NAME

mount_nfs – mount remote NFS resources

SYNOPSIS

mount [**–F nfs**] [**–r**] [**–m**] [**–o** *specific_options*] [**–O**] *resource* | *mount_point*

DESCRIPTION

mount attaches a named *resource* to the file system hierarchy at the pathname location *mount_point*, which must already exist. If *mount_point* has any contents prior to the **mount** operation, the contents remain hidden until the *resource* is once again unmounted.

If the resource is listed in the **/etc/vfstab** file, the command line can specify either *resource* or *mount_point*, and **mount** will consult **/etc/vfstab** for more information. If the **–F** option is omitted, **mount** takes the file system type from **/etc/vfstab**.

mount maintains a table of mounted file systems in **/etc/mnttab**, described in **mnttab**(4).

OPTIONS

–r Mount the specified file system read-only.

–m Do not append an entry to the **/etc/mnttab** table of mounted file systems

–o *specific_options*
Set file system specific options according to a comma-separated list chosen from words below.

rw | **ro** *resource* is mounted read-write or read-only. The default is **rw**.

suid | **nosuid** Setuid execution allowed or disallowed. The default is **suid**.

remount If a file system is mounted read-only, remounts the file system read-write.

bg | **fg** If the first attempt fails, retry in the background, or, in the foreground. The default is **fg**.

quota Enables **quota**(1M) to check whether the user is over quota on this file system; if the file system has quotas enabled on the server, quotas will still be checked for operations on this file system.

noquota Prevent **quota**(1M) from checking whether the user exceeded the quota on this file system; if the file system has quotas enabled on the server, quotas will still be checked for operations on this file system.

retry=*n* The number of times to retry the mount operation. The default is **10000**.

vers=<*NFS version number***>**
By default, the version of NFS protocol used between the client and the server is the highest one available on both systems. If the NFS server does not support NFS Version 3, then the NFS mount will use NFS Version 2

proto=<*netid***>** **<***netid***>** is a value of network_id field from entry in the **/etc/netconfig** file. By default, the transport protocol used for the NFS mount will be first available connection oriented transport supported on both the client and the server. If no connection oriented transport is found, then the first available connectionless transport is used. This default behavior can be overridden with the **proto=<***netid***>** option.

port=*n* The server IP port number. The default is **NFS_PORT**.

grpid By default, the GID associated with a newly created file will obey the System V semantics; that is, the GID is set to the effective GID of the calling process. This behavior may be overridden on a per-directory basis by setting the set-GID bit of the parent directory; in this case, the GID of a newly created file is set to the GID of the parent directory (see **open**(2) and **mkdir**(2)). Files created on file systems that are mounted with the **grpid** option will obey BSD semantics independent of whether the set-GID bit of the parent directory is set; that is, the GID is unconditionally inherited from that of the parent directory.

rsize=*n* Set the read buffer size to *n* bytes. The default value is **32768**.

wsize=*n* Set the write buffer size to *n* bytes. The default value is **32768**.

timeo=*n* Set the NFS timeout to *n* tenths of a second. The default value is **11**.

retrans=*n* Set the number of NFS retransmissions to *n*. The default value is **5**.

soft | hard Return an error if the server does not respond, or continue the retry request until the server responds. The default value is **hard**.

intr | nointr Allow (do not allow) keyboard interrupts to kill a process that is hung while waiting for a response on a hard-mounted file system. The default is **intr**.

secure Use DES authentication for NFS transactions.

posix Request POSIX.1 semantics for the file system. Requires a mount version 2 **mountd**(1M) on the server.

kerberos Use Kerberos authentication for NFS transactions.

noac Suppress attribute caching.

acdirmax=*n* Hold cached attributes for no more than *n* seconds after directory update. The default value is **60**.

acdirmin=*n* Hold cached attributes for at least *n* seconds after directory update. The default value is **30**.

acregmax=*n* Hold cached attributes for no more than *n* seconds after file modification. The default value is **60**.

acregmin=*n* Hold cached attributes for at least *n* seconds after file modification. The default value is **3**.

actimeo=*n* Set *min* and *max* times for regular files and directories to *n* seconds.

–O Overlay mount. Allow the file system to be mounted over an existing mount point, making the underlying file system inaccessible. If a mount is attempted on a pre-existing mount point without setting this flag, the mount will fail, producing the error “device busy”.

NFS FILE SYSTEMS

Background vs. Foreground

File systems mounted with the **bg** option indicate that **mount** is to retry in the background if the server's mount daemon (**mountd**(1M)) does not respond. **mount** retries the request up to the count specified in the **retry=***n* option. Once the file system is mounted, each NFS request made in the kernel waits **timeo=***n* tenths of a second for a response. If no response arrives, the time-out is multiplied by **2** and the request is retransmitted. When the number of retransmissions has reached the number specified in the **retrans=***n* option, a file system mounted with the **soft** option returns an error on the request; one mounted with the **hard** option prints a warning message and continues to retry the request.

Hard vs. Soft

File systems that are mounted read-write or that contain executable files should always be mounted with the **hard** option. Applications using **soft** mounted file systems may incur unexpected I/O errors.

Authenticated Requests

The server may require authenticated NFS requests from the client. Either **secure** or **kerberos** authentication may be required.

File Attributes

To improve NFS read performance, files and file attributes are cached. File modification times get updated whenever a write occurs. However, file access times may be temporarily out-of-date until the cache gets refreshed.

The attribute cache retains file attributes on the client. Attributes for a file are assigned a time to be flushed. If the file is modified before the flush time, then the flush time is extended by the time since the last modification (under the assumption that files that changed recently are likely to change soon). There is a minimum and maximum flush time extension for regular files and for directories. Setting **actimeo=***n* sets flush time to *n* seconds for both regular files and directories.

EXAMPLES

To mount an NFS file system:

example# mount serv:/usr/src /usr/src

To mount an NFS file system readonly with no suid privileges:

example# mount –r –o nosuid serv:/usr/src /usr/src

To mount an NFS file system over Version 2, with the UDP transport:

example# mount -o vers=2,proto=udp serv:/usr/src /usr/src

FILES

/etc/mnttab	table of mounted file systems
/etc/dfs/fstypes	default distributed file system type
/etc/vfstab	table of automatically mounted resources

SEE ALSO

mountall(1M), **mountd**(1M), **quota**(1M), **mkdir**(2), **mount**(2), **open**(2), **umount**(2), **mnttab**(4)

NOTES

A NFS server should not attempt to mount its own filesystems (see **lofs**(7FS)).

If the directory on which a file system is to be mounted is a symbolic link, the file system is mounted on *the directory to which the symbolic link refers*, rather than being mounted on top of the symbolic link itself.

SunOS 4.X used the **biod** maintenance procedure to perform parallel read-ahead and write-behind on NFS clients. SunOS 5.X obsoleted **biod** with multi-threaded processing, which transparently performs parallel read-ahead and write-behind.