS

SIGSETPASK(2) currently is not supported by Digital Equipment Corporation.

SIGSTACK(2)

NAME

sigstack – set and/or get signal stack context

SYNTAX

```
#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.

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

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

DIAGNOSTICS

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)

STATUS

SIGSTACK(2) currently is not supported by Digital Equipment Corporation.

NAME

sigvec – software signal facilities

SYNTAX

```
#include <signal.h>

struct sigvec {
    int      (*sv_handler)();
    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

SIGVEC(2)

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.

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

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

DIAGNOSTICS

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

- [EFAULT] Either *vec* or *ovec* points to memory which is not a valid part of the process address space.
- [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), 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

SIGVEC(2)

Subscript-range	SIGFPE	FPE_SUBRNG_TRAP
Floating overflow fault	SIGFPE	FPE_FLTOVE_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_PRIVIN_FAULT
Customer-reserved instr.	SIGEMT	
Reserved operand	SIGILL	ILL_RESOP_FAULT
Reserved addressing	SIGILL	ILL_RESAD_FAULT
Trace pending	SIGTRAP	
Bpt instruction	SIGTRAP	
Compatibility-mode	SIGILL	hardware supplied code
Chme	SIGSEGV	
Chms	SIGSEGV	
Chmu	SIGSEGV	

STATUS

SIGVEC(2) currently is not supported by Digital Equipment Corporation.

NAME

socket – create an endpoint for communication

SYNTAX

```
#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

SOCKET(2)

may also be transmitted as described in *send(2)* and received as described in *recv(2)*.

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

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

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

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

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

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

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

DIAGNOSTICS

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”.

STATUS

SOCKET(2) currently is not supported by Digital Equipment Corporation.

SOCKETPAIR(2)

NAME

socketpair – create a pair of connected sockets

SYNTAX

```
#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.

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

DIAGNOSTICS

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)

STATUS

SOCKETPAIR(2) currently is not supported by Digital Equipment Corporation.

NAME

stat, *lstat*, *fstat* – get file status

SYNTAX

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

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

```
stat(path, buf)
```

```
char *path;
```

```
struct stat *buf;
```

```
lstat(path, buf)
```

```
char *path;
```

```
struct stat *buf;
```

```
fstat(fd, buf)
```

```
int fd;
```

```
struct stat *buf;
```

DESCRIPTION

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

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

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

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

```
struct stat {
    dev_t      st_dev;      /* device inode resides on */
    ino_t      st_ino;     /* this inode's number */
    u_short    st_mode;    /* protection */
    short      st_nlink;   /* number or hard links to the file */
    short      st_uid;     /* user-id of owner */
    short      st_gid;     /* group-id of owner */
    dev_t      st_rdev;    /* the device type, for inode that is device */
    off_t      st_size;    /* total size of file */
    time_t     st_atime;   /* file last access time */
    int        st_spare1;
    time_t     st_mtime;   /* file last modify time */
    int        st_spare2;
    time_t     st_ctime;   /* file last status change time */
    int        st_spare3;
    long       st_blksize; /* optimal blocksize for file system i/o ops */
};
```

```

        long      st_blocks;    /* actual number of blocks allocated */
        long      st_spare4[2];
    };

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

st_mtime    Time when data was last modified. It is not set by changes of owner, group, link count, or mode. Changed by the following system calls: mknod(2), utimes(2), write(2).

st_ctime    Time when file status was last changed. It is set both both by writing and changing the i-node. Changed by the following system calls: chmod(2), chown(2), link(2), mknod(2), unlink(2), utimes(2), write(2).

```

The status information word *st mode* has bits:

```

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

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

DIAGNOSTICS

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

- [ENOTDIR] A component of the path prefix is not a directory.
- [EPERM] The pathname contains a character with the high-order bit set.
- [ENOENT] The pathname was too long.
- [ENOENT] The named file does not exist.
- [EACCES] Search permission is denied for a component of the path prefix.

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

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

[EBADF] *Fildes* is not a valid open file descriptor.

[EFAULT] *Buf* points to an invalid address.

[ELOOP] Too many symbolic links were encountered in translating the pathname.

SEE ALSO

`chmod(2)`, `chown(2)`, `utimes(2)`

RESTRICTIONS

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

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)`).

STATUS

STAT(2) currently is not supported by Digital Equipment Corporation.

SWAPON(2)

NAME

swapon – add a swap device for interleaved paging/swapping

SYNTAX

```
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)

RESTRICTIONS

There is no way to stop swapping on a disk so that the pack may be dismounted.

STATUS

SWAPON(2) currently is not supported by Digital Equipment Corporation.

NAME

symlink – make symbolic link to a file

SYNTAX

```
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.

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.

DIAGNOSTICS

The symbolic link is made unless one or more of the following are true:

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

SEE ALSO

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

STATUS

SYMLINK(2) currently is not supported by Digital Equipment Corporation.

SYNC(2)

NAME

`sync` – update super-block

SYNTAX

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

SEE ALSO

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

STATUS

SYNC(2) currently is not supported by Digital Equipment Corporation.

NAME

syscall – indirect system call

SYNTAX

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)*).

RESTRICTIONS

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

STATUS

SYSCALL(2) currently is not supported by Digital Equipment Corporation.

TRUNCATE(2)

NAME

`truncate` – truncate a file to a specified length

SYNTAX

`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.

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.

DIAGNOSTICS

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)`

RESTRICTIONS

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.

STATUS

TRUNCATE(2) currently is not supported by Digital Equipment Corporation.

NAME

umask – set file creation mode mask

SYNTAX

```
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.

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

SEE ALSO

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

STATUS

UMASK(2) currently is not supported by Digital Equipment Corporation.

UNLINK(2)

NAME

unlink – remove directory entry

SYNTAX

```
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.

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

DIAGNOSTICS

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] *Path* 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)

STATUS

UNLINK(2) currently is not supported by Digital Equipment Corporation.

NAME

`utimes` – set file times

SYNTAX

```
#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.

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

DIAGNOSTICS

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 *times* is NULL and write access is denied.
- [EROFS] The file system containing the file is mounted read-only.
- [EFAULT] *Tvp* points outside the process’s allocated address space.
- [ELOOP] Too many symbolic links were encountered in translating the pathname.

SEE ALSO

`stat(2)`

STATUS

UTIMES(2) currently is not supported by Digital Equipment Corporation.

VFORK(2)

NAME

vfork – spawn new process in a virtual memory efficient way

SYNTAX

```
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*.

RESTRICTIONS

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.

STATUS

VFORK(2) currently is not supported by Digital Equipment Corporation.

NAME

vhangup – virtually “hangup” the current control terminal

SYNTAX

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.

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

SEE ALSO

init (8)

STATUS

VHANGUP(2) currently is not supported by Digital Equipment Corporation.

WAIT(2)

NAME

`wait`, `wait3` – wait for process to terminate

SYNTAX

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

pid = wait(status)
int pid;
union wait *status;

pid = wait(0)
int pid;

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

pid = wait3(status, options, rusage)
int pid;
union wait *status;
int options;
struct rusage *rusage;
```

DESCRIPTION

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

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

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

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

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.

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.

DIAGNOSTICS

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)

STATUS

WAIT(2) currently is not supported by Digital Equipment Corporation.

WRITE(2)

NAME

write, writev – write on a file

SYNTAX

```
write(d, buf, nbytes)  
int d;  
char *buf;  
int nbytes;  
  
#include <sys/types.h>  
#include <sys/uio.h>  
  
writev(d, iov, iovectlen)  
int d;  
struct iovec *iov;  
int iovectlen;
```

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.

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

DIAGNOSTICS

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

- | | |
|----------|--|
| [EBADF] | <i>D</i> is not a valid descriptor open for writing. |
| [EPIPE] | An attempt is made to write to a pipe that is not open for reading by any process. |
| [EPIPE] | An attempt is made to write to a socket of type SOCK_STREAM which is not connected to a peer socket. |
| [EFBIG] | An attempt was made to write a file that exceeds the process’s file size limit or the maximum file size. |
| [EFAULT] | Part of <i>iov</i> or data to be written to the file points outside the process’s allocated address space. |

SEE ALSO

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

STATUS

WRITE(2) currently is not supported by Digital Equipment Corporation.

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

INTRO(3)

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	error.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
dbm.3	dbm.3x	data base subroutines
delete	dbm.3x	data base subroutines
dfrac	fmin.3f	return extreme values
dfimax	fmin.3f	return extreme values
dfimax	range.3f	return extreme values
dfimin	fmin.3f	return extreme values
dfimin	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
endfsent	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
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

INTRO(3)

<code>fdate</code>	<code>fdate.3f</code>	return date and time in an ASCII string
<code>feof</code>	<code>ferror.3s</code>	stream status inquiries
<code>ferror</code>	<code>ferror.3s</code>	stream status inquiries
<code>fetch</code>	<code>dbm.3x</code>	data base subroutines
<code>fflush</code>	<code>fclose.3s</code>	close or flush a stream
<code>frac</code>	<code>fmin.3f</code>	return extreme values
<code>ffs</code>	<code>bstring.3</code>	bit and byte string operations
<code>fgetc</code>	<code>getc.3f</code>	get a character from a logical unit
<code>fgetc</code>	<code>getc.3s</code>	get character or word from stream
<code>fgets</code>	<code>gets.3s</code>	get a string from a stream
<code>fileno</code>	<code>ferror.3s</code>	stream status inquiries
<code>firstkey</code>	<code>dbm.3x</code>	data base subroutines
<code>fimax</code>	<code>fmin.3f</code>	return extreme values
<code>fimax</code>	<code>range.3f</code>	return extreme values
<code>fimin</code>	<code>fmin.3f</code>	return extreme values
<code>fimin</code>	<code>range.3f</code>	return extreme values
<code>floor</code>	<code>floor.3m</code>	absolute value, floor, ceiling functions
<code>flush</code>	<code>flush.3f</code>	flush output to a logical unit
<code>fork</code>	<code>fork.3f</code>	create a copy of this process
<code>fpecent</code>	<code>trpffe.3f</code>	trap and repair floating point faults
<code>fprintf</code>	<code>printf.3s</code>	formatted output conversion
<code>fputc</code>	<code>putc.3f</code>	write a character to a fortran logical unit
<code>fputc</code>	<code>putc.3s</code>	put character or word on a stream
<code>fputs</code>	<code>puts.3s</code>	put a string on a stream
<code>fread</code>	<code>fread.3s</code>	buffered binary input/output
<code>free</code>	<code>malloc.3</code>	memory allocator
<code>frexp</code>	<code>frexp.3</code>	split into mantissa and exponent
<code>fscanf</code>	<code>scanf.3s</code>	formatted input conversion
<code>fseek</code>	<code>fseek.3f</code>	reposition a file on a logical unit
<code>fseek</code>	<code>fseek.3s</code>	reposition a stream
<code>fstat</code>	<code>stat.3f</code>	get file status
<code>ftell</code>	<code>fseek.3f</code>	reposition a file on a logical unit
<code>ftell</code>	<code>fseek.3s</code>	reposition a stream
<code>ftime</code>	<code>time.3c</code>	get date and time
<code>fwrite</code>	<code>fread.3s</code>	buffered binary input/output
<code>gamma</code>	<code>gamma.3m</code>	log gamma function
<code>gcvt</code>	<code>ecvt.3</code>	output conversion
<code>gerror</code>	<code>perror.3f</code>	get system error messages
<code>getarg</code>	<code>getarg.3f</code>	return command line arguments
<code>getc</code>	<code>getc.3f</code>	get a character from a logical unit
<code>getc</code>	<code>getc.3s</code>	get character or word from stream
<code>getchar</code>	<code>getc.3s</code>	get character or word from stream
<code>getcwd</code>	<code>getcwd.3f</code>	get pathname of current working directory
<code>getdiskbyname</code>	<code>getdisk.3x</code>	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
getprotobynumber	getprotoent.3n	get protocol entry
getprotoent	getprotoent.3n	get protocol entry
getpw	getpw.3	get name from uid
getpwent	getpwent.3	get password file entry
getpwnam	getpwent.3	get password file entry
getpwuid	getpwent.3	get password file entry
gets	gets.3s	get a string from a stream
getservbyname	getservent.3n	get service entry
getservbyport	getservent.3n	get service entry
getservent	getservent.3n	get service entry
getuid	getuid.3f	get user or group ID of the caller
getw	getc.3s	get character or word from stream
getwd	getwd.3	get current working directory pathname
gmtime	ctime.3	convert date and time to ASCII
gmtime	time.3f	return system time
gtty	stty.3c	set and get terminal state (defunct)
hostnm	hostnm.3f	get name of current host
htonl	byteorder.3n	convert values between host and network byte order
htons	byteorder.3n	convert values between host and network byte order
hypot	hypot.3m	Euclidean distance
iargc	getarg.3f	return command line arguments
idate	idate.3f	return date or time in numerical form
ierrno	perror.3f	get system error messages

INTRO(3)

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
lnblnk	index.3f	tell about character objects
loc	loc.3f	return the address of an object
localtime	ctime.3	convert date and time to ASCII
log	exp.3m	exponential, logarithm, power, square root
log10	exp.3m	exponential, logarithm, power, square root
long	long.3f	integer object conversion

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

INTRO(3)

rewind	fseek.3s	reposition a stream
rewinddir	directory.3	directory operations
rexec	rexec.3x	return stream to a remote command
rindex	index.3f	tell about character objects
rindex	string.3	string operations
rresvport	rcmd.3x	routines for returning a stream to a remote command
ruserok	rcmd.3x	routines for returning a stream to a remote command
scandir	scandir.3	scan a directory
scanf	scanf.3s	formatted input conversion
seekdir	directory.3	directory operations
setbuf	setbuf.3s	assign buffering to a stream
setbuffer	setbuf.3s	assign buffering to a stream
setegid	setuid.3	set user and group ID
seteuid	setuid.3	set user and group ID
setfsent	getfsent.3x	get file system descriptor file entry
setgid	setuid.3	set user and group ID
setgrent	getgrent.3	get group file entry
sethostent	gethostent.3n	get network host entry
setjmp	setjmp.3	non-local goto
setkey	crypt.3	DES encryption
setlinebuf	setbuf.3s	assign buffering to a stream
setnetent	getnetent.3n	get network entry
setprotoent	getprotoent.3n	get protocol entry
setpwent	getpwent.3	get password file entry
setrgid	setuid.3	set user and group ID
setruid	setuid.3	set user and group ID
setservent	getservent.3n	get service entry
setstate	random.3	better random number generator
setuid	setuid.3	set user and group ID
short	long.3f	integer object conversion
signal	signal.3	simplified software signal facilities
signal	signal.3f	change the action for a signal
sin	sin.3m	trigonometric functions
sinh	sinh.3m	hyperbolic functions
sleep	sleep.3	suspend execution for interval
sleep	sleep.3f	suspend execution for an interval
space	plot.3x	graphics interface
sprintf	printf.3s	formatted output conversion
sqrt	exp.3m	exponential, logarithm, power, square root
srand	rand.3c	random number generator
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
ttyname	ttyname.3	find name of a terminal
ttyslot	ttyname.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

INTRO(3)

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

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

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	fmin.3f	return extreme values
dfmax	fmin.3f	return extreme values
dfmin	fmin.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
ffrac	fmin.3f	return extreme values
fgetc	getc.3f	get a character from a logical unit
flmax	fmin.3f	return extreme values
flmin	fmin.3f	return extreme values
flush	flush.3f	flush output to a logical unit
fork	fork.3f	create a copy of this process
fpectx	trpfpe.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
lnblk	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

INTRO(3F)

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
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
twrite	topen.3f	f77 tape I/O
unlink	unlink.3f	remove a directory entry
wait	wait.3f	wait for a process to terminate

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

INTRO(3N)

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

stdio — standard buffered input/output package

SYNTAX

```
#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*, *ferror*, *fileno*.

SEE ALSO

open(2), *close*(2), *read*(2), *write*(2), *fread*(3S), *fseek*(3S), *f**(3S)

DIAGNOSTICS

The value **EOF** is returned uniformly to indicate that a **FILE** pointer has not been initialized with *fopen*, input (output) has been attempted on an output (input) stream, or a **FILE** pointer designates corrupt or otherwise unintelligible **FILE** data.

For purposes of efficiency, this implementation of the standard library has been changed to line buffer output to a terminal by default and attempts to do this transparently by flushing the output whenever a *read*(2) from the standard input is necessary. This is almost always transparent, but may cause confusion or malfunctioning of programs which use standard i/o routines but use *read*(2) themselves to read from the standard input.

In cases where a large amount of computation is done after printing part of a line on an output terminal, it is necessary to *fflush*(3S) the standard output before going off and computing so that the output will appear.

LIST OF FUNCTIONS

<i>Name</i>	<i>Appears on Page</i>	<i>Description</i>
clearerr	ferror.3s	stream status inquiries
fclose	fclose.3s	close or flush a stream
fdopen	fopen.3s	open a stream
feof	ferror.3s	stream status inquiries
ferror	ferror.3s	stream status inquiries
fflush	fclose.3s	close or flush a stream
fgetc	getc.3s	get character or word from stream
fgets	gets.3s	get a string from a stream
fileno	ferror.3s	stream status inquiries
fopen	fopen.3s	open a stream
fprintf	printf.3s	formatted output conversion
fputc	putc.3s	put character or word on a stream
fputs	puts.3s	put a string on a stream
fread	fread.3s	buffered binary input/output
freopen	fopen.3s	open a stream
fscanf	scanf.3s	formatted input conversion
fseek	fseek.3s	reposition a stream
ftell	fseek.3s	reposition a stream
fwrite	fread.3s	buffered binary input/output
getc	getc.3s	get character or 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

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*, *rresuport*, *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
endfsent	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

INTRO(3X)

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

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.

RESTRICTIONS

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

STATUS

ABORT(3) currently is not supported by Digital Equipment Corporation.

ABORT(3F)

NAME

abort – terminate abruptly with memory image

SYNTAX

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)

STATUS

ABORT(3F) currently is not supported by Digital Equipment Corporation.

NAME

abs – integer absolute value

SYNTAX

abs(i)

int i;

DESCRIPTION

Abs returns the absolute value of its integer operand.

SEE ALSO

floor(3M) for *fabs*

RESTRICTIONS

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.

STATUS

ABS(3) currently is not supported by Digital Equipment Corporation.

ACCESS(3F)

NAME

access – determine accessibility of a file

SYNTAX

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)

RESTRICTIONS

Pathnames can be no longer than MAXPATHLEN as defined in *<sys/param.h>*.

STATUS

ACCESS(3F) currently is not supported by Digital Equipment Corporation.

NAME

alarm – schedule signal after specified time

SYNTAX

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)

STATUS

ALARM(3C) currently is not supported by Digital Equipment Corporation.

ALARM(3F)

NAME

alarm – execute a subroutine after a specified time

SYNTAX

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)

RESTRICTIONS

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.

STATUS

ALARM(3F) currently is not supported by Digital Equipment Corporation.

NAME

assert – program verification

SYNTAX

#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.

STATUS

ASSERT(3X) currently is not supported by Digital Equipment Corporation.

ATOF(3)

NAME

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

SYNTAX

double *atof*(*nptr*)

char **nptr*;

atoi(*nptr*)

char **nptr*;

long *atol*(*nptr*)

char **nptr*;

DESCRIPTION

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

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

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

DIAGNOSTICS

If an overflow occurs, *atof* returns the largest value possible.

If an underflow occurs, *atof* returns a 0 value.

SEE ALSO

scanf(3S)

STATUS

ATOF(3) is supported by Digital Equipment Corporation.

NAME

bessel functions – of two kinds for integer orders

SYNTAX

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

BESSEL (3F)

SEE ALSO

j0(3M), perror(3F)

STATUS

BESSEL (3F) currently is not supported by Digital Equipment Corporation.

NAME

bit – and, or, xor, not, rshift, lshift bitwise functions

SYNTAX

(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.

STATUS

BIT(3F) currently is not supported by Digital Equipment Corporation.

BSTRING(3)

NAME

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

SYNTAX

***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.

RESTRICTIONS

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

STATUS

BSTRING(3) currently is not supported by Digital Equipment Corporation.

NAME

htonl, htons, ntohl, ntohs – convert values between host and network byte order

SYNTAX

```
#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)*

RESTRICTIONS

The VAX handles bytes in the reverse from most everyone else.

