

# *NFS*

The NFS module implements the client side of the Network File System, using RPC and running on top of UDP. The process requires no configuration. However, it does assume the presence of both a DNS and a NIS module in the system.

## **Caveat**

In its current implementation, the NFS code does not follow symbolic links automatically, because of the various possibilities of remote-mount points across systems. Also, although the *DEV\_LIST* message is supported and returns a list of mount points, the device attributes of the mount points are not returned in the reply.

## **Process Information**

Prototype Name	nfs
Link Order	does not matter
Process Name	“nfs”

## **Process Operation**

The module has only a main process. It accepts the *Standard* messageset for dataflow communications and the *Files* messageset for filing system operations.

## **Protocol Design**

The following table shows how the NFS protocol elements are mapped to the corresponding shared library routines and ROME messages.

command	level	stdlib	message	ROME routine
GETATTR	[2,3]	<i>stat</i>	FSLIB_ATTR_GET	<i>fslib_get_attr</i>
SETATTR	[2,3]	<i>chown, chmod</i>	FSLIB_ATTR_SET	<i>fslib_set_attr</i>
ROOT	[2]			__ <sup>1</sup>
LOOKUP	[2,3]			<i>fopen</i> <sup>2</sup>
ACCESS	[3]	<i>access</i>		__ <sup>3</sup>
READLINK	[2,3]	<i>readlink</i>	FSLIB_READLINK	<i>readlink</i>
READ	[2,3]	<i>read</i>	<i>Standard</i>	<i>fread etc.</i>
WRITECACHE	[2]			__ <sup>4</sup>
WRITE	[2,3]	<i>write</i>	<i>Standard</i>	<i>fwrite etc.</i>
CREATE	[2,3]	<i>open</i>	OPEN	<i>fopen</i> <sup>2</sup>
REMOVE	[2,3]	<i>unlink</i>	FSLIB_DELETE	<i>delete</i>
RENAME	[2,3]	<i>rename</i>	FSLIB_RENAME	<i>rename</i>
LINK	[2,3]	<i>link</i>	FSLIB_LINK	<i>link</i>
SYMLINK	[2,3]	<i>symlink</i>	FSLIB_LINK	<i>link</i>
MKDIR	[2,3]	<i>mkdir</i>	FSLIB_MKDIR	<i>mkdir</i>
RMDIR	[2,3]	<i>rmdir</i>	FSLIB_RMDIR	<i>rmdir</i>
READDIR	[2,3]		FSLIB_READDIR	<i>readdir</i> <sup>5</sup>
READDIRPLUS	[3]		FSLIB_READDIR	<i>readdir</i> <sup>5</sup>
STATFS	[3]		FSLIB_DEV_ATTR	<i>fslib_get_device_attr</i> <sup>6</sup>
FSSTAT	[3]		FSLIB_DEV_ATTR	<i>fslib_get_device_attr</i> <sup>6</sup>
FSINFO	[3]		FSLIB_DEV_ATTR	<i>fslib_get_device_attr</i> <sup>6</sup>
MKNOD	[3]	<i>mknod</i>		__ <sup>7</sup>
PATHCONF	[3]	<i>fstat64</i>	FSLIB_READDIR	<i>readdir</i> <sup>5</sup>
COMMIT	[3]	<i>sync</i>		__ <sup>8</sup>

## Notes

1. The ROOT operation is obsolete and has been replaced by the *MOUNT* protocol, which is invoked as needed during *OPEN* processing.
2. The NFS server is (almost) stateless with no explicit *OPEN* or *CLOSE* operations. The *LOOKUP* and *CREATE* operations are performed during *OPEN* processing, according to the read or write options on the *OPEN* message.
3. There is no direct support for the *access* routine, the information can be derived from the directory-level operations.
4. The *WRITECACHE* operation is obsolete.
5. The generic *readdir* operation is a superset of the level-2 and level-3 functionality.
6. The generic device interrogation interface combines these three routines.
7. The current ROME implementation does not support the creation of special device nodes on the remote machine
8. The NFS level-3 caching operations are not supported.

## Authorisation

Although some operations (for example reading the export list from a server) can be performed without any authentication, more NFS RPCs require authorisation. In the current implementation of the NFS client, the intention is that the applications all run under the same *UNIX* userid and group, these values must be set globally before any of the NFS operations are invoked. This is done by calling the *nfs\_authorize* routine.

