

# Distributed File Systems Design

Paul Krzyzanowski  
Rutgers University

September 14, 2012

## 1 Introduction

Presently, our most common exposure to distributed systems that exemplify some degree of transparency is through distributed file systems. We'd like remote files to look and feel just like local ones.

A file system is responsible for the organization, storage, retrieval, naming, sharing, and protection of files. File systems provide directory services, which convert a file name (possibly a hierarchical one) into an internal identifier (e.g. inode, FAT index). They contain a representation of the file data itself and methods for accessing it (read/write). The file system is responsible for controlling access to the data and for performing low-level operations such as buffering frequently-used data and issuing disk I/O requests.

Our goals in designing a distributed file system are to present certain degrees of transparency to the user and the system.

**access transparency** Clients are unaware that files are distributed and can access them in the same way as local files are accessed.

**location transparency** A consistent name space exists encompassing local as well as remote files. The name of a file does not give it location.

**concurrency transparency** All clients have the same view of the state of the file system. This means that if one process is modifying a file, any other processes on the same system or remote systems that are accessing the files will see the modifications in a coherent manner.

**failure transparency** The client and client programs should operate correctly after a server failure.

**heterogeneity** File service should be provided across different hardware and operating system platforms.

**scalability** The file system should work well in small environments (1 machine, a dozen machines) and also scale gracefully to huge ones (hundreds through tens of thousands of systems).

**replication transparency** To support scalability, we may wish to replicate files across multiple servers. Clients should be unaware of this.

**migration transparency** Files should be able to move around without the client's knowledge.

**support fine-grained distribution of data** To optimize performance, we may wish to locate individual objects near the processes that use them.

**tolerance for network partitioning** The entire network or certain segments of it may be unavailable to a client during certain periods (e.g. disconnected operation of a laptop). The file system should be tolerant of this.

## 2 Distributed file system concepts

A **file service** is a specification of what the file system offers to clients. A **file server** is the implementation of a file service and runs on one or more machines.

A **file** itself contains a name, data, and attributes (such as owner, size, creation time, access rights). An immutable file <sup>1</sup> is one that, once created, cannot be changed. Immutable files are easy to cache and to replicate across servers since their contents are guaranteed to remain unchanged.

Two forms of protection are generally used in distributed file systems, and they are essentially the same techniques that are used in single-processor non-networked systems:

**capabilities** Each user is granted a ticket (capability) from some trusted source for each object to which it has access. The capability specifies what kinds of access are allowed.

**access control lists** Each file has a list of users associated with it and access permissions per user. Multiple users may be organized into an entity known as a group.

### 2.1 File service types

To provide a remote system with file service, we will have to select one of two models of operation. One of these is the **upload/download** model. In this model, there are two fundamental operations: *read file* transfers an entire file from the server to the requesting client, and *write file* copies the file back to the server. It is a simple model and efficient in that it provides local access to the file when it is being used. Three problems are evident. It can be wasteful if the client needs access to only a small amount of the file data. It can be problematic if the client doesn't have enough space to cache the entire file. Finally, what happens if others need to modify the same file? The second model is a remote access model. The file service provides remote operations such as *open*, *close*, *read bytes*, *write bytes*, *get attributes*, etc. The file system itself runs on servers. The drawback in this approach is the servers are accessed for the duration of file access rather than once to download the file and again to upload it.

Another important distinction in providing file service is that of understanding the difference between directory service and file service. A directory service, in the context of file systems, maps human-friendly textual names for files to their internal locations, which can be used by the file service. The file service itself provides the file interface (this is mentioned above). Another component of file distributed file systems is the client module. This is the client-side interface for file and directory service. It provides a local file system interface to client software (for example, the VFS layer of a UNIX/Linux kernel).

---

<sup>1</sup>The Bullet server on the Amoeba operating system is an example of a system that uses immutable files.

## 2.2 Naming issues

In designing a distributed file service, we should consider whether all machines (and processes) should have the exact same view of the directory hierarchy. We might also wish to consider whether the name space on all machines should have a global root directory (a.k.a. “super root”) so that files can be accessed as, for example, `//server/path`. This is a model that was adopted by the Apollo Domain System, an early distributed file system, and more recently by the web community in the construction of a uniform resource locator (URL).

In considering our goals in name resolution, we must distinguish between location transparency and location independence. By location transparency we mean that the path name of a file gives no hint to where the file is located. For instance, we may refer to a file as `//server1/dir/file`. The server (server) can move anywhere without the client caring, so we have location transparency. However, if the file moves to server2 things will not work. If we have location independence, the files can be moved without their names changing. Hence, if machine or server names are embedded into path names we do not achieve location independence.

It is desirable to have access transparency, so that applications and users can access remote files just as they access local files. To facilitate this, the remote file system name space should be syntactically consistent with the local name space. One way of accomplishing this is by redefining the way files are named and require an explicit syntax for identifying remote files. This can cause legacy applications to fail and user discontent (users will have to learn a new way of naming their files). An alternate solution is to use a file system mounting mechanism to overlay portions of another file system over a node in a local directory structure. Mounting is used in the local environment to construct a uniform name space from separate file systems (which reside on different disks or partitions) as well as incorporating special-purpose file systems into the name space (e.g. `/proc` on many UNIX systems allows file system access to processes). A remote file system can be mounted at a particular point in the local directory tree. Attempts to access files and directories under that node will be directed to the driver for that file system.

