

True incremental Backup System[®]



Multiple Level Synthetic Backup Consolidation for AFS

Overview

Terminology/Backup Model

Network Backups

- . Initial Full
- . True incremental

Mixed Mode Backup

- . Network Synthetic Cumulative Incremental

Synthetic Backup

- . Full
- . Cumulative Incremental
- . Partial Cumulative Incremental

Workload Comparisons

Future Development

Terminology (General)

Primary Storage: Data being protected. An individual AFS volume. Some or all of the volumes within an AFS Cell. Multiple AFS Cells.

Backup Volume: A full or partial copy of Primary Storage data. For an AFS volume this is a “vos dump” image.

Secondary Storage: Backup media (disk or tape) where Backup Volumes are stored.

Backup System: System that creates Backup Volumes of Primary Storage data and stores the results in Secondary Storage.

Backup Cache: Hard drive staging area where Backup Volumes are generated before being moved or copied to Secondary Storage.

Disk Library: Extension of a Backup Cache that allows a Backup Volume to be retained on disk (nearline) in addition to or instead of tape.

Terminology (Types of Backup Volumes)

Full (level 0): A complete copy of an AFS volume as it existed at a point in time


Cumulative Incremental (level n): Contains the changes to an AFS volume since the most recent backup at level (n-1) or lower.

Differential (level 1): The changes to an AFS volume since the most recent Full Backup.

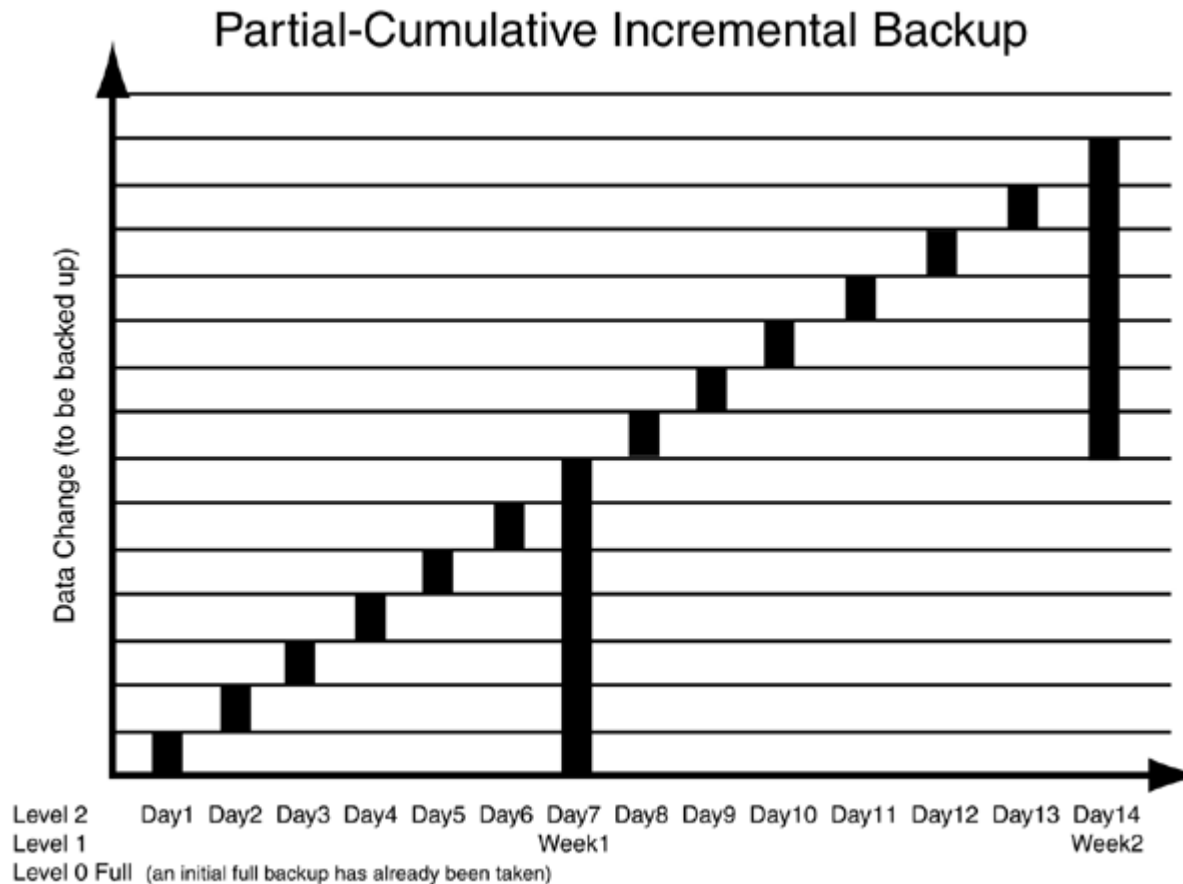
Partial Cumulative Incremental (level n): The changes to an AFS volume since the most recent backup at level (n) or lower.

