




Tivoli Storage, IBM Software Group

# Using IBM Tivoli Storage Manager in a large scale environment

Frank Mueller ([muellerfrank@de.ibm.com](mailto:muellerfrank@de.ibm.com))  
NSC'08  
October 2008

# Tivoli Storage Manager

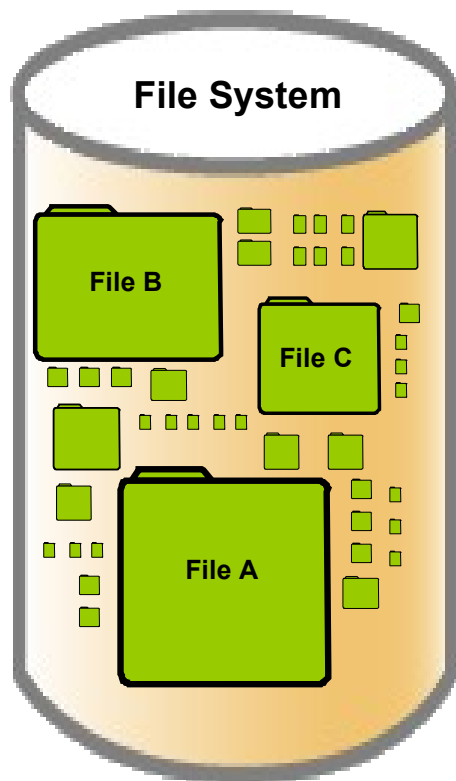


1. BA Client: Backup or Archive files to a TSM Server
2. BA Client: Restore or Retrieve files to a TSM Server
1. HSM Client: Migrate files to a TSM Server
2. HSM Client: Recall files to a TSM Server

# Agenda

- Introduction
- **What are the difficulties for a Backup Software?**
- *TSM BA client Backup Techniques*
- *TSM HSM clients - Overview*
- *TSM for Space Management – Lifting the Limit*
- *TSM for Space Management – the ILM solution*
- *Questions and Answers*

# What are the difficulties for a Backup Software?



Main Backup software is able to handle large (Peta Byte) file system

Performance to backup large files (TB or GB) depends on our lab environment.

To backup small files depends you need good techniques from your backup software.

# Backup Software challenges



Genomics & Drug  
Discovery



Check & Document  
Imaging



Digital Audio &  
Video Files



Medical Imaging &  
Patient Records



CAD & CAM Designs

# HSM – essential part of green ILM solutions



Reduce electrical  
power consumption  
through migration  
to offline media



Transparent  
Archive & Recall



Optimization of  
data storage  
through policy based  
media selection



Protection & Recovery  
through close backup  
integration

# Agenda

- Introduction
- What are the difficulties for a Backup Software?
- **TSM BA client Backup Techniques**
- *TSM HSM clients - Overview*
- *TSM for Space Management – Lifting the Limit*
- *TSM for Space Management – the ILM solution*
- *Questions and Answers*

## What is covered in this part of the presentation


- The various backup techniques offered by TSM Backup-Archive Client:
  - What you get vs. what you give up with different technologies and variations
  - Decision points
- Provide qualitative advice on evaluating different technologies and variations
  - What tools are available today to help with these decisions



## What is NOT covered in this part of the presentation

- How to implement:
  - Not a tutorial in option syntax
  - No dsm.opt or dsm.sys examples!
- Other backup or archive considerations
- Recovery or retrieve considerations

# Progressive Incremental Backup



1. Client queries server for current view of file system
2. Server returns list of active versions for entire file system
3. Client scans and compares with local file system
4. Client backs up new and changed files


# Progressive Incremental

- a.k.a Full Progressive Incremental or Full Incremental or “Incremental Forever”
- What you get
  - Single-instance store
  - Easier restore
    - No concept of applying incremental or differential changes to full
- When to use this
  - Default technology in TSM
  - Hot spots:
    - Memory
      - Number of active backup versions
      - Length of file name
    - Run time
      - Biggest factor is file system scan time

## Progressive Incremental Derivatives

- Memory efficient backup
- Memory efficient disk caching
- Virtual mount points
- Incremental by date
- File lists
- Journal-based backups

# Memory Efficient




- Client queries server for current view of first directory
- Server returns list of files in directory (db query)
- Client scans and compares with local file system
- Client backs up new and changed files
- Client queries server for current view of next directory, etc.