## URL Processing

Handling file names is particularly complicated in NFS because of the lack of state and the various possibilities for mounting and accessing remote filing systems. The URL passed on the *OPEN* message contains a full path to a file, for example:

nfs://anlicus.nec-lab.com/export/home/ljf/max

or

nfs://~ljf/max

which must be converted into an NFS handle for subsequent operations. The first stage is to resolve usernames into file-system mount points. Firstly the *passwd.byname* map is used to retrieve the user's home directory. Then the *auto\_home* map resolves the automount point to a host and filing system. After these operations, all filenames are in the first format given above, with an explicit hostname and path.

The next stage is to find the appropriate mount point on the remote filesystem. If this is the first time a file has been accessed from this host, the *MOUNTPROG* service is used to read the list of exported filing systems on that machine, and the returned list is cached for future reference. Once the list is known, it is scanned to find a matching directory prefix with the supplied path.

If the directory has not previously been accessed the *MNT* procedure of the mount protocol is used to retrieve the handle for the mount point. This leaves a path and terminal filename relative to a known file handle on the remote system. For example:

nfs://anlicus.nec-lab.com/export/home/ljf/max = File-Handle + /ljf/ + max

The path is resolved into a sequence of File Handles using the *LOOKUP* procedure of the main *NFS* protocol, and checking that each returned handle is to a directory. Finally, the last component is resolved relative to its containing directory. If this file does not exist, and the *OPEN* was for writing, a new file is made with the *CREATE* procedure on the remote system, otherwise an error is returned.

The expected result of this processing is to generate a File-Handle for the destination file and have a datapath open to the remote NFS server for subsequent operations.

## Messages

**CLOSE** messages caused any buffered output to be written to the remote device, and the data structures associated with the file are freed.

**FETMBLK/GETMBLK** messages cause internally-buffered data to be transferred to the application's buffer or the *READ* procedure to be invoked remotely to transfer data to the local machine. *GETMBLK* replies contain the returns from remote procedure calls and are handled within the RPC layer, as described in the RPC module documentation.

**FLUSH** messages force any buffered output to be written to the remote file.

**FSLIB\_ATTR\_GET** messages return the filing system attributes for the URL passed in the *name* parameter of the message.

FSLIB\_ATTR\_SET messages set the attributes for the URL passed in the *name* parameter. Currently only the *mode*, *uid* and *gid* parameters are updated.

FSLIB\_DELETE messages first locate the file handle using the same method as for *OPEN* and then issue the *REMOVE* procedure on that handle.

FSLIB\_DEV\_ATTR messages return the device attributes for the file system specified in the *devspec* parameter. The values returned are for the mount point on which the file resides.

FSLIB\_DEV\_LIST messages return a list of the mount points for the hostname specified in the *devspec* parameter.

FSLIB\_LINK messages create either hard or symbolic links in the filing system.

FSLIB\_MKDIR messages first locate the parent directory file handle using the same method as for *OPEN* and then issue the *MKDIR* procedure on the final component.

FSLIB\_READDIR messages use the *READDIR* procedure to retrieve directory entries from the remote directory. The ‘cookie’ value is saved between calls to allow multiple calls to retrieve a list of entries in sequence.

FSLIB\_RENAME messages use the *RENAME* procedure to move the file. Renaming a file across different directories on the same device is supported.

FSLIB\_RMDIR messages first locate the parent directory file handle using the same method as for *OPEN* and then issue the *RMDIR* procedure on the final component.

FSLIB\_SEEK messages set the file position which will be used to read or write data on subsequent calls.

FSLIB\_TELL messages return the current file position.

FSLIB\_TRUNCATE messages set the size of the file to the current file position by updating the file size attribute directly.

OPEN messages are handled as described in the ‘URL’ section above.

OUTMBLK/PUTMBLK messages invoke the *WRITE* procedure to transfer data to the remote file.

NEWMBLK messages return buffers for applications to use

RETMBLK messages just free the associated buffers.

TIMEOUT messages are used to trigger re-transmissions at the RPC layer for requests originating from this process.

## Shared Library Macros and Routines

### **nfs\_authorize**

```
void nfs_authorize(  
    char *host,  
    uint uid,  
    uint gid)
```

The *nfs\_authorize* routine sets the *UNIX* authorisation for subsequent NFS calls. It is assumed that only one userid is active for all processes within the system.