

Introduction to AFS IMSA Intersession 2003

AFS Odds and Ends

Brian Sebby, IMSA '96

Copyright 2003 by Brian Sebby, data@imsa.edu . Copies of these slides are available at <http://www.sebby.org/afs/> .

.backup volumes

- When the `vos backup` command is run, it creates a static backup volume for the volume you backed up.
- This volume can be explicitly accessed by appending the string `“.backup”` to the volume name. As in our previous example, the backup volume of `“user.bob”` would be `“user.bob.backup”`.
- You can use this explicit name to mount the backup volume in your AFS space. This is commonly done to give users a link to their backup volume in their home directories; you could run a command like:

```
% fs mkm oldfiles user.bob.backup
```
- This allows users to access their old data (usually as it appeared yesterday), which can let the system administrator avoid performing restores if files are deleted accidentally.
- You can use `vos exa` on `“.backup”` volumes to find information about them, such as when the backup was created.

Odds and Ends

- The following slides are about a wide variety of topics, some of which were skipped earlier, some of which are described in further detail, and a few that are just for your information.
- Please suggest any other topics you'd like me to add to this.
- As always, for the most complete information about AFS, please read through the various manuals in the documents collection at <http://www.openafs.org/doc/index.htm> .
- All of the slides I have prepared are available at <http://www.sebby.org/afs/> .

Volume naming conventions

- Because AFS cells tend to have a large number of volumes, it is useful to have a logical way of creating volume names to make them easier to manage and remember.
- For instance, if you have a number of user home directories, you may want to create a volume to contain links to those home directories, as well as separate volumes for each user.
- For the main user volume, you could name it `“user”` and mount it as `/afs/<cellname>/user/` .
- Then, you could create user home volumes by using the format `“user.<login name>”`. For instance, the user `“bob”` would have a home volume called `“user.bob”`, which would be mounted under `/afs/<cellname>/user/bob` .
- You could do similar conventions for other types of volumes. For instance, if you had a number of applications in AFS with individual volumes, you could create an `“appl”` volume, and volumes such as `“appl.perl-5.6”`, `“appl.mozilla-1.3”`, etc.

.readonly volumes

- When a volume is created, three sequential volume IDs are assigned to it. These refer to the Read-Write, Read-Only, and Backup volumes. However, the Read-Only and Backup volumes do not get created until the volume is replicated or `vos backup` is run, respectively.
- When you access a volume that is mounted with its normal volume name (say, `“user.bob”`) and it is replicated, you will automatically go to the read-only version of the volume unless you go through a read-write mount point.
- However, if you want to explicitly refer to the Read-Only volume, you can refer to the volume with the string `“.readonly”` appended to it. In this example, you could explicitly refer to the read-only version as `“user.bob.readonly”`.
- This could be used if you wish to create a mount point that always points to a read-only volume. You can also use commands like `vos exa` on the explicit `“.readonly”` volume to get information about when it was last replicated, etc.

The vos backupsys command

- One handy command for creating backup volumes is the `vos backupsys` command. The usage is:

```
% vos backupsys -h
Usage: vos backupsys [-prefix <common prefix on
volume(s)>+] [-server <machine name>] [-partition
<partition name>] [-exclude] [-xprefix <negative
prefix on volume(s)>+] [-dryrun] [-cell <cell name>]
[-noauth] [-localauth] [-verbose] [-help]
```
- By specifying a prefix, server, or partition, you can create backup volumes for volumes that match the strings you listed. For instance, to backup all volumes that begin with `“user”` you could type:

```
vos backupsys -prefix user ,
```

and to backup all volumes on `machine1`, you could type:

```
vos backupsys -server machine1 .
```
- This is often used in `vos` cron jobs with the `-localauth` flag to automate the creation of backup volumes to use when performing backups of your AFS cell.

