
Oxford University
TSM Symposium 2003
New TSM Features:
Is it all Good?
STORServer
Backup Appliance
Kelly J. Lipp
STORServer, Inc.


Oxford University
TSM Symposium 2003
Agenda
TSM 5.x New Features - General
Newer Client Features
Journal Engine
Subfile Backup
Group Backup
Open File Engine
Multi-Stream to Tape (Simultaneous Copypools)
SAN Backup Techniques
STORServer
Backup Appliance
2

Oxford University TSM Symposium 2003

Considering New Features

- Are they right for you?
- Just because they built it doesn't mean you should use it!
- Evaluate your environment to determine which features makes sense for you.
- Implement and test

STORServer® Backup Appliance

3

Oxford University TSM Symposium 2003

Backup Client Journal Engine

- Separate service tracks file changes dynamically
- Client database logs file changes between backups
- When to use it:
  - When a client has a lot of files
  - Lengthy scan process is unacceptable
- When NOT to use it:
  - Small number of files

STORServer® Backup Appliance

4

Backup Client Subfile Backup

- Move only the bytes of a file that have changed rather than the entire file
- When to use it:
  - Wide area network implementations
  - Laptop Backups
- When Not to use it:
  - Server backup
  - Database backup

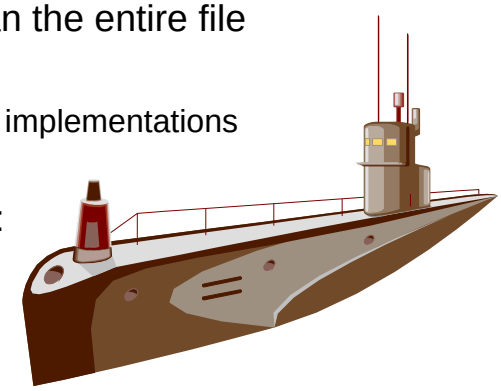

5

Backup Client Group Backup

- Allows creation of a logical entity of files spanning multiple filespaces
- Manage this group of files as an entity
- Sort of a backup “archive”
- When to use it:
  - When an application spans multiple drives
  - To capture a number of disparate files as a single entity
  - When an archive option doesn’t quite get it
- When NOT to use it:
  - When a regular backup will do

6

Oxford University TSM Symposium 2003

Backup Client Open File Backup

- Separate service (Logical Volume Snapshot)
- When to use it:
  - When quiescing applications is not an option
  - When you must back up an open file
  - When an agent does not exist for the application
- When NOT to use it:
  - When an agent exists

STORServer® Backup Appliance

7

Oxford University TSM Symposium 2003

Simultaneous Writes to Copy Pools

- Define stgpool poolname copystgpool= to direct TSM to multi-stream the data
- Up to 10 copystgpools can be defined
- When to use it:
  - Large objects: Oracle, SQL, Exchange
  - To reduce daily processing time
- When NOT to use it:
  - Small objects
  - Insufficient tape resources

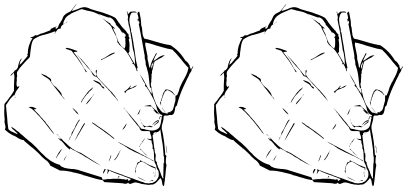

STORServer® Backup Appliance

8

SAN Backups

- Fallicy: SAN Backups are faster than LAN backups.
- Fallicy: SAN Backups are better than LAN backups.
- Truth: SAN Backups are more complicated and certainly more expensive than LAN Backups!
- Gigabit LANs are as fast as SANs
- Bottleneck is not the Wire!

STORServer® Backup Appliance

9

LAN Free Backups

- Clients have visibility to the drives via a SAN
- TSM Server controls library, drive and tape allocation
- Data moves from disk to the client to the tape
- Metadata moves via the LAN

STORServer® Backup Appliance

10

Server Free Backups

- Clients have visibility to the drives via the SAN
- Third Party Datamover exists in the SAN
  - Limited to IBM SAN Data Gateway
- TSM Client and TSM Server set up the paths
- Datamover is directed to move logical blocks from disk to tape.
- Data does NOT move through the client
- Metadata is moved via the LAN

11

SAN Backups

- When to use it:
  - When SAN infrastructure exists, and
  - Large object backup is necessary
- When NOT to use it:
  - When LAN Backups will do (which is most of the time)
  - When cost is an issue

12

Bottom Line

- Use New TSM Features when they make sense to you.
- Determine what you hope to gain by implementing a new feature.
- Research the ADSM.ORG list and APAR list for possible feature issues
- Test and ensure the feature is doing what you expect it to do

STORServer® Backup Appliance

13