

Distributed File Systems

Distributed File Systems (DFS)

- Used by the A&T Unix and Windows systems.
- Files can be stored on a file server or on peer systems.
- Reduces the mass storage required to make large data sets available to large numbers of users.
- Allows centralized access and update of files.
- Allows diskless workstations to function.

Goals

- A Distributed File System should look to a client like a conventional, centralized file system.
- Enforce access rights and security.
 - Prevent viewing or modification as the file is transmitted over the network.
- Be fault tolerant.
- Be scalable.

Andrew File System (AFS)

- Used at A&T.
- Can run on PCs or Unix systems.
- Requires client computers to have a disk.
- Copies files to the client computer's hard drive.
- Accessing a file from a hard drive is faster than accessing it over a network.
- Network traffic is not necessary during the use of the file.

AFS Design Assumptions

- Most files are small (<10K).
- Most files are used only by one user.
- Reads occur about 6 times more often than writes.
- Most writes are by the file owner. Few files have shared writes.
- Files are referenced in bursts (temporal locality).

AFS Operation

- A request is sent to the file server when a file or directory is opened.
- The server copies the file over the network to the hard drive of the client.
- The user application accesses the file from the local hard drive.
- When the file is closed, it is copied back to the server if it has been updated.

AFS Caching

- When a file is closed, the client keeps a copy of the file.
- If the file is opened again, the server copies the file to the client only if it has been updated since the client last closed the file.

Windows Network Stack

I/O Manager	
File System	Redirector
Disk manager	Transport layer
Disk driver	Network layer
Disk controller	NDIS
	Network device driver
	NIC

Windows File Sharing

1. When an application reads a remote file, the redirector sends a request to the server.
 2. The server reads the requested block from the file and sends it to the client.
 3. The redirector in the client returns the data to the application as if it was read from a local disk.
- The server disk organization is transparent to the client system.

DFS Access Techniques

- File caching
 - Entire file is copied to the client upon opening.
 - Used by the Andrew File System
- Block based access
 - Blocks of the file are transferred from the server to the client as they would be read from the disk.
 - Used by Windows and the Network File System.

Semantics of Sharing

- **Unix semantics** - every read of a file sees the effects of all previous writes. It is possible for clients to share the location pointer into a file.
- **Session semantics** - writes to an open file are immediately visible to local clients but invisible to remote clients. Once a file is closed all changes are visible to new opens.
- **Immutable shared file semantics** - Files cannot be changed.

Naming

- Naming is a mapping between logical files and their physical location.
- **location transparency** - name doesn't tell where the file is located.
- **location independence** - name doesn't change if file is moved or viewed from another client.


Naming Schemes

- `\\server\local_name`
 - not location transparent or independent
- Mount remote system on local directory
 - Remote file system attached as sub-tree of local file system.
 - Names may vary depending on mount point
- Single tree on all systems
 - Local files can be exceptions.

Unix `mount` Command

- Logically attaches a remote file system to a node in the local directory structure.

```
mount -t nfs serv:/usr/src /my/localdir
```



Fault Tolerance

- Failure transparency - continue even if client, server or net temporarily fails
- Stateless servers
 - Server does not keep information about open files.
 - Clients must supply full information to access a file.
 - Restart of server does not affect client
- Servers can be replicated.

Replication

- Files can be copied on multiple servers for fault tolerance or performance.
- Frequently used read-only data can be located at nearby servers for quick access.
- Replication should be transparent to users
- If users can make updates to the same data on multiple servers, consistency problems can occur.
- Updating replicating servers can be immediate or periodic.

Security

- More important with a distributed file system than with a local disk.
- Data has to be protected as it traverses the network.
- Clients must verify their identity.

Distributed Databases

- A remote file system is different than a remote database.
- A database system may perform processing on the data before sending a result to the client.