To summarize, our naming options are:

- machine and path naming (*machine:path*, */machine/path*).
- mount remote file systems onto the local directory hierarchy (merging the local and remote name spaces).
- provide a single name space which looks the same on all machines.

The first two of these options are relatively easy to implement.

## 2.3 Types of names

When we talk about file names, we refer to symbolic names (for example, `server.c`). These names are used by people (users or programmers) to refer to files. Another “name” is the identifier used by the system internally to refer to a file. We can think of this as a **binary name** (more precisely, as an address). On most POSIX file systems, this would be the device number and inode number.

Directories provide a mapping from symbolic names to file addresses (binary names). Typically, one symbolic name maps to one file address. If multiple symbolic names map onto one binary name, these are called *hard links*. On inode-based file systems (e.g., most

UNIX systems), hard links must exist within the same device since the address (inode) is unique only on that device. On Windows systems, they are not supported because file attributes are stored with the name of the file. Having two symbolic names refer to the same data will cause problems in synchronizing file attributes (how would you locate other files that point to this data?). A hack to allow multiple names to refer to the same file (whether its on the same device or a different device) is to have the symbolic name refer to a single file address but that file may have an attribute to tell the system that its contents contain a symbolic file name that should be dereferenced. Essentially, this adds a level of indirection: access a file which contains another file name, which references the file attributes and data. These files are known as symbolic links. Finally, it is possible for one symbolic name to refer to multiple file addresses. This doesn't make much sense on a local system<sup>2</sup>, but can be useful on a networked file system to provide fault tolerance or enable the system to use the file address which is most efficient.

## 2.4 Semantics of file sharing

The analysis of file sharing semantics is that of understanding how files behave. For instance, on most systems, if a read follows a write, the read of that location will return the values just written. If two writes occur in succession, the following read will return the results of the last write. File systems that behave this way are said to observe sequential semantics.

Sequential semantics can be achieved in a distributed system if there is only one server and clients do not cache data. This can cause performance problems since clients will be going to the server for every file operation (such as single-byte reads). The performance problems can be alleviated with client caching. However, now if the client modifies its cache and another client reads data from the server, it will get obsolete data. Sequential semantics no longer hold.

One solution is to make all the writes write-through to the server. This is inefficient and does not solve the problem of clients having invalid copies in their cache. To solve this, the server would have to notify all clients holding copies of the data.

Another solution is to relax the semantics. We will simply tell the users that things do not work the same way on the distributed file system as they did on the local file system. The new rule can be "changes to an open file are initially visible only to the process (or machine) that modified it." These are known as **session semantics**.

Yet another solution is to make all the files immutable. That is, a file cannot be open for modification, only for reading or creating. If we need to modify a file, we'll create a completely new file under the old name. Immutable files are an aid to replication but they do not help with changes to the file's contents (or, more precisely, that the old file is obsolete because a new one with modified contents succeeded it). We still have to contend with the issue that there may be another process reading the old file. It's possible to detect that a file has changed and start failing requests from other processes.

A final alternative is to use atomic transactions. To access a file or a group of files, a process first executes a *begin transaction* primitive to signal that all future operations will be executed indivisibly. When the work is completed, an *end transaction* primitive is executed.

---

<sup>2</sup>It really does make sense in a way. In the late 1980's, David Korn created a file system that allowed multiple directories to be mounted over the same directory node. Several operating systems later adopted this technique and called it *union mounts*. It was a core aspect of the Plan 9 operating system, where it allowed one to build a fully custom name space and avoid the need for a PATH environment variable in searching for executables. The executables are always found in /bin. The name is resolved by searching through the file systems (directories) mounted on that node in a last-mounted, first-searched order.

If two or more transactions start at the same time, the system ensures that the end result is as if they were run in some sequential order. All changes have an all or nothing property.

## 2.5 File usage patterns

It's clear that we cannot have the best of all worlds and that compromises will have to be made between the desired semantics and efficiency (efficiency encompasses increasing client performance, reducing network traffic, and minimizing server load). To design or select a suitable distributed file system, it is important to understand the usage patterns within a file system.

A comprehensive study was made by Satyanarayanan in 1981 which showed the following use patterns. Bear in mind that this is in the days before people kept vast collections of audio and video files.

- Most files are under 10K bytes in size. This suggests that it may be feasible to transfer entire files (a simpler design). However, a file system should still be able to support large files.
- Most files have short lifetimes (many files are temporary files created by editors and compilers). It may be a good idea to keep these files local and see if they will be deleted soon.
- Few files are shared. While sharing is a big issue theoretically, in practice it's hardly done. We may choose to accept session semantics and not worry too much about the consequences.
- Files can be grouped into different classes, with each class exhibiting different properties: system binaries:
  - read/only, widely distributed. These are good candidates for replication.
  - compiler and editor temporary files :
  - short, unshared, disappear quickly. We'd like to keep these local if possible.
  - mailboxes (email): not shared, frequently updated. We don't want to replicate these.
  - ordinary data files: these may be shared.

## 2.6 System design issues

### Name resolution

In looking up the pathname of a file (e.g. via the `namei` function in the UNIX kernel), we may choose to evaluate a pathname a component at a time. For example, for a pathname `aaa/bbb/ccc`, we would perform a remote lookup of `aaa`, then another one of `bbb`, and finally one of `ccc`). Alternatively, we may pass the rest of the pathname to the remote machine as one lookup request once we find that a component is remote. The drawback of the latter scheme is (a) the remote server may be asked to walk up the tree by processing `..` (parent node) components and reveal more of its file system than it wants and (b) other components cannot be mounted underneath the remote tree on the local system. Because of this, component at a time evaluation is generally favored but it has performance problems (a lot more messages). We may choose to keep a local cache of component resolutions.