# Memory Efficient Backup

- What you get
  - Smaller backup memory footprint
- What you don't get
  - Memory savings in directories with large number of objects
- What you give up
  - More network chit-chat
  - Increased backup run time
- When to use this
  - Incremental backups run out of memory
  - Can get rough estimate on memory usage

# Memory Efficient Disk Caching



- Client queries server for current view of file system
- Server returns list of files for entire file system; client stores inventory on disk rather than in memory
- Client scans and compares with local file system
- Client backs up files

# Memory Efficient Disk Caching

- How it works
  - Same as progressive incremental except ...
  - B-A client can store TSM Server db queries on disk instead of in memory
- What you get
  - Even smaller memory footprint if regular memoryefficient backup doesn't reduce memory
- What you give up
  - Processing time (takes a bit longer to store and retrieve information from disk)
- When to use this
  - Incremental backup running out of memory and memoryefficient backup not sufficient Need individual file restore
  - Journal based backup technology is not available on your platform (note you might need this just to get the first backup done for journal based backups).



# Virtual Mount Point




- Large file system logically partitioned
- Managed as separate file spaces on TSM Server

# Virtual Mount Point

- How this works
  - Instead of mapping entire file system to single namespace on TSM Server (file space), user can define
- What you get
  - Progressive incremental only has to work against subset of file system thus reducing memory requirements
- What you don't get
  - Relief from directories with large number of objects
- What you give up
  - Centralized management of file system on TSM Server
- When to use this
  - Large, balanced (number of directories and objects) UNIX and Linux file systems can be efficiently divided

# Incremental by Date




- Client queries server last backup of entire file system
- Server returns timestamp of last backup (end) of entire file system
- Client scans and compares with local file system
- Client backs up files

# Incremental By Date

- What you get
  - Reduced time determining which files have changed
  - Removing processing time on server to query database
  - Removing network traffic to communicate results
- What you give up
  - Flexibility on scope of backup operation; you must be doing backups of the entire file system
  - File backups where changes do not affect date (attribute, mode, ACL)
    - Rename, copy, move, security change
  - Expiration of deleted files
  - Policy rebinding
  - Entire file system still needs to be scanned
  - No free pass for having client and server clocks out of synch
  - May not be able to use when client and server are in different time zones.
- When to use this
  - **Not meeting backup windows and ...**
  - File system changes purely additive or changes, e.g., code repository or data that is basically “additive” (think of data like your family photos)
- Can be effective when doing weekly (or periodic) full incremental backups
- Read Chapter 9 in the User’s Guide for more critical information on this option

# File List Backup




1. Application creates list of files for backup
2. Application passes list of files to client
3. Client backs up files (SELECTIVE backup)

## File lists

- What you get
  - Selective backup eliminates the query of the TSM Server db and scan of local file system
- What you give up
  - Time – you have to find some mechanism to feed the file list
  - Implicit file specifications – file list will not accept wild cards or directory recursion
- When to use this
  - **Not meeting backup windows and ...**
  - File system changes are predictable, e.g., an application is creating the changes and it would be easy to interface into the application to get that list of changes

# Journal Based Backup




1. Journal monitors file system for changes
2. Client queries journal for file system changes
3. Client backs up files

# Journal Based Backup

- What you get
  - Time to determine which files have changed is greatly reduced
- What you don't get
  - Relief from doing a full progressive (full) incremental backup against a file system on some periodic basis:
    - Initial backup to enable the journal
    - If a discrepancy is found between journal and TSM Server database
    - If the velocity of changing files is high enough to cause a notification buffer overrun
    - If the journal service is stopped and restarted
    - In general, it is a good idea to do periodic progressive incremental backups
- When to use this
  - **Not meeting backup window and ...**
  - Small number of files (< 1,000,000) and small number of changes between backups (<1,000,000)
  - Large number of objects (<10,000,000) with 10-15% change rate try JBB (should emphasize a low change velocity, too... large numbers of changes over a short timeframe can cause notification buffer overflows)
  - Large number of objects + large number of changes will stress JBB ... not a good fit