STATUS

BYTEORDER(3N) currently is not supported by Digital Equipment Corporation.

CHDIR(3F)

NAME

chdir — change default directory

SYNTAX

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)

RESTRICTIONS

Pathnames can be no longer than MAXPATHLEN as defined in *<sys/param.h>*.

Use of this function may cause **inquire** by unit to fail.

STATUS

CHDIR(3F) currently is not supported by Digital Equipment Corporation.

NAME

chmod – change mode of a file

SYNTAX

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)

RESTRICTIONS

Pathnames can be no longer than MAXPATHLEN as defined in *<sys/param.h>*.

STATUS

CHMOD(3F) currently is not supported by Digital Equipment Corporation.

CRYPT(3)

NAME

crypt, *setkey*, *encrypt* – DES encryption

SYNTAX

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

```
char *key, *salt;
```

```
encrypt(block)
```

```
char *block;
```

DESCRIPTION

Crypt is the password encryption routine.

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 hashing 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 *encrypt* entry is 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 hashing algorithm using the key set by *crypt*.

SEE ALSO

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

RESTRICTIONS

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

STATUS

CRYPT(3) is supported by Digital Equipment Corporation.

NAME

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

SYNTAX

```
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 \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.

CTIME(3)

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)

RESTRICTIONS

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

STATUS

CTIME(3) currently is not supported by Digital Equipment Corporation.

NAME

`isalpha`, `isupper`, `islower`, `isdigit`, `isalnum`, `isspace`, `ispunct`, `isprint`, `isctrl`, `isascii` — character classification macros

SYNTAX

```
#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)).

isalpha *c* is a letter

isupper *c* is an upper case letter

islower *c* is a lower case letter

isdigit *c* is a digit

isalnum *c* is an alphanumeric character

isspace *c* is a space, tab, carriage return, newline, or formfeed

ispunct *c* is a punctuation character (neither control nor alphanumeric)

isprint *c* is a printing character, code 040(8) (space) through 0176 (tilde)

isctrl *c* is a delete character (0177) or ordinary control character (less than 040).

isascii *c* is an ASCII character, code less than 0200

SEE ALSO

`ascii`(7)

STATUS

CTYPE(3) currently is not supported by Digital Equipment Corporation.

CURSES(3X)

NAME

curses – screen functions with “optimal” cursor motion

SYNTAX

`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)`

FUNCTIONS

<code>addch(ch)</code>	add a character to <i>stdscr</i>
<code>addstr(str)</code>	add a string to <i>stdscr</i>
<code>box(win,vert,hor)</code>	draw a box around a window
<code>crmode()</code>	set cbreak mode
<code>clear()</code>	clear <i>stdscr</i>
<code>clearok(scr,boolf)</code>	set clear flag for <i>scr</i>
<code>clrtoebot()</code>	clear to bottom on <i>stdscr</i>
<code>clrtoeol()</code>	clear to end of line on <i>stdscr</i>
<code>delch()</code>	delete a character
<code>deleteln()</code>	delete a line
<code>delwin(win)</code>	delete <i>win</i>
<code>echo()</code>	set echo mode
<code>endwin()</code>	end window modes
<code>erase()</code>	erase <i>stdscr</i>
<code>getch()</code>	get a char through <i>stdscr</i>
<code>getcap(name)</code>	get terminal capability <i>name</i>
<code>getstr(str)</code>	get a string through <i>stdscr</i>
<code>gettmode()</code>	get tty modes
<code>getyx(win,y,x)</code>	get (y,x) co-ordinates
<code>inch()</code>	get char at current (y,x) co-ordinates
<code>initscr()</code>	initialize screens
<code>insch(c)</code>	insert a char
<code>insertln()</code>	insert a line
<code>leaveok(win,boolf)</code>	set leave flag for <i>win</i>
<code>longname(termbuf,name)</code>	get long name from <i>termbuf</i>
<code>move(y,x)</code>	move to (y,x) on <i>stdscr</i>
<code>mvcur(lasty,lastx,newy,newx)</code>	actually move cursor
<code>newwin(lines,cols,begin y,begin x)</code>	create a new window

<code>nl()</code>	set newline mapping
<code>nocrmode()</code>	unset cbreak mode
<code>noecho()</code>	unset echo mode
<code>nonl()</code>	unset newline mapping
<code>noraw()</code>	unset raw mode
<code>overlay(win1,win2)</code>	overlay win1 on win2
<code>overwrite(win1,win2)</code>	overwrite win1 on top of win2
<code>printw(fmt,arg1,arg2,...)</code>	printf on <i>stdscr</i>
<code>raw()</code>	set raw mode
<code>refresh()</code>	make current screen look like <i>stdscr</i>
<code>resetty()</code>	reset tty flags to stored value
<code>savetty()</code>	stored current tty flags
<code>scanw(fmt,arg1,arg2,...)</code>	scanf through <i>stdscr</i>
<code>scroll(win)</code>	scroll <i>win</i> one line
<code>scrollok(win,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>wclrto bot(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>	printf on <i>win</i>
<code>wrefresh(win)</code>	make screen look like <i>win</i>
<code>wscanw(win,fmt,arg1,arg2,...)</code>	scanf 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>

STATUS

CURSES(3X) currently is not supported by Digital Equipment Corporation.

NAME

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

SYNTAX

```
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*.

RESTRICTIONS

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.

STATUS

DBM(3X) currently is not supported by Digital Equipment Corporation.

DIRECTORY(3)

NAME

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

SYNTAX

```
#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)

STATUS

DIRECTORY(3) is supported by Digital Equipment Corporation.

ECVT(3)

NAME

ecvt, *fcvt*, *gevt* – output conversion

SYNTAX

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 **gevt*(value, ndigit, buf)

double value;

char *buf;

DESCRIPTION

*Ecv*t 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.

*Fcv*t 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*.

*Gcv*t 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)

RESTRICTIONS

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

STATUS

ECVT(3) currently is not supported by Digital Equipment Corporation.

NAME

end, *etext*, *edata* – last locations in program

SYNTAX

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)*

STATUS

END(3) currently is not supported by Digital Equipment Corporation.

ETIME(3F)

NAME

etime, *dtime* – return elapsed execution time

SYNTAX

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)

STATUS

ETIME(3F) currently is not supported by Digital Equipment Corporation.

NAME

`execl`, `execv`, `execle`, `execlp`, `execvp`, `exec`, `exece`, `execx`, `environ` — execute a file

SYNTAX

```

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[];

execx(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 *execx* 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.

EXECL(3)

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.

RESTRICTIONS

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.

STATUS

EXECL(3) currently is not supported by Digital Equipment Corporation.

NAME

`exit` – terminate a process after flushing any pending output

SYNTAX

```
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)`

STATUS

EXIT(3) currently is not supported by Digital Equipment Corporation.

EXIT(3F)

NAME

exit – terminate process with status

SYNTAX

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)

STATUS

EXIT(3F) currently is not supported by Digital Equipment Corporation.

NAME

exp, *log*, *log10*, *pow*, *sqrt* – exponential, logarithm, power, square root

SYNTAX

```
#include <math.h>
```

```
double exp(x)
```

```
double x;
```

```
double log(x)
```

```
double x;
```

```
double log10(x)
```

```
double x;
```

```
double pow(x, y)
```

```
double x, y;
```

```
double sqrt(x)
```

```
double x;
```

DESCRIPTION

Exp returns the exponential function of x .

Log returns the natural logarithm of x ; *log10* returns the base 10 logarithm.

Pow returns x^y .

Sqrt returns the square root of x .

SEE ALSO

hypot(3M), *sinh*(3M), *intro*(3M)

DIAGNOSTICS

Exp and *pow* return a huge value when the correct value would overflow; *errno* is set to ERANGE. *Pow* returns 0 and sets *errno* to EDOM when the second argument is negative and non-integral and when both arguments are 0.

Log returns 0 when x is zero or negative; *errno* is set to EDOM.

Sqrt returns 0 when x is negative; *errno* is set to EDOM.

STATUS

EXP(3M) currently is not supported by Digital Equipment Corporation.

FCLOSE(3S)

NAME

fclose, *fflush* – close or flush a stream

SYNTAX

```
#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.

STATUS

FCLOSE(3S) is supported by Digital Equipment Corporation.

NAME

fdate – return date and time in an ASCII string

SYNTAX

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)

STATUS

FDATE(3F) currently is not supported by Digital Equipment Corporation.

FERROR(3S)

NAME

`ferror`, `feof`, `clearerr`, `fileno` – stream status inquiries

SYNTAX

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

```
feof(stream)
```

```
FILE *stream;
```

```
ferror(stream)
```

```
FILE *stream
```

```
clearerr(stream)
```

```
FILE *stream
```

```
fileno(stream)
```

```
FILE *stream;
```

DESCRIPTION

Feof returns non-zero when end of file is read on the named input *stream*, otherwise zero.

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

Clrerr resets the error indication on the named *stream*.

Fileno returns the integer file descriptor associated with the *stream*, see *open*(2).

These functions are implemented as macros; they cannot be redeclared.

SEE ALSO

fopen(3S), *open*(2)

STATUS

FERROR(3S) is supported by Digital Equipment Corporation.

NAME

fmin, *fmax*, *frac*, *dflmin*, *dflmax*, *dfrac*, *inmax* – return extreme values

SYNTAX

function *fmin*()

function *fmax*()

function *frac*()

double precision function *dflmin*()

double precision function *dflmax*()

double precision function *dfrac*()

function *inmax*()

DESCRIPTION

Functions *fmin* and *fmax* return the minimum and maximum positive floating point values respectively. Functions *dflmin* and *dflmax* return the minimum and maximum positive double precision floating point values. Function *inmax* returns the maximum positive integer value.

The functions *frac* and *dfrac* 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

STATUS

FLMIN(3F) currently is not supported by Digital Equipment Corporation.

FLOOR(3M)

NAME

fabs, *floor*, *ceil* – absolute value, floor, ceiling functions

SYNTAX

```
#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)

STATUS

FLOOR(3M) currently is not supported by Digital Equipment Corporation.

NAME

flush - flush output to a logical unit

SYNTAX

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)

STATUS

FLUSH(3F) currently is not supported by Digital Equipment Corporation.

FOPEN(3S)

NAME

fopen, *freopen*, *fdopen* – open a stream

SYNTAX

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

```
FILE *fopen(filename, type)
```

```
char *filename, *type;
```

```
FILE *freopen(filename, type, stream)
```

```
char *filename, *type;
```

```
FILE *stream;
```

```
FILE *fdopen(fildes, type)
```

```
char *type;
```

DESCRIPTION

Fopen opens the file named by *filename* and associates a stream with it. *Fopen* returns a pointer to be used to identify the stream in subsequent operations.

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

"r" open for reading

"w" create for writing

"a" append: open for writing at end of file, or create for writing

In addition, each *type* may be followed by a '+' to have the file opened for reading and writing. "r+" positions the stream at the beginning of the file, "w+" creates or truncates it, and "a+" positions it at the end. Both reads and writes may be used on read/write streams, with the limitation that an *fseek*, *rewind*, or reading an end-of-file must be used between a read and a write or vice-versa.

Freopen substitutes the named file in place of the open *stream*. It returns the original value of *stream*. The original stream is closed. *Freopen* is typically used to attach the preopened constant names, **stdin**, **stdout**, **stderr**, to specified files.

Fdopen associates a stream with a file descriptor obtained from *open*, *dup*, *creat*, or *pipe*(2). The *type* of the stream must agree with the mode of the open file.

SEE ALSO

open(2), *fclose*(3)

DIAGNOSTICS

Fopen and *freopen* return the pointer **NULL** if *filename* cannot be accessed.

RESTRICTIONS

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.

STATUS

FOPEN(3S) is supported by Digital Equipment Corporation.

FORK(3F)

NAME

fork — create a copy of this process

SYNTAX

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 *fork/exec* can be performed using *system(3F)*.

FILES

/usr/lib/libU77.a

SEE ALSO

fork(2), wait(3F), kill(3F), system(3F), perror(3F)

STATUS

FORK(3F) currently is not supported by Digital Equipment Corporation.

NAME

fread, *fwrite* – buffered binary input/output

SYNTAX

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

```
fread(ptr, sizeof(*ptr), nitems, stream)  
char *ptr; unsigned nitems, sizeof(*ptr)  
FILE *stream;
```

```
fwrite(ptr, sizeof(*ptr), nitems, stream)  
char *ptr; unsigned nitems, sizeof(*ptr)  
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.

STATUS

FREAD(3S) is supported by Digital Equipment Corporation.

FREXP(3)

NAME

frexp, *ldexp*, *modf* – split into mantissa and exponent

SYNTAX

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 $value = x * 2^n$ indirectly through *eptr*.

Ldexp returns the quantity $value * 2^{exp}$.

Modf returns the positive fractional part of *value* and stores the integer part indirectly through *iptr*.

STATUS

FREXP(3) currently is not supported by Digital Equipment Corporation.

NAME

fseek, *ftell* – reposition a file on a logical unit

SYNTAX

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)

STATUS

FSEEK(3F) currently is not supported by Digital Equipment Corporation.

FSEEK(3S)

NAME

`fseek`, `ftell`, `rewind` – reposition a stream

SYNTAX

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

```
fseek(stream, offset, ptrname)
```

```
FILE *stream;
```

```
long offset;
```

```
long ftell(stream)
```

```
FILE *stream;
```

```
rewind(stream)
```

```
FILE *stream;
```

DESCRIPTION

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

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 and is 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.

STATUS

FSEEK(3S) is supported by Digital Equipment Corporation.

NAME

gamma – log gamma function

SYNTAX

```
#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.

STATUS

GAMMA(3M) currently is not supported by Digital Equipment Corporation.

GETARG(3F)

NAME

getarg, *iargc* – return command line arguments

SYNTAX

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 0*th* 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)

STATUS

GETARG(3F) currently is not supported by Digital Equipment Corporation.

NAME

getc, fgetc – get a character from a logical unit

SYNTAX

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 f77 I/O error code. See *perror(3F)*.

FILES

/usr/lib/libU77.a

SEE ALSO

getc(3S), *intro(2)*, *perror(3F)*

STATUS

GETC(3F) currently is not supported by Digital Equipment Corporation.

GETC(3S)

NAME

getc, *getchar*, *fgetc*, *getw* — get character or word from stream

SYNTAX

```
#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*.

RESTRICTIONS

Because it is implemented as a macro, *getc* treats a stream argument with side effects incorrectly. In particular, '*getc(*f++)*;' doesn't work as expected.

STATUS

GETC(3S) is supported by Digital Equipment Corporation.

NAME

getcwd – get pathname of current working directory

SYNTAX

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)

RESTRICTIONS

Pathnames can be no longer than MAXPATHLEN as defined in *<sys/param.h>*.

STATUS

GETCWD(3F) currently is not supported by Digital Equipment Corporation.

GETDISKBYNAME(3X)

NAME

getdiskbyname – get disk description by its name

SYNTAX

```
#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) 2/6/83  */

/*
 * 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)

STATUS

GETDISKBYNAME(3X) currently is not supported by Digital Equipment Corporation.

NAME

`getenv` – value for environment name

SYNTAX

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)*

STATUS

GETENV(3) is supported by Digital Equipment Corporation.

GETENV(3F)

NAME

getenv – get value of environment variables

SYNTAX

subroutine *getenv* (**ename**, **evaluate**)
character*(*) **ename**, **evaluate**

DESCRIPTION

Getenv searches the environment list (see *environ(7)*) for a string of the form *ename=value* and returns *value* in *evaluate* if such a string is present, otherwise fills *evaluate* with blanks.

FILES

/usr/lib/libU77.a

SEE ALSO

environ(7), *execve(2)*

STATUS

GETENV(3F) currently is not supported by Digital Equipment Corporation.

NAME

getfsent, *getfsspec*, *getfsfile*, *getfstype*, *setfsent*, *endfsent* – get file system descriptor file entry

SYNTAX

```
#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)

GETFSENT(3X)

DIAGNOSTICS

Null pointer (0) returned on EOF or error.

RESTRICTIONS

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

STATUS

GETFSENT(3X) currently is not supported by Digital Equipment Corporation.

NAME

getgrent, getgrgid, getgrnam, setgrent, endgrent – get group file entry

SYNTAX

```
#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.

```
/*      grp.h      4.1      83/05/03 */

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)

GETGREN(3)

DIAGNOSTICS

A null pointer (0) is returned on EOF or error.

RESTRICTIONS

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

STATUS

GETGREN(3) currently is not supported by Digital Equipment Corporation.

NAME

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

SYNTAX

```
#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.

GETHOSTENT(3N)

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.

RESTRICTIONS

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.

STATUS

GETHOSTENT(3N) currently is not supported by Digital Equipment Corporation.

NAME

getlog – get user's login name

SYNTAX

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)

STATUS

GETLOG(3F) currently is not supported by Digital Equipment Corporation.

GETLOGIN(3)

NAME

getlogin – get login name

SYNTAX

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.

RESTRICTIONS

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

STATUS

GETLOGIN(3) currently is not supported by Digital Equipment Corporation.

NAME

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

SYNTAX

```
#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.

GETNETENT (3N)

FILES

/etc/networks

SEE ALSO

networks(5)

DIAGNOSTICS

Null pointer (0) returned on EOF or error.

RESTRICTIONS

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.

STATUS

GETNETENT (3N) currently is not supported by Digital Equipment Corporation.

NAME

getpass – read a password

SYNTAX

```
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)

RESTRICTIONS

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

STATUS

GETPASS(3) currently is not supported by Digital Equipment Corporation.

GETPID(3F)

NAME

getpid – get process id

SYNTAX

integer function getpid()

DESCRIPTION

Getpid returns the process ID number of the current process.