## Should servers maintain state?

A **stateless system** is one in which the client sends a request to a server, the server carries it out, and returns the result. Between these requests, no client-specific information is stored on the server. A **stateful system** is one where information about client connections is maintained on the server.

In a stateless system:

- Each request must be complete – the file has to be fully identified and any offsets specified.
- Fault tolerance: if a server crashes and then recovers, no state was lost about client connections because there was no state to maintain.
- No remote open/close calls are needed (they only serve to establish state).
- No wasted server space per client.
- No limit on the number of open files on the server; they aren't "open" – the server maintains no per-client state.
- No problems if the client crashes. The server does not have any state to clean up.

On a stateful system:

- requests are shorter (less info to send).
- better performance in processing the requests.
- idempotency works; cache coherence is possible.
- file locking is possible; the server can keep state that a certain client is locking a file (or portion thereof).

## Caching

We can employ caching to improve system performance. There are four places in a distributed system where we can hold data:

1. On the server's disk
2. In a cache in the server's memory
3. In the client's memory
4. On the client's disk

The first two places are not an issue since any interface to the server can check the centralized cache. It is in the last two places that problems arise and we have to consider the issue of cache consistency. Several approaches may be taken:

**write-through** What if another client reads its own cached copy? All accesses would require checking with the server first (adds network congestion) or require the server to maintain state on who has what files cached. Write-through also does not alleviate congestion on writes.

**delayed writes** Data can be buffered locally (where consistency suffers) but files can be updated periodically. A single bulk write is far more efficient than lots of little writes every time any file contents are modified. Unfortunately the semantics become ambiguous.

**write on close** This is admitting that the file system uses session semantics.

**centralized control** Server keeps track of who has what open in which mode. We would have to support a stateful system and deal with signaling traffic.

### 3 Distributed File Systems: case studies

It is clear that compromises have to be made in a practical design. For example, we may trade-off file consistency for decreased network traffic. We may choose to use a connectionless protocol to enable clients to survive a server crash gracefully but sacrifice file locking. This section examines a few distributed file systems.

#### 3.1 Network File System (NFS)

Sun's Network File System (NFS) is one of the earliest distributed file systems, is still widely used, and is the *de facto* standard network file system on various flavors of UNIX, Linux and BSD and is natively supported in Apple's OS X. We will look at its early design to understand what the designers were trying to do and why certain decisions were made. The design goals of NFS were:

- Any machine can be a client and/or a server.
- NFS must support diskless workstations (that are booted from the network). Diskless workstations were Sun's major product line.
- Heterogeneous systems should be supported: clients and servers may have different hardware and/or operating systems. Interfaces for NFS were published to encourage the widespread adoption of NFS.
- High performance: try to make remote access as comparable to local access through caching and read-ahead.

From a transparency point of view NFS offers:

**Access transparency** Remote (NFS) files are accessed through normal system calls; On POSIX systems, the protocol is implemented under the VFS layer.

**Location transparency** The client adds remote file systems to its local name space via mount. File systems must be exported at the server. The user is unaware of which directories are local and which are remote. The location of the mount point in the local system is up to the client's administrator.

Failure transparency : NFS is stateless; UDP is used as a transport. If a server fails, the client retries.

**Performance transparency** Caching at the client will be used to improve performance

**No migration transparency** The client mounts machines from a server. If the resource moves to another server, the client must know about the move.

**No support for all Unix file system features** NFS is stateless, so stateful operations such as file locking are a problem. All UNIX file system controls may not be available.

**Devices?** Since NFS had to support diskless workstations, where every file is remote, remote device files had to refer to the client's local devices. Otherwise there would be no way to access local devices in a diskless environment.

## 3.2 NFS protocols

The NFS client and server communicate over remote procedure calls (Sun/ONC RPC) using two protocols: the **mounting protocol** and the **directory and file access protocol**. The mounting protocol is used to request a access to an exported directory (and the files and directories within that file system under that directory). The directory and file access protocol is used for accessing the files and directories (e.g. read/write bytes, create files, etc.). The use of RPC's external data representation (XDR) allows NFS to communicate with heterogeneous machines. The initial design of NFS ran only with remote procedure calls over UDP. This was done for two reasons. The first reason is that UDP is somewhat faster than TCP but does not provide error correction (the UDP header provides a checksum of the data and headers). The second reason is that UDP does not require a connection to be present. This means that the server does not need to keep per-client connection state and there is no need to reestablish a connection if a server was rebooted.

The lack of UDP error correction is remedied in the fact that remote procedure calls have built-in retry logic. The client can specify the maximum number of retries (default is 5) and a timeout period. If a valid response is not received within the timeout period the request is re-sent. To avoid server overload, the timeout period is then doubled. The retry continues until the limit has been reached. This same logic keeps NFS clients fault-tolerant in the presence of server failures: a client will keep retrying until the server responds.

### Mounting protocol

The client sends the pathname to the server and requests permission to access the contents of that directory. If the name is valid and exported (stored in `/etc/exports` on Linux and BSD systems and in `/etc/dfs/sharetab` on System V release 4/SunOS 5.x), the server returns a file handle to the client. This file handle contains all the information needed to identify the file on the server: *{file system type, disk ID, inode number, security info}*.

