

Lecture 4

Andrew File System

(AFS)

Richard J. Orgass
Heinz School
Carnegie Mellon University

Carnegie Mellon

1

References

- www.openafs.org
- Open AFS Documentation
 - AFS Administrator's Guide
 - AFS Administration Reference
 - AFS Quick Start Guide for UNIX
 - AFS Quick Start Guide for Windows NT/2000
 - AFS Users Guide
 - IBM AFS 3.6 Release Notes
- Kerberos Authentication
 - web.mit.edu/kerberos/www/
 - <http://web.mit.edu/kerberos/www/dialogue.html>
 - <http://www.ornl.gov/~jar/HowToKerb.html>

Carnegie Mellon

2

Cells

- Administratively Independent Site
- Administrator decisions
 - configuration
 - number of server machines
 - number of server processes
 - number of client machines
 - number and identity of users
 - etc.,
- Cells combine local files into a global filespace
 - access is transparent
 - user's don't know anything about file's location to access it
 - path name of file is all that is needed
- Demo list of cells and finding files in them

Carnegie Mellon

3

Volumes

- Enables distributing files across many machines while maintaining uniform namespace.
- Container for a set of related files
 - different sizes
 - smaller than a disk partition
- Can be replicated to different machines
 - read-only clones
 - location independence of clones
- Used for
 - improved performance
 - requests for files to multiple servers
 - popular volumes available in many locations
 - off site backup
- Relatively small size makes them easy to move from one partition or machine to another

Carnegie Mellon

4

Volumes -- 2 --

- Each volume corresponds logically to a directory in the file tree
 - keeps together
 - all data that makes up the files in the directory
 - a user's volume contains all of his files
- Unit of resource management
 - space quota is associated with each volume
 - number of replicas determines maximum number of read only users

Carnegie Mellon

5

Namespace and Access

- Provides
 - transparent access to all the files in a cell's file space
 - User doesn't know location of file on file server(s)
 - User doesn't know name and location of file servers
 - Uniform Namespace
 - Same path name to a file in ALL AFS cells
 - ▲ recall demo with cells slide

Carnegie Mellon

6

Mount Points

- A volume is mounted on a directory in the local file system.
 - Local directory is the root directory of the file system in the volume
 - The local directory is also the mount point

Carnegie Mellon

7

Replication

- Making copy or clone of a source read/write volume
- Placing copy on one or more additional file server machines in a cell
- Increases availability of contents
 - if one file server fails, another can serve the pages
 - users still have access when a server goes down

Carnegie Mellon

8

Cache Manager

- Files are cached on local machines to increase the speed of access to commonly used files
- Most frequently used files remain in cache, less used files are flushed to make space for new files
- Consistency of files is dealt with by means of callbacks.
- Callback
 - promise by a file server to a cache manager to inform the latter when a change is made to any of the data delivered by the file server:
 - When a writeable copy of a file is delivered a callback is sent with it. If the source version of the file is changed by another user, the file server breaks the callback associated with the file to tell the cache manager that it must update the cached copy
 - When a file server delivers a file from a read-only volume to the Cache Manager, the file server sends along a callback associated with the entire volume.

Carnegie Mellon

9

AFS Server Processes and CM

- File Server
 - delivers data files from the file server to local workstations as requested
 - stores files again when the user saves any changes to the files
- Basic Overseer Server (BOS Server)
 - ensures that the other server processes on its server machine are running correctly as much of the time as possible
- Authentication Server
 - helps ensure that communications on the network are secure
 - verifies user identities at login
 - provides facilities through which participants in transactions prove their identity to one another (mutually authenticate)
 - maintains authentication databases

Carnegie Mellon

10

Server Processes and CM -- 2

- Protection Server
 - helps users control who has access to their files and directories
 - users grant access to several other users at once by putting them in a group entry in the protection database
- Volume Server
 - all types of volume manipulation
 - helps administrator move volumes from one server machine to another to balance the workload among various machines
- Volume Location Server (VL Server)
 - maintains volume location database
 - records the location of volumes as the move from file server machine to file server machine
 - key to transparent file access for users

Carnegie Mellon

11

Server Processes and CM -- 2 --

- Update Server
 - distributes new versions of AFS
 - server process software
 - configuration information
 - critical to stable systems performance
- Backup Server
 - maintains the backup database
 - stores information related to the backup system
 - enables administrator to back up data from volumes to tape
 - data can be restored from tape if it is lost from the file system
- Salvager
 - not a server process in the sense others are
 - runs only after the file server or volume server fails
 - repairs any inconsistencies caused by the failure
 - administrator can run it directly

Carnegie Mellon

12

Server Processes and CM -- 3 --

- Network Time Protocol Daemon (NTPD)
 - not an AFS server process per say
 - plays a vital role in AFS Operation
 - synchronizes the internal clock on a file server machine with those on other machines
- Cache Manager
 - resides on an AFS client rather than a file server machine
 - not a process but part of the kernel on AFS client machines
 - communicates with AFS server processes
 - main responsibilities
 - retrieve files for application programs running on the client
 - maintain the files in the cache