



ROYAL INSTITUTE
OF TECHNOLOGY

AFS Backups

Mattias Amnefelt

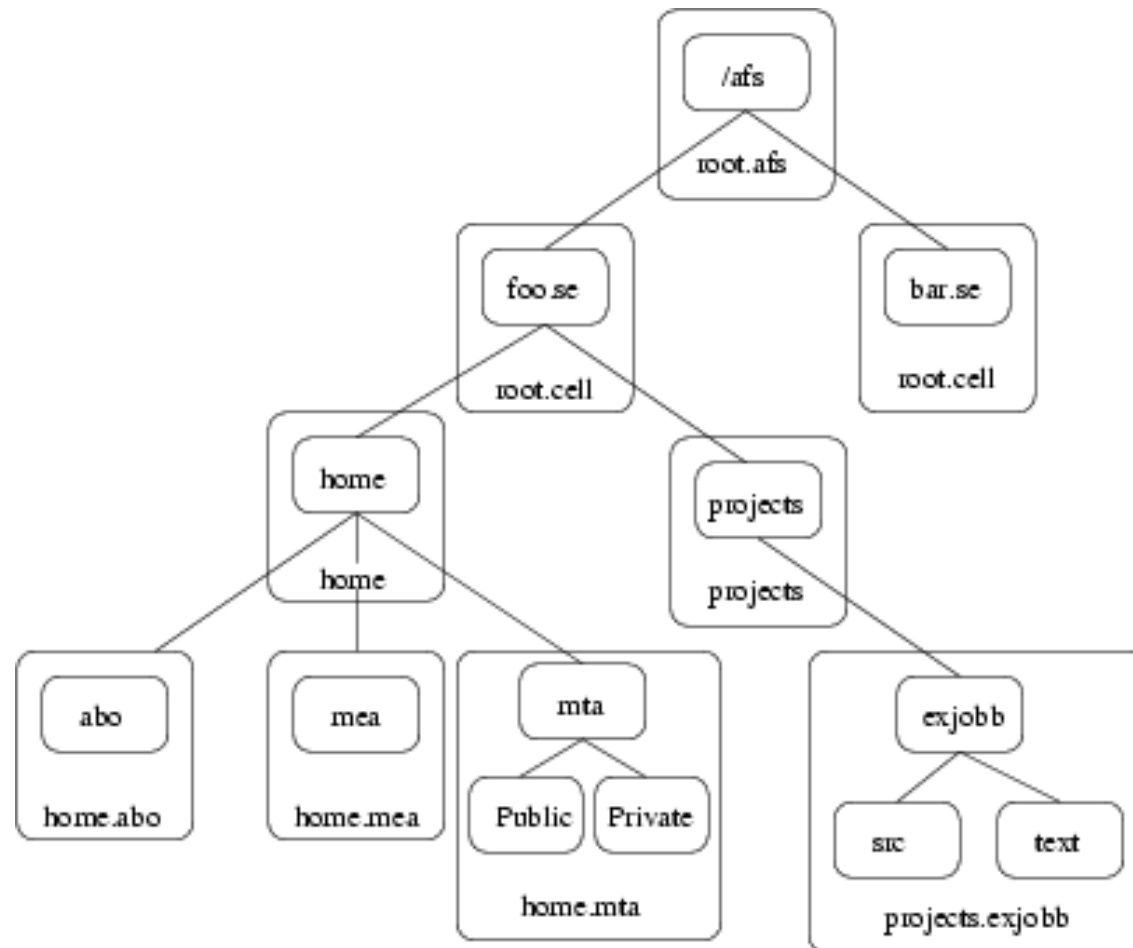
KTH

mattiasa@kth.se

AFS Structure



ROYAL INSTITUTE
OF TECHNOLOGY



Types of backups



ROYAL INSTITUTE
OF TECHNOLOGY

- Backups are taken for two distinct reasons
 - Crash recovery – when a disk or server crashes, the latest data needs to be restored
 - Retrieval of old data – if a file has been accidentally removed or overwritten, it may need to be restored from an older backup
- Backups in AFS can be taken in two different modes
 - Volume backups
 - File level backups

Backup volumes

- OldFiles makes restoration of data on a users behalf fairly unnecessary.
 - A snapshot of each volume as it was during the last backup is kept on the fileserver.
 - The user can go into this volume and restore files if they are accidentally removed.
 - This eases the load on backup administrators a lot.