True incremental: The changes to an AFS volumes since the most recent Backup.


Differential (level 1): The changes to an AFS volume since the most recent Full Backup



Partial Cumulative Incremental (level n): The changes to an AFS volume since the most recent backup at level (n) or lower.



True incremental: The changes to an AFS volumes since the most recent Backup



Terminology (Backup Methods)

Network: Backup Volumes are generated by copying data directly from Primary Storage. The Backup System communicates directly with the AFS Cell to obtain data.

Sythetic: Backup Volumes are generated from data already existing in the Backup Cache and Secondary Storage. No communication with the AFS Cell is required.

Need more info?

www.teradactyl.com/Products/BackupsDefs.html

Backup Model (Centralized Backup System)

Dedicated backup server hardware

- Minimize the impact of backups on the cell and network
- Maximize the amount of data a backup server can support
- Use backup hardware for other backups (UNIX/Windows)

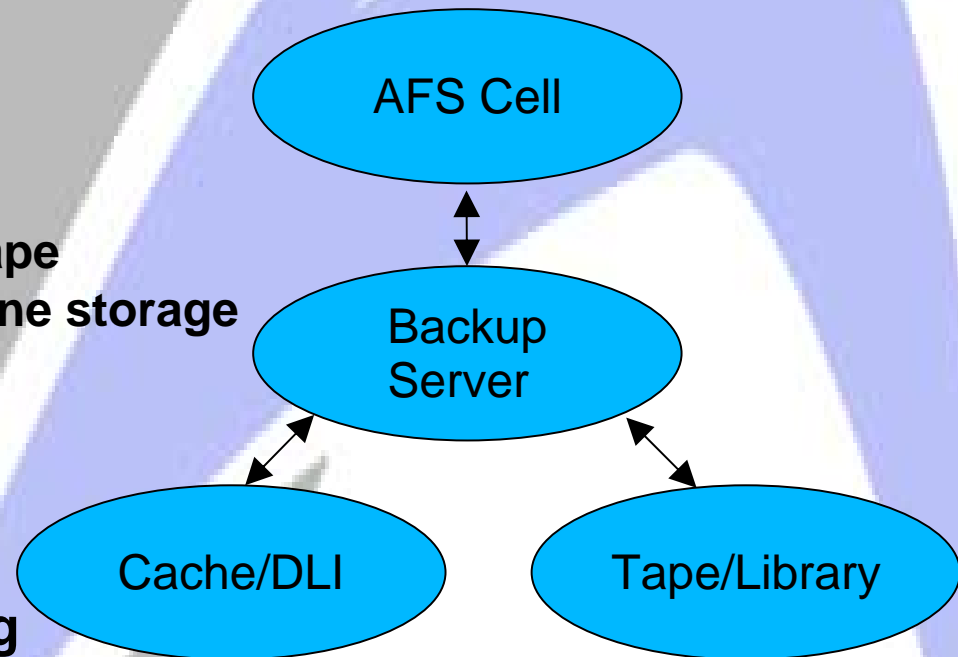
Backup Cache

- Enables parallel backups
- Streams backup volumes to tape
- Optional Disk Library for nealine storage

AFS Volume relocation

Need more info?

“Considerations when Choosing a Backup System for AFS”



Network Full Backups

Used to take an initial complete copy of each AFS volume from a cell

Three basic steps (order of operations is important!)

- . AFS Volume identification
- . Generation of updated .backup volume
- . Full backup of .backup volume

AFS Volume identification (afsgen)

- . VLDB volume listing
- . Fileserver volume listing
- . Detects duplicate volumes
- . Reports Off-line volumes
- . Reports inconsistencies between vldb and file server states
- . Generates volume listing with server location and last update

Network Full Backups

Generation of updated .backup volume

- . Run vos backupsys at scheduled time
- . Let TiBS do it

Full backup of .backup volume

```
vos dump -time 0 -id rw_vol.backup | afssrv -t last_update(rw_vol/afsgen)
```

What does afssrv do?