Mounting an NFS file system is accomplished by parsing the path name, contacting the remote machine for a file handle, and creating an in-memory vnode at the mount point. A vnode points to an inode for a local UNIX file or, in the case of NFS, an rnode. The rnode contains specific information about the state of the file from the point of view of the client. Two forms of mounting are supported:

**static** In this case, file systems are mounted with the mount command (generally during system boot).

**automounting** One problem with static mounting is that if a client has a lot of remote resources mounted, boot-time can be excessive, particularly if any of the remote systems are not responding and the client keeps retrying. Another problem is that each



machine has to maintain its own name space. If an administrator wants all machines to have the same name space, this can be an administrative headache. To combat these problems the automounter was introduced.

The automounter allows mounts and unmounts to be performed in response to client requests. A set of remote directories is associated with a local directory. None are mounted initially. the first time any of these is referenced, the operating system sends a message to each of the servers. The first reply wins and that file system gets mounted (it is up to the administrator to ensure that all file systems are the same). To configure this, the automounter relies on mapping files that provide a mapping of client pathname to the server file system. These maps can be shared to facilitate providing a uniform naming space to a number of clients.

## Directory and file access protocol

Clients send RPC messages to the server to manipulate files and directories. A file is accessed by performing a lookup remote procedure call. This returns a file handle and attributes. It is *not* like an *open* in that no information is stored in any system tables on the server. After that, the handle may be passed as a parameter for other functions. For example, a *read(handle, offset, count)* function will read *count* bytes from location *offset* in the file referred to by *handle*.

The entire directory and file access protocol is encapsulated in sixteen functions<sup>3</sup>. These are:

Function	Description
<i>null</i>	no-operation but ensure that connectivity exists
<i>lookup</i>	lookup the file name in a directory
<i>create</i>	create a file or a symbolic link
<i>remove</i>	remove a file from a directory
<i>rename</i>	rename a file or directory
<i>read</i>	read bytes from a file
<i>write</i>	write bytes to a file
<i>link</i>	create a link to a file
<i>symlink</i>	create a symbolic link to a file
<i>readlink</i>	read the data in a symbolic link (do not follow the link)
<i>mkdir</i>	create a directory
<i>rmdir</i>	remove a directory
<i>readdir</i>	read from a directory
<i>getattr</i>	get attributes about a file or directory (type, access and modify times, and access permissions)
<i>setattr</i>	set file attributes
<i>statfs</i>	get information about the remote file system

<sup>3</sup>These functions are present in versions 2 and 3 of the NFS protocol. Version 3 added six more functions and even more were added in version 4.

## Accessing files

Files are accessed through conventional system calls (thus providing access transparency). If you recall conventional UNIX systems, a hierarchical pathname is dereferenced to the file location with a kernel function called *namei*. This function maintains a reference to a current directory, looks at one component and finds it in the directory, changes the reference to that directory, and continues until the entire path is resolved. At each point in traversing this pathname, it checks to see whether the component is a mount point, meaning that name resolution should continue on another file system. In the case of NFS, it continues with remote procedure calls to the server hosting that file system.

Upon realizing that the rest of the pathname is remote, *namei* will continue to parse one component of the pathname at a time to ensure that references to .. (dot-dot, the parent directory) and to symbolic links become local if necessary. Each component is retrieved via a remote procedure call which performs an NFS lookup. This procedure returns a file handle. An in-memory *rnode* is created and the VFS layer in the file system creates a *vnode* to point to it.

The application can now issue read and write system calls. The file descriptor in the user's process will reference the in-memory *vnode* at the VFS layer, which in turn will reference the in-memory *rnode* at the NFS level which contains NFS-specific information, such as the file handle. At the NFS level, NFS *read*, *write*, etc. operations may now be performed, passing the file handle and *local* state (such as file offset) as parameters. No information is maintained on the server between requests; it is a *stateless* system.

The RPC requests have the user ID and group ID number sent with them. This is a security hole that may be stopped by turning on RPC encryption.

## Performance

NFS performance was usually found to be slower than accessing local files because of the network overhead. To improve performance, reduce network congestion, and reduce server load, file data is cached at the client. Entire pathnames are also cached at the client to improve performance for directory lookups.

**server caching** Server caching is automatic at the server in that the same buffer cache is used as for all other files on the server. The difference for NFS-related writes is that they are all **write-through** to avoid unexpected data loss if the server dies.

**client caching** The goal of client caching is to reduce the amount of remote operations. Three forms of information are cached at the client: file data, file attribute information, and pathname bindings. NFS caches the results of *read*, *readlink*, *getattr*, *lookup*, and *readdir* operations. The danger with caching is that **inconsistencies may arise**. NFS tries to avoid inconsistencies (and/or increase performance) with:

- *validation*: if caching one or more blocks of a file, save a time stamp. When a file is opened or if the server is contacted for a new data block, compare the last modification time. If the remote modification time is more recent, invalidate the cache.
- Validation is performed every three seconds on open files.
- Cached data blocks are assumed to be valid for three seconds.
- Cached directory blocks are assumed to be valid for thirty seconds.

- Whenever a page is modified, it is marked dirty and scheduled to be written (asynchronously). The page is flushed when the file is closed.

Transfers of data are done in large chunks; the default is 8K bytes. As soon as a chunk is received, the client immediately requests the next 8K-byte chunk. This is known as **read-ahead**. The assumption is that most file accesses are sequential and we might as well fetch the next block of data while we're working on our current block, anticipating that we'll likely need it. This way, by the time we do, it will either be there or we don't have to wait too long for it since it is on its way.

## Problems

The biggest problem with NFS is file consistency. The caching and validation policies do not guarantee session semantics.

