

# **AFS Usage and Backups using TiBS at Fermilab**

Presented by Kevin Hill

# Agenda

- History and current usage of AFS at Fermilab
- About Teradactyl
- How TiBS (True Incremental Backup System) and TeraMerge works

# AFS History at FERMI

- Original Mission
  - Pilot program started on FNALU (the central Unix cluster)
  - Used for general purpose storage
  - Many multi-flavored “smaller” servers
  - FNALU to be Restricted to physics applications
  - AFS seen as a method to share physics data site-wide and to other HEP labs that were beginning to use AFS

# AFS History at FERMI

- Quickly the mission changed
  - Growth of FNALU general user base due to migration off VMS
    - FNALU became a general purpose UNIX cluster with a focus on Physics analysis
  - Need for more CPU changed the system spec to a few large multi-cpu machine

# The Last Few Years

- Usage increase
  - Introduction of Linux a large factor
  - Introduction of a stable Windows client
    - provides a method for file sharing between Windows and UNIX
  - Many more central web servers hosted from AFS
  - Scientific users finally appreciate a single login environment and password

# Statistics

	6/2002	3/2003
AFS Users	3427	3773 (+10%)
AFS Groups	640	696 (+9%)
Experiments/ Projects	33	33
Clusters	9	12* (+33%)

# Statistics

	6/2002	3/2003
Servers	11	11* (Many infrastructure changes)
Storage	2.6 TB RAID5 (1.6TB on SAN)	4.2TB RAID 5 (all on SAN)
Network	100Mb/s Switched	100Mb/s Switched

# Statistics

	6/2002	3/2003
Disk in Use	1.2TB	1.5TB (+25%)
Volumes Served	11766	12564 (+7%)



# AFS Currently

- Sparc Servers
- SAN Based Storage
- 3 Database Servers
- 7 File Servers
- 1 Backup Server
- Running OpenAFS

# Current Usage

- 4.19 TB total storage
- 1.42 TB in use
- 25841 Volumes
- 3667 AFS users

# Backups

- Original backup system
  - 5 DLT 7000's
  - Operator assisted tape mounts
  - Scripts using native afs backup utilities
  - Window approaching 40+ hours for Full backup

# Backups

- Current System
  - 4 AIT3 drives
  - 120 slot tape library
  - Using TiBS Software from Teradactyl
  - 4-5 hour backup window
    - Now includes over 180 non-afs client machines
  - Full over network rarely needed

# About Teradactyl®

- Founded in 1999
- Located in Albuquerque, New Mexico
- Company develops advanced backup and recovery solutions
- Support for AFS®, Linux®, Mac® OS, UNIX®, and Windows®
- Primary product line: True incremental Backup System® (TiBS)
- Patented TeraMerge® technology
- Primary business focus is education and government.
- DOE sole source for AFS backup solution at Fermi National Accelerator Lab.


# Example: 3 level backup (4 week cycle)

Level 0: Full backup every four weeks

Level 1: Weekly incremental backup

Level 2: Daily incremental backup

Typical network bandwidth over a four week period:



# What's wrong with the traditional approach?

## Tradeoff between backup and restore

1. Fewer levels make restores faster, but increase the impact of backups on networks and clients.
2. More levels reduce the impact of backups, but increase restore time, complexity, and risk to tape failure.

## Multiple copies of unchanging data is inefficient

Periodic full backups are taken when only a percentage of the data has changed. Incremental backups continue to take copies of changed data even after it has already been backed up.

# TeraMerge

- Minimizes the time required for daily network backups
  - TiBS clients only send file changes since the last successful backup.
- Produces lower level backups without network intervention
  - TiBS reuses data to generate new backup volumes. This process eliminates the need to copy files again from backup clients.
- Creates full backup volumes entirely offline
  - A TiBS server uses the most recent backup data to create a current full backup image.