# Image Backup



1. Client sends logical block image of file system to server


# Image Backup

- No mapping between file and block
- Off-line and static backups vs. on-line backups
- What you get
  - Faster backups
  - No scan time to determine what has changed
  - Faster overall data movement
- What you give up
  - Ability to restore individual files directly from TSM Server (we know, we know...)
    - You would need to restore image to alternate location and pull data directly using Explorer or equivalent.

# Image Backup


- Image backup
  - Backup at logical volume level
  - Ability to take image + incremental backups to provide single file restore
- Off-line image backup
  - Volume mounted read-only
  - Available for AIX, Windows, Linux x86, Solaris and HP
  - Exploited best by flashcopy operations
- On-line image (dynamic) backup
  - Volume left on-line
  - Fuzzy backup of changing data
- On-line image backup using snapshot
  - Volume left on-line
  - Image backup taken at single point-in-time (“crash consistent”)
  - Available for Windows and Linux x86

# Image + Incremental by Date



- *dsmc backup image* sends logical block image of file system to server
- Upon subsequent *dsmc backup image -mode=incremental*, client queries server last backup of entire file system
- Server returns timestamp of last backup of entire file system
- Client scans and compares with local file system
- Client creates transactions of files for backup

## Image + Incremental by Date Restore Scenario




1. Client request incremental image restore
2. Server returns the base image
3. Server returns additional files which need to be applied to the base image to satisfy the recovery point

## Image + Incremental By Date


- What you get
  - Benefit of image backup plus some individual file protection of files that are changing
- What you give up
  - See Incremental by Date
  - No reconciliation of deleted files

## Image + Progressive Incremental Backup



- *dsmc backup image* sends logical block image of file system to server
- Upon subsequent *dsmc incremental*, client queries server for current view of file system
- Server returns list of files for entire file system
- Client scans and compares with local file system
- Client backs up changed files

## Image + Progressive Incremental Restore Scenario



- Client requests incremental image restore
- Server returns the base image
- Server returns additional files which need to be applied to the base image to satisfy the recovery point
- Server optionally returns the list of files which need to be deleted from the base image to satisfy the recovery point



# Image + Progressive Incremental


- **What you get**

- Benefit of image backup restore for improving RTO (Recovery Time Objective)
- Benefit of progressive incremental backups for better RPO (Recovery Point Objective)

- **What you give up**

- Additional time to take periodic image backups
- Storage space at TSM Server

## TSM Backup triggered via mmbackup from GPFS



- Full backup
  - GPFS create a file list (using internal snapshot functionality)  
(mmbackup device -n configfile [-t **full** | incremental] [-s sortDir])
- Incremental backup
  - GPFS create a diff file list (using internal snapshot functionality)  
(mmbackup device -n configfile [-t full | **incremental**] [-s sortDir])

## TSM Backup triggered via mmbackup from GPFS

### ■ What you get

- Reduced time determining which files have changed
- Removing processing time on server to query database
- Removing network traffic to communicate results
- Smaller memory footprint

### ■ What you give up

- Flexibility on scope of backup operation; you must be doing backups of the entire file system
- Changes in the ACL's trigger a backup (attribute, mode, ACL)
  - Rename, copy, move, security change
- Settings in the TSM management class will be ignored
- Expiration of deleted files
- Entire file system still needs to be scanned

### ■ When to use this


- Not meeting backup windows

# Agenda

- Introduction
- What are the difficulties for a Backup Software?
- TSM BA client Backup Techniques
- **TSM HSM clients - Overview**
- *TSM for Space Management – Lifting the Limit*
- *TSM for Space Management – the ILM solution*
- *Questions and Answers*