FILES

/usr/lib/libU77.a

SEE ALSO

getpid(2)

STATUS

GETPID(3F) currently is not supported by Digital Equipment Corporation.

NAME

getprotoent, getprotobynumber, getprotobyname, setprotoent, endprotoent — get protocol entry

SYNTAX

```
#include <netdb.h>

struct protoent *getprotoent()

struct protoent *getprotobyname(name)
char *name;

struct protoent *getprotobynumber(proto)
int proto;

setprotoent(stayopen)
int stayopen

endprotoent()
```

DESCRIPTION

Getprotoent, *getprotobyname*, and *getprotobynumber* each return a pointer to an object with the following structure containing the broken-out fields of a line in the network protocol data base, */etc/protocols*.

```
struct protoent {
    char    *p_name;          /* official name of protocol */
    char    **p_aliases;     /* alias list */
    long    p_proto;        /* protocol number */
};
```

The members of this structure are:

p_name The official name of the protocol.

p_aliases A zero terminated list of alternate names for the protocol.

p_proto The protocol number.

Getprotoent reads the next line of the file, opening the file if necessary.

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

Endprotoent closes the file.

Getprotobyname and *getprotobynumber* sequentially search from the beginning of the file until a matching protocol name or protocol number is found, or until EOF is encountered.

FILES

/etc/protocols

GETPROTOENT(3N)

SEE ALSO

protocols(5)

DIAGNOSTICS

Null pointer (0) returned on EOF or error.

RESTRICTIONS

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.

STATUS

GETPROTOENT(3N) currently is not supported by Digital Equipment Corporation.

NAME

`getpw` – get name from uid

SYNTAX

```
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.

STATUS

GETPW(3C) currently is not supported by Digital Equipment Corporation.

GETPWENT(3)

NAME

getpwent, *getpwuid*, *getpwnam*, *setpwent*, *endpwent* – get password file entry

SYNTAX

```
#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.

```
/*      pwd.h      4.1      83/05/03 */

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.

RESTRICTIONS

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

STATUS

GETPWENT(3) currently is not supported by Digital Equipment Corporation.

GETS(3S)

NAME

gets, *fgets* – get a string from a stream

SYNTAX

```
#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.

RESTRICTIONS

Gets deletes a newline, while *fgets* keeps it.

STATUS

GETS(3S) is supported by Digital Equipment Corporation.

NAME

getservent, getservbyport, getservbyname, setservent, endservent – get service entry

SYNTAX

```
#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.

GETSERVENT(3N)

FILES

/etc/services

SEE ALSO

getprotoent(3N), services(5)

DIAGNOSTICS

Null pointer (0) returned on EOF or error.

RESTRICTIONS

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

STATUS

GETSERVENT(3N) currently is not supported by Digital Equipment Corporation.

NAME

getuid, getgid – get user or group ID of the caller

SYNTAX

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)

STATUS

GETUID(3F) currently is not supported by Digital Equipment Corporation.

GETWD(3)

NAME

getwd – get current working directory pathname

SYNTAX

```
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.

DIAGNOSTICS

Getwd returns zero and places a message in *pathname* if an error occurs.

RESTRICTIONS

Getwd may fail to return to the current directory if an error occurs.

Pathnames can be no longer than MAXPATHLEN as defined in <sys/param.h>.

STATUS

GETWD(3) is supported by Digital Equipment Corporation.

NAME

hostnm -- get name of current host

SYNTAX

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)

STATUS

HOSTNM(3F) currently is not supported by Digital Equipment Corporation.

HYPOT(3M)

NAME

hypot, cabs – Euclidean distance

SYNTAX

```
#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*

STATUS

HYPOT(3M) currently is not supported by Digital Equipment Corporation.

NAME

idate, *itime* – return date or time in numerical form

SYNTAX

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 ≥ 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)

STATUS

IDATE(3F) currently is not supported by Digital Equipment Corporation.

INDEX(3F)

NAME

index, rindex, lnblnk, len – tell about character objects

SYNTAX

(intrinsic) function index (string, substr)
character*(*) string, substr

integer function rindex (string, substr)
character*(*) string, substr

function lnblnk (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.

Lnblnk 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

STATUS

INDEX(3F) currently is not supported by Digital Equipment Corporation.

NAME

`inet_addr`, `inet_network`, `inet_ntoa`, `inet_makeaddr`, `inet_lnaof`, `inet_netof` – Internet address manipulation routines

SYNTAX

```
#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.

INET(3N)

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.

STATUS

INET(3N) currently is not supported by Digital Equipment Corporation.

NAME

`initgroups` – initialize group access list

SYNTAX

```
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.

RESTRICTIONS

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*.

STATUS

INITGROUPS(3X) currently is not supported by Digital Equipment Corporation.

INSQUE(3)

NAME

insque, *remque* – insert/remove element from a queue

SYNTAX

```
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 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.

STATUS

INSQUE(3) currently is not supported by Digital Equipment Corporation.

NAME

ioinit – change f77 I/O initialization

SYNTAX

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 fortran 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.

IOINIT(3F)

If the argument *verbose* 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”

RESTRICTIONS

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.

STATUS

IOINIT(3F) currently is not supported by Digital Equipment Corporation.

NAME

`j0, j1, jn, y0, y1, yn` – Bessel functions

SYNTAX

```
#include <math.h>
```

```
double j0(x)
```

```
double x;
```

```
double j1(x)
```

```
double x;
```

```
double jn(n, x)
```

```
double x;
```

```
double y0(x)
```

```
double x;
```

```
double y1(x)
```

```
double x;
```

```
double yn(n, x)
```

```
double x;
```

DESCRIPTION

These functions calculate Bessel functions of the first and second kinds for real arguments and integer orders.

DIAGNOSTICS

Negative arguments cause `y0`, `y1`, and `yn` to return a huge negative value and set `errno` to `EDOM`.

STATUS

`JO(3M)` currently is not supported by Digital Equipment Corporation.

KILL (3F)

NAME

kill – send a signal to a process

SYNTAX

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)

STATUS

KILL (3F) currently is not supported by Digital Equipment Corporation.

NAME

lib2648 – subroutines for the HP 2648 graphics terminal

SYNTAX

```
#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 *ttyinit()* 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 *outstr*.

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.

aon() Turn the alphanumeric display on.

areaclear(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 *outstr*.

outstr(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 *fprintf/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()

Set the display to draw lines in normal (solid) mode. (This is reversed by inverse video mode.)

setxor()

Set the display to draw lines in exclusive or mode.

sync() Force all accumulated output to be displayed on the screen. This should be followed by *fflush*(*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 *fflush*(*stdout*) automatically.

togvid()

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()

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()

Set inverse video mode.

vidnorm()

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()

Turn zoom off. This forces the screen to zoom level 1 without affecting the current internal zoom number.

zoomon()

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)

STATUS

LIB2648(3X) currently is not supported by Digital Equipment Corporation.

LINK(3F)

NAME

link – make a link to an existing file

SYNTAX

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)

RESTRICTIONS

Pathnames can be no longer than MAXPATHLEN as defined in *<sys/param.h>*.

STATUS

LINK(3F) currently is not supported by Digital Equipment Corporation.

NAME

loc — return the address of an object

SYNTAX

function loc (arg)

DESCRIPTION

The returned value will be the address of *arg*.

FILES

/usr/lib/libU77.a

STATUS

LOC(3F) currently is not supported by Digital Equipment Corporation.

LONG(3F)

NAME

long, short – integer object conversion

SYNTAX

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

STATUS

LONG(3F) currently is not supported by Digital Equipment Corporation.

NAME

`malloc`, `free`, `realloc`, `calloc`, `alloca` – memory allocator

SYNTAX

```
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.

MALLOC(3)

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.

RESTRICTIONS

When *realloc* returns 0, the block pointed to by *ptr* may be destroyed.

Currently, the allocator is unsuitable for direct use in a large virtual environment where many small blocks are kept, since it keeps all allocated and freed blocks on a circular list. Just before more memory is allocated, all allocated and freed blocks are referenced.

Alloca is machine dependent.

STATUS

MALLOC(3) is supported by Digital Equipment Corporation.

NAME

mktemp – make a unique file name

SYNTAX

```
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)

STATUS

MKTEMP(3) currently is not supported by Digital Equipment Corporation.

MONITOR(3)

NAME

monitor, monstartup, moncontrol – prepare execution profile

SYNTAX

monitor(lowpc, highpc, buffer, bufsize, nfunc)

int (*lowpc)(), (*highpc)();

short buffer[];

monstartup(lowpc, highpc)

int (*lowpc)(), (*highpc)();

moncontrol(mode)

DESCRIPTION

There are two different forms of monitoring available: An executable program created by:

```
cc -p ...
```

automatically includes calls for the *prof*(1) monitor and includes an initial call to its start-up routine *monstartup* with default parameters; *monitor* need not be called explicitly except to gain fine control over profil buffer allocation. An executable program created by:

```
cc -pg ...
```

automatically includes calls for the *gprof*(1) monitor.

Monstartup is a high level interface to *profil*(2). *Lowpc* and *highpc* specify the address range that is to be sampled; the lowest address sampled is that of *lowpc* and the highest is just below *highpc*. *Monstartup* allocates space using *sbrk*(2) and passes it to *monitor* (see below) to record a histogram of periodically sampled values of the program counter, and of counts of calls of certain functions, in the buffer. Only calls of functions compiled with the profiling option `-p` of *cc*(1) are recorded.

To profile the entire program, it is sufficient to use

```
extern etext();
...
monstartup((int) 2, etext);
```

Etext lies just above all the program text, see *end*(3).

To stop execution monitoring and write the results on the file *mon.out*, use

```
monitor(0);
```

then *prof*(1) can be used to examine the results.

Moncontrol is used to selectively control profiling within a program. This works with either *prof*(1) or *gprof*(1) type profiling. When the program starts, profiling begins. To stop the collection of histogram ticks and call counts use *moncontrol*(0); to resume the collection of histogram ticks and call counts use *moncontrol*(1). This allows the cost of particular operations to be measured. Note that an output file will be produced upon program exit regardless of the state of *moncontrol*.

Monitor is a low level interface to *profil(2)*. *Lowpc* and *highpc* are the addresses of two functions; *buffer* is the address of a (user supplied) array of *bufsize* short integers. At most *nfunc* call counts can be kept. For the results to be significant, especially where there are small, heavily used routines, it is suggested that the buffer be no more than a few times smaller than the range of locations sampled. *Monitor* divides the buffer into space to record the histogram of program counter samples over the range *lowpc* to *highpc*, and space to record call counts of functions compiled with the **-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)*

STATUS

MONITOR(3) currently is not supported by Digital Equipment Corporation.

NICE(3C)

NAME

nice - set program priority

SYNTAX

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)

STATUS

NICE(3C) currently is not supported by Digital Equipment Corporation.

NAME

nlist – get entries from name list

SYNTAX

```
#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.

STATUS

NLIST(3) currently is not supported by Digital Equipment Corporation.

PAUSE(3C)

NAME

pause – stop until signal

SYNTAX

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.

DIAGNOSTICS

Pause always returns:

[EINTR] The call was interrupted, i.e., always returns -1.

SEE ALSO

kill(2), select(2), sigpause(2)

STATUS

PAUSE(3C) currently is not supported by Digital Equipment Corporation.

NAME

`perror`, `sys_errlist`, `sys_nerr` – system error messages

SYNTAX

```
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)*

STATUS

PERROR(3) is supported by Digital Equipment Corporation.

PERROR(3F)

NAME

`perror`, `error`, `ierrno` – get system error messages

SYNTAX

subroutine `perror` (`string`)
character*(*) `string`

subroutine `error` (`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*

RESTRICTIONS

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"

STATUS

PERROR(3F) currently is not supported by Digital Equipment Corporation.

PLOT(3X)

NAME

plot: openpl, erase, label, line, circle, arc, move, cont, point, linemod, space, closepl – graphics interface

SYNTAX

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)*

STATUS

PLOT(3X) currently is not supported by Digital Equipment Corporation.

NAME

popen, *pclose* – initiate I/O to/from a process

SYNTAX

```
#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.

RESTRICTIONS

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 *fflush*, see *fclose*(3).

Popen always calls *sh*, never calls *cs**h*.

STATUS

POPEN(3) currently is not supported by Digital Equipment Corporation.

PRINTF(3S)

NAME

printf, fprintf, sprintf – formatted output conversion

SYNTAX

```
#include <stdio.h>

printf(format [, arg ] ... )
char *format;

fprintf(stream, format [, arg ] ... )
FILE *stream;
char *format;

char *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 **Ox(OX)** 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

- dox** 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).

PRINTF(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)

RESTRICTIONS

Very wide fields (>128 characters) fail.

STATUS

PRINTF(3S) is supported by Digital Equipment Corporation.

NAME

psignal, sys_siglist – system signal messages

SYNTAX

```
psignal(sig, s)
unsigned sig;
char *s;
char *sys_siglist[];
```

DESCRIPTION

Psignal produces a short message on the standard error file describing the indicated signal. First the argument string *s* is printed, then a colon, then the name of the signal and a newline. Most usefully, the argument string is the name of the program which incurred the signal. The signal number should be from among those found in *<signal.h>*.

To simplify variant formatting of signal names, the vector of message strings *sys siglist* is provided; the signal number can be used as an index in this table to get the signal name without the newline. The define *NSIG* defined in *<signal.h>* is the number of messages.

SEE ALSO

sigvec(2), perror(3)

STATUS

PSIGNAL(3) is supported by Digital Equipment Corporation.

PUTC(3F)

NAME

`putc`, `fputc` – write a character to a fortran logical unit

SYNTAX

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)`

STATUS

PUTC(3F) currently is not supported by Digital Equipment Corporation.

NAME

`putc`, `putchar`, `fputc`, `putw` – put character or word on a stream

SYNTAX

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

```
int putc(c, stream)
```

```
char c;
```

```
FILE *stream;
```

```
putchar(c)
```

```
fputc(c, stream)
```

```
FILE *stream;
```

```
putw(w, stream)
```

```
FILE *stream;
```

DESCRIPTION

Putc appends the character *c* to the named output *stream*. It returns the character written.

Putchar(c) is defined as *putc(c, stdout)*.

Fputc behaves like *putc*, but is a genuine function rather than a macro.

Putw appends word (that is, **int**) *w* to the output *stream*. It returns the word written.

Putw neither assumes nor causes special alignment in the file.

SEE ALSO

`fopen(3S)`, `fclose(3S)`, `getc(3S)`, `puts(3S)`, `printf(3S)`, `fread(3S)`

DIAGNOSTICS

These functions return the constant **EOF** upon error. Since this is a good integer, `ferror(3S)` should be used to detect *putw* errors.

RESTRICTIONS

Because it is implemented as a macro, *putc* treats a stream argument with side effects incorrectly. In particular, '`putc(c, *f++)`;' doesn't work as expected.

STATUS

PUTC(3S) is supported by Digital Equipment Corporation.

PUTS(3S)

NAME

puts, *fputs* – put a string on a stream

SYNTAX

```
#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*

RESTRICTIONS

Puts appends a newline, while *fputs* does not.

STATUS

PUTS(3S) is supported by Digital Equipment Corporation.

NAME

qsort — quicker sort

SYNTAX

```
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)

STATUS

QSORT(3) currently is not supported by Digital Equipment Corporation.

QSORT(3F)

NAME

qsort — quick sort

SYNTAX

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)

STATUS

QSORT(3F) currently is not supported by Digital Equipment Corporation.

NAME

rand, srand – random number generator

SYNTAX

srand(seed)

int seed;

rand()

DESCRIPTION

The newer random(3) should be used in new applications; rand remains for compatilby.

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)

STATUS

RAND(3C) currently is not supported by Digital Equipment Corporation.

RAND(3F)

NAME

rand, *drand*, *irand* – return random values

SYNTAX

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)

RESTRICTIONS

The algorithm returns a 31 bit quantity on the VAX.

STATUS

RAND(3F) currently is not supported by Digital Equipment Corporation.

NAME

random, srandom, initstate, setstate – better random number generator; routines for changing generators

SYNTAX

```
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 $16*(2^{31}-1)$.

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*.

RANDOM(3)

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)

STATUS

RANDOM(3) currently is not supported by Digital Equipment Corporation.

NAME

rcmd, *rresvport*, *ruserok* – routines for returning a stream to a remote command

SYNTAX

```
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

RCMD(3X)

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)

STATUS

RCMD(3X) currently is not supported by Digital Equipment Corporation.

NAME

`re_comp`, `re_exec` – regular expression handler

SYNTAX

```
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 \ ).
```

STATUS

REGEX(3) currently is not supported by Digital Equipment Corporation.

RENAME(3F)

NAME

rename – rename a file

SYNTAX

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)

RESTRICTIONS

Pathnames can be no longer than MAXPATHLEN as defined in <sys/param.h>.

STATUS

RENAME(3F) currently is not supported by Digital Equipment Corporation.

NAME

`rexec` — return stream to a remote command

SYNTAX

```
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. 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 a auxiliary channel to a control process will be setup, and a descriptor for it will be placed in **fd2p*. The control process will return diagnostic output from the command (unit 2) on this channel, and will also accept bytes on this channel as being UNIX signal numbers, to be forwarded to the process group of the command. If *fd2p* is 0, then the **`stderr`** (unit 2 of the remote command) will be made the same as the **`stdout`** and no provision is made for sending arbitrary signals to the remote process, although you may be able to get its attention by using out-of-band data.

SEE ALSO

rcmd(3X), *rexecd(8C)*

STATUS

REXEC(3X) currently is not supported by Digital Equipment Corporation.

SCANDIR(3)

NAME

`scandir` – scan a directory

SYNTAX

```
#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.

STATUS

SCANDIR(3) currently is not supported by Digital Equipment Corporation.

NAME

`scanf`, `fscanf`, `sscanf` – formatted input conversion

SYNTAX

```
#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

SCANF(3S)

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.

RESTRICTIONS

The success of literal matches and suppressed assignments is not directly determinable.

STATUS

SCANF(3S) is supported by Digital Equipment Corporation.

SETBUF(3S)

NAME

setbuf, setbuffer, setlinebuf – assign buffering to a stream

SYNTAX

```
#include <stdio.h>

setbuf(stream, buf)
FILE *stream;
char *buf;

setbuffer(stream, buf, size)
FILE *stream;
char *buf;
int size;

setlinebuf(stream)
FILE *stream;
```

DESCRIPTION

The three types of buffering available are unbuffered, block buffered, and line buffered. When an output stream is unbuffered, information appears on the destination file or terminal as soon as written; when it is block buffered many characters are saved up and written as a block; when it is line buffered characters are saved up until a newline is encountered or input is read from stdin. *Flush* (see *fclose(3S)*) may be used to force the block out early. Normally all files are block buffered. A buffer is obtained from *malloc(3)* upon the first *getc* or *putc(3S)* on the file. If the standard stream **stdout** refers to a terminal it is line buffered. The standard stream **stderr** is always unbuffered.