1. Checks volume time stamps (as gmt seconds)
 $last_update(rw_vol/afsgen) < creation(rw_vol.backup)$
2. Creates a Backup Volume in the Backup Cache (backup0.dat)
3. Creates an online file lookup database (backup.fldb)
4. Saves state information (backup.tdb)
 $BV_Time = last_update(rw_vol/afsgen)$
 $BV_ParentTime = BV_Time$
5. Signals a Pending Backup job (remove_cache=1)

Transfer to Secondary Storage

All new Backup Volumes are created in the Backup Cache

Secondary Storage is defined as Media Pools

- . Store data to disk or tape
- . Store data to disk and tape
- . Optional tape mirroring for offsite

Disk Only (tibsdld)

- . Uses virtual tape number and offset
- . Creates a subdirectory in the the backup cache
/cache/tibsdld/pool_type/pool_name/tapenumber/tapeoffset
(e.g. /cache/tibsdld/full/Full/1/1)
- . Hard links to data in the backup cache

Tape Only (tibstaped)

- . Copies data from the backup cache to tape (tape block size)
- . Writes data to two tapes simultaneously for tape mirroring
- . Backup Volumes may span multiple tape volumes

Transfer to Secondary Storage

Disk and Tape (tibstaped)

- . Writes data to tape first
- . Performs Disk Library updates
`/cache/tibsd/ pool_type/poll_name/tapenumber/tapeoffset`

Additional Cleanup

- . Copy file lookup database from the backup cache to online media databases
- . If the `remove_cache` flag is set, then remove the vos dump image from the backup cache and copy key state information to a fast cache recovery location

Network True incremental Backups

Used to take changes for each updated AFS volume since the last backup from the .backup volume.

Three basic steps (order of operations is important!)

Volume identification (afsgen)

Generation of updated .backup volumes

True incremental backup of .backup volumes

True incremental backups

```
vos dump -time BV_Time(cache) -id rw_vol.backup |  
afssrv -t last_update(rw_vol/afsgen) -p BV_Time(cache)
```

NOTE: If `last_update(cache) == last_update(rw_vol/afsgen)`
then no backup is required!

Network True incremental Backups

What does afssrv do?

1. Checks volume time stamps (as gmt seconds)
`last_update(rw_vol/afsgen) < create(rw_vol.backup)`
2. Creates a Backup Volume in the Backup Cache (backup0.dat)
3. Creates a file lookup database (backup.fldb)
4. Saves state information (backup.tdb)
`BV_Time = last_update(rw_vol/afsgen)`
`BV_ParentTime = BV_Time (previous cache)`
5. Signals a Pending Backup job (remove_cache=?)
Based on this backup level's media pool definition
cumulative incremental: `remove_cache=0`
partial cumulative incremental: `remove_cache=1`

Network Synthetic Cumulative Incremental

A mixed mode backup with both a network and synthetic backup component.

Four basic steps (order of operations is important!)

Volume identification (afsgen)

Generation of updated .backup volumes

True incremental backup of .backup volumes

**Synthetic backup consolidation with the previous
cache Backup volume**

True incremental Backup component

Transfers all directories and files to the new Backup Volume

Tracks file backup tags that represent unchanging file data

Network Synthetic Cumulative Incremental

Synthetic Backup Consolidation component

Opens the previous vos dump image in the Backup Cache

Skips all directory information

Transfers only files matching tags from True incremental

Appends remaining tags to the new Backup Volume (backup0.dat)

Saves state information (backup.tdb)

BV_Time=last_update(rw_vol/afsgen)

BV_ParentTime=BV_ParentTime(previous cache)

Signals a Pending Backup job (remove_cache=0)

Synthetic Full Backups

Generates an updated Full Backup using data from the Backup Cache and Secondary Storage.

Basic Steps:

- 1. Scan the current Cumulative Incremental Backup in the Backup Cache.**
Transfer all directories and files to the new Backup Volume
Remove and track any tags
- 2. Scan the previous Full Backup Volume**
Tranfers only files matching tags found in the cache Backup Volume
Verfies that there are no outstanding tags (backup0.dat)
- 3. Saves state information (backup.tdb)**
BV_Time = last_update(previous cache)
BV_ParentTime = BV_Time
- 4. Signals a Pending Backup job (remove_cache=1)**

Advantages of Synthetic Backups

No communication with the cell

No impact on the network

Can be run outside of network backup windows

Allows backup servers to run 24/7

Constant restore process continuously validates data

Built in redundancy, each new synthetic backup is a function of one or more backup volumes. Failed volumes can be regenerated w/o need mirror data.

Synthetic Cumulative Incremental Backups (level n)

Generates an updated Cumulative Incremental Backup using data from the Backup Cache and Secondary Storage.

Basic Steps:

1. Scan the current Cumulative Incremental Backup in the Backup Cache.
Transfer all directories and files to the new Backup Volume
Remove and track any tags
2. Scan the previous, level n, Cumulative Incremental Backup Volume
Tranfers only files matching tags from the cache Backup Volume
Appends remaining tags to the new Backup Volume (backup0.dat)
3. Saves state information (backup.tdb)
BV_Time = last_update(cache)
BV_ParentTime = BV_ParentTime (previous level n volume)
4. Signals a Pending Backup job
If the next lower level (n-1) backup is scheduled keep data in the backup cache, otherwise remove

Partial Cumulative Incremental Backups

Partial Cumulative Incremental (level n): The changes to an AFS volume since the most recent backup at level (n) or lower.

