



| Tivoli Storage, IBM Software Group

Facts and News about TSM HSM: Where are we and where will we go?

| Stefan Bender
Oxford University TSM Symposium
September 2007

Disclaimer

- This presentation describes potential future enhancements to the IBM Tivoli Storage Manager family of products
- All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only
- Information in this presentation does not constitute a commitment to deliver the described enhancements or to do so in a particular timeframe
- IBM reserves the right to change product plans, features, and delivery schedules according to business needs and requirements
- This presentation uses the following designations regarding availability of potential product enhancements
 - Planned 5.5: Planned for delivery in TSM v5.5 (2007)
 - Next Release Candidate: Candidate for delivery in the next release after v5.5
 - Future Candidate: Candidate for delivery in future release

Presentation Contents

This presentation covers details on IBM Tivoli Storage Manager for Space Management and on IBM Tivoli Storage Manager HSM for Windows

It provides a brief overview of the products, its functions and planned improvements for 5.5, and what is coming next.

Presenter Information

Stefan Bender



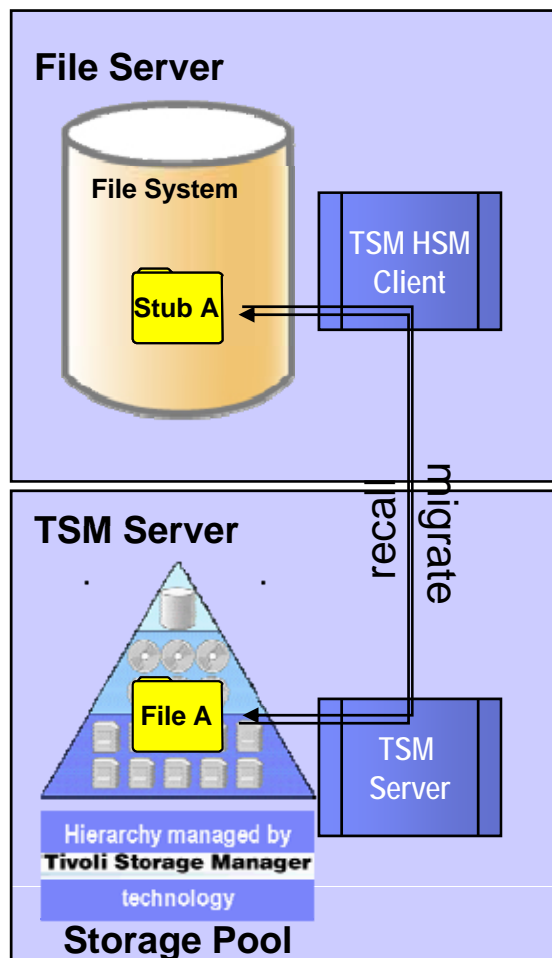
- IBMer since 2000, TSMer since 2002
- Located in Mainz, Germany
- IBM Tivoli Storage Manager Development of the HSM clients
- Recent Projects:
 - Reconciliation and Backup Integration for HSM for Windows

Agenda

- TSM HSM clients - Overview
- TSM HSM for Windows – Closing the Gap
- TSM for Space Management – Lifting the Limit
- TSM for Space Management – Entering new Territories
- TSM HSM clients - Futures
- Questions and Answers


TSM HSM clients - Overview

What is 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

TSM HSM clients - Overview

TSM HSM 5.4 is available in two flavors:

- **IBM Tivoli Storage Manager for Space Management**
 - AIX 5.3
 - JFS2
 - GPFS 2.3, GPFS 3.1
 - Linux: SLES 9, SLES 10 (32 Bit only), RHEL 4, GPFS 2.3 & GPFS 3.1
 - Solaris 9,10 VxFS 4
 - HP-UX 11i V2 JFS
- **IBM Tivoli Storage Manager HSM for Windows**
 - Windows 2003 Server & Enterprise Server (32 Bit)
 - Windows 2003 Server & Enterprise Server R2 (32 Bit)
 - all NTFS V5 (full support for Reparse Points and Sparse file needed)