Setbuf is used after a stream has been opened but before it is read or written. The character array *buf* is used instead of an automatically allocated buffer. If *buf* is the constant pointer **NULL**, input/output will be completely unbuffered. A manifest constant **BUFSIZ** tells how big an array is needed:

```
char buf[BUFSIZ];
```

Setbuffer, an alternate form of *setbuf*, is used after a stream has been opened but before it is read or written. The character array *buf* whose size is determined by the *size* argument is used instead of an automatically allocated buffer. If *buf* is the constant pointer **NULL**, input/output will be completely unbuffered.

Setlinebuf is used to change *stdout* or *stderr* from block buffered or unbuffered to line buffered. Unlike *setbuf* and *setbuffer* it can be used at any time that the file descriptor is active.

A file can be changed from unbuffered or line buffered to block buffered by using *freopen* (see *fopen(3S)*). A file can be changed from block buffered or line buffered to unbuffered by using *freopen* followed by *setbuf* with a buffer argument of **NULL**.

SEE ALSO

fopen(3S), getc(3S), putc(3S), malloc(3), fclose(3S), puts(3S), printf(3S), fread(3S)

RESTRICTIONS

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.

STATUS

SETBUF(3S) is supported by Digital Equipment Corporation.

SETJMP(3)

NAME

setjmp, longjmp – non-local goto

SYNTAX

```
#include <setjmp.h>
```

```
setjmp(env)  
jmp buf env;
```

```
longjmp(env, val)  
jmp buf env;
```

```
_setjmp(env)  
jmp_buf env;
```

```
_longjmp(env, val)  
jmp_buf env;
```

DESCRIPTION

These routines are useful for dealing with errors and interrupts encountered in a low-level subroutine of a program.

Setjmp saves its stack environment in *env* for later use by *longjmp*. It returns value 0.

Longjmp restores the environment saved by the last call of *setjmp*. It then returns in such a way that execution continues as if the call of *setjmp* had just returned the value *val* to the function that invoked *setjmp*, which must not itself have returned in the interim. All accessible data have values as of the time *longjmp* was called.

Setjmp and *longjmp* save and restore the signal mask *sigmask(2)*, while *_setjmp* and *_longjmp* manipulate only the C stack and registers.

SEE ALSO

sigvec(2), sigstack(2), signal(3)

RESTRICTIONS

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.

STATUS

SETJMP(3) currently is not supported by Digital Equipment Corporation.

NAME

setuid, seteuid, setruid, setgid, setegid, setrgid – set user and group ID

SYNTAX

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.

STATUS

SETUID(3) currently is not supported by Digital Equipment Corporation.

SIGNAL(3C)

NAME

signal – simplified software signal facilities

SYNTAX

```
#include <signal.h>

(*signal(sig, func))()
void (*func)();
```

DESCRIPTION

Signal is a simplified interface to the more general *sigvec(2)* facility.

A signal is generated by some abnormal event, initiated by a user at a terminal (quit, interrupt, stop), by a program error (bus error, etc.), by request of another program (kill), or when a process is stopped because it wishes to access its control terminal while in the background (see *tty(4)*). Signals are optionally generated when a process resumes after being stopped, when the status of child processes changes, or when input is ready at the control terminal. Most signals cause termination of the receiving process if no action is taken; some signals instead cause the process receiving them to be stopped, or are simply discarded if the process has not requested otherwise. Except for the SIGKILL and SIGSTOP signals, the *signal* call allows signals either to be ignored or to cause an interrupt to a specified location. The following is a list of all signals with names as in the include file *<signal.h>*:

SIGHUP	1	hangup
SIGINT	2	interrupt
SIGQUIT	3*	quit
SIGILL	4*	illegal instruction
SIGTRAP	5*	trace trap
SIGIOT	6*	IOT instruction
SIGEMT	7*	EMT instruction
SIGFPE	8*	floating point exception
SIGKILL	9	kill (cannot be caught or ignored)
SIGBUS	10*	bus error
SIGSEGV	11*	segmentation violation
SIGSYS	12*	bad argument to system call
SIGPIPE	13	write on a pipe with no one to read it
SIGALRM	14	alarm clock
SIGTERM	15	software termination signal
SIGURG	16•	urgent condition present on socket
SIGSTOP	17†	stop (cannot be caught or ignored)
SIGTSTP	18†	stop signal generated from keyboard
SIGCONT	19•	continue after stop
SIGCHLD	20•	child status has changed
SIGTTIN	21†	background read attempted from control terminal
SIGTTOU	22†	background write attempted to control terminal
SIGIO	23•	i/o is possible on a descriptor (see <i>fcntl(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 *vfork*(2) the child inherits all signals. *Execve*(2) resets all caught signals to the default action; ignored signals remain ignored.

RETURN VALUE

The previous action is returned on a successful call. Otherwise, -1 is returned and *errno* is set to indicate the error.

ERRORS

Signal will fail and no action will take place if one of the following occur:

- [EINVAL] *Sig* is not a valid signal number.
- [EINVAL] An attempt is made to ignore or supply a handler for SIGKILL or SIGSTOP.
- [EINVAL] An attempt is made to ignore SIGCONT (by default SIGCONT is ignored).

SEE ALSO

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

NOTES (VAX-11)

The handler routine can be declared:

```
handler(sig, code, scp)
```

Here *sig* is the signal number, into which the hardware faults and traps are mapped as defined below. Code is a parameter which is either a constant as given below or, for compatibility mode faults, the code provided by the hardware. *Scp* is a pointer to the *struct*

SIGNAL(3C)

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_PRIVIN_FAULT
Customer-reserved instr.	SIGEMT	
Reserved operand	SIGILL	ILL_RESOP_FAULT
Reserved addressing	SIGILL	ILL_RESAD_FAULT
Trace pending	SIGTRAP	
Bpt instruction	SIGTRAP	
Compatibility-mode	SIGILL	hardware supplied code
Chme	SIGSEGV	
Chms	SIGSEGV	
Chmu	SIGSEGV	

STATUS

SIGNAL(3C) is supported by Digital Equipment Corporation.

NAME

signal — change the action for a signal

SYNTAX

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.

STATUS

SIGNAL(3F) currently is not supported by Digital Equipment Corporation.

SIN(3M)

NAME

sin, *cos*, *tan*, *asin*, *acos*, *atan*, *atan2* – trigonometric functions

SYNTAX

```
#include <math.h>

double sin(x)
double x;

double cos(x)
double x;

double asin(x)
double x;

double acos(x)
double x;

double atan(x)
double x;

double atan2(x, y)
double x, y;
```

DESCRIPTION

Sin, *cos* and *tan* return trigonometric functions of radian arguments. The magnitude of the argument should be checked by the caller to make sure the result is meaningful.

Asin returns the arc sin in the range $-\pi/2$ to $\pi/2$.

Acos returns the arc cosine in the range 0 to π .

Atan returns the arc tangent of x in the range $-\pi/2$ to $\pi/2$.

Atan2 returns the arc tangent of x/y in the range $-\pi$ to π .

DIAGNOSTICS

Arguments of magnitude greater than 1 cause *asin* and *acos* to return value 0; *errno* is set to EDOM. The value of *tan* at its singular points is a huge number, and *errno* is set to ERANGE.

RESTRICTIONS

The value of *tan* for arguments greater than about $2^{*}31$ is unreliable.

STATUS

SIN(3M) currently is not supported by Digital Equipment Corporation.

NAME

sinh, *cosh*, *tanh* – hyperbolic functions

SYNTAX

```
#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.

STATUS

SINH(3M) currently is not supported by Digital Equipment Corporation.

SLEEP(3)

NAME

sleep – suspend execution for interval

SYNTAX

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)

STATUS

SLEEP(3) currently is not supported by Digital Equipment Corporation.

NAME

sleep – suspend execution for an interval

SYNTAX

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)

STATUS

SLEEP(3F) currently is not supported by Digital Equipment Corporation.

STAT(3F)

NAME

stat, lstat, fstat – get file status

SYNTAX

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), perror(3F), time(3F)

RESTRICTIONS

Pathnames can be no longer than MAXPATHLEN as defined in *<sys/param.h>*.

STATUS

STAT(3F) currently is not supported by Digital Equipment Corporation.

NAME

strcat, strncat, strcmp, strncmp, strcpy, strncpy, strlen, index, rindex – string operations

SYNTAX

```
#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 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.

STRING(3)

STATUS

STRING(3) is supported by Digital Equipment Corporation.

NAME

`stty`, `gty` – set and get terminal state (defunct)

SYNTAX

```
#include <sgtty.h>
```

```
stty(fd, buf)
```

```
int fd;
```

```
struct sgttyb *buf;
```

```
gty(fd, buf)
```

```
int fd;
```

```
struct sgttyb *buf;
```

DESCRIPTION

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

Stty sets the state of the terminal associated with *fd*. *Gty* retrieves the state of the terminal associated with *fd*. To set the state of a terminal the call must have write permission.

The *stty* call is actually “`ioctl(fd, TIOCSETP, buf)`”, while the *gty* call is “`ioctl(fd, TIOCGETP, buf)`”. See *ioctl(2)* and *tty(4)* for an explanation.

DIAGNOSTICS

If the call is successful 0 is returned, otherwise -1 is returned and the global variable *errno* contains the reason for the failure.

SEE ALSO

`ioctl(2)`, `tty(4)`

STATUS

STTY(3C) currently is not supported by Digital Equipment Corporation.

SWAB(3)

NAME

swab -- swap bytes

SYNTAX

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.

STATUS

SWAB(3) currently is not supported by Digital Equipment Corporation.

NAME

syslog, openlog, closelog – control system log

SYNTAX

```
#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 */dev/log*, it writes on */dev/console* instead and returns *-1*.

Closelog can be used to close the log file.

EXAMPLES

```
syslog(LOG_SALERT, "who: internal error 23");
```

```
openlog("serverftp", LOG_PID);
syslog(LOG_INFO, "Connection from host %d", CallingHost);
```

SEE ALSO

syslog(8)

STATUS

SYSLOG(3) currently is not supported by Digital Equipment Corporation.

SYSTEM(3)

NAME

system — issue a shell command

SYNTAX

```
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.

STATUS

SYSTEM(3) is supported by Digital Equipment Corporation.

NAME

system – execute a UNIX command

SYNTAX

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)

RESTRICTIONS

String can not be longer than NCARGS-50 characters, as defined in *<sys/param.h>*.

STATUS

SYSTEM(3F) currently is not supported by Digital Equipment Corporation.

TERMCAP(3X)

NAME

tgetent, tgetnum, tgetflag, tgetstr, tgoto, tputs – terminal independent operation routines

SYNTAX

```
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 */etc/termcap*. This can speed up entry into programs that call *tgetent*, as well as to help debug new terminal descriptions or to make one for your terminal if you can't write the file */etc/termcap*.

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)

STATUS

TERMCAP(3X) currently is not supported by Digital Equipment Corporation.

TIME(3C)

NAME

time, ftime – get date and time

SYNTAX

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)

STATUS

TIME(3C) currently is not supported by Digital Equipment Corporation.

NAME

time, *ctime*, *ltime*, *gmtime* – return system time

SYNTAX

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)

STATUS

TIME (3F) currently is not supported by Digital Equipment Corporation.

TIMES(3C)

NAME

times – get process times

SYNTAX

```
#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)`

STATUS

TIMES(3C) currently is not supported by Digital Equipment Corporation.

NAME

topen, tclose, tread, twrite, trewin, tskipf, tstate – f77 tape I/O

SYNTAX

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, eoff, eotf, tcsr)

integer tlu, fileno, recno, tcsr

logical errf, eoff, 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 *perror*(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

TOPEN(3F)

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 *tlu*.) 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*, *eoff*, 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 *tcsr* 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)*

STATUS

TOPEN(3F) currently is not supported by Digital Equipment Corporation.

NAME

traper — trap arithmetic errors

SYNTAX

integer function traper (mask)

DESCRIPTION

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)

STATUS

TRAPER(3F) currently is not supported by Digital Equipment Corporation.

TRAPOV(3F)

NAME

trapov – trap and repair floating point overflow

SYNTAX

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 (REV 7) 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 "dfimax()".

FILES

/usr/lib/libF77.a

SEE ALSO

trpfpe(3F), signal(3F), range(3F)

RESTRICTIONS

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.

STATUS

TRAPOV(3F) currently is not supported by Digital Equipment Corporation.

NAME

trpfpe, fpecnt — trap and repair floating point faults

SYNTAX

subroutine trpfpe (numesg, rtnval)
double precision rtnval

integer function fpecnt ()

common /fpeflt/ fperr
logical fperr

DESCRIPTION

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 “dfimax()”.

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.

Fpecnt returns the number of faults since the last call to *trpfpe*.

The logical value in the common block labelled *fpeflt* will be set to **.true.** each time a fault occurs.

FILES

/usr/lib/libF77.a

SEE ALSO

signal(3F), range(3F)

RESTRICTIONS

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.

STATUS

TRPFPE(3F) currently is not supported by Digital Equipment Corporation.

TTYNAM(3F)

NAME

ttynam, *isatty* – find name of a terminal port

SYNTAX

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'.

STATUS

TTYNAM(3F) currently is not supported by Digital Equipment Corporation.

NAME

`ttyname`, `isatty`, `ttyslot` – find name of a terminal

SYNTAX

`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.

RESTRICTIONS

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

STATUS

TTYNAME(3) currently is not supported by Digital Equipment Corporation.

UNGETC(3S)

NAME

`ungetc` – push character back into input stream

SYNTAX

```
#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.

STATUS

UNGETC(3S) is supported by Digital Equipment Corporation.

NAME

unlink – remove a directory entry

SYNTAX

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), filsys(5), perror(3F)

RESTRICTIONS

Pathnames can be no longer than MAXPATHLEN as defined in *<sys/param.h>*.

STATUS

UNLINK(3F) currently is not supported by Digital Equipment Corporation.

UTIME(3C)

NAME

utime – set file times

SYNTAX

```
#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)

STATUS

UTIME(3C) currently is not supported by Digital Equipment Corporation.

NAME

`valloc` – aligned memory allocator

SYNTAX

```
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.

RESTRICTIONS

Vfree isn't implemented.

STATUS

VALLOC(3) currently is not supported by Digital Equipment Corporation.

VARARGS(3)

NAME

varargs – variable argument list

SYNTAX

```
#include <varargs.h>

function(va_alist)
va_dcl
va_list pvar;
va_start(pvar);
f = va_arg(pvar, type);
va_end(pvar);
```

DESCRIPTION

This set of macros provides a means of writing portable procedures that accept variable argument lists. Routines having variable argument lists (such as *printf(3)*) that do not use varargs are inherently nonportable, since different machines use different argument passing conventions.

va_alist is used in a function header to declare a variable argument list.

va_dcl is a declaration for **va_alist**. Note that there is no semicolon after **va_dcl**.

va_list is a type which can be used for the variable *pvar*, which is used to traverse the list. One such variable must always be declared.

va_start(pvar) is called to initialize *pvar* to the beginning of the list.

va_arg(pvar, type) will return the next argument in the list pointed to by *pvar*. *Type* is the type the argument is expected to be. Different types can be mixed, but it is up to the routine to know what type of argument is expected, since it cannot be determined at run-time.

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_alist)
va_dcl
{
    va_list ap;
    char *file;
    char *args[100];
    int argno = 0;

    va_start(ap);
    file = va_arg(ap, char *);
    while (args[argno++] = va_arg(ap, char *))
        ;
}
```

```
    va_end(ap);  
    return execv(file, args);  
}
```

RESTRICTIONS

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.

STATUS

VARARGS(3) currently is not supported by Digital Equipment Corporation.

VLIMIT(3C)

NAME

`vlimit` – control maximum system resource consumption

SYNTAX

`#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 `cs(1)`.

The system refuses to extend the data or stack space when the limits would be exceeded in the normal way; a *break* call fails if the data space limit is reached, or the process is killed when the stack limit is reached (since the stack cannot be extended, there is no way to send a signal!).

A file i/o operation which would create a file which is too large will cause a signal SIGXFSZ to be generated, this normally terminates the process, but may be caught. When the cpu time limit is exceeded, a signal SIGXCPU is sent to the offending process; to allow it time to process the signal it is given 5 seconds grace by raising the cpu time limit.

SEE ALSO

`cs(1)`

RESTRICTIONS

If LIM_NORAISE is set, then no grace should be given when the cpu time limit is exceeded.

STATUS

VLIMIT(3C) currently is not supported by Digital Equipment Corporation.

VTIMES(3C)

NAME

vtimes – get information about resource utilization

SYNTAX

```
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_untime;          /* user time (*HZ) */  
    int    vm_stime;           /* system time (*HZ) */  
    /* divide next two by untime+stime to get averages */  
    unsigned vm_idrssi;       /* integral of d+s rss */  
    unsigned vm_ixrssi;       /* integral of text rss */  
    int    vm_maxrss;         /* maximum rss */  
    int    vm_majflt;         /* major page faults */  
    int    vm_minflt;         /* minor page faults */  
    int    vm_nswap;          /* number of swaps */  
    int    vm_inblk;          /* block reads */  
    int    vm_oublk;          /* block writes */  
};
```

The *vm_untime* 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_untime+vm_stime* would be the 60. *vm_idrssi* integrates data and stack segment usage, while *vm_ixrssi* integrates text segment usage. *vm_maxrss* reports the maximum instantaneous sum of the text+data+stack core-resident page count.

The *vm_majflt* field gives the number of page faults which resulted in disk activity; the *vm_minflt* field gives the number of page faults incurred in simulation of reference bits; *vm_nswap* is the number of swaps which occurred. The number of file system input/output events are reported in *vm_inblk* and *vm_oublk*. These numbers account only for real i/o; data supplied by the caching mechanism is charged only to the first process to read or write the data.

SEE ALSO

time(2), wait3(2)

STATUS

VTIMES(3C) currently is not supported by Digital Equipment Corporation.

WAIT(3F)

NAME

wait – wait for a process to terminate

SYNTAX

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)*

STATUS

WAIT(3F) currently is not supported by Digital Equipment Corporation.