NFS assumes that clocks between machines are synchronized and performs no clock synchronization between client and server. One place where this hurts is in distributed software development environments. A program such as `make`, which compares times of files (such as object and source) to determine whether to regenerate them, can either fail or give confusing results.

Because of its stateless design, open with append mode cannot be guaranteed to work. You can open a file, get the attributes (size), and then write at that offset, but you'll have no assurance that somebody else did not write to that location after you received the attributes. In that case your write will overwrite the other once since it will go to the old end-of-file byte offset.

Also because of its stateless nature, file locking cannot work. File locking implies that the server keeps track of which processes have locks on the file. Sun's solution to this was to provide a separate process (a lock manager). This introduces *state* to NFS.

One common programming practice under UNIX file systems for manipulating temporary data in files is to open a temporary file and then remove it from the directory. The name is gone, but the data persists because you still have the file open. Under NFS, the server maintains no state about remotely opened files and removing a file will cause the file to disappear. Since legacy applications depended on this, Sun's solution was to create a special hack for UNIX: if the same process that has a file open attempts to delete it, it is instead moved to a temporary name and deleted on close. It's not a perfect solution, but it works well.

Permission bits might change on the server and disallow future access to a file. Since NFS is stateless, it has to check access permissions each time it receives an NFS request. With local file systems, once access is granted initially, a process can continue accessing the file even if permissions change.

By default, no data is encrypted and Unix-style authentication is used (used ID, group ID). NFS supports two additional forms of authentication: Diffie-Hellman and Kerberos. However, data is never encrypted and user-level software should be used to encrypt files if this is necessary.

## More fixes

The original version of NFS was released in 1985, with version 2 released around 1988. In 1992, NFS was enhanced to version 3 (SunOS 5.5). This version is still supported,

although version 4 was introduced around 2003. Several changes were added to enhance the performance of the system in version 3:

1. NFS was enhanced to support TCP. UDP caused more problems over wide-area networks than it did over LANs because of errors. To combat that and to support larger data transfer sizes, NFS was modified to support TCP as well as UDP. To minimize connection setup, all traffic can be multiplexed over one TCP connection.
2. NFS always relied on the system buffer cache for caching file data. The buffer cache is often not very large and useful data was getting flushed because of the size of the cache. Sun introduced a caching file system, CacheFS, that provides more caching capability by using the disk. Memory is still used as before, but a local disk is used if more cache space is needed. Data can be cached in chunks as large as 64K bytes and entire directories can be cached.
3. NFS was modified to support asynchronous writes. If a client needed to send several write requests to a server, it would send them one after another. The server would respond to a request only after the data was flushed to the disk. Now multiple writes can be collected and sent as an aggregate request to the server. The server does not have to ensure that the data is on stable storage (disk) until it receives a commit request from the client.
4. File attributes are returned with each remote procedure call now. The overhead is slight and saves clients from having to request file attributes separately (which was a common operation).
5. Version 3 allows 64-bit rather than the old 32-bit file offsets (supporting file sizes over 18 million terabytes).
6. An enhanced lock manager was added to provide monitored locks. A status monitor monitors hosts with locks and informs a lock manager of a system crash. If a server crashes, the status monitor reinstates locks on recovery. If a client crashes, all locks from that client are freed on the server.

Version 4, described in RFC 3530, is a major enhancement to NFS and the original stateless model of the system is essentially gone now. A few of the major additions in version 4 were:

**Single pseudo file system export:** Instead of exporting a set of directories, a server can now export a single pseudo file system that appears as one directory but is created from multiple components.

**Compound RPC:** The system is still based on ONC RPC but supports the use of *compound RPC* reduce network latency. Compound RPCs allow one to issue several disparate requests in one message.

**Connection-oriented transport:** NFS now requires the use of TCP.

**Access control lists and server objects:** NFS now supports arbitrary server objects beyond files and directories. This was added primarily to support Windows systems. Support for control lists was also added.

**File locking and consistency:** NFS still has a weak consistency model for general file access but now supports mandatory as well as advisory file locking. The NFS server can also direct the client on what data it can cache.

**Security:** Strong security is now required. One of three mechanisms are supported and negotiated at mount time: Kerberos, LIPKEY, and SPK-3. Identification is sent as username strings instead of user ID numbers.

### 3.3 Andrew File System (AFS)

The goal of the Andrew File System (from Carnegie Mellon University, then a product of Transarc Corp., and now part of the Transarc division of IBM and available via the IBM public license) was to support information sharing on a large scale (thousands to 10000+ users). There were several incarnations of AFS, with the first version being available around 1984, AFS-2 in 1986, and AFS-3 in 1989).

The assumptions about file usage were:

- most files are small
- reads are much more common than writes
- most files are read/written by one user
- files are referenced in bursts (locality principle). Once referenced, a file will probably be referenced again.

From these assumptions, the original goal of AFS was to use **whole file serving** on the server (send an entire file when it is opened) and **whole file caching** on the client (save the entire file onto a local disk). To enable this mode of operation, the user would have a cache partition on a local disk devoted to AFS. If a file was updated then the file would be written back to the server when the application performs a *close*. The local copy would remain cached at the client.

