

Distributed File Systems

File Characteristics

From Andrew File System work:

- most files are small—transfer files rather than disk blocks?
- reading more common than writing
- most access is sequential
- most files have a short lifetime—lots of applications generate temporary files (such as a compiler).
- file sharing (involving writes) is unusual—argues for client caching
- processes use few files
- files can be divided into classes—handle “system” files and “user” files differently.

Distributed File Systems

Primarily look at three traditional distributed file systems as we look at issues.

1. File Transfer Protocol (FTP). Motivation is to provide file sharing (not a distributed file system). 1970s.

Connect to a remote machine and interactively send or fetch an arbitrary file. FTP deals with authentication, listing a directory contents, ascii or binary files, etc.

Typically, a user connecting to an FTP server must specify an account and password. Often, it is convenient to set up a special account in which no password is needed. Such systems provide a service called *anonymous* FTP where userid is “anonymous” and password is typically user email address. Now typically handled via HTTP.

2. Sun’s Network File System (NFS). Motivated by wanting to extend a Unix file system to a distributed environment. Easy file sharing and compatability with existing systems. Mid-1980’s.

Stateless in that servers do not maintain state about clients. RPC calls supported:

- searching for a file within a directory
- reading a set of directory entries
- manipulating links and directories
- accessing file attributes
- reading/writing file data

Latest version of NFS (version 4) introduces some amount of state.

3. Andrew File System (AFS). Research project at CMU in 1980s. Company called Transarc, acquired by IBM. Primary motivation was to build a scalable distributed file system. Look at pictures.

Other file systems:

1. CODA: AFS spin-off at CMU. Disconnection and fault recovery.
2. Sprite: research project at UCB in 1980’s. To build a distributed Unix system.
3. Echo. Digital SRC.
4. Amoeba Bullet File Server: Tanenbaum research project.
5. xFs: serverless file system—file system distributed across multiple machines. Research project at UCB.
- 6.

Traditional Distributed File System Issues

Naming

How are files named? Access independent? Is the name location independent?

- FTP. location and access dependent.
- NFS. location dependent through client mount points. Largely transparent for ordinary users, but the same remote file system could be mounted differently on different machines. Access independent. See Fig 9-3. Has automount feature for file systems to be mounted on demand. All clients could be configured to have same naming structure.
- AFS. location independent. Each client has the same look within a *cell*. Have a cell at each site. See Fig 13-15.

Migration

Can files be migrated between file server machines? What must clients be aware of?

- FTP. Sure, but end-user must be aware.
- NFS. Must change mount points on the client machines.
- AFS. On a per-volume (collection of files managed as a single unit) basis.

Directories

Are directories and files handled with the same or a different mechanism?

- FTP. Directory listing handled as remote command.
- NFS. Unix-like.
- AFS. Unix-like.

Amoeba has separate mechanism for directories and files.

Sharing Semantics

What type of file sharing semantics are supported if two processes accessing the same file?

Possibilities:

- Unix semantics – every operation on a file is instantly visible to all processes.
- session semantics – no changes are visible to other processes until the file is closed.
- immutable files – files cannot be changed (new versions must be created)
- FTP. User-level copies. No support.
- NFS. Mostly Unix semantics.
- AFS. Session semantics.

Immutable files in Amoeba.

Caching

What, if any, file caching is supported?

Possibilities:

- write-through – all changes made on client are immediately written through to server
- write-back – changes made on client are cached for some amount of time before being written back to server.
- write-on-close – one type of write-back where changes are written on close (matches session semantics).

- FTP. None. User maintains own copy (whole file)
- NFS. File attributes (inodes) and file data blocks are cached separately. Cached attributes are validated with the server on file open.

Version 3: Uses read-ahead and delayed writes from client cache. Time-based at block level. New/changed files may not be visible for 30 seconds. Neither Unix nor session semantics. Non-deterministic semantics as multiple processes can have the same file open for writing.

Version 4: (like AFS) Client must flush modified file contents back to the server on close of file at client. Server can also *delegate* a file to a client so that the client can handle all requests for the file without checking with the server. However, server must now maintain state about open delegations and recall (with a callback) a delegation if the file is needed on another machine.

- AFS. File-level caching with callbacks (explain). Session semantics. Concurrent sharing is not possible.

Locking

Does the system support locking of files?

- FTP. N/A.
- NFS. Has mechanism, but external to NFS in v3. Internal to file system in version 4.
- AFS. Does support.

Replication/Reliability

Is file replication/reliability supported and how?

- FTP. No.
- NFS. minimal support in version 4.
- AFS. For read-only volumes within a cell. For example binaries and system libraries.

Scalability

Is the system scalable?

- FTP. Yes. Millions of users.
- NFS. Not so much. 10-100s
- AFS. Better than NFS, keep traffic away from file servers. 1000s.

Homogeneity

Is hardware/software homogeneity required?

- FTP. No.
- NFS. No.
- AFS. No.

File System Interface

Is the application interface compatible to Unix or is another interface used?

- FTP. Separate.
- NFS. The same.
- AFS. The same.

Security

What security and protection features are available to control access?

- FTP. Account/password authorization.
- NFS. RPC Unix authentication. Version 4 uses RPCSEC_GSS, a general security framework that can use proven security mechanisms such as Kerberos.
- AFS. Unix permissions for files, access control lists for directories. CODA has secure RPC implementation.

State/Stateless

Do file system servers maintain state about clients?

- FTP. No.
- NFS. No. In Version 4 servers maintains state about delegations and file locking.
- AFS. Yes.

Papers

“Elephant: The File System that Never Forgets”

<http://www.cs.wpi.edu/~cs535/s08/santry:osdi99/>

“A Comparison of File System Workloads”

<http://www.cs.wpi.edu/~cs535/s08/roselli:usenix00/>

“The Google File System”

<http://www.cs.wpi.edu/~cs535/s08/ghemawat:sosp03.pdf>

“Shark: Scaling File Servers via Cooperative Caching”

<http://www.cs.wpi.edu/~cs535/s08/annapureddy:nsdi05.pdf>

“Do Incentives Build Robustness in BitTorrent?”

<http://www.cs.wpi.edu/~cs535/s08/piatek:nsdi07.pdf>