NAME

a.out – assembler and link editor output

SYNTAX

```
#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).

A.OUT(5)

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.

A.OUT(5)

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 `r_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)`

STATUS

A.OUT(5) currently is not supported by Digital Equipment Corporation.

NAME

acct – execution accounting file

SYNTAX

```
#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          6.1          83/07/29*/

/*
 * Accounting structures;
 * these use a comp_t type which is a 3 bits base 8
 * exponent, 13 bit fraction “floating point” number.
 */
typedef u_short comp_t;

struct acct
{
    char      ac_comm[10]; /* Accounting command name */
    comp_t    ac_utime;    /* 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 */

#ifdef KERNEL
struct acct      acctbuf;
struct inode     *acctp;
#endif
```

ACCT(5)

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)

STATUS

ACCT(5) currently is not supported by Digital Equipment Corporation.

NAME

aliases – aliases file for sendmail

SYNTAX

`/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.

RESTRICTIONS

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.

STATUS

ALIASES(5) currently is not supported by Digital Equipment Corporation.

AR(5)

NAME

`ar` – archive (library) file format

SYNTAX

```
#include <ar.h>
```

DESCRIPTION

The archive command `ar` combines several files into one. Archives are used mainly as libraries to be searched by the link-editor `ld`.

A file produced by `ar` has a magic string at the start, followed by the constituent files, each preceded by a file header. The magic number and header layout as described in the include file are:

```
/*      ar.h      4.183/05/03*/

#define ARMAG "!<arch>\n"
#define SARMAG 8

#define ARFMAG "'\n"

struct ar_hdr {
    char    ar_name[16];
    char    ar_date[12];
    char    ar_uid[6];
    char    ar_gid[6];
    char    ar_mode[8];
    char    ar_size[10];
    char    ar_fmag[2];
};
```

The name is a blank-padded string. The `ar_fmag` field contains ARFMAG to help verify the presence of a header. The other fields are left-adjusted, blank-padded numbers. They are decimal except for `ar_mode`, which is octal. The date is the modification date of the file at the time of its insertion into the archive.

Each file begins on a even (0 mod 2) boundary; a new-line is inserted between files if necessary. Nevertheless the size given reflects the actual size of the file exclusive of padding.

There is no provision for empty areas in an archive file.

The encoding of the header is portable across machines. If an archive contains printable files, the archive itself is printable.

SEE ALSO

`ar(1)`, `ld(1)`, `nm(1)`

RESTRICTIONS

File names lose trailing blanks.

) **STATUS**

AR(5) currently is not supported by Digital Equipment Corporation.

CORE(5)

NAME

core – format of memory image file

SYNTAX

```
#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)*

STATUS

CORE(5) currently is not supported by Digital Equipment Corporation.

NAME

dir – format of directories

SYNTAX

```
#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

/*
```

DIR(5)

* 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)

STATUS

DIR(5) currently is not supported by Digital Equipment Corporation.

NAME

disktab – disk description file

SYNTAX

```
#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

DISKTAB(5)

SEE ALSO

newfs(8), diskpart(8)

STATUS

DISKTAB(5) currently is not supported by Digital Equipment Corporation.

NAME

dump, dumpdates – incremental dump format

SYNTAX

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

DESCRIPTION

Tapes used by *dump* and *restore*(8) contain:

- a header record
- two groups of bit map records
- a group of records describing directories
- a group of records describing files

The format of the header record and of the first record of each description as given in the include file *<dumprest.h>* is:

```
#define NTREC      10
#define MLEN       16
#define MSIZ      4096
```

```
#define TS_TAPE      1
#define TS_INODE     2
#define TS_BITS      3
#define TS_ADDR      4
#define TS_END       5
#define TS_CLRI      6
#define MAGIC        (int) 60011
#define CHECKSUM     (int) 84446
```

```
struct  spcl {
    int          c_type;
    time_t      c_date;
    time_t      c_ddate;
    int         c_volume;
    daddr_t     c_tapea;
    ino_t       c_inumber;
    int         c_magic;
    int         c_checksum;
    struct      dinode      c_dinode;
    int         c_count;
    char        c_addr[BSIZE];
} spcl;
```

```
struct  idates {
```

DUMP(5)

```
    char          id_name[16];
    char          id_incno;
    time_t        id_ddate;
};
```

```
#define DUMPOUTFMT "%-16s %c %s"      /* for printf */
/* name, incno, ctime(date) */
#define DUMPINFMT  "%16s %c %[\n]|n"  /* inverse for scanf */
```

NTREC is the number of 1024 byte records in a physical tape block. MLEN is the number of bits in a bit map word. MSIZ is the number of bit map words.

The TS_ entries are used in the c_type field to indicate what sort of header this is. The types and their meanings are as follows:

TS_TAPE Tape volume label
TS_INODE A file or directory follows. The 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)*

STATUS

DUMP(5) currently is not supported by Digital Equipment Corporation.

NAME

fs, inode – format of file system volume

SYNTAX

```
#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_sbkn0;            /* addr of super-block in filesys */
    daddr_t fs_cbkn0;            /* offset of cyl-block in filesys */
    daddr_t fs_ibkn0;            /* offset of inode-blocks in filesys */
    daddr_t fs_dbkn0;            /* offset of first data after cg */
    long fs_cgoffset;            /* cylinder group offset in cylinder */
    long fs_cgmask;              /* used to calc mod fs_ntrak */
    time_t fs_time;              /* last time written */
    long fs_size;                 /* number of blocks in fs */
    long fs_dsize;                /* number of data blocks in fs */
    long fs_ncg;                  /* number of cylinder groups */
    long fs_bsize;                /* size of basic blocks in fs */
    long fs_fsize;                /* size of frag blocks in fs */
    long fs_frag;                 /* number of frags in a block in fs */
    /* these are configuration parameters */
    long fs_minfree;              /* minimum percentage of free blocks */
    long fs_rotdelay;            /* num of ms for optimal next block */
    long fs_rps;                  /* disk revolutions per second */
    /* these fields can be computed from the others */
    long fs_bmask;                /* "blkoff" calc of blk offsets */
    long fs_fmask;                /* "fragoff" calc of frag offsets */
    long fs_bshift;               /* "lblkno" calc of logical blkno */
    long fs_fshift;               /* "numfrags" calc number of frags */
    /* these are configuration parameters */
    long fs_maxcontig;            /* max number of contiguous blks */
    long fs_maxbpg;              /* max number of blks per cyl group */
    /* these fields can be computed from the others */
```

```

long    fs_fragshift;           /* block to frag shift */
long    fs_fsbtodb;            /* fsbtodb and dbtfsb shift constant */
long    fs_sbsize;             /* actual size of super block */
long    fs_csmask;             /* csum block offset */
long    fs_csshift;            /* csum block number */
long    fs_nindir;             /* value of NINDIR */
long    fs_inopb;              /* value of INOPB */
long    fs_nspf;               /* value of NSPF */
long    fs_sparecon[6];        /* reserved for future constants */
/* sizes determined by number of cylinder groups and their sizes */
daddr_t fs_csaddr;             /* blk addr of cyl grp summary area */
long    fs_cssize;             /* size of cyl grp summary area */
long    fs_cgsize;             /* cylinder group size */
/* these fields should be derived from the hardware */
long    fs_ntrak;              /* tracks per cylinder */
long    fs_nsect;              /* sectors per track */
long    fs_spc;                /* sectors per cylinder */
/* this comes from the disk driver partitioning */
long    fs_ncyl;               /* cylinders in file system */
/* these fields can be computed from the others */
long    fs_cpg;                /* cylinders per group */
long    fs_ipg;                /* inodes per group */
long    fs_fpg;                /* blocks per group * fs_frag */
/* this data must be re-computed after crashes */
struct  csum fs_cstotal; /* cylinder summary information */
/* these fields are cleared at mount time */
char    fs_fmod;               /* super block modified flag */
char    fs_clean;              /* file system is clean flag */
char    fs_ronly;              /* mounted read-only flag */
char    fs_flags;              /* currently unused flag */
char    fs_fsmnt[MAXMNTLEN];   /* name mounted on */
/* these fields retain the current block allocation info */
long    fs_cgrotor;            /* last cg searched */
struct  csum *fs_csp[MAXCSBUFS]; /* list of fs_cs info buffers */
long    fs_cpc;                /* cyl per cycle in postbl */
short   fs_postbl[MAXCPG][NRPOS]; /* head of blocks for each rotation */
long    fs_magic;              /* magic number */
u_char  fs_rotbl[1];           /* list of blocks for each rotation */
/* actually longer */
};

```

Each disk drive contains some number of file systems. A file system consists of a number of cylinder groups. Each cylinder group has inodes and data.

A file system is described by its super-block, which in turn describes the cylinder groups. The super-block is critical data and is replicated in each cylinder group to protect against catastrophic loss. This is done at file system creation time and the critical super-block data does not change, so the copies need not be referenced further unless disaster strikes.

Addresses stored in inodes are capable of addressing fragments of 'blocks'. File system blocks of at most size MAXBSIZE can be optionally broken into 2, 4, or 8 pieces, each of which is addressable; these pieces may be DEV_BSIZE, or some multiple of a DEV_BSIZE unit.

Large files consist of exclusively large data blocks. To avoid undue wasted disk space, the last data block of a small file is allocated as only as many fragments of a large block as are necessary. The file system format retains only a single pointer to such a fragment, which is a piece of a single large block that has been divided. The size of such a fragment is determinable from information in the inode, using the "blksize(fs, ip, lbn)" macro.

The file system records space availability at the fragment level; to determine block availability, aligned fragments are examined.

The root inode is the root of the file system. Inode 0 can't be used for normal purposes and historically bad blocks were linked to inode 1, thus the root inode is 2 (inode 1 is no longer used for this purpose, however numerous dump tapes make this assumption, so we are stuck with it). The *lost+found* directory is given the next available inode when it is initially created by *mkfs*.

fs_minfree gives the minimum acceptable percentage of file system blocks which may be free. If the freelist drops below this level only the super-user may continue to allocate blocks. This may be set to 0 if no reserve of free blocks is deemed necessary, however severe performance degradations will be observed if the file system is run at greater than 90% full; thus the default value of *fs_minfree* is 10%.

Empirically the best trade-off between block fragmentation and overall disk utilization at a loading of 90% comes with a fragmentation of 4, thus the default fragment size is a fourth of the block size.

Cylinder group related limits: Each cylinder keeps track of the availability of blocks at different rotational positions, so that sequential blocks can be laid out with minimum rotational latency. NRPOS is the number of rotational positions which are distinguished. With NRPOS 8 the resolution of the summary information is 2ms for a typical 3600 rpm drive.

fs_rotdelay gives the minimum number of milliseconds to initiate another disk transfer on the same cylinder. It is used in determining the rotationally optimal layout for disk blocks within a file; the default value for *fs_rotdelay* is 2ms.

Each file system has a statically allocated number of inodes. An inode is allocated for each NBPI bytes of disk space. The inode allocation strategy is extremely conservative.

MAXIPG bounds the number of inodes per cylinder group, and is needed only to keep the structure simpler by having the only a single variable size element (the free bit map). 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. 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>*.

STATUS

FS(5) currently is not supported by Digital Equipment Corporation.

FSTAB(5)

NAME

`fstab` – static information about the filesystems

SYNTAX

```
#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)`

STATUS

FSTAB(5) currently is not supported by Digital Equipment Corporation.

GETTYTAB(5)

NAME

gettytab -- terminal configuration data base

SYNTAX

/etc/gettytab

DESCRIPTION

Gettytab is a simplified version of the *termcap(5)* data base used to describe terminal lines. The initial terminal login process *getty(8)* accesses the *gettytab* file each time it starts, allowing simpler reconfiguration of terminal characteristics. Each entry in the data base is used to describe one class of terminals.

There is a default terminal class, *default*, that is used to set global defaults for all other classes. (That is, the *default* entry is read, then the entry for the class required is used to override particular settings.)

CAPABILITIES

Refer to *termcap(5)* for a description of the file layout. The *default* column below lists defaults obtained if there is no entry in the table obtained, nor one in the special *default* table.

Name	Type	Default	Description
ap	bool	false	terminal uses any parity
bd	num	0	backspace delay
bk	str	0377	alternate end of line character (input break)
cb	bool	false	use crt backspace mode
cd	num	0	carriage-return delay
ce	bool	false	use crt erase algorithm
ck	bool	false	use crt kill algorithm
cl	str	NULL	screen clear sequence
co	bool	false	console - add \n after login prompt
ds	str	^Y	delayed suspend character
ec	bool	false	leave echo OFF
ep	bool	false	terminal uses even parity
er	str	^?	erase character
et	str	^D	end of text (EOF) character
ev	str	NULL	initial enviroment
f0	num	unused	tty mode flags to write messages
f1	num	unused	tty mode flags to read login name
f2	num	unused	tty mode flags to leave terminal as
fd	num	0	form-feed (vertical motion) delay
fl	str	^O	output flush character
hc	bool	false	do NOT hangup line on last close
he	str	NULL	hostname editing string
hn	str	hostname	hostname
ht	bool	false	terminal has real tabs
ig	bool	false	ignore garbage characters in login name

im	str	NULL	initial (banner) message
in	str	^C	interrupt character
is	num	unused	input speed
kl	str	^U	kill character
lc	bool	false	terminal has lower case
lm	str	login:	login prompt
ln	str	^V	“literal next” character
lo	str	/bin/login	program to exec when name obtained
nd	num	0	newline (line-feed) delay
nl	bool	false	terminal has (or might have) a newline character
nx	str	default	next table (for auto speed selection)
op	bool	false	terminal uses odd parity
os	num	unused	output speed
pc	str	\0	pad character
pe	bool	false	use printer (hard copy) erase algorithm
pf	num	0	delay between first prompt and following flush (seconds)
ps	bool	false	line connected to a MICOM port selector
qu	str	^\ ^	quit character
rp	str	^R	line retype character
rw	bool	false	do NOT use raw for input, use cbreak
sp	num	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 enviroment)
ub	bool	false	do unbuffered output (of prompts etc)
uc	bool	false	terminal is known upper case only
we	str	^W	word erase character
xc	bool	false	do NOT echo control chars as ^X
xf	str	^S	XOFF (stop output) character
xn	str	^Q	XON (start output) character

If no line speed is specified, speed will not be altered from that which prevails when *getty* is entered. Specifying an input or output speed 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.

GETTYTAB(5)

Delays are specified in milliseconds, the nearest possible delay available in the tty driver will be used. Should greater certainty be desired, delays with values 0, 1, 2, and 3 are interpreted as choosing that particular delay algorithm from the driver.

The **cl** screen clear string may be preceded by a (decimal) number of milliseconds of delay required (a la termcap). This delay is simulated by repeated use of the pad character **pc**.

The initial message, and login message, **im** and **lm** may include the character sequence **%h** to obtain the hostname. (**%%** obtains a single **'%'** character.) The hostname is normally obtained from the system, but may be set by the **hn** table entry. In either case it may be edited with **he**. The **he** string is a sequence of characters, each character that is neither **'@'** nor **'#'** is copied into the final hostname. A **'@'** in the **he** string, causes one character from the real hostname to be copied to the final hostname. A **'#'** in the **he** string, causes the next character of the real hostname to be skipped. Surplus **'@'** and **'#'** characters are ignored.

When **getty** execs the login process, given in the **lo** string (usually **"/bin/login"**), it will have set the environment to include the terminal type, as indicated by the **tt** string (if it exists). The **ev** string, can be used to enter additional data into the environment. It is a list of comma separated strings, each of which will presumably be of the form *name=value*.

If a non-zero timeout is specified, with **to**, then **getty** will exit within the indicated number of seconds, either having received a login name and passed control to *login*, or having received an alarm signal, and exited. This may be useful to hangup dial in lines.

Output from *getty* is even parity unless **op** is specified. **Op** may be specified with **ap** to allow any parity on input, but generate odd parity output. Note: this only applies while *getty* is being run, terminal driver limitations prevent a more complete implementation. *Getty* does not check parity of input characters in *RAW* mode.

SEE ALSO

termcap(5), getty(8).

RESTRICTIONS

Since some users insist on changing their default special characters, it is wise to define 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.

Currently *login*(1) destroys the environment, so there is no point setting it in *gettytab*.

STATUS

GETTYTAB(5) currently is not supported by Digital Equipment Corporation.

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)*

RESTRICTIONS

The *passwd(1)* command won't change the passwords.

STATUS

GROUP(5) currently is not supported by Digital Equipment Corporation.

HOSTS(5)

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)

STATUS

HOSTS(5) currently is not supported by Digital Equipment Corporation.

NAME

mtab – mounted file system table

SYNTAX

```
#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)

STATUS

MTAB(5) currently is not supported by Digital Equipment Corporation.

NETWORKS(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)

STATUS

NETWORKS(5) currently is not supported by Digital Equipment Corporation.

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)

STATUS

PASSWD(5) currently is not supported by Digital Equipment Corporation.

PHONES(5)

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)`

STATUS

PHONES(5) currently is not supported by Digital Equipment Corporation.

NAME

plot – graphics interface

DESCRIPTION

Files of this format are produced by routines described in *plot(3X)*, and are interpreted for various devices by commands described in *plot(1G)*. A graphics file is a stream of plotting instructions. Each instruction consists of an ASCII letter usually followed by bytes of binary information. The instructions are executed in order. A point is designated by four bytes representing the x and y values; each value is a signed integer. The last designated point in an **l**, **m**, **n**, or **p** instruction becomes the ‘current point’ for the next instruction.

Each of the following descriptions begins with the name of the corresponding routine in *plot(3X)*.

- m** move: The next four bytes give a new current point.
- n** cont: Draw a line from the current point to the point given by the next four bytes. See *plot(1G)*.
- p** point: Plot the point given by the next four bytes.
- l** line: Draw a line from the point given by the next four bytes to the point given by the following four bytes.
- t** label: Place the following ASCII string so that its first character falls on the current point. The string is terminated by a newline.
- a** arc: The first four bytes give the center, the next four give the starting point, and the last four give the end point of a circular arc. The least significant coordinate of the end point is used only to determine the quadrant. The arc is drawn counter-clockwise.
- c** circle: The first four bytes give the center of the circle, the next two the radius.
- e** erase: Start another frame of output.
- f** linemod: Take the following string, up to a newline, as the style for drawing further lines. The styles are ‘dotted,’ ‘solid,’ ‘longdashed,’ ‘shortdashed,’ and ‘dotdashed.’ Effective only in *plot 4014* and *plot ver*.
- s** space: The next four bytes give the lower left corner of the plotting area; the following four give the upper right corner. The plot will be magnified or reduced to fit the device as closely as possible.

Space settings that exactly fill the plotting area with unity scaling appear below for devices supported by the filters of *plot(1G)*. The upper limit is just outside the plotting area. In every case the plotting area is taken to be square; points outside may be displayable on devices whose face isn’t square.

```
4014    space(0, 0, 3120, 3120);
ver     space(0, 0, 2048, 2048);
300, 300s space(0, 0, 4096, 4096);
450     space(0, 0, 4096, 4096);
```

PLOT(5)

SEE ALSO

plot(1G), plot(3X), graph(1G)

STATUS

PLOT(5) currently is not supported by Digital Equipment Corporation.

NAME

printcap – printer capability data base

SYNTAX

/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 lp 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 lp is a tty, clear flag bits (sgtty.h)
ff	str	“\f”	string to send for a form feed
fo	bool	false	print a form feed when device is opened
fs	num	0	like ‘fc’ but set bits
gf	str	NULL	graph data filter (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

PRINTCAP(5)

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

STATUS

PRINTCAP(5) currently is not supported by Digital Equipment Corporation.

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)

STATUS

PROTOCOLS(5) currently is not supported by Digital Equipment Corporation.

REMOTE(5)

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 *x* 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 “*tip*300”, where 300 is the baud rate with which the connection is to be made. When the *cu* interface is used, entries of the form “*cu*300” 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.
- di** (str) Disconnect message sent to the host when a disconnect is requested by the user.
- du** (bool) This host is on a dial-up line.
- dv** (str) UNIX device(s) to open to establish a connection. If this file refers to a terminal line, *tip*(1C) attempts to perform an exclusive open on the device to insure only one user at a time has access to the port.
- el** (str) Characters marking an end-of-line. The default is NULL. “~” 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)*

STATUS

REMOTE(5) currently is not supported by Digital Equipment Corporation.

SERVICES(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)

STATUS

SERVICES(5) currently is not supported by Digital Equipment Corporation.

NAME

stab – symbol table types

SYNTAX

```
#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 */
```

STAB(5)

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

```

STAB(5)

SEE ALSO

as(1), ld(1), dbx(1), a.out(5)

RESTRICTIONS

Sdb assumes that a symbol of type N_GSYM with name *name* is located at address *_name*.