## HSM – in the TSM way of definition

- HSM automatically migrates files from on-line to less expensive near-line storage, based on policy settings
- Small *stub files* are retained on disk, appearing as the original files, thus ensuring transparency of HSM for user applications
- Stub files contain information pointing to corresponding entry in TSM server DB
- Migrated files are recalled automatically back to disk as required



# The TSM for Space Management toolbox



- Main Space Management modules:
  - Module for monitoring the file system (Monitor Daemon)
  - Module for searching candidates (Scout Daemon)
  - Modules for migrating files (Migrate & Automigrate)
  - Module for recalling files (Recall Daemon)
  - Module for synchronizing file system and TSM Server (Reconcile)
  - Helper Modules (e.g. special ls and df commands)

## Space Management version 5.5 – lifting the limits

- With version 5.5 HSM has become the ability to manage file systems with more than 100 million files.
- Major parts of HSM are completely redesigned
  - Module for searching candidates (Scout Daemon)
  - Module for synchronizing file system and TSM Server (Reconcile)

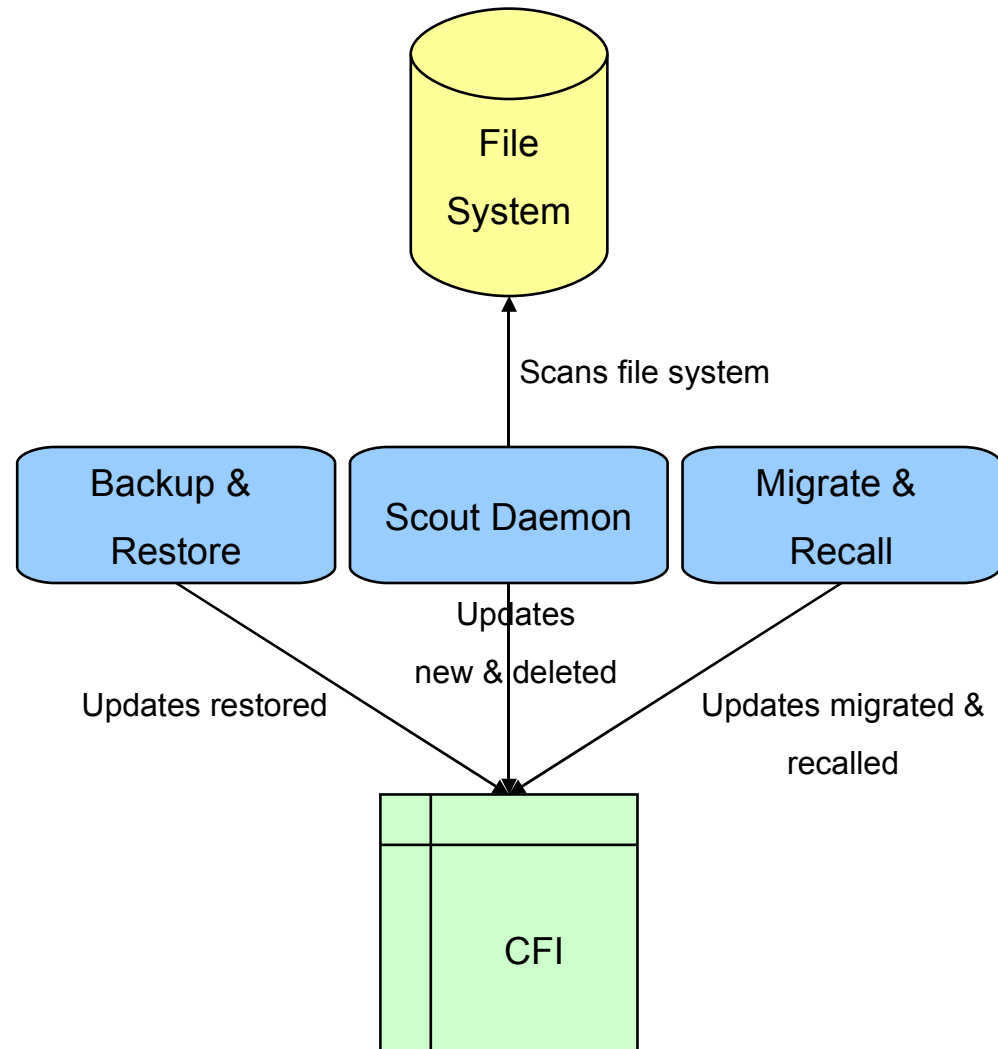
# Agenda

- Introduction
- What are the difficulties for a Backup Software?
- TSM BA client Backup Techniques
- TSM HSM clients - Overview
- **TSM for Space Management – Lifting the Limit**
- *TSM for Space Management – the ILM solution*
- *Questions and Answers*