TSM HSM clients - Overview

Major differences between TSM HSM on UNIX and Windows

UNIX HSM

- Capacity oriented approach (classic HSM)
- Automatic migration when a high threshold is reached
- Avoids out-of-space condition
- Command line

Windows HSM

- File value oriented approach (more like ILM)
- Granular migration jobs define exact Placement of individual files
- Does not avoid out-of-space condition
- GUI

TSM HSM clients - Overview

Major differences between TSM HSM on UNIX and Windows

UNIX HSM

- Uses Space Management Pool on TSM Server
- DMAPI
- TSM internal Interface for Client/Server Communication

Windows HSM

- Uses Archive Pool on TSM Server
- File System Filter / Reparse Points / Sparse Files
- TSM API used for Client/Server Communication

TSM HSM for Windows – Closing the Gap

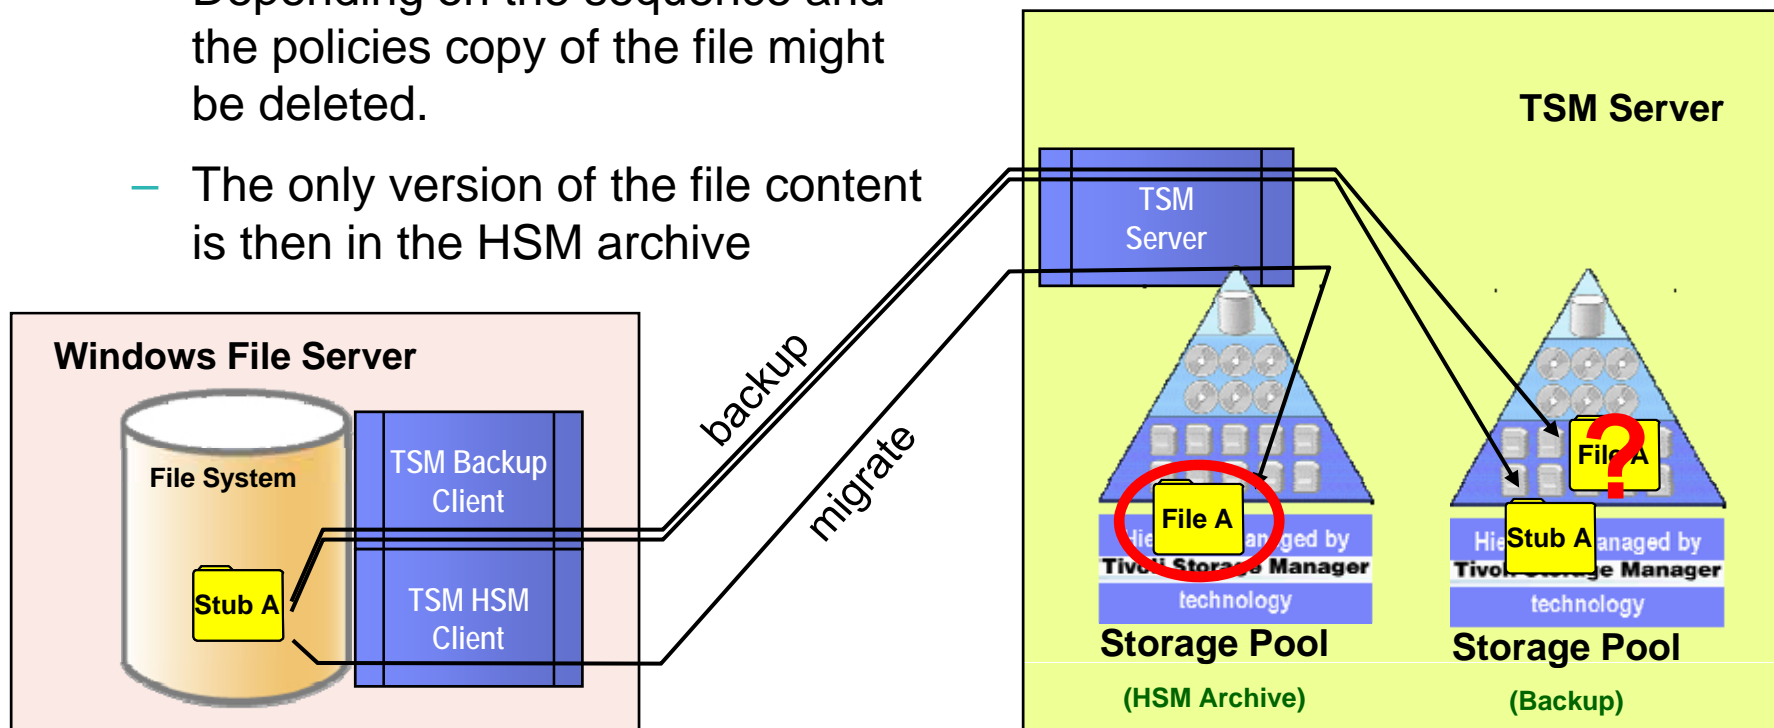
Two major functional Improvements planned for 5.5:

- Better Integration of Backup & HSM
- Reconciliation

TSM HSM for Windows – Closing the Gap

Integration of Backup & HSM (Current Situation)

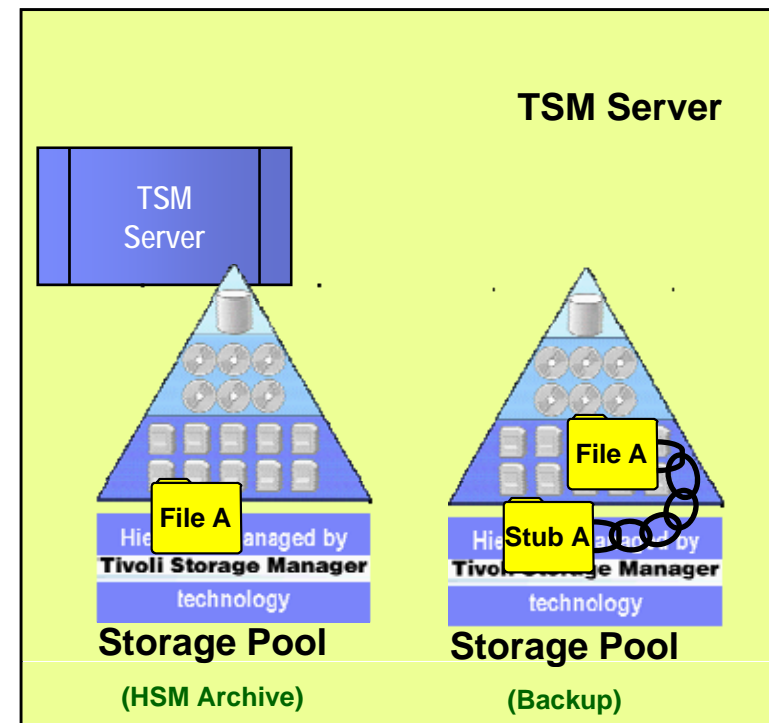

- After subsequent backup and migration, a copy of the file and the stub is in the Backup Pool.
- Depending on the sequence and the policies copy of the file might be deleted.
- The only version of the file content is then in the HSM archive



TSM HSM for Windows – Closing the Gap

Integration of Backup & HSM (Improvements planned for 5.5)

- Still have one copy of the file and the stub in the Backup Pool, ...
- ... now both are tied together.
- The copy of the file will not be expired until copy of the stub expires.
- On restore either file or stub can be recreated.



TSM HSM for Windows – Closing the Gap

Integration of Backup & HSM (5.5 Details)

- Each stub will always have an associated copy of the file in the backup pool
- If an appropriate copy of the file does not exist on backup of a stub, the file will be recalled prior to the backup
- Archive and Selective Backup will always recall stubs
- On restore either the stub or the file can be recreated
- By default the stub is restored, but only if the stub is accessible
- New B/A client options are provided to control the backup/restore
- A plug-in is used for the interaction between B/A and HSM
- No inline copy of migrated files
- TSM server for HSM and Backup can be different