STATUS

STAB(5) currently is not supported by Digital Equipment Corporation.

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

TAR(5)

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)

RESTRICTIONS

Names or linknames longer than NAMSIZ produce error reports and cannot be dumped.

STATUS

TAR(5) currently is not supported by Digital Equipment Corporation.

NAME

termcap – terminal capability data base

SYNTAX

/etc/termcap

DESCRIPTION

Termcap is a data base describing terminals, used, *e.g.*, by *vi*(1) and *curses*(3X). Terminals are described in *termcap* by giving a set of capabilities which they have, and by describing how operations are performed. Padding requirements and initialization sequences are included in *termcap*.

Entries in *termcap* consist of a number of ‘:’ separated fields. The first entry for each terminal gives the names which are known for the terminal, separated by ‘|’ characters. The first name is always 2 characters long and is used by older version 6 systems which store the terminal type in a 16 bit word in a systemwide data base. The second name given is the most common abbreviation for the terminal, and the last name given should be a long name fully identifying the terminal. The second name should contain no blanks; the last name may well contain blanks for readability.

CAPABILITIES

(P) indicates padding may be specified

(P*) indicates that padding may be based on no. lines affected

Name Type Pad? Description

ae	str	(P)	End alternate character set
al	str	(P*)	Add new blank line
am	bool		Terminal has automatic margins
as	str	(P)	Start alternate character set
bc	str		Backspace if not ^H
bs	bool		Terminal can backspace with ^H
bt	str	(P)	Back tab
bw	bool		Backspace wraps from column 0 to last column
CC	str		Command character in prototype if terminal settable
cd	str	(P*)	Clear to end of display
ce	str	(P)	Clear to end of line
ch	str	(P)	Like cm but horizontal motion only, line stays same
cl	str	(P*)	Clear screen
cm	str	(P)	Cursor motion
co	num		Number of columns in a line
cr	str	(P*)	Carriage return, (default ^M)
cs	str	(P)	Change scrolling region (vt100), like cm
cv	str	(P)	Like ch but vertical only.
da	bool		Display may be retained above
dB	num		Number of millisec of bs delay needed
db	bool		Display may be retained below
dC	num		Number of millisec of cr delay needed

TERMCAP (5)

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 ce \r\n (Delta Data)
xs	bool	Standout not erased by writing over it (HP 264?)
xt	bool	Tabs are destructive, magic so char (Telery 1061)

A Sample Entry

The following entry, which describes the Concept-100, is among the more complex entries in the *termcap* file as of this writing. (This particular concept entry is outdated, and is used as an example only.)

```
c1|c100|concept100:is= \EU \Ef \E7 \E5 \E8 \E1 \ENH \EK \E \200 \Eo& \200:\
:al=3* \E R:am:bs:cd=16* \E C:ce=16 \E S:c1=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 ****, and the sequences **\n \r \t \b \f** give a newline, return, tab, backspace and formfeed. Finally, characters may be given as three octal digits after a ****, and the characters **^** and **** may be given as **^** and ****. If it is necessary to place a **:** in a capability it must be escaped in octal as **\072**. If it is necessary to place a null character in a string capability it must be encoded as **\200**. The routines which deal with *termcap* use C strings, and strip the high bits of the output very late so that a **\200** comes out as a **\000** would.

Preparing Descriptions

We now outline how to prepare descriptions of terminals. The most effective way to prepare a terminal description is by imitating the description of a similar terminal in *termcap* and to build up a description gradually, using partial descriptions with *ex* to check that they are correct. Be aware that a very unusual terminal may expose deficiencies in the ability of the *termcap* file to describe it or bugs in *ex*. To easily test a new terminal description you can set the environment variable **TERMCAP** to a pathname of a file containing the description you are working on and the editor will look there rather than in */etc/termcap*. **TERMCAP** can also be set to the *termcap* entry itself to avoid reading the file when starting up the editor. (This only works on version 7 systems.)

Basic capabilities

The number of columns on each line for the terminal is given by the **co** numeric capability. If the terminal is a CRT, then the number of lines on the screen is given by the **li** capability. If the terminal wraps around to the beginning of the next line when it reaches the right margin, then it should have the **am** capability. If the terminal can clear its screen, then this is given by the **cl** string capability. If the terminal can backspace, then it should have the **bs** capability, unless a backspace is accomplished by a character other than **^H** (ugh) in which case you should give this character as the **bc** string capability. If it overstrikes

(rather than clearing a position when a character is struck over) then it should have the **os** capability.

A very important point here is that the local cursor motions encoded in *termcap* are undefined at the left and top edges of a CRT terminal. The editor will never attempt to backspace around the left edge, nor will it attempt to go up locally off the top. The editor assumes that feeding off the bottom of the screen will cause the screen to scroll up, and the **am** capability tells whether the cursor sticks at the right edge of the screen. If the terminal has switch selectable automatic margins, the *termcap* file usually assumes that this is on, i.e. **am**.

These capabilities suffice to describe hardcopy and “glass-tty” terminals. Thus the model 33 teletype is described as

```
t3|33|tty33:co#72:os
```

while the Lear Siegler ADM-3 is described as

```
clladm3|3|lsi adm3:am:bs:cl=~Z:li#24:co#80
```

Cursor addressing

Cursor addressing in the terminal is described by a **cm** string capability, with *printf*(3S) like escapes **%x** in it. These substitute to encodings of the current line or column position, while other characters are passed through unchanged. If the **cm** string is thought of as being a function, then its arguments are the line and then the column to which motion is desired, and the **%** encodings have the following meanings:

```
%d      as in printf, 0 origin
%2       like %2d
%3       like %3d
%.       like %c
%+x      adds x to value, then %.
%>xy    if value > x adds y, no output.
%r       reverses order of line and column, no output
%i       increments line/column (for 1 origin)
%%       gives a single %
%n       exclusive or row and column with 0140 (DM2500)
%B       BCD (16*(x/10)) + (x%10), no output.
%D       Reverse coding (x-2*(x%16)), no output. (Delta Data).
```

Consider the HP2645, which, to get to row 3 and column 12, needs to be sent `\E&a12c03Y` padded for 6 milliseconds. Note that the order of the rows and columns is inverted here, and that the row and column are printed as two digits. Thus its **cm** capability is “**cm=6\E&%r%2c%2Y**”. The Microterm ACT-IV needs the current row and column sent preceded by a **^T**, with the row and column simply encoded in binary, “**cm=~T%.%.**”. Terminals which use “**%.%**” need to be able to backspace the cursor (**bs** or **bc**), and to move the cursor up one line on the screen (**up** introduced below). This is necessary because it is not always safe to transmit `\t`, `\n` **^D** and `\r`, as the system may change or discard them.

A final example is the LSI ADM-3a, which uses row and column offset by a blank character, thus “`cm=\E=% + % +`”.

Cursor motions

If the terminal can move the cursor one position to the right, leaving the character at the current position unchanged, then this sequence should be given as **nd** (non-destructive space). If it can move the cursor up a line on the screen in the same column, this should be given as **up**. If the terminal has no cursor addressing capability, but can home the cursor (to very upper left corner of screen) then this can be given as **ho**; similarly a fast way of getting to the lower left hand corner can be given as **ll**; this may involve going up with **up** from the home position, but the editor will never do this itself (unless **ll** does) because it makes no assumption about the effect of moving up from the home position.

Area clears

If the terminal can clear from the current position to the end of the line, leaving the cursor where it is, this should be given as **ce**. If the terminal can clear from the current position to the end of the display, then this should be given as **cd**. The editor only uses **cd** from the first column of a line.

Insert/delete line

If the terminal can open a new blank line before the line where the cursor is, this should be given as **al**; this is done only from the first position of a line. The cursor must then appear on the newly blank line. If the terminal can delete the line which the cursor is on, then this should be given as **dl**; this is done only from the first position on the line to be deleted. If the terminal can scroll the screen backwards, then this can be given as **sb**, but just **al** suffices. If the terminal can retain display memory above then the **da** capability should be given; if display memory can be retained below then **db** should be given. These let the editor understand that deleting a line on the screen may bring non-blank lines up from below or that scrolling back with **sb** may bring down non-blank lines.

Insert/delete character

There are two basic kinds of intelligent terminals with respect to insert/delete character which can be described using *termcap*. The most common insert/delete character operations affect only the characters on the current line and shift characters off the end of the line rigidly. Other terminals, such as the Concept 100 and the Perkin Elmer Owl, make a distinction between typed and untyped blanks on the screen, shifting upon an insert or delete only to an untyped blank on the screen which is either eliminated, or expanded to two untyped blanks. You can find out which kind of terminal you have by clearing the screen and then typing text separated by cursor motions. Type “`abc def`” using local cursor motions (not spaces) between the “`abc`” and the “`def`”. Then position the cursor before the “`abc`” and put the terminal in insert mode. If typing characters causes the rest of the line to shift rigidly and characters to fall off the end, then your terminal does not distinguish between blanks and untyped positions. If the “`abc`” shifts over to the “`def`” which then move together around the end of the current line and onto the next as you insert, you have the second type of terminal, and should give the capability **in**, which stands for “insert

null”.

If your terminal does something different and unusual then you may have to modify the editor to get it to use the insert mode your terminal defines. We have seen no terminals which have an insert mode not falling into one of these two classes.

The editor can handle both terminals which have an insert mode, and terminals which send a simple sequence to open a blank position on the current line. Give as **im** the sequence to get into insert mode, or give it an empty value if your terminal uses a sequence to insert a blank position. Give as **ei** the sequence to leave insert mode (give this, with an empty value also if you gave **im** so). Now give as **ic** any sequence needed to be sent just before sending the character to be inserted. Most terminals with a true insert mode will not give **ic**, terminals which send a sequence to open a screen position should give it here. (Insert mode is preferable to the sequence to open a position on the screen if your terminal has both.) If post insert padding is needed, give this as a number of milliseconds in **ip** (a string option). Any other sequence which may need to be sent after an insert of a single character may also be given in **ip**.

It is occasionally necessary to move around while in insert mode to delete characters on the same line (e.g. if there is a tab after the insertion position). If your terminal allows motion while in insert mode you can give the capability **mi** to speed up inserting in this case. Omitting **mi** will affect only speed. Some terminals (notably Datamedia's) must not have **mi** because of the way their insert mode works.

Finally, you can specify delete mode by giving **dm** and **ed** to enter and exit delete mode, and **dc** to delete a single character while in delete mode.

Highlighting, underlining, and visible bells

If your terminal has sequences to enter and exit standout mode these can be given as **so** and **se** respectively. If there are several flavors of standout mode (such as inverse video, blinking, or underlining – half bright is not usually an acceptable “standout” mode unless the terminal is in inverse video mode constantly) the preferred mode is inverse video by itself. If the code to change into or out of standout mode leaves one or even two blank spaces on the screen, as the TVI 912 and Teleray 1061 do, then **ug** should be given to tell how many spaces are left.

Codes to begin underlining and end underlining can be given as **us** and **ue** respectively. If the terminal has a code to underline the current character and move the cursor one space to the right, such as the Microterm Mime, this can be given as **uc**. (If the underline code does not move the cursor to the right, give the code followed by a nondestructive space.)

Many terminals, such as the HP 2621, automatically leave standout mode when they move to a new line or the cursor is addressed. Programs using standout mode should exit standout mode before moving the cursor or sending a newline.

If the terminal has a way of flashing the screen to indicate an error quietly (a bell replacement) then this can be given as **vb**; it must not move the cursor. If the terminal should be placed in a different mode during open and visual modes of *ex*, this can be given as **vs** and **ve**, sent at the start and end of these modes respectively. These can be used to change,

e.g., from a underline to a block cursor and back.

If the terminal needs to be in a special mode when running a program that addresses the cursor, the codes to enter and exit this mode can be given as **ti** and **te**. This arises, for example, from terminals like the Concept with more than one page of memory. If the terminal has only memory relative cursor addressing and not screen relative cursor addressing, a one screen-sized window must be fixed into the terminal for cursor addressing to work properly.

If your terminal correctly generates underlined characters (with no special codes needed) even though it does not overstrike, then you should give the capability **ul**. If overstrikes are erasable with a blank, then this should be indicated by giving **eo**.

Keypad

If the terminal has a keypad that transmits codes when the keys are pressed, this information can be given. Note that it is not possible to handle terminals where the keypad only works in local (this applies, for example, to the unshifted HP 2621 keys). If the keypad can be set to transmit or not transmit, give these codes as **ks** and **ke**. Otherwise the keypad is assumed to always transmit. The codes sent by the left arrow, right arrow, up arrow, down arrow, and home keys can be given as **kl**, **kr**, **ku**, **kd**, and **kh** respectively. If there are function keys such as f0, f1, ..., f9, the codes they send can be given as **k0**, **k1**, ..., **k9**. If these keys have labels other than the default f0 through f9, the labels can be given as **l0**, **l1**, ..., **l9**. If there are other keys that transmit the same code as the terminal expects for the corresponding function, such as clear screen, the *termcap* 2 letter codes can be given in the **ko** capability, for example, “:ko=cl,ll,sf,sb:”, which says that the terminal has clear, home down, scroll down, and scroll up keys that transmit the same thing as the cl, ll, sf, and sb entries.

The **ma** entry is also used to indicate arrow keys on terminals which have single character arrow keys. It is obsolete but still in use in version 2 of vi, which must be run on some minicomputers due to memory limitations. This field is redundant with **kl**, **kr**, **ku**, **kd**, and **kh**. It consists of groups of two characters. In each group, the first character is what an arrow key sends, the second character is the corresponding vi command. These commands are **h** for **kl**, **j** for **kd**, **k** for **ku**, **l** for **kr**, and **H** for **kh**. For example, the mime would be **:ma=^Kj^Zk^Xl**: indicating arrow keys left (^H), down (^K), up (^Z), and right (^X). (There is no home key on the mime.)

Miscellaneous

If the terminal requires other than a null (zero) character as a pad, then this can be given as **pc**.

If tabs on the terminal require padding, or if the terminal uses a character other than ^I to tab, then this can be given as **ta**.

Hazeltine terminals, which don't allow “” characters to be printed should indicate **hz**. Datamedia terminals, which echo carriage-return linefeed for carriage return and then ignore a following linefeed should indicate **nc**. Early Concept terminals, which ignore a linefeed immediately after an **am** wrap, should indicate **xn**. If an erase-eol is required to

get rid of standout (instead of merely writing on top of it), **xs** should be given. Teleray terminals, where tabs turn all characters moved over to blanks, should indicate **xt**. Other specific terminal problems may be corrected by adding more capabilities of the form **xx**.

Other capabilities include **is**, an initialization string for the terminal, and **if**, the name of a file containing long initialization strings. These strings are expected to properly clear and then set the tabs on the terminal, if the terminal has settable tabs. If both are given, **is** will be printed before **if**. This is useful where **if** is `/usr/lib/tabset/std` but **is** clears the tabs first.

Similar Terminals

If there are two very similar terminals, one can be defined as being just like the other with certain exceptions. The string capability **tc** can be given with the name of the similar terminal. This capability must be *last* and the combined length of the two entries must not exceed 1024. Since *termlib* routines search the entry from left to right, and since the **tc** capability is replaced by the corresponding entry, the capabilities given at the left override the ones in the similar terminal. A capability can be canceled with **xx@** where **xx** is the capability. For example, the entry

```
hnl2621nl: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)`

RESTRICTIONS

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.

STATUS

TERMCAP(5) currently is not supported by Digital Equipment Corporation.

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)

STATUS

TP(5) currently is not supported by Digital Equipment Corporation.

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)`

STATUS

TTYS(5) currently is not supported by Digital Equipment Corporation.

TTYTYPE(5)

NAME

ttytype – data base of terminal types by port

SYNTAX

/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)

RESTRICTIONS

Some lines are merely known as "dialup" or "plugboard."

STATUS

TTYTYPE(5) currently is not supported by Digital Equipment Corporation.

NAME

types – primitive system data types

SYNTAX

```
#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;
```

TYPES(5)

```
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)*

STATUS

TYPES(5) currently is not supported by Digital Equipment Corporation.

NAME

utmp, wtmp – login records

SYNTAX

```
#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)*

STATUS

UTMP(5) currently is not supported by Digital Equipment Corporation.

UUENCODE(5)

NAME

uuencode – format of an encoded uuencode file

DESCRIPTION

Files output by *uuencode(1C)* consist of a header line, followed by a number of body lines, and a trailer line. *Uudecode(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

uuencode(1C), uudecode(1C), uuseend(1C), uuencode(1C), mail(1)

STATUS

UUENCODE(5) currently is not supported by Digital Equipment Corporation.

NAME

vfont – font formats for the Benson-Varian or Versatec

SYNTAX

/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)

VFONT(5)

STATUS

VFONT(5) currently is not supported by Digital Equipment Corporation.