## Lifting the limits - the scout daemon

- Based on a self updating metadata system
- Complete File Index (CFI) stores metadata information used for candidates selection
- Scout daemon scans the file system to find new and deleted files
- TSM tools automatically update CFI during migrate recall and restore activities

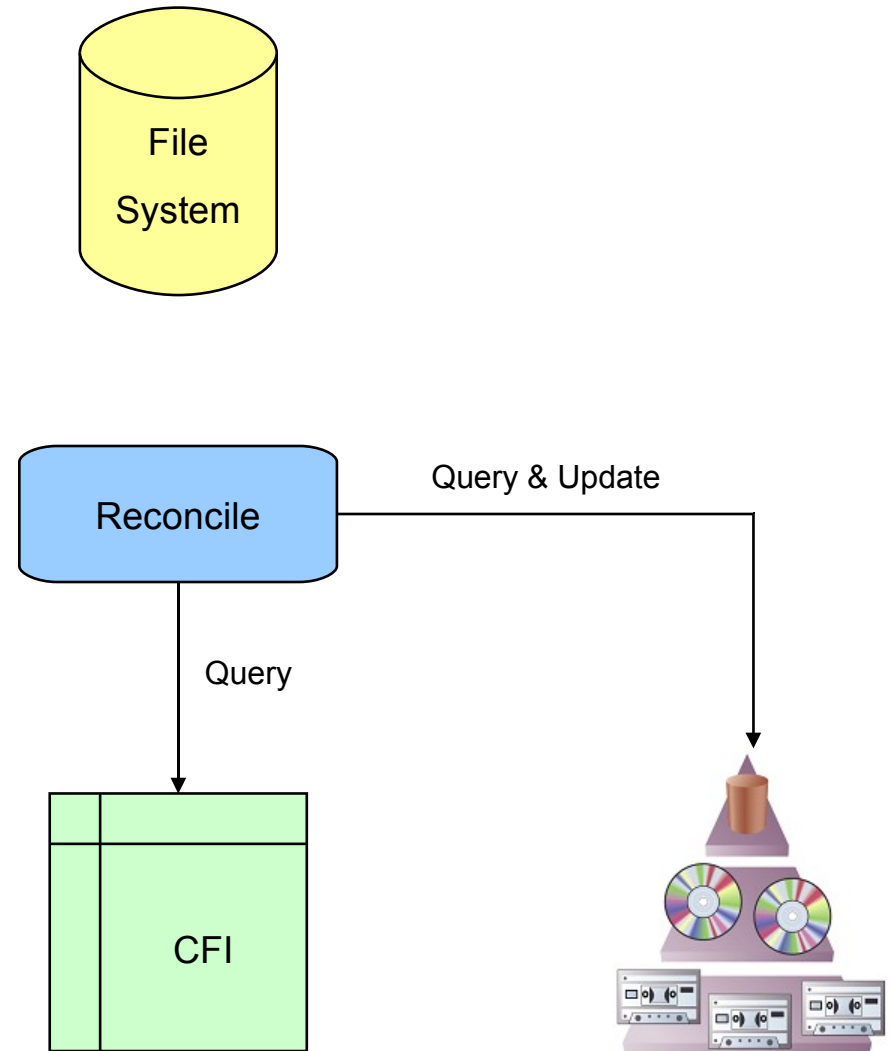


## Lifting the limits - the scout daemon

- If threshold for file system was reached scout daemons internal search engine queries the CFI for best candidates
  - Pre-migrated files
  - Biggest and oldest files
- The migration of pre-migrated files is very fast. HSM is able to free space in a small time frame