#### Implementation

The client's machine has one disk partition devoted to the AFS cache (for example, 100M bytes, or whatever the client can spare). The client software manages this cache in an LRU (least recently used) manner and the clients communicate with a set of trusted servers. Each server presents a **location-transparent** hierarchical file name space to its clients. On the server, each physical disk partition contains files and directories that can be grouped into one or more volumes. A volume is nothing more than an administrative unit of organization (e.g., a user's home directory, a local source tree). Each volume has a directory structure (a rooted hierarchy of files and directories) and is given a name and ID. Servers are grouped into administrative entities called *cells*. ***A cell is a collection of servers, administrators, clients, and users. Each cell is autonomous but cells may cooperate and present users with one uniform name space\****. The goal is that *every client* will see the same name space (by convention, under a directory */afs*). Listing the directory */afs* shows the participating cells (e.g., */afs/mit.edu*).

Internally, each file and directory is identified by three 32-bit numbers:

**volume ID:** This identifies the volume to which the object belongs. The client caches the binding between volume ID and server, but the server is responsible for maintaining the bindings.

**vnode ID:** This is the “handle” (vnode number) that refers to the file on a particular server and disk partition (volume).

**uniqueifier:** This is a unique number to ensure that the same vnode IDs are not reused.

Each server maintains a copy of a database that maps a volume number to its server. If the client request is incorrect (because a volume moved to a different server), the server forwards the request. This provides AFS with migration transparency: volumes may be moved between servers without disrupting access.

Communication in AFS is with RPCs via UDP. Access control lists are used for protection; UNIX file permissions are ignored. The granularity of access control is directory based; the access rights apply to all files in the directory. Users may be members of groups and access rights specified for a group. Kerberos is used for authentication.

## Cache coherence

The server copies a file to the client and provides a **callback promise**: *it will notify the client when any other process modifies the file.*

When a server gets an update from a client, it notifies all the clients by sending a **callback** (via RPC). Each client that receives the callback then invalidates the cached file. If a client that had a file cached was down, on restart, it contacts the server with the timestamps of each cached file to decide whether to invalidate the file. Note that if a process has a file open, it can continue using it, even if it has been invalidated in the cache. Upon close, the contents will still be propagated to the server. There is no further mechanism for coherency. AFS presents **session semantics**.

Under AFS, read-only files may be replicated on multiple servers.

Whole file caching is not feasible for very large files, so AFS caches files in 64K byte chunks (by default) and directories in their entirety. File modifications are propagated only on close. Directory modifications are propagated immediately.

AFS does not support byte-range file locking. Advisory file locking (query to see whether a file has a lock on it) is supported.

## AFS Summary

AFS demonstrates that whole file (or large chunk) caching offers dramatically reduced loads on servers, creating an environment that scales well. The AFS file system provides a uniform name space from all workstations, unlike NFS, where the client mounts each NFS file system at a client-specific location (the name space is uniform only under the /afs directory, however). Establishing the same view of the file name space from each client is easier than with NFS. This enables users to move to different workstations and see the same view of the file system.

Access permission is handled through control lists per directory, but there is no per-file access control. Workstation/user authentication is performed via the Kerberos authentication protocol using a trusted third party (more on this in the security section).

A limited form of replication is supported. Replicating read-only (and read-mostly at your own risk) files can alleviate some performance bottlenecks for commonly accessed files (e.g. password files, system binaries).

### 3.4 Coda

Coda is a descendent of AFS, also created at CMU (c. 1990–1992). Its goals are:

- Provide better support for replication of file volumes than offered by AFS. AFS' limited form (read-only volumes) of replication will be a limiting factor in scaling the system. We would like to support widely shared read/write files.
- Provide constant data availability in disconnected environments through hoarding (user-directed caching). This requires logging updates on the client and reintegration when the client is reconnected to the network. Such a scheme will support the mobility of PCs.
- Improve fault tolerance. Failed servers and network problems shouldn't seriously inconvenience users.

To achieve these goals, AFS was modified in two substantial ways:

1. File volumes can be replicated to achieve higher throughput of file access operations and improve fault tolerance.
2. The caching mechanism was extended to enable disconnected clients to operate.

Volumes can be replicated to group of servers. The set of servers that can host a particular volume is the **volume storage group (VSG)** for that volume. In identifying files and directories, a client no longer uses a volume ID as AFS did, but instead uses a replicated volume ID. The client performs a one-time lookup to map the replicated volume ID to a list of servers and local volume IDs. This list is cached for efficiency. Read operations can take place from any of these servers to distribute the load. A write operation has to be multicast to all available servers. Since some servers may be inaccessible at a particular point in time, a client may be able to access only a subset of the VSG. This subset is known as the **Available Volume Storage Group**, or **AVSG**.

Since some volume servers may be inaccessible, special treatment is needed to ensure that clients do not read obsolete data. Each file copy has a version stamp. Before fetching a file, a client requests version stamps for that file from all available servers. If some servers are found to have old versions, the client initiates a **resolution process** which tries to automatically resolve differences (administrative intervention may be required if the process finds problems that it cannot fix). Resolution is only initiated by the client. The process is handled entirely by the servers.

#### Disconnected operation

If a client's AVSG is empty, then the client is operating in a **disconnected operation** mode. If a file is not cached locally and is needed, nothing can be done: the system simply retries access and fails. For writes, however, the client does not report a failure of an update. Instead, the client logs the update locally in a Client Modification Log (CML). The user is oblivious to this. On reconnection, a process of **reintegration** with the server(s) commences to bring the server up to date. The CML is played back (the log playback is optimized so that only the latest changes are sent). The system tries to resolve conflicts automatically. This is not always possible (for example, someone may have modified the

same parts of the file on a server while our client was disconnected). In cases where conflicts arise, user intervention is required to reconcile the differences.