NAME

vgrindefs – vgrind’s language definition data base

SYNTAX

/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.

```
Clc:  :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 *x* 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)

STATUS

VGRINDEF5(5) currently is not supported by Digital Equipment Corporation.

NAME

miscellaneous – miscellaneous useful information pages

DESCRIPTION

This section contains miscellaneous documentation, mostly in the area of text processing macro packages for *troff*(1).

ascii	map of ASCII character set
environ	user environment
eqnchar	special character definitions for eqn
hier	file system hierarchy
mailaddr	mail addressing description
man	macros to typeset manual pages
me	macros for formatting papers
ms	macros for formatting manuscripts
term	conventional names for terminals

ASCII(7)

NAME

ascii – map of ASCII character set

SYNTAX

cat /usr/pub/ascii

DESCRIPTION

Ascii is a map of the ASCII character set, to be printed as needed. It contains:

000 nul	001 soh	002 stx	003 etx	004 eot	005 enq	006 ack	007 bel
010 bs	011 ht	012 nl	013 vt	014 np	015 cr	016 so	017 si
020 dle	021 dc1	022 dc2	023 dc3	024 dc4	025 nak	026 syn	027 etb
030 can	031 em	032 sub	033 esc	034 fs	035 gs	036 rs	037 us
040 sp	041 !	042 "	043 #	044 \$	045 %	046 &	047 '
050 (051)	052 *	053 +	054 ,	055 -	056 .	057 /
060 0	061 1	062 2	063 3	064 4	065 5	066 6	067 7
070 8	071 9	072 :	073 ;	074 <	075 =	076 >	077 ?
100 @	101 A	102 B	103 C	104 D	105 E	106 F	107 G
110 H	111 I	112 J	113 K	114 L	115 M	116 N	117 O
120 P	121 Q	122 R	123 S	124 T	125 U	126 V	127 W
130 X	131 Y	132 Z	133 [134 \	135]	136 ^	137
140 `	141 a	142 b	143 c	144 d	145 e	146 f	147 g
150 h	151 i	152 j	153 k	154 l	155 m	156 n	157 o
160 p	161 q	162 r	163 s	164 t	165 u	166 v	167 w
170 x	171 y	172 z	173 {	174	175 }	176 ~	177 del

00 nul	01 soh	02 stx	03 etx	04 eot	05 enq	06 ack	07 bel
08 bs	09 ht	0a nl	0b vt	0c np	0d cr	0e so	0f si
10 dle	11 dc1	12 dc2	13 dc3	14 dc4	15 nak	16 syn	17 etb
18 can	19 em	1a sub	1b esc	1c fs	1d gs	1e rs	1f us
20 sp	21 !	22 "	23 #	24 \$	25 %	26 &	27 '
28 (29)	2a *	2b +	2c ,	2d -	2e .	2f /
30 0	31 1	32 2	33 3	34 4	35 5	36 6	37 7
38 8	39 9	3a :	3b ;	3c <	3d =	3e >	3f ?
40 @	41 A	42 B	43 C	44 D	45 E	46 F	47 G
48 H	49 I	4a J	4b K	4c L	4d M	4e N	4f O
50 P	51 Q	52 R	53 S	54 T	55 U	56 V	57 W
58 X	59 Y	5a Z	5b [5c \	5d]	5e ^	5f
60 `	61 a	62 b	63 c	64 d	65 e	66 f	67 g
68 h	69 i	6a j	6b k	6c l	6d m	6e n	6f o
70 p	71 q	72 r	73 s	74 t	75 u	76 v	77 w
78 x	79 y	7a z	7b {	7c	7d }	7e ~	7f del

FILES

/usr/pub/ascii

STATUS

ASCII(7) currently is not supported by Digital Equipment Corporation.

ENVIRON(7)

NAME

environ — user environment

SYNTAX

extern char **environ;

DESCRIPTION

An array of strings called the ‘environment’ is made available by *execve*(2) when a process begins. By convention these strings have the form ‘*name=value*’. The following names are used by various commands:

PATH The sequence of directory prefixes that *sh*, *time*, *nice*(1), etc., apply in searching for a file known by an incomplete path name. The prefixes are separated by ‘.’. *Login*(1) sets `PATH=:/usr/ucb:/bin:/usr/bin`.

HOME A user’s login directory, set by *login*(1) from the password file *passwd*(5).

TERM The kind of terminal for which output is to be prepared. This information is used by commands, such as *nroff* or *plot*(1G), which may exploit special terminal capabilities. See */etc/termcap* (*termcap*(5)) for a list of terminal types.

SHELL The file name of the users login shell.

TERMCAP

The string describing the terminal in **TERM**, or the name of the termcap file, see *termcap*(5), *termcap*(3X).

EXINIT A startup list of commands read by *ex*(1), *edit*(1), and *vi*(1).

USER The login name of the user.

PRINTER The name of the default printer to be used by *lpr*(1), *lpq*(1), and *lprm*(1).

Further names may be placed in the environment by the *export* command and ‘*name=value*’ arguments in *sh*(1), or by the *setenv* command if you use *csh*(1). Arguments may also be placed in the environment at the point of an *execve*(2). It is unwise to conflict with certain *sh*(1) variables that are frequently exported by ‘.profile’ files: MAIL, PS1, PS2, IFS.

SEE ALSO

csh(1), *ex*(1), *login*(1), *sh*(1), *execve*(2), *system*(3), *termcap*(3X), *termcap*(5)

STATUS

ENVIRON(7) currently is not supported by Digital Equipment Corporation.

NAME

eqnchar — special character definitions for eqn

SYNTAX

eqn /usr/pub/eqnchar [files] | **troff** [options]

neqn /usr/pub/eqnchar [files] | **nroff** [options]

DESCRIPTION

Eqnchar contains *troff* and *nroff* character definitions for constructing characters that are not available on the Graphic Systems typesetter. These definitions are primarily intended for use with *eqn* and *neqn*. It contains definitions for the following characters

<i>ciplus</i>	⊕			<i>square</i>	□
<i>citimes</i>	⊗	<i>langle</i>	<	<i>circle</i>	○
<i>wig</i>	~	<i>rangle</i>	>	<i>blot</i>	■
<i>-wig</i>	≡	<i>hbar</i>	ℏ	<i>bullet</i>	●
<i>> wig</i>	⋈	<i>ppd</i>	⊥	<i>prop</i>	⋈
<i>< wig</i>	⋈	<i><-></i>	↔	<i>empty</i>	∅
<i>= wig</i>	≡	<i><=></i>	↔	<i>member</i>	∈
<i>star</i>	*	<i><</i>	⋈	<i>nomem</i>	∉
<i>bigstar</i>	*	<i>></i>	⋈	<i>cup</i>	∪
<i>=dot</i>	⋅	<i>ang</i>	∟	<i>cap</i>	∩
<i>orsign</i>	∨	<i>rang</i>	┌	<i>incl</i>	⊆
<i>andsign</i>	∧	<i>3dot</i>	⋮	<i>subset</i>	⊂
<i>=del</i>	≠	<i>thf</i>	⋮	<i>supset</i>	⊃
<i>oppA</i>	∇	<i>quarter</i>	¼	<i>!subset</i>	⊄
<i>oppE</i>	∇	<i>3quarter</i>	¾	<i>!supset</i>	⊅
<i>angstrom</i>	Å	<i>degree</i>	°		

FILES

/usr/pub/eqnchar

SEE ALSO

troff(1), eqn(1)

STATUS

EQNCHAR(7) currently is not supported by Digital Equipment Corporation.

HIER(7)

NAME

hier -- file system hierarchy

DESCRIPTION

The following gives a quick tour through the **root file system**. Listed are the major directory hierarchy and representative files.

/ root directory for root file system

/bin directory for utility programs (see also **/usr/bin**)

as
assembler

cc
C compiler executive (see also **/lib/ccom**, **/lib/cpp**, **/lib/c2**)

csh
C shell

.

.

.

/dev directory for devices (4)

MAKEDEV
shell script to create special files

MAKEDEV.local
site specific part of **MAKEDEV**

console
main console, *tty*(4)

hp*
disks, *hp*(4)

rhp*
raw disks, *hp*(4)

ra*
UNIBUS disks *ra*(4)

tty*
terminals, *tty*(4)

.

.

.

/etc directory for maintenance utilities and data

cron
clock daemon, *cron*(8)

disktab
disk characteristics and partition tables, *disktab*(5)

dump
dump program, *dump*(8)

dumpdatesdump history for *dump*(8)**fstab**file system configuration table, *fstab*(5)**getty**part of *login*, *getty*(8)**group**group file, *group*(5)**hosts**host name to network address mapping file, *hosts*(5)**init**parent of all processes, *init*(8)**motd**message-of-the-day file, *login*(1)**mount**mount program, *mount*(8)**mtab**mounted file table, *mtab*(5)**networks**network name to network number mapping file, *networks*(5)**passwd**password file, *passwd*(5)**phones**phone numbers for remote hosts, *phones*(5)**protocols**name to number mapping file, *protocols*(5)**rc**

shell script to bring the system to multi-user mode

rc.localsite dependent portion of *rc***remote**names and description of remote hosts for *tip*(1C), *remote*(5)**services**network services definition file, *services*(5)**termcap**description of terminal capabilities, *termcap*(5)**ttys**properties of terminals, *ttys*(5)**ttytype**terminal type table, *ttytype*(5).
.
.**/lib**directory object libraries (see also */usr/lib*)

ccom
C compiler proper
cpp
C preprocessor
c2
C code improver
libc.a
system calls and standard I/O (2,3,3S)

.
.
.

/lost+found
directory for connecting detached files for *fsck*(8)
/sys
symbolic link, normally to /usr/sys
/tmp
directory for temporary files (see also /usr/tmp)
e*
used by *ed*(1)
ctm*
used by *cc*(1)
.
.
.
/usr
general purpose directory, normally on which the /usr file system is mounted
(see description below)
/vmunix kernel image

The following gives a quick tour through the **/usr file system**. Listed are the major directory hierarchy and representative files.

/usr
root directory for /usr file system
/usr/adm
directory for administrative information
crash
directory for crash dumps
vmcore.?,vmunix.?
crash dump files
lpacct
line printer accounting, *lpr*(1)
messages
hardware error messages
tracct

phototypesetter accounting, *troff*(1)
vaacct, vpacct
 varian and versatec accounting: *vpr*(1), *vtroff*(1), *pac*(8)
wtmp
 login history, *utmp*(5)
 .
 .
 .

/usr/bin directory for utility programs (keeps /bin small)
/usr/dict directory for word lists
spellhist
 history file, *spell*(1)
words
 word list, *look*(1)
 .
 .
 .

/usr/doc directories containing files for the Vol.2 documentation
as
 assembler manual
c C manual
 .
 .
 .

/usr/games
 directory for games
hangman
 hangman game
lib
 library directory for games
 .
 .
 .

/usr/guest
 directory for guest accounts

/usr/include
 directory for standard #include files
a.out.h
 object file layout, *a.out*(5)
math.h

- math (3M)
- stdio.h**
 - standard I/O, *intro*(3S)
- sys**
 - symbolic link to /sys/h (system generation #include files)
 - .
 - .
 - .

/usr/lib directory for object libraries (keeps /lib small)

- atrun**
 - system scheduler, *at*(1)
- crontab**
 - system clock daemon table
- font**
 - directory for *nroff*(1) and *troff*(1) fonts
- lint**
 - directory for utility files for *lint*(1)
- tmac**
 - directory for *troff*(1) macros
- units**
 - directory of conversion tables for *units*(1)
- uucp**
 - directory for *uucp*(1C) programs and data
 - .
 - .
 - .

/usr/man directory for unformatted and preformatted man pages

- man1**
 - directory for section 1 (unformatted)
- man2**
 - directory for section 2 (unformatted)
- man3**
 - directory for section 3 (unformatted)
 - .
 - .
 - .
- cat1**
 - directory for section 1 (preformatted)
- cat2**
 - directory for section 2 (preformatted)
- cat3**
 - directory for section 3 (preformatted)

```

)
    .
    .
    .
/usr/mdec
    directory for ULTRIX-32 boot files
/usr/msgs
    directory for messages, msgs(1)
/usr/new directory for binaries of new versions of programs
/usr/preserve
    directory for editor temp files preserved after crashes/hangups
/usr/pub directory for binaries of user programs
/usr/skel
    directory for sample user startup files
    .cshrc
        startup file for csh(1)
    .login
        login startup file for csh(1)
    .mailrc
        startup file for mail(1)
    .profile
        startup file for sh(1)
    .project
        lists information used by finger(1)
/usr/spool
    directory for delayed execution files
    at
        directory used by at(1)
    lpd
        directory used by lpr(1)
        lock
            present when line printer is active
        cf*
            copy of file to be printed, if necessary
        df*
            daemon control file, lpd(8)
        tf*
            transient control file (exists while lpr is working)
    mail
        directory of mailboxes for mail(1)
        name

```

- mail file for user *name*
- name.lock*
 - lock file (exists while *name* is receiving mail)
- uucp**
 - directory for work files and staging area for *uucp*(1C)
 - LOGFILE**
 - summary log
- /usr/src** directory for generic sources
 - usr.bin**
 - directory for user sources
 - troff**
 - directory for nroff and troff sources
 - term**
 - directory of description files for new printers
- /usr/sys** directory for system files
 - BINARY**
 - directory for system object files, *make*(1)
 - cassette**
 - directory of files for boot cassette
 - conf**
 - directory of configuration files, *config*(8)
 - data**
 - directory for drive partition tables
 - floppy**
 - directory of files for floppy disk
 - h**
 - directory for system #include files
 - mdec**
 - directory of headers for 11/750 boot blocks
 - net**
 - directory for general network files
 - netimp**
 - directory for IMP network files
 - netinet**
 - directory for DARPA internet network files
 - netpup**
 - directory for PUP network files
 - stand**
 - directory for standalone boot binaries
 - sys**
 - directory for machine dependent system files
 - vax**
 - directory for VAX specific system files

vaxif

directory of network interface drivers for the VAX

vaxmba

directory of drivers for devices on the MASSBUS

vaxuba

directory of drivers for devices on the UNIBUS

/usr/tmp symbolic link to /tmp

SEE ALSO

apropos(1), finger(1), find(1), grep(1), ls(1), ncheck(8), whatis(1), whereis(1), which(1)

STATUS

HIER(7) is supported by Digital Equipment Corporation.

MAILADDR(7)

NAME

mailaddr – mail addressing description

DESCRIPTION

Mail addresses are based on the ARPANET protocol listed at the end of this manual page. These addresses are in the general format

user@domain

where a domain is a hierarchical dot separated list of subdomains. For example, the address

eric@monet.Berkeley.ARPA

is normally interpreted from right to left: the message should go to the ARPA name tables (which do not correspond exactly to the physical ARPANET), then to the Berkeley gateway, after which it should go to the local host monet. When the message reaches monet it is delivered to the user “eric”.

Unlike some other forms of addressing, this does not imply any routing. Thus, although this address is specified as an ARPA address, it might travel by an alternate route if that was more convenient or efficient. For example, at Berkeley the associated message would probably go directly to monet over the Ethernet rather than going via the Berkeley ARPANET gateway.

Abbreviation. Under certain circumstances it may not be necessary to type the entire domain name. In general anything following the first dot may be omitted if it is the same as the domain from which you are sending the message. For example, a user on “calder.Berkeley.ARPA” could send to “eric@monet” without adding the “.Berkeley.ARPA” since it is the same on both sending and receiving hosts.

Certain other abbreviations may be permitted as special cases. For example, at Berkeley ARPANET hosts can be referenced without adding the “.ARPA” as long as their names do not conflict with a local host name.

Compatibility. Certain old address formats are converted to the new format to provide compatibility with the previous mail system. In particular,

host:user

is converted to

user@host

to be consistent with the *rcp*(1C) command.

Also, the syntax:

host!user

is converted to:

user@host.UUCP

This is normally converted back to the “host!user” form before being sent on for compatibility with older UUCP hosts.

The current implementation is not able to route messages automatically through the UUCP network. Until that time you must explicitly tell the mail system which hosts to send your message through to get to your final destination.

Case Distinctions. Domain names (i.e., anything after the “@” sign) may be given in any mixture of upper and lower case with the exception of UUCP hostnames. Most hosts accept any mixture of case in user names, with the notable exception of MULTICS sites.

Differences with ARPA Protocols. Although the UNIX addressing scheme is based on the ARPA mail addressing protocols, there are some significant differences.

At the time of this writing the only “top level” domain defined by ARPA is the “.ARPA” domain itself. This is further restricted to having only one level of host specifier. That is, the only addresses that ARPA accepts at this time must be in the format “user@host.ARPA” (where “host” is one word). In particular, addresses such as:

eric@monet.Berkeley.ARPA

are not currently legal under the ARPA protocols. For this reason, these addresses are converted to a different format on output to the ARPANET, typically:

eric%monet@Berkeley.ARPA

Route-addr. Under some circumstances it may be necessary to route a message through several hosts to get it to the final destination. Normally this routing is done automatically, but sometimes it is desirable to route the message manually. An address that shows these relays are termed “route-addr.” These use the syntax:

<@hosta,@hostb:user@hostc>

This specifies that the message should be sent to hosta, from there to hostb, and finally to hostc. This path is forced even if there is a more efficient path to hostc.

Route-addr occur frequently on return addresses, since these are generally augmented by the software at each host. It is generally possible to ignore all but the “user@host” part of the address to determine the actual sender.

Postmaster. Every site is required to have a user or user alias designated “postmaster” to which problems with the mail system may be addressed.

CSNET. Messages to CSNET sites can be sent to “user.host@UDel-Relay”.

SEE ALSO

mail(1), sendmail(8); Crocker, D. H., *Standard for the Format of Arpa Internet Text Messages*, RFC822.

STATUS

MAILADDR(7) currently is not supported by Digital Equipment Corporation.

MAN(7)

NAME

man – macros to typeset manual

SYNTAX

nroff **-man** file ...

troff **-man** file ...

DESCRIPTION

These macros are used to lay out pages of this manual. A skeleton page may be found in the file `/usr/man/man0/xx`.

Any text argument *t* may be zero to six words. Quotes may be used to include blanks in a ‘word’. If *text* is empty, the special treatment is applied to the next input line with *text* to be printed. In this way `.I` may be used to italicize a whole line, or `.SM` followed by `.B` to make small bold letters.

A prevailing indent distance is remembered between successive indented paragraphs, and is reset to default value upon reaching a non-indented paragraph. Default units for indents *i* are ens.

Type font and size are reset to default values before each paragraph, and after processing font and size setting macros.

These strings are predefined by **-man**:

`*R` ‘®’, ‘(Reg)’ in *nroff*.

`*S` Change to default type size.

FILES

`/usr/lib/tmac/tmac.an`

`/usr/man/man0/xx`

SEE ALSO

`troff(1)`, `man(1)`

RESTRICTIONS

Relative indents don’t nest.

REQUESTS

Request	Causes Break	If no Argument	Explanation
<code>.B t</code>	no	<code>t=n.t.l.*</code>	Text <i>t</i> is bold.
<code>.BI t</code>	no	<code>t=n.t.l.</code>	Join words of <i>t</i> alternating bold and italic.
<code>.BR t</code>	no	<code>t=n.t.l.</code>	Join words of <i>t</i> alternating bold and Roman.
<code>.DT</code>	no	<code>.5i li...</code>	Restore default tabs
<code>.HP i</code>	no	<code>i=p.i.</code>	Set prevailing indent to <i>i</i> . Begin paragraph with hanging indent.
<code>.I t</code>	no	<code>t=n.t.l.*</code>	Text <i>t</i> is italic.
<code>.IB t</code>	no	<code>t=n.t.l.</code>	Join words of <i>t</i> alternating italic and bold.

.IP <i>x i</i>	yes	<i>x</i> =""	Same as .TP with tag <i>x</i> .
.IR <i>t</i>	no	<i>t</i> =n.t.l.	Join words of <i>t</i> alternating italic and Roman.
.LP	yes	—	Same as .PP
.PD <i>d</i>	no	<i>d</i> =.4v	Interparagraph distance is <i>d</i> .
.PP	yes	—	Begin paragraph. Set prevailing indent to .5i.
.RE	yes	—	End of relative indent. Set prevailing indent to amount of starting .RS
.RB <i>t</i>	no	<i>t</i> =n.t.l.	Join words of <i>t</i> alternating Roman and bold.
.RI <i>t</i>	no	<i>t</i> =n.t.l.	Join words of <i>t</i> alternating Roman and italic.
.RS <i>i</i>	no	<i>i</i> =p.i.	Start relative indent, move left margin distance <i>i</i> . Set prevailing indent to .5i for nested indents.
.SH <i>t</i>	yes	<i>t</i> =n.t.l.	Subhead
.SM <i>t</i>	no	<i>t</i> =n.t.l.	Text <i>t</i> is small.
.TH <i>n c s v m</i>	yes	—	Begin page named <i>n</i> of chapter <i>c</i> ; <i>x</i> is extra commentary, e.g., 'local', for page foot center; <i>v</i> alters page foot left, e.g. '4th Berkeley Distribution'; <i>m</i> alters page head center, e.g., 'Brand X Programmer's Manual'. Set prevailing indent and tabs to .5i.
.TP <i>i</i>	yes	<i>i</i> =p.i.	Set prevailing indent to <i>i</i> . Begin indented paragraph with hanging tag given by next test line. If tag doesn't fit, place it on separate line.

* n.t.l. = next text line; p.i. = prevailing indent

STATUS

MAN(7) currently is not supported by Digital Equipment Corporation.

ME(7)

NAME

me – macros for formatting papers

SYNTAX

nroff **-me** [options] file ...

troff **-me** [options] file ...

DESCRIPTION

This package of *nroff* and *troff* macro definitions provides a canned formatting facility for technical papers in various formats. When producing 2-column output on a terminal, filter the output through *col(1)*.

The macro requests are defined below. Many *nroff* and *troff* requests are unsafe in conjunction with this package, however these requests may be used with impunity after the first *.pp*:

