



October 12-14, 2004

NFSv4 In Depth

Spencer Shepler

Senior Staff Engineer

Sun Microsystems

spencer.shepler@sun.com



October 12-14, 2004


NFSv4

- Brief review of operations
- Solaris client structure and its use of the NFSv4 operations
- Interesting errors



October 12-14, 2004

Many to One





October 12-14, 2004

NFSv4 Operations

- COMPOUND allows for flexibility
- Client implementations are just starting to take advantage of this
- Expect more behavioral changes as time passes and minor version work is done



October 12-14, 2004

NFSv4 Operations

- ACCESS
- CLOSE
- COMMIT
- CREATE
- DELEGPURGE
- DELEGRETURN
- GETATTR
- GETFH
- LINK
- LOCK
- LOCKT
- LOCKU
- LOOKUP
- LOOKUPP
- NVERIFY
- OPEN
- OPENATTR
- OPEN_CONFIRM
- OPEN_DOWNGRADE
- PUTFH
- PUTPUBFH
- PUTROOTFH
- READ
- READDIR
- READLINK
- RENAME
- RESTOREFH
- SAVEFH
- SECINFO
- SETATTR
- SETCLIENTID
- SETCLIENTID_CONFIRM
- VERIFY
- WRITE
- RELEASE_LOCKOWNER



October 12-14, 2004

Ex: mount server:/test/dir

- Client generates this COMPOUND
 - PUTROOFH
 - GETFH
 - LOOKUP(test)
 - GETFH
 - GETATTR
 - SECINFO(dir)
 - LOOKUP(dir)
 - GETFH
 - GETATTR



October 12-14, 2004

Ex: mount server:/test/dir

- GETATTR is for these attributes
 - SYMLINK_SUPPORT
 - LINK_SUPPORT
 - FH_EXPIRE_TYPE
 - TYPE
 - SUPPORTED_ATTRS
 - FSID



October 12-14, 2004

Ex: mount server:/test/dir

- SETCLIENTID
 - Contains callback info (addr/program #)
 - Returns clientid used for OPEN
- SETCLIENTID_CONFIRM



October 12-14, 2004

Ex: Is file

- Is generates this COMPOUND
 - PUTFH
 - SAVEFH
 - LOOKUP (file)
 - GETFH
 - GETATTR
 - RESTOREFH
 - NVERIFY
 - GETATTR



October 12-14, 2004

Ex: Is directory

- Is of directory generates these **COMPOUNDS**
 - Getattr (of containing directory/cache chk)
 - Lookup (directory)
 - Access
 - Getattr (of containing directory – open)
 - Readdir (with full attributes if first readdir)



October 12-14, 2004

Ex: ls directory

- Readdir COMPOUND
 - PUTFH
 - READDIR
 - (first time with attrs)
 - SIZE, CHANGE, TYPE, RDATTR_ERROR, FSID, FILEID,
 - FILEHANDLE, OWNER_GROUP, OWNER, NUMLINKS, MODE, TIME_ACCESS, SPACE_USED, RAWDEV, MOUNTED_ON_FILEID, TIME_MODIFY, TIME_METADATA
 - LOOKUPP
 - GETFH
 - GETATTR



October 12-14, 2004

user/group mapping

- user@domain / group@domain
- Use of user-level nfsmapid
 - /etc/default/nfs NFSMAPID_DOMAIN
 - _nfsv4idmapdomain TXT RR
 - Local “domain”



October 12-14, 2004

Ex: open(2)

- Open COMPOUND
 - PUTFH
 - OPEN
 - GETFH
 - GETATTR
- Server may ask for OPEN_CONFIRM
- Stateid is returned
- Delegation may be returned
- Type and stateid



October 12-14, 2004

Ex: read/write

- read/write COMPOUND
 - PUTFH
 - READ
- The operation carries a stateid
 - Delegation Stateid
 - Lock Stateid
 - Open Stateid
 - Special Stateid
 -



October 12-14, 2004

Ex: close (2)

- Flush of modified data
- Close COMPOUND
 - PUTFH
 - GETATTR
 - CLOSE
- If applicable, delegreturn most likely to come later



October 12-14, 2004

Recovery events

- Server reboot
 - NFS4ERR_STALE_CLIENTID
 - NFS4ERR_STALE_STATEID
- Network partition
 - NFS4ERR_EXPIRED



October 12-14, 2004

Interrupted open/close

- Client must ensure that the operation occurs at the server and clean up if applicable



October 12-14, 2004

Server Namespace

- Shares are managed the same as before
- Client must traverse namespace from root to target directory
- With MOUNT protocol, some intermediate permission checking avoided
- Clients may not traverse entire namespace with single mount
- NFS4ERR_WRONGSEC



October 12-14, 2004

RENEW

- NFS4ERR_CB_PATH_DOWN
- NFS4ERR_ADMIN_REVOKED
- NFS4ERR_LEASE_MOVED



October 12-14, 2004

Errors that are interesting

- NFS4ERR_GRACE
- NFS4ERR_NO_GRACE
- NFS4ERR_DELAY
- NFS4ERR_FHEXPIRED
- NFS4ERR_STALE
- NFS4ERR_MOVED
- NFS4ERR_WRONGSEC



October 12-14, 2004

NFSv4 In Depth

Spencer Shepler

Senior Staff Engineer

Sun Microsystems

spencer.shepler@sun.com