AFS Databases and Log Files

- The AFS databases are stored in the directory `/usr/afs/db` on the database servers.
- Files that end in `.DB0` are the actual databases, and the files that end with `.DBSYS1` are transaction files.
- The file names for the various databases are:
kaserver.[DB0|DBSYS1] – kaserver (Authentication DB)
prdb.[DB0|DBSYS1] – ptserver (Protection DB)
vldb.[DB0|DBSYS1] – vlserver (Volume Location DB)
bdb.[DB0|DBSYS1] – buserver (Backup DB)
- These should be backed up like any other Unix file.
- Log files are stored in `/usr/afs/logs`. The files are:
BosLog – log file for bossserver, FileLog – log file for fileserver,
VolSerLog – log file for volsrver, SalvageLog – log file for salvager
AuthLog – log file for kaserver, VLLog – log file for vlserver,
BackupLog – log file for buserver

Avoid Native AFS Backups

- The native commands to perform backups of AFS are outlined in the *AFS Administrator's Guide*.
- However, if you have access to Veritas NetBackup or another backup service that can back up AFS, I recommend you use that.
- The AFS backup suite is a convoluted series of commands that require you to have template files for your tape drive, run one command to find which tape to use, another to restore the data, and a third to actually control the tape drive.
- In addition, AFS backups are only on the volume level, so you must restore an entire volume (which will get restored with a new extension, mount that volume, and then extract what you want and remove the restored volume, even if you only want to restore one file.
- However, as backup solutions that can handle AFS tend to be expensive commercial products, the native backup utilities may be needed. It is more important to actually have backups of your data, even if it is difficult to do so.

Ubik

- The AFS databases stay in sync with each other via a database protocol known as Ubik.
- At any given time, one database machine is designated the sync site, and the other machines are secondary sites. The sync site has a read-write copy of the databases, while the other sites have read-only copies.
- When a change is made to a database, the sync site receives the change, changes its database, then informs the secondary sites. Each change increments the version number of the database.
- To decide which site is the sync site, the DB servers have an election, and each votes for the lowest IP address that it can contact.
- A server must have a strict majority of servers voting for it to be elected the sync site. This is why you want three database servers – if one machine goes down, the other two can still elect a new sync site as they will have a majority. When the machine comes back it will resync its databases.

AFS Ports

- The following are the UDP ports that AFS uses for communication:
7000 filesrver 7001 cache manager callback service
7002 ptserver 7003 vlserver (vldb)
7004 kaserver 7005 volsrver (volume management)
7007 bossrver 7008 upsrver
7009 AFS/NFS Translator rmtsys remote pioctl
7020 AFS backup coordinator 7021 AFS backup buserver
7025-7032 AFS backup tape controllers 7101 xstat
2106 fs monitor port (read by venusmon)
Next available port pts, kas, fs, klog, etc...
- Knowing these ports may be useful to use AFS in a firewall environment.
- See my slides about "AFS and Firewalls" at <http://www.sebby.org/afs/> for more information about this subject.

Using IP addresses in ACLs

- You can use IP addresses in ACLs to limit access to directories to certain machines.
- To do this, you have to create an entry in the protection database for the IP address or subnet address that you want to use.
- To create an entry for the IP address 123.123.123.123, you would type:
`pts createuser 123.123.123.123`
- To create an entry for the subnet 123.123.123.*, you would type:
`pts createuser 123.123.123.0`
- You could then add these new protection entities to ACLs or groups. For instance, you could create a group of all IP ranges at your site.
- Users connecting from that IP address or IP range would then have the rights associated with that ACL for that directory.
- Note that they do not need to be authenticated to AFS to have those rights as long as they are connecting from that IP address.

The Future of AFS

- Many things are in store for the future of AFS. Now that it is an open source project, new changes and innovations are starting to be introduced.
- Kerberos v5 will become the standard authentication method in future releases of AFS. You can download packages today to use a Kerberos server instead of the kaserver, but the old kaserver is still the default.
- New features such as disconnected operations are coming. This will allow you to use AFS when you are not connected to a network.
- Keep looking at <http://www.openafs.org> for more informaton.
- You can subscribe to the openafs-info@openafs.org mailing list to ask questions and keep up to date.