## Lifting the limits – the reconcile

- Reconciliation synchronizes the HSM managed file systems with the TSM Server DB
- The CFI is an image of the file system in metadata format
- The Reconciliation uses the CFI to do fast queries instead of long running file system scans
  - Increased performance
  - No limitation in number of files

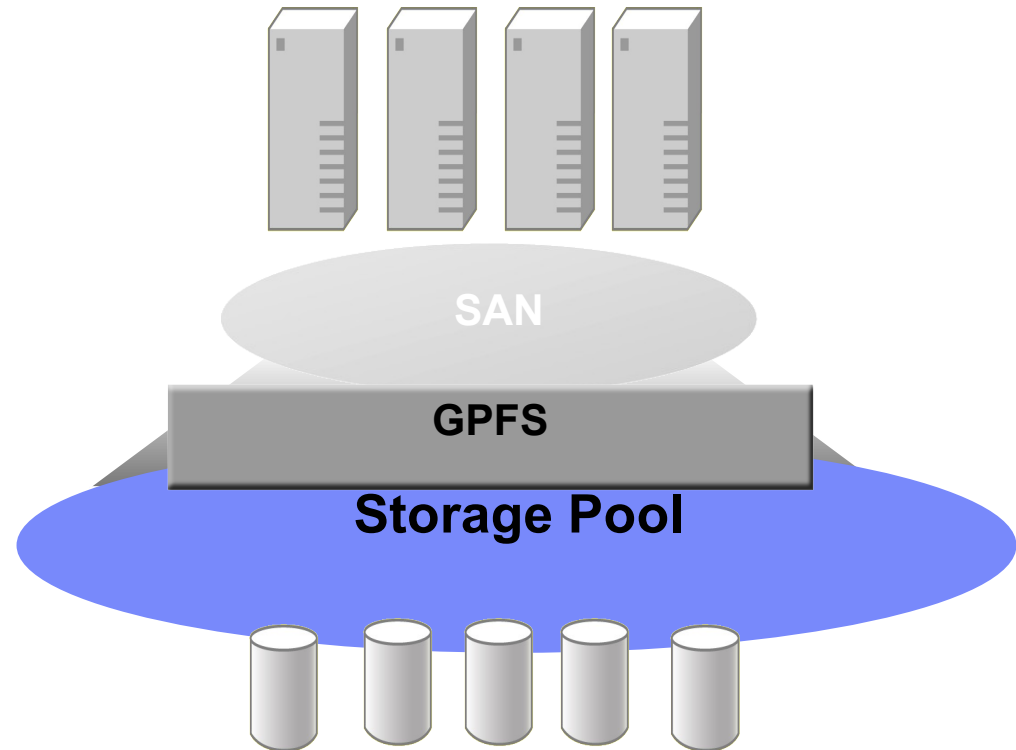


# Agenda

- Introduction
- What are the difficulties for a Backup Software?
- TSM BA client Backup Techniques
- TSM HSM clients - Overview
- TSM for Space Management – Lifting the Limit
- **TSM for Space Management – the ILM solution**
- *Questions and Answers*


## Excursion - GPFS without storage pools

- The entire file system is striped across all discs in one pool
- Seperate pools for data and metadata possible



## Excursion - GPFS with storage pools

- One name space across all pools
- Files in the same directory can be in different storage pools
- Files are placed in storage pools at creation time using policies
- Files can be moved between pools for policy reasons