```
.bp      begin new page
.br      break output line here
.sp n    insert n spacing lines
.ls n    (line spacing) n=1 single, n=2 double space
.na      no alignment of right margin
.ce n    center next n lines
.ul n    underline next n lines
.sz +n   add n to point size
```

Output of the *eqn*, *neqn*, *refer*, and *tbl(1)* preprocessors for equations and tables is acceptable as input.

FILES

/usr/lib/tmac/tmac.e

/usr/lib/me/*

SEE ALSO

eqn(1), *troff(1)*, *refer(1)*, *tbl(1)*

-me Reference Manual, Eric P. Allman

Writing Papers with Nroff Using *-me*

REQUESTS

In the following list, “initialization” refers to the first *.pp*, *.lp*, *.ip*, *.np*, *.sh*, or *.uh* macro.

This list is incomplete; see *The -me Reference Manual* for interesting details.

Request	Initial Value	Cause Break	Explanation
<i>.(c</i>	-	yes	Begin centered block
<i>.(d</i>	-	no	Begin delayed text
<i>.(f</i>	-	no	Begin footnote
<i>.(l</i>	-	yes	Begin list
<i>.(q</i>	-	yes	Begin major quote
<i>.(x x</i>	-	no	Begin indexed item in index <i>x</i>
<i>.(z</i>	-	no	Begin floating keep

.jc	-	yes	End centered block
.jd	-	yes	End delayed text
.jf	-	yes	End footnote
.jl	-	yes	End list
.jq	-	yes	End major quote
.jx	-	yes	End index item
.jz	-	yes	End floating keep
.++ <i>m H</i>	-	no	Define paper section. <i>m</i> defines the part of the paper, and can be C (chapter), A (appendix), P (preliminary, e.g., abstract, table of contents, etc.), B (bibliography), RC (chapters renumbered from page one each chapter), or RA (appendix renumbered from page one).
.+c <i>T</i>	-	yes	Begin chapter (or appendix, etc., as set by .++). <i>T</i> is the chapter title.
.1c	1	yes	One column format on a new page.
.2c	1	yes	Two column format.
.EN	-	yes	Space after equation produced by <i>eqn</i> or <i>neqn</i> .
.EQ <i>x y</i>	-	yes	Precede equation; break out and add space. Equation number is <i>y</i> . The optional argument <i>x</i> may be <i>I</i> to indent equation (default), <i>L</i> to left-adjust the equation, or <i>C</i> to center the equation.
.TE	-	yes	End table.
.TH	-	yes	End heading section of table.
.TS <i>x</i>	-	yes	Begin table; if <i>x</i> is <i>H</i> table has repeated heading.
.ac <i>A N</i>	-	no	Set up for ACM style output. <i>A</i> is the Author's name(s), <i>N</i> is the total number of pages. Must be given before the first initialization.
.b <i>x</i>	no	no	Print <i>x</i> in boldface; if no argument switch to boldface.
.ba + <i>n</i>	0	yes	Augments the base indent by <i>n</i> . This indent is used to set the indent on regular text (like paragraphs).
.bc	no	yes	Begin new column
.bi <i>x</i>	no	no	Print <i>x</i> in bold italics (nofill only)
.bx <i>x</i>	no	no	Print <i>x</i> in a box (nofill only).
.ef 'x'y'z' ""	no	no	Set even footer to <i>x y z</i>
.eh 'x'y'z' ""	no	no	Set even header to <i>x y z</i>
.fo 'x'y'z' ""	no	no	Set footer to <i>x y z</i>
.hx	-	no	Suppress headers and footers on next page.
.he 'x'y'z' ""	no	no	Set header to <i>x y z</i>
.hl	-	yes	Draw a horizontal line
.i <i>x</i>	no	no	Italicize <i>x</i> ; if <i>x</i> missing, italic text follows.
.ip <i>x y</i>	no	yes	Start indented paragraph, with hanging tag <i>x</i> . Indentation is <i>y</i> ens (default 5).
.lp	yes	yes	Start left-blocked paragraph.
.lo	-	no	Read in a file of local macros of the form <i>.*x</i> . Must be given before initialization.
.np	1	yes	Start numbered paragraph.
.of 'x'y'z' ""	no	no	Set odd footer to <i>x y z</i>
.oh 'x'y'z' ""	no	no	Set odd header to <i>x y z</i>

ME(7)

.pd	-	yes	Print delayed text.
.pp	no	yes	Begin paragraph. First line indented.
.r	yes	no	Roman text follows.
.re	-	no	Reset tabs to default values.
.sc	no	no	Read in a file of special characters and diacritical marks. Must be given before initialization.
.sh <i>n x</i>	-	yes	Section head follows, font automatically bold. <i>n</i> is level of section, <i>x</i> is title of section.
.sk	no	no	Leave the next page blank. Only one page is remembered ahead.
.sz <i>+n</i>	10p	no	Augment the point size by <i>n</i> points.
.th	no	no	Produce the paper in thesis format. Must be given before initialization.
.tp	no	yes	Begin title page.
.u <i>x</i>	-	no	Underline argument (even in <i>troff</i>). (Nofill only).
.uh	-	yes	Like .sh but unnumbered.
.xp <i>x</i>	-	no	Print index <i>x</i> .

STATUS

ME(7) currently is not supported by Digital Equipment Corporation.

NAME

ms — text formatting macros

SYNTAX

nroff -ms [options] file ...

troff -ms [options] file ...

DESCRIPTION

This package of *nroff* and *troff* macro definitions provides a formatting facility for various styles of articles, theses, and books. When producing 2-column output on a terminal or lineprinter, or when reverse line motions are needed, filter the output through *col*(1). All external -ms macros are defined below. Many *nroff* and *troff* requests are unsafe in conjunction with this package. However, the first four requests below may be used with impunity after initialization, and the last two may be used even before initialization:

```
.bp    begin new page
.br    break output line
.sp n  insert n spacing lines
.ce n  center next n lines
.ls n  line spacing: n=1 single, n=2 double space
.na    no alignment of right margin
```

Font and point size changes with `\f` and `\s` are also allowed; for example, “`\f Iword \fR`” will italicize *word*. Output of the *tbl*, *eqn*, and *refer*(1) preprocessors for equations, tables, and references is acceptable as input.

FILES

/usr/lib/tmac/tmac.x

/usr/lib/ms/x.???

SEE ALSO

eqn(1), *refer*(1), *tbl*(1), *troff*(1)

REQUESTS

Macro Name	Initial Value	Break? Reset?	Explanation
.AB <i>x</i>	—	y	begin abstract; if <i>x</i> =no don't label abstract
.AE	—	y	end abstract
.AI	—	y	author's institution
.AM	—	n	better accent mark definitions
.AU	—	y	author's name
.B <i>x</i>	—	n	embolden <i>x</i> ; if no <i>x</i> , switch to boldface
.B1	—	y	begin text to be enclosed in a box
.B2	—	y	end boxed text and print it
.BT	date	n	bottom title, printed at foot of page
.BX <i>x</i>	—	n	print word <i>x</i> in a box
.CM	if t	n	cut mark between pages
.CT	—	y,y	chapter title: page number moved to CF (TM only)
.DA <i>x</i>	if n	n	force date <i>x</i> at bottom of page; today if no <i>x</i>

.DE	-	y	end display (unfilled text) of any kind
.DS <i>x y</i>	I	y	begin display with keep; <i>x</i> =I,L,C,B; <i>y</i> =indent
.ID <i>y</i>	8n,5i	y	indented display with no keep; <i>y</i> =indent
.LD	-	y	left display with no keep
.CD	-	y	centered display with no keep
.BD	-	y	block display; center entire block
.EF <i>x</i>	-	n	even page footer <i>x</i> (3 part as for .tl)
.EH <i>x</i>	-	n	even page header <i>x</i> (3 part as for .tl)
.EN	-	y	end displayed equation produced by <i>eqn</i>
.EQ <i>x y</i>	-	y	break out equation; <i>x</i> =L,I,C; <i>y</i> =equation number
.FE	-	n	end footnote to be placed at bottom of page
.FP	-	n	numbered footnote paragraph; may be redefined
.FS <i>x</i>	-	n	start footnote; <i>x</i> is optional footnote label
.HD	undef	n	optional page header below header margin
.I <i>x</i>	-	n	italicize <i>x</i> ; if no <i>x</i> , switch to italics
.IP <i>x y</i>	-	y,y	indented paragraph, with hanging tag <i>x</i> ; <i>y</i> =indent
.IX <i>x y</i>	-	y	index words <i>x y</i> and so on (up to 5 levels)
.KE	-	n	end keep of any kind
.KF	-	n	begin floating keep; text fills remainder of page
.KS	-	y	begin keep; unit kept together on a single page
.LG	-	n	larger; increase point size by 2
.LP	-	y,y	left (block) paragraph.
.MC <i>x</i>	-	y,y	multiple columns; <i>x</i> =column width
.ND <i>x</i>	if t	n	no date in page footer; <i>x</i> is date on cover
.NH <i>x y</i>	-	y,y	numbered header; <i>x</i> =level, <i>x</i> =0 resets, <i>x</i> =S sets to <i>y</i>
.NL	10p	n	set point size back to normal
.OF <i>x</i>	-	n	odd page footer <i>x</i> (3 part as for .tl)
.OH <i>x</i>	-	n	odd page header <i>x</i> (3 part as for .tl)
.P1	if TM	n	print header on 1st page
.PP	-	y,y	paragraph with first line indented
.PT	- -	n	page title, printed at head of page
.PX <i>x</i>	-	y	print index (table of contents); <i>x</i> =no suppresses title
.QP	-	y,y	quote paragraph (indented and shorter)
.R	on	n	return to Roman font
.RE	5n	y,y	retreat: end level of relative indentation
.RP <i>x</i>	-	n	released paper format; <i>x</i> =no stops title on 1st page
.RS	5n	y,y	right shift: start level of relative indentation
.SH	-	y,y	section header, in boldface
.SM	-	n	smaller; decrease point size by 2
.TA	8n,5n	n	set tabs to 8n 16n ... (nroff) 5n 10n ... (troff)
.TC <i>x</i>	-	y	print table of contents at end; <i>x</i> =no suppresses title
.TE	-	y	end of table processed by <i>tbl</i>
.TH	-	y	end multi-page header of table
.TL	-	y	title in boldface and two points larger

.TM	off	n	UC Berkeley thesis mode
.TS <i>x</i>	-	y,y	begin table; if <i>x</i> =H table has multi-page header
.UL <i>x</i>	-	n	underline <i>x</i> , even in <i>troff</i>
.UX <i>x</i>	-	n	UNIX; trademark message first time; <i>x</i> appended
.XA <i>x y</i>	-	y	another index entry; <i>x</i> =page or no for none; <i>y</i> =indent
.XE	-	y	end index entry (or series of .IX entries)
.XP	-	y,y	paragraph with first line exdented, others indented
.XS <i>x y</i>	-	y	begin index entry; <i>x</i> =page or no for none; <i>y</i> =indent
.1C	on	y,y	one column format, on a new page
.2C	-	y,y	begin two column format
.]-	-	n	beginning of <i>refer</i> reference
.[0	-	n	end of unclassifiable type of reference
.[N	-	n	N= 1:journal-article, 2:book, 3:book-article, 4:report

REGISTERS

Formatting distances can be controlled in `-ms` by means of built-in number registers. For example, this sets the line length to 6.5 inches:

```
.nr LL 6.5i
```

Here is a table of number registers and their default values:

Name	Register Controls	Takes Effect	Default
PS	point size	paragraph	10
VS	vertical spacing	paragraph	12
LL	line length	paragraph	6i
LT	title length	next page	same as LL
FL	footnote length	next .FS	5.5i
PD	paragraph distance	paragraph	1v (if n), .3v (if t)
DD	display distance	displays	1v (if n), .5v (if t)
PI	paragraph indent	paragraph	5n
QI	quote indent	next .QP	5n
FI	footnote indent	next .FS	2n
PO	page offset	next page	0 (if n), ~1i (if t)
HM	header margin	next page	1i
FM	footer margin	next page	1i
FF	footnote format	next .FS	0 (1, 2, 3 available)

When resetting these values, make sure to specify the appropriate units. Setting the line length to 7, for example, will result in output with one character per line. Setting FF to 1 suppresses footnote superscripting; setting it to 2 also suppresses indentation of the first line; and setting it to 3 produces an .IP-like footnote paragraph.

Here is a list of string registers available in `-ms`; they may be used anywhere in the text:

Name	String's Function
*Q	quote (" in <i>nroff</i> , " in <i>troff</i>)
*U	unquote (" in <i>nroff</i> , " in <i>troff</i>)
*-	dash (-- in <i>nroff</i> , — in <i>troff</i>)
*(MO	month (month of the year)
*(DY	day (current date)
**	automatically numbered footnote
*'	acute accent (before letter)
*'	grave accent (before letter)
*	circumflex (before letter)
*,	cedilla (before letter)
*:	umlaut (before letter)
*	tilde (before letter)

When using the extended accent mark definitions available with `.AM`, these strings should come after, rather than before, the letter to be accented.

RESTRICTIONS

Floating keeps and regular keeps are diverted to the same space, so they cannot be mixed together with predictable results.

STATUS

MS(7) currently is not supported by Digital Equipment Corporation.

TERM(7)

NAME

`term` – conventional names for terminals

DESCRIPTION

Certain commands use these terminal names. They are maintained as part of the shell environment (see `sh(1)`, `environ(7)`).

<code>adm3a</code>	Lear Seigler Adm-3a
<code>2621</code>	Hewlett-Packard HP262? series terminals
<code>hp</code>	Hewlett-Packard HP264? series terminals
<code>c100</code>	Human Designed Systems Concept 100
<code>h19</code>	Heathkit H19
<code>mime</code>	Microterm mime in enhanced ACT IV mode
<code>1620</code>	DIABLO 1620 (and others using HyType II)
<code>300</code>	DASI/DTC/GSI 300 (and others using HyType I)
<code>33</code>	TELETYPE® Model 33
<code>37</code>	TELETYPE Model 37
<code>43</code>	TELETYPE Model 43
<code>735</code>	Texas Instruments TI735 (and TI725)
<code>745</code>	Texas Instruments TI745
<code>dumb</code>	terminals with no special features
<code>dialup</code>	a terminal on a phone line with no known characteristics
<code>network</code>	a terminal on a network connection with no known characteristics
<code>4014</code>	Tektronix 4014
<code>vt52</code>	Digital Equipment Corp. VT52

The list goes on and on. Consult `/etc/termcap` (see `termcap(5)`) for an up-to-date and locally correct list.

Commands whose behavior may depend on the terminal either consult `TERM` in the environment, or accept arguments of the form `-Tterm`, where `term` is one of the names given above.

SEE ALSO

`stty(1)`, `tabs(1)`, `plot(1G)`, `sh(1)`, `environ(7)` `ex(1)`, `clear(1)`, `more(1)`, `ul(1)`, `tset(1)`, `termcap(5)`, `termcap(3X)`, `ttytype(5)`
`troff(1)` for `nroff`

STATUS

`TERM(7)` currently is not supported by Digital Equipment Corporation.

HOW TO ORDER ADDITIONAL DOCUMENTATION

DIRECT TELEPHONE ORDERS

In Continental USA
and Puerto Rico
call **800-258-1710**

In Canada
call **800-267-6146**

In New Hampshire,
Alaska or Hawaii
call **603-884-6660**

DIRECT MAIL ORDERS (U.S. and Puerto Rico*)

DIGITAL EQUIPMENT CORPORATION
P.O. Box CS2008
Nashua, New Hampshire 03061

DIRECT MAIL ORDERS (Canada)

DIGITAL EQUIPMENT OF CANADA LTD.
940 Belfast Road
Ottawa, Ontario, Canada K1G 4C2
Attn: A&SG Business Manager

INTERNATIONAL

DIGITAL EQUIPMENT CORPORATION
A&SG Business Manager
c/o Digital's local subsidiary
or approved distributor

Internal orders should be placed through the Software Distribution Center (SDC), Digital Equipment Corporation, Northboro, Massachusetts 01532

*Any prepaid order from Puerto Rico must be placed
with the Local Digital Subsidiary:
809-754-7575

Reader's Comments

Note: This form is for document comments only. DIGITAL will use comments submitted on this form at the company's discretion. If you require a written reply and are eligible to receive one under Software Performance Report (SPR) service, submit your comments on an SPR form.

Did you find this manual understandable, usable, and well-organized? Please make suggestions for improvement. _____

Did you find errors in this manual? If so, specify the error and the page number.

Please indicate the type of user/reader that you most nearly represent.

- Assembly language programmer
- Higher-level language programmer
- Occasional programmer (experienced)
- User with little programming experience
- Student programmer
- Other (please specify) _____

Name _____ Date _____

Organization _____

Street _____

City _____ State _____ Zip Code _____
or _____
Country _____

-----Do Not Tear - Fold Here and Tape-----

digital



No Postage
Necessary
if Mailed in the
United States



BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO 33 MAYNARD MASS.

POSTAGE WILL BE PAID BY ADDRESSEE

Documentation Manager
ULTRIX-32™ Documentation Group
MK02-1/H10
Continental Blvd.
Merrimack, N.H.
03054

-----Do Not Tear - Fold Here and Tape-----

Cut Along Dotted Line

Notes:

Notes:

Notes:

Notes:

Notes:

Notes:

Notes:

Notes:

digital™