Does not require consolidation with previous level n backup

True incremental vs. Network Synthetic Cumulative Incremental

No need to read read/append previous Backup Cache Volume

Typically about ½ the size of a Network Synthetic Cumulative Incremental

The Backup Cache will always be empty after a True incremental Backup

True incremental Backups MUST be stored in the Disk Library

Recommended that they also be stored to tape

Synthetic Partial Cumulative Incremental Backups

Generates an updated Partial Cumulative Incremental Backup using data from the Backup Cache and Secondary Storage.

Does not require consolidation with previous level n backup

Different behavior depending on next higher level backup (n+1)

If level n+1 backup is a Cumulative Incremental Backup:

No consolidations are required

All data required for the level n backup is currently in the backup cache

Signals a Pending Backup job

If the next lower level (n-1) backup is scheduled keep data in the backup cache, otherwise remove it

Synthetic Partial Cumulative Incremental Backups

If level n+1 Backup is a Partial Cumulative Incremental Backup:

All level n+1 backups generated since the most recent level n or lower backup must be consolidated with any data in the backup cache to form a temporary level n+1 Cumulative Incremental Backup

Signals a Pending Backup job

If the next lower level (n-1) backup is scheduled keep data in the backup cache, otherwise remove it

This same temporary consolidation is required for a level n Synthetic Cumulative Incremental Backup and for Synthetic Full Backups.

Synthetic Partial Cumulative Incremental Backups

Advantages over Cumulative Incremental Backups

Smaller average backup volume sizes allow more nearline storage

Improves data recovery time from disk

Reduces server processing of network and mid-level incremental backups

Reduces tape costs for network and mid-level incremental backups

Allows same backup hardware to support more data

Backup server

Tape devices

Tape libraries

Caveats

Requires more disk space: Keep enough disk backups at each level to ensure that all backup consolidations occur from disk

Older restores may require additional tapes: As a rule of thumb, the older the data restore request, the longer the user is willing to wait.

Avg. Daily Data Transfer (Network/Cell)

Schedule	Traditional Systems	Synthetic Cumulative	Synthetic Partial Cumulative
1 Level (Daily)	100%	N/A	N/A
2 Level (Weekly)	15%	.6%	.6%
3 Level (Monthly)	5%	.6%	.6%
4 Level (Quarterly)	2.5%	.6%	.6%

Avg. Daily New Backup Generation (Backup Server)

Schedule	Traditional Systems	Synthetic Cumulative	Partial Cumulative*
1 Level (Daily)	100%	N/A	N/A
2 Level (Weekly)	15%	15%	15%
3 Level (Monthly)	5%	5%	4%
4 Level (Quarterly)	2.5%	2.5%	1.8%

* partially based on predicted results

Example Schedules (Network Backup Window/Offsite)

3 Level Backup

Level 0 Monthly Synthetic Full (mirror tape)

**Level 1 Weekly Synthetic Cumulative Incremental
(mirror tape)**

Level 2 Network True incremental (disk/tape)

Small disk library for 1-2 weeks of True incremental Backups

Copy of Week/Full tapes sent offsite

True incremental tapes sent offsite

Optionally mirror True incremental Backups for better onsite recovery.

Example Schedules (Backup/Recovery Performance)

4 Level Backups

Level 0 Quarterly Synthetic Full (disk/tape)

**Level 1 Monthly Synthetic Partial Cumulative Incremental
(disk/tape)**

**Level 2 Weekly Synthetic Parital Cumulative Incremental
(disk/tape)**

Level 3 Network True incremental (disk/tape)

Very large disk library (300%+ Primary Storage Size)

Small network backup window

All Synthetic Backups produced on disk

Short term restores (within a month) performed from disk

Example Schedules (Archive Tape Cost/Disk Space)

4 Level Backups

Level 0 Archive Quarterly Synthetic Full (disk/tape)

Level 1 Archive Monthly Synthetic Partial Cumulative Incremental (disk/tape)

Level 2 Weekly Synthetic Cumulative Incremental (tape)

Level 3 Network Synthetic Cumulative Incremental (tape)

Archive Monthly backups retained in the disk library long enough to generate new Synthetic Full Backups

Save \$1000s in archive tape costs vs. Monthly Cumulative Incremental

Older restores may require multiple Monthly tapes

Archive stability enhanced by mirroring Monthly tapes

Future Work relevant for AFS Backup and Recovery

Lazy Synthetic Consolidation

**Reduce Disk Library size requirements
Allow consolidations to occur on older data
(not just the most recent backups)**

Restores

**Single file/sub-directory restore
Restore data outside of AFS (restorevol)
GUI**

True incremental Backup System®



Thank you for your time!