## HSM - part of the GPFS ILM solution

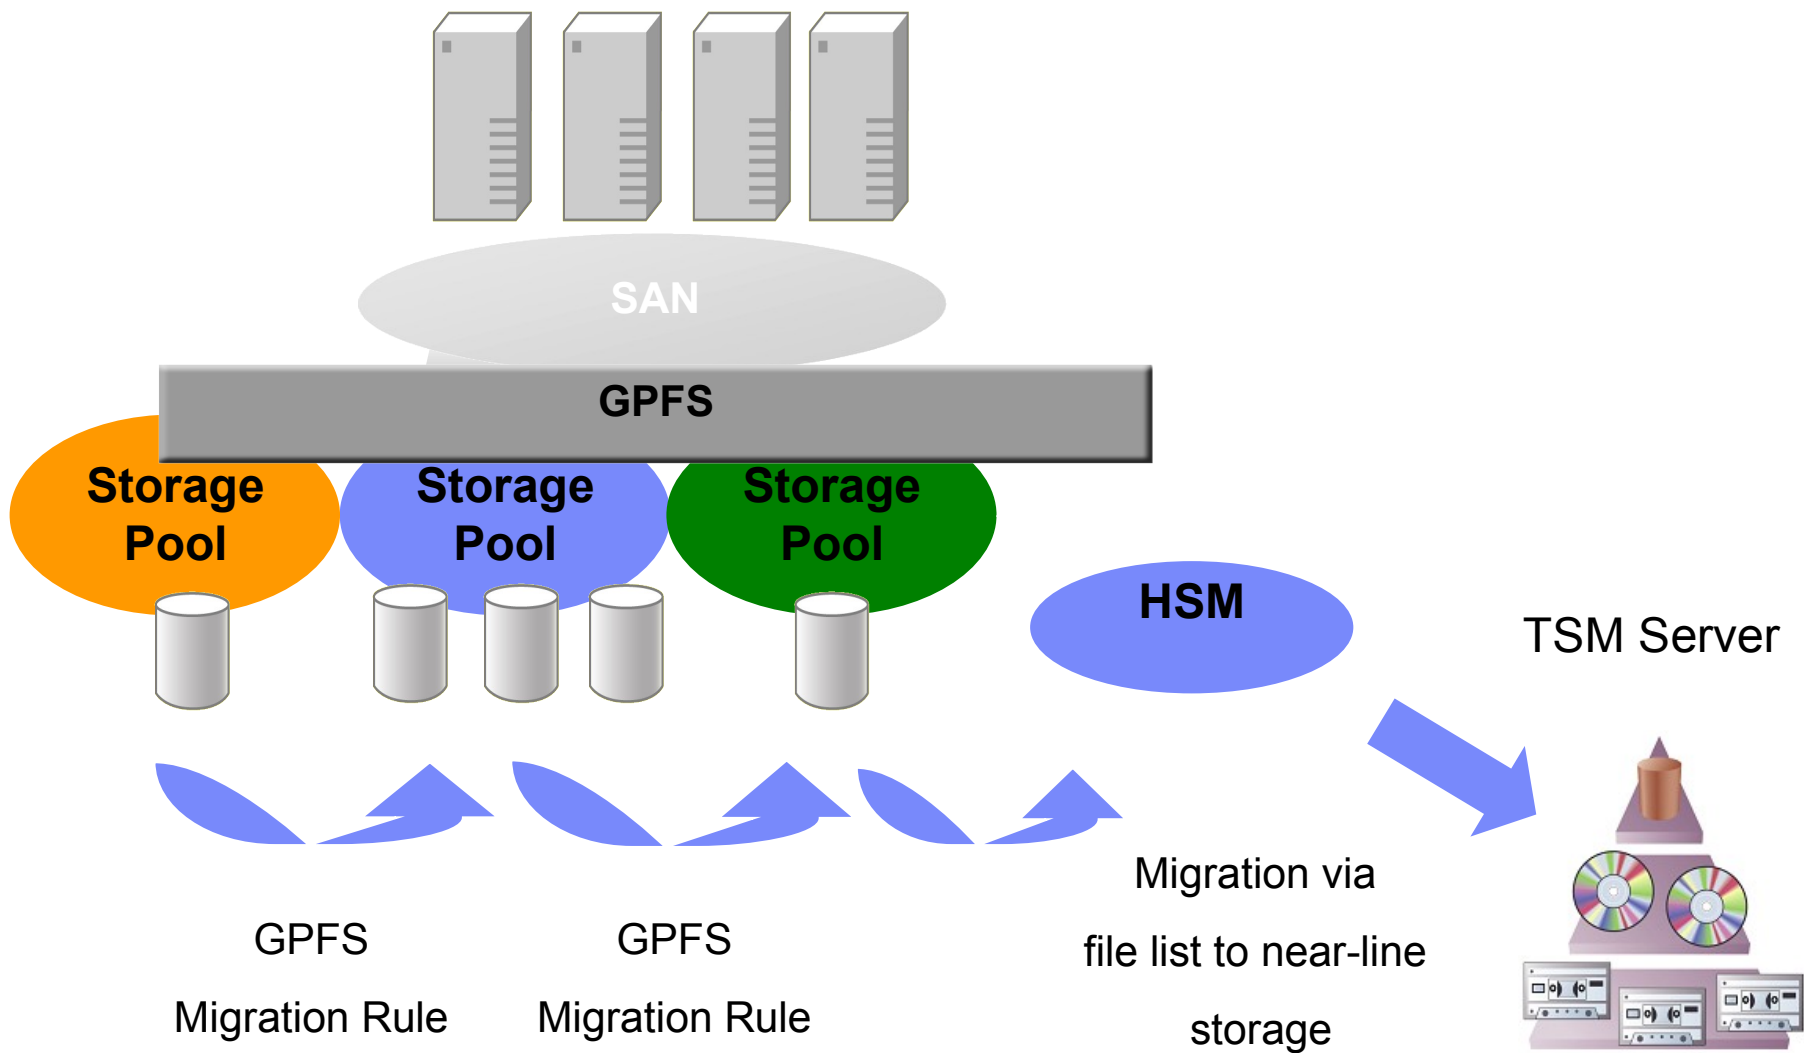
- GPFS has officially announced that TSM HSM is part of their ILM strategy.
  - TSM HSM enhances GPFS to the ability to do ILM to tape.
- With GPFS 3.2 and TSM HSM 5.5 the customer can use HSM as an external GPFS pool.
- Migration / automigration can be defined in GPFS through policies.