## Incremental Network Merge Backup



This process copies changes since the last successful network backup (Full or Incremental) from a backup client volume into the backup server cache. These changes are merged with previous incremental data to create a new incremental backup volume. The new volume represents all of the current changes since the last lower level backup. Incremental Network Merge Backups can be taken in parallel from several backup clients to a backup server disk. As each backup completes, the data is streamed efficiently to tape from the backup server cache. By reusing data in the backup server disk cache, this process reduces the workload that daily backups place on networks and backup clients.

## Flush Merge Backup



This process takes cumulative data from a higher level Flush Merge or Incremental backup in the backup server disk cache and merges it with the previous Flush Merge backup at this level. The merged data is sent to the backup server cache. As individual backup volumes are completed they are efficiently streamed to tape. If there are no pending lower level backup, the data is removed from the backup cache once it has been written to tape. The Flush Merge process may be performed in parallel from multiple tape devices. Additional levels of Flush Merge backups allow a backup server to support larger amounts of data.

## Full Merge Backup



This process takes cumulative data from a higher level Flush Merge or Incremental Backup in the backup server disk cache and merges it with the previous Full Backup tape volume. The merged data is sent to the backup server cache. As individual backups are completed they are efficiently streamed to tape. The data is removed from the backup cache once it has been written to tape. This process may be performed in parallel from multiple tape devices.



## The key advantages of TeraMerge®

- Reduce loads on networks and backup clients
- Reduce number of tapes required for restores
- Reduces average restore times
- Recovers from single tape failures on lower level backups
- Current incremental data is mirrored in the cache
- Lower tape cost than tape mirroring and tape striping.

# TiBS Backup Features

- TiBS only needs to backup new and changed files – No periodic network full backups are required or recommended!
- Automatic verification of tapes as they are merged
- Backup sub-directories and not just entire file systems
- Single file backup capability
- Omit rules to eliminate unnecessary data backup
- Windows registry and security data
- UNIX special files are supported by TiBS
- Disk caching allows parallel backup processing
- TeraMerge of lower level off-line backups
- Automated generation of mirrored volumes
- Status notification can be sent to different administrators based on configuration
- Flexible and customizable reporting

# TiBS Restore Features

- Efficient single pass restore
- Restores only versions of files required
- Restore incremental data from disk cache
- Search for any file on any tape
- Recover backup server from tape
- Redirect restore to any other TiBS client
- Data checksums ensure data is restored correctly
- Can restore data while continuing with backups
- Does not restore data which was intentionally deleted

# TiBS Archive Features

- Long-term storage managed by site tape retention policies.
- Tape pools can have individual tape retention policies.
- Original files can be retained or manually deleted by backup administrator.
- A reference of every file is automatically maintained by the TiBS File Lookup Database (FLDB) for easy location.
- Archive tapes can be managed in separate tape library for easy retrieval.
- Tape scan utilities can check integrity of stored tape volumes over time.
- Archive tapes are produced by taking a final incremental from a backup client and then generating a final consolidated full backup from data already on the backup server.

# TiBS Centralized Solution for AFS

Generates new full and lower level backups offline

Detects and reports corrupted volumes and orphaned vnodes

Reports discrepancies between vldb and file servers

Uses UNIX backup client to backup vldb and file server critical information

Online file lookup database (CMU-ECE 20GB since Jan 2000)

Only processes volumes which have been updated/created on a daily basis

Supports both IBM-AFS and OpenAFS



# OS Support

- Supports a wide variety of Linux OS clients.
- Backup Servers available in Linux and Solaris operating systems.
- Backup Server supports AFS, Mac OS, UNIX, and Windows clients.
- Backup Servers are compatible with the automated tape library interface.
- Tape support for AIT, DLT, SDLT, LTO and other popular formats.
- Clients available for Windows, Linux, Solaris, SGI, and MacOS

# Conclusions

- TiBS allows decent frequency of Full Backups without the network overhead
- Produces real full tapes
- Easy to use command line tools
- Needs GUI tools
- File Database very handy
- <http://www.teradactyl.com/>
  - Lots of FAQ's and stuff...