

NAME

nfsuserd - load user and group information into the kernel for NFSv4 services plus support manage-gids for all NFS versions

SYNOPSIS

nfsuserd [-**domain** *domain_name*] [-**usertimeout** *minutes*] [-**usermax** *max_cache_size*] [-**verbose**] [-**force**] [-**manage-gids**] [*num_servers*]

DESCRIPTION

nfsuserd loads user and group information into the kernel for NFSv4. For Kerberized NFSv4 mounts, it must be running on both client(s) and server for correct operation. For non-Kerberized NFSv4 mounts, this daemon must be running unless all client(s) plus the server are configured to put uid/gid numbers in the owner and owner_group strings.

It also provides support for manage-gids and must be running on the server if this is being used for any version of NFS.

Upon startup, it loads the machine's DNS domain name, plus timeout and cache size limit into the kernel. It then preloads the cache with group and user information, up to the cache size limit and forks off *num_servers* (default 4) children which are the servers that service requests from the kernel for cache misses. The master is there for the sole purpose of terminating the servers. To stop the **nfsuserd**, send a SIGUSR1 to the master.

The following options are available:

-domain *domain_name*

This option allows you to override the default DNS domain name, which is acquired by taking either the suffix on the machine's hostname or, if that name is not a fully qualified host name, the canonical name as reported by `getaddrinfo(3)`.

-usertimeout *minutes*

Overrides the default timeout for cache entries, in minutes. The longer the time out, the better the performance, but the longer it takes for replaced entries to be seen. If your user/group database management system almost never re-uses the same names or id numbers, a large timeout is recommended. The default is 1 minute.

-usermax *max_cache_size*

Overrides the default upper bound on the cache size. The larger the cache, the more kernel memory is used, but the better the performance. If your system can afford the memory use, make this the sum of the number of entries in your group and password databases. The default

is 200 entries.

-verbose

When set, the server logs a bunch of information to syslog.

-force This flag option must be set to restart the daemon after it has gone away abnormally and refuses to start, because it thinks nfsuserd is already running.

-manage-gids

This flag enables manage-gids for the NFS server nfsd(8). When this is enabled, all NFS requests using AUTH_SYS authentication take the uid from the RPC request and uses the group list for that uid provided by getgrouplist(3) on the server instead of the list of groups provided in the RPC authenticator. This can be used to avoid the 16 group limit for AUTH_SYS.

num_servers

Specifies how many servers to create (max 20). The default of 4 may be sufficient. You should run enough servers, so that ps(1) shows almost no running time for one or two of the servers after the system has been running for a long period. Running too few will have a major performance impact, whereas running too many will only tie up some resources, such as a process table entry and swap space.

SEE ALSO

getgrent(3), getgrouplist(3), getpwent(3), nfsv4(4), group(5), passwd(5), nfsd(8)

HISTORY

The **nfsuserd** utility was introduced with the NFSv4 experimental subsystem in 2009.

BUGS

The **nfsuserd** use getgrent(3), getgrouplist(3) and getpwent(3) library calls to resolve requests and will hang if the servers handling those requests fail and the library functions don't return. See group(5) and passwd(5) for more information on how the databases are accessed.