## HSM - part of the GPFS ILM solution

- Candidates selection is identically to GPFS native rule based pool to pool migration
- The GPFS policy engine is able to scan 1 million files in about 10 seconds and deliver migration candidates to HSM.
- In a GPFS cluster environment it is possible to handle more than 100 million files with this approach.



# HSM - part of the GPFS ILM solution



# HSM classical concept vs. GPFS ILM solution

## HSM driven management (classic)

- HSM monitors the file system and initiates migration. Automigration migrates data when high threshold is reached
- Scout Daemon delivers candidates based on the CFI
- HSM modules performing automated migration until the low threshold was reached

## GPFS driven management (GPFS ILM)

- GPFS policy daemon monitors the file system
- GPFS policy daemon scans the file system and delivers candidates
- GPFS calls HSM migration module with a list of candidates for migration

## HSM & GPFS - advantages of the close integration

- HSM enhances GPFS to the ability to do ILM to tape
- Policy based candidates selection instead of size and age criteria only
- Performance
  - Fast candidates selection
  - Distributed migrate from different cluster nodes

# Agenda

- Introduction
- What are the difficulties for a Backup Software?
- TSM BA client Backup Techniques
- TSM HSM clients - Overview
- TSM for Space Management – Lifting the Limit
- TSM for Space Management – the ILM solution
- **Questions and Answers**



# Questions ?

## References (TSM for Space Management)

- TSM for Space Management on the Web (also Demo available):
  - <http://www-306.ibm.com/software/tivoli/products/storage-mgr-space/>
- TSM for Space Management Client Manuals:
  - <http://www.elink.ibm.com/public/applications/publications/cgibin/pbi.cgi?CTY=US&FNC=SRX&PBL=GC32-0794-03>
- TSM for Space Management platform support information:
  - <http://www-306.ibm.com/software/tivoli/products/storage-mgr/platforms.html>
- TSM for Space Management IBM Tivoli Enterprise Space Management Console:
  - <http://www-306.ibm.com/software/tivoli/resource-center/storage/code-ent-con-stor-mgr-space.jsp>
- TSM for Space Management and HSM for Windows white paper (Update 08/08):
  - <http://www-1.ibm.com/support/docview.wss?uid=swg27002498>