

NAME

read – read from a file descriptor

SYNOPSIS

```
#include <unistd.h>
```

```
ssize_t read(int fd, void *buf, size_t count);
```

DESCRIPTION

read() attempts to read up to *count* bytes from file descriptor *fd* into the buffer starting at *buf*.

If *count* is zero, **read()** returns zero and has no other results. If *count* is greater than `SSIZE_MAX`, the result is unspecified.

RETURN VALUE

On success, the number of bytes read is returned (zero indicates end of file), and the file position is advanced by this number. It is not an error if this number is smaller than the number of bytes requested; this may happen for example because fewer bytes are actually available right now (maybe because we were close to end-of-file, or because we are reading from a pipe, or from a terminal), or because **read()** was interrupted by a signal. On error, `-1` is returned, and *errno* is set appropriately. In this case it is left unspecified whether the file position (if any) changes.

ERRORS**EINTR**

The call was interrupted by a signal before any data was read.

EAGAIN

Non-blocking I/O has been selected using **O_NONBLOCK** and no data was immediately available for reading.

EIO I/O error. This will happen for example when the process is in a background process group, tries to read from its controlling tty, and either it is ignoring or blocking `SIGTTIN` or its process group is orphaned. It may also occur when there is a low-level I/O error while reading from a disk or tape.

EISDIR

fd refers to a directory.

EBADF

fd is not a valid file descriptor or is not open for reading.

EINVAL

fd is attached to an object which is unsuitable for reading.

EFAULT

buf is outside your accessible address space.

Other errors may occur, depending on the object connected to *fd*. POSIX allows a **read** that is interrupted after reading some data to return `-1` (with *errno* set to `EINTR`) or to return the number of bytes already read.

CONFORMING TO

SVr4, SVID, AT&T, POSIX, X/OPEN, BSD 4.3

RESTRICTIONS

On NFS file systems, reading small amounts of data will only update the time stamp the first time, subsequent calls may not do so. This is caused by client side attribute caching, because most if not all NFS clients leave atime updates to the server and client side reads satisfied from the client's cache will not cause atime updates on the server as there are no server side reads. UNIX semantics can be obtained by disabling client side attribute caching, but in most situations this will substantially increase server load and decrease performance.

Many filesystems and disks were considered to be fast enough that the implementation of **O_NONBLOCK** was deemed unnecessary. So, **O_NONBLOCK** may not be available on files and/or disks.

SEE ALSO

close(2), fcntl(2), ioctl(2), lseek(2), readdir(2), readlink(2), select(2), write(2), fread(3), readv(3)