To further support disconnected operation, it is desirable to cache all the files that will be needed for work to proceed when disconnected and keep them up to date even if they are not being actively used. To do this, Coda supports a **hoard database** that contains a list of these “important” files. The hoard database is constructed both by monitoring a user’s file activity and allowing a user to explicitly specify files and directories that should be present on the client. The client frequently asks the server to send updates if necessary (that is, when it receives a callback).

### 3.5 Distributed File System (DFS), also known as AFS v3

DFS is the file system that is part of the Open Group’s (formerly the Open Software Foundation or OSF) Distributed Computing Environment (DCE) and is the third version of AFS. Like AFS, it assumes that:

- most file accesses are sequential
- most file lifetimes are short
- the majority of accesses are whole-file transfers
- the majority of accesses are to small files

With these assumptions, the conclusion is that file caching can reduce network traffic and server load. Since the studies on file usage in the early and mid 1980s, it was noted that file throughput per user has increased dramatically and that typical file sizes became much larger. However, disk capacity and network throughput also became much larger.

DFS implements a strong consistency model (unlike AFS) with Unix semantics supported. This means that a read will return the effects of all writes that precede it. Cache consistency under DFS is maintained by the use of **tokens**.

A *token* is a guarantee from the server that a client can perform certain operations on the cached file. The server will revoke a token if another client attempts a conflicting operation. A server grants and revokes tokens. It will grant any number of read tokens to clients but as soon as one client requests write access, the server will revoke all outstanding read and write tokens and issue a single write token to the requestor. This token scheme makes long term caching possible (which is not possible under NFS). Caching is in units of chunk sizes that range from 8K to 256K bytes. Caching is both in client memory and on the disk. DFS also employs read-ahead (similar to NFS) to attempt to bring additional chunks off the file to the client before they are needed.

DFS is integrated with DCE security services. File protection is via access control lists (ACL) and all communication between client and server is via authenticated remote procedure calls.

### 3.6 Server Message Block (SMB)

SMB is a protocol for sharing files, devices, and communication abstractions (such as named pipes or mailslots). It was created by Microsoft and Intel in 1987 and evolved over the years.

SMB is a client-server request-response protocol. Servers make file systems and other resources available to clients and clients access the shared file systems (and printers) from



the servers. The protocol is connection-oriented and initially required either Microsoft's IPX/SPX or NetBIOS over either TCP/IP or NetBEUI (these are session-layer APIs). A typical session proceeds as follows:

1. Client sends a *negprot* SMB to the server. This is a protocol negotiation request.
2. The server responds with a *version number* of the protocol (and version-specific information, such as a maximum buffer size and naming information).
3. The client *logs on* (if required) by sending a *sesssetupX* SMB, which includes a username and password. It receives an acknowledgement or failure from the server. If successful, the server sends back a user ID (UID) of the logged-on user. This UID must be submitted with future requests.
4. The client can now connect to a tree. It sends a *tcon* (or *tconX*) SMB with the network name of the shared resource to request access to the resource. The server responds with a *tree identifier* (TID) that the client will use for all future requests for that resource.
5. Now the client can send *open*, *read*, *write*, *close* SMBs.

Machine naming was restricted to 15-character NetBIOS names if NetBEUI or TCP/IP are used. Since SMB was designed to operate in a small local-area network, clients find out about resources either by being configured to know about the servers in their environment or by having each server periodically broadcast information about its presence. Clients would listen for these broadcasts and build lists of servers. This is fine in a LAN environment but does not scale to wide-area networks (e.g. a TCP/IP environment with multiple subnets or networks). To combat this deficiency, Microsoft introduced browse servers and the Windows Internet Name Service (WINS).

The SMB security model has two levels:

1. **Share level** Protection is applied per “share” (resource). Each share can have a password. The client needs to know that password to be able to access all files under that share. This was the only security model in early versions of SMB and became the default under Windows 95 and future systems.
2. **User level** Protection is applied to individual files in each share based on user access rights. A user (client) must log into a server and be authenticated. The client is then provided with a UID which must be presented for all future accesses.

### 3.7 CIFS (SMB evolves)

SMB continued to evolve and eventually Microsoft made the protocol public. CIFS is a version of SMB based on the public protocol and continued evolution.

It is based on the server message block protocol and draws from concepts in AFS and DFS. SMB shunned excessive client-side caching for fears of inconsistency but it really is useful both for client performance and to alleviate load from servers so coherent caching was a key addition to the protocol. To support wide-area (slow) networks, CIFS allows multiple requests to be combined into a single message to minimize round-trip latencies. The obsolete requirement for NetBIOS or NetBEUI have been dropped. CIFS is transport-independent but requires a reliable connection-oriented message-stream transport. Sharing is in units of

directory trees or individual devices. For access, the client sends authentication information to the server (name and password). The granularity of authorization is up to the server (e.g., individual files or an entire directory tree).

The caching mechanism is one of the more interesting aspects of CIFS. Network load is reduced if the amount of times that the client informs the server of changes is minimized. This also minimizes the load on servers. An extreme optimization leads to session semantics but a goal in CIFS is to provide coherent caching. Client caching is safe if any number of clients are reading data. Read-ahead operations (prefetch) are also safe as long as other clients are only reading the file. Write-behind (delayed writes) is safe only if a single client is accessing the file. None of these optimizations is safe if multiple clients are writing the file. In that case, operations will have to go directly to the server.