ROYAL INSTITUTE
OF TECHNOLOGY

Volume backups

- Uses the AFS dump protocol
 - Same as for volume transfers between file servers
- Has two modes
 - Full - includes all data in the volume
 - Incremental – taken against a full, or another incremental



ROYAL INSTITUTE
OF TECHNOLOGY

Volume backups



ROYAL INSTITUTE
OF TECHNOLOGY

- Dumps include
 - Directory structure
 - File data
 - File system metadata, such as quota and ACLs
- A dump becomes a single file on the dump client, which can then be stored to secondary storage.

File level backups



ROYAL INSTITUTE
OF TECHNOLOGY

- Uses the file system client to access file data.
- Stores files to the backup system in the same way as local backups.
- Does not automatically include file system metadata, which must be archived separately.
- Is often done by mounting each volume in a separate area of the file system and then backing up all files in that volume w/o traversing mount points.

Method comparison



ROYAL INSTITUTE
OF TECHNOLOGY

- Volume backups
 - + The way AFS was designed to be backed up.
 - + Includes all metadata
 - Can waste space because pruning is done on volume basis
 - Must restore an entire volume to restore a file
- File level backups
 - + Can save space, if your backup system is intelligent
 - + Can restore single files
 - Cannot use standard tools
 - Must archive metadata separately
 - Depends on the file system client

AFS backup software



ROYAL INSTITUTE
OF TECHNOLOGY

- OpenAFS (and Transarc AFS) comes with a backup toolkit which include
 - Backup db run on the AFS database servers
 - Tape control software run on the tape server
 - Backup command client used to initiate backups
- It is a fairly competent backup system, but it is difficult to automate.

ITE Backup example



ROYAL INSTITUTE
OF TECHNOLOGY

- At the IT-Department at KTH we have a IBM Tivoli Storage Management System (TSM)
- We use volume backups
 - Using vos dump we write the volume to a file on the backup server
 - This file is then inserted into TSM as an archived file.
 - We only use TSM for storage and bypass all of TSMs automatic pruning.

ITE Backup example



ROYAL INSTITUTE
OF TECHNOLOGY

- Some characteristics of our backups
 - Full backups are taken every 60 days – if needed
 - Incrementals are taken against these full backups – if needed
 - If the sum of all incrementals are larger than a new full would be, then a full is taken
- Cleaning
 - Always keep the latest two fulls.
 - Always keep the latest 30 incrementals
 - After that, time is divided into ranges of exponentially increasing size
 - The newest full in each range is kept
 - Gives one full 2, 4, 8, 16... months old

AFS with amanda



ROYAL INSTITUTE
OF TECHNOLOGY

- Cornell University has developed a toolkit for integrating AFS backups with the Amanda backup software.
- It uses volume level backups and stores the volumes to an Amanda server

File level backups



ROYAL INSTITUTE
OF TECHNOLOGY

- Stanford is one AFS-user I know does file level backups.
- I don't know enough about these to have anything to say.
- Perhaps someone in the audience does?

References



ROYAL INSTITUTE
OF TECHNOLOGY

- The OpenAFS backup software is introduced in the AFS Administrators Guide and Reference available at <http://www.openafs.org/>
- KTH's afs-tsm is available at <ftp://ftp.e.kth.se/pub/tsm/afs-tsm/>
- Cornell's afs-amanda is available at <ftp://ftp.ccmr.cornell.edu/pub/amanda-afs/>

Questions, comments, rants?



ROYAL INSTITUTE
OF TECHNOLOGY

Mattias Amnefelt

KTH

mattiasa@kth.se