NAME

nfsdcltrack - NFSv4 Client Tracking Callout Program

SYNOPSIS

nfsdcltrack [-d] [-f] [-s stable storage dir] <command> <args...>

DESCRIPTION

nfsdcltrack is the NFSv4 client tracking callout program. It is not necessary to install this program on machines that are not acting as NFSv4 servers.

When a network partition is combined with a server reboot, there are edge conditions that can cause the server to grant lock reclaims when other clients have taken conflicting locks in the interim. A more detailed explanation of this issue is described in RFC 3530, section 8.6.3 and in RFC 5661, section 8.4.3.

In order to prevent these problems, the server must track a small amount of per-client information on stable storage. This program provides the userspace piece of that functionality. When the kernel needs to manipulate the database that stores this info, it will execute this program to handle it.

OPTIONS

-d, **--debug**
Enable debug level logging.

-f, **--foreground**
Log to stderr instead of syslog.

-s *storagedir*, **--storagedir**=*storage_dir*
Directory where stable storage information should be kept. The default value is */var/lib/nfs/nfsdcltrack*.

COMMANDS

nfsdcltrack requires a command for each invocation. Supported commands are:

init Initialize the database. This command requires no argument.

create
Create a new client record (or update the timestamp on an existing one). This command requires a hex-encoded nfs_client_id4 as an argument.

remove
Remove a client record from the database. This command requires a hex-encoded nfs_client_id4 as an argument.

check
Check to see if a nfs_client_id4 is allowed to reclaim. This command requires a hex-encoded nfs_client_id4 as an argument.

gracedone
Remove any unreclaimed client records from the database. This command requires a epoch boot time as an argument.

LEGACY TRANSITION MECHANISM

The Linux kernel NFSv4 server has historically tracked this information on stable storage by manipulating information on the filesystem directly, in the directory to which */proc/fs/nfsd/nfsv4recoverydir* points. If the kernel passes the correct information, then nfsdcltrack can use it to allow a seamless transition from the old client tracking scheme to the new one.

On a **check** operation, if there is no record of the client in the database, nfsdcltrack will look to see if the **NFSDCLTRACK_LEGACY_RECDIR** environment variable is set. If it is, then it will fetch that value and see if a directory exists by that name. If it does, then the check operation will succeed and the directory will be removed.

On a **gracedone** operation, nfsdcltrack will look to see if the **NFSDCLTRACK_LEGACY_TOPDIR** environment variable is set. If it is, then it will attempt to clean out that directory prior to exiting.

Note that this transition is one-way. If the machine subsequently reboots back into an older kernel that does

not support the nfsdcltrack upcall then the clients will not be able to recover their state.

NOTES

This program requires a kernel that supports the nfsdcltrack usermodehelper upcall. This support was first added to mainline kernels in 3.8.

AUTHORS

nfsdcltrack was developed by Jeff Layton <jlayton@redhat.com>.