To support this behavior, the server grants **opportunistic locks (oplocks)** to a client for each file that it is accessing. This is a slight modification of the token granting scheme used in DFS. An oplock takes one of the following forms:

**exclusive oplock:** Tells the client that it is the only one with the file open (for write).

Local caching, read-ahead, and write-behind are allowed. The server must receive an update upon file close. If someone else opens the file, the server has the previous client break its oplock. The client must send the server any lock and write data and acknowledge that it no longer has the lock.

**level II oplock:** Allows multiple clients to have the same file open as long as none are writing to the file. It tells the client that there are multiple concurrent clients, none of whom have modified the file (read access). Local caching of reads as well as read-ahead are allowed. All other operations must be sent directly to the server.

**batch oplock:** Allows the client to keep the file open on the server even if a local process that was using it has closed the file. A client requests a batch oplock if it expects that programs may behave in a way that generates a lot of traffic (accessing the same file over and over). This oplock tells the client that it is the only one with the file open. All operations may be done on cached data and data may be cached indefinitely.

**no oplocks:** Tells the client that other clients may be writing data to the file: all requests other than reads must be sent to the server. Read operations may work from a local cache only if the byte range was locked by the client.

The server has the right to asynchronously send a message to the client changing the oplock. For example, a client may be granted an exclusive oplock initially since nobody else was accessing the file. Later on, when another client opened the same file for read, the oplock was changed to a level II oplock. When another client opened the file for writing, both of the earlier clients were sent a message revoking their oplocks.

## 4 References

- The NFS Distributed File Service, A White Paper from SunSoft, 1994 Sun Microsystems, Inc.
- RFC 1094: NFS: Network File System Protocol Specification, Sun Microsystems, March 1989

- Sun’s Network File System<sup>4</sup>, Operating Systems: Three Easy Pieces, Remzi H. Arpaci-Dusseu and Andrea C. Arpaci-Dusseu
- S. Shepler et al, RFC 3530<sup>5</sup>, Network Working Group, IETF, April 2003

## AFS

- John H. Howard, An Overview of the Andrew File System<sup>6</sup>, Carnegie Mellon University, 1988

OpenAFS Project [] [OpenAFS Project]: <http://www.openafs.org> “Open AFS project page” class=external

- Remzi H. Arpaci-Dusseu and Andrea C. Arpaci-Dusseu, The Andrew File System (AFS)<sup>7</sup>, Operating Systems: Three Easy Pieces.
- AFS at Carnegie Mellon<sup>8</sup>

## CIFS/SMB

- MS-CIFS<sup>9</sup>: Common Internet File System (CIFS) Protocol, Microsoft Developer Network, Microsoft, v20, October 2012. (782 pages!)
- Richard Sharpe, Just what is SMB?<sup>10</sup>, [www.samba.org](http://www.samba.org), v1.2, October 2002.
- Server Message Block article<sup>11</sup>, Wikipedia
- Opportunistic Locks (Windows)<sup>12</sup>, Windows Dev Center, Microsoft, October 2012

## Coda

- Coda File System<sup>13</sup>, Project web page, Carnegie Mellon University.
- Peter J. Braam, The Coda Distributed File System<sup>14</sup>, Linux Journal, June 1998
- Mahadev Satyanarayanan, James J. Kistler, Puneet Kumar, Maria E. Okasaki, Ellen H. Siegel, David C. Steere, Coda: A Highly Available File System for a Distributed Workstation Environment<sup>15</sup>, IEEE Transactions on Computers, Vol. 39, No 4, April 1990.

---

<sup>4</sup><http://pages.cs.wisc.edu/~remzi/OSTEP/dist-nfs.pdf>

<sup>5</sup><http://www.ietf.org/rfc/rfc3530.txt>

<sup>6</sup><http://ra.adm.cs.cmu.edu/anon/usr/ftp/itc/CMU-ITC-062.pdf>

<sup>7</sup><http://pages.cs.wisc.edu/~remzi/OSTEP/dist-afs.pdf>

<sup>8</sup><http://www.cs.cmu.edu/~help/afs/index.html>

<sup>9</sup>[http://msdn.microsoft.com/en-us/library/ee442092\(v=prot.20\).aspx](http://msdn.microsoft.com/en-us/library/ee442092(v=prot.20).aspx)

<sup>10</sup><http://www.samba.org/cifs/docs/what-is-smb.html>

<sup>11</sup>[http://en.wikipedia.org/wiki/Server\\_Message\\_Block](http://en.wikipedia.org/wiki/Server_Message_Block)

<sup>12</sup>[http://msdn.microsoft.com/en-us/library/windows/desktop/aa365433\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/windows/desktop/aa365433(v=vs.85).aspx)

<sup>13</sup><http://www.coda.cs.cmu.edu>

<sup>14</sup><http://www.cs.cmu.edu/afs/cs/project/coda-www/ResearchWebPages/docdir/lj98.pdf>

<sup>15</sup><http://www.cs.cmu.edu/afs/cs/project/coda-www/ResearchWebPages/docdir/tcc90.pdf>

## DCE

- Brad Curtis Johnson, A Distributed Computing Environment Framework: An OSF Perspective<sup>16</sup>, June 1991
- Norbert Leser, Towards a Worldwide Distributed File System<sup>17</sup>, Open Software Foundation, September 1990.
- Agustin Mena III and Carl Burnett, *Performance Characteristics of the DCE Distributed File Service*, IBM Corp.

---

<sup>16</sup><http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.51.4971>

<sup>17</sup><http://www.opengroup.org/dce/info/papers/dev-dce-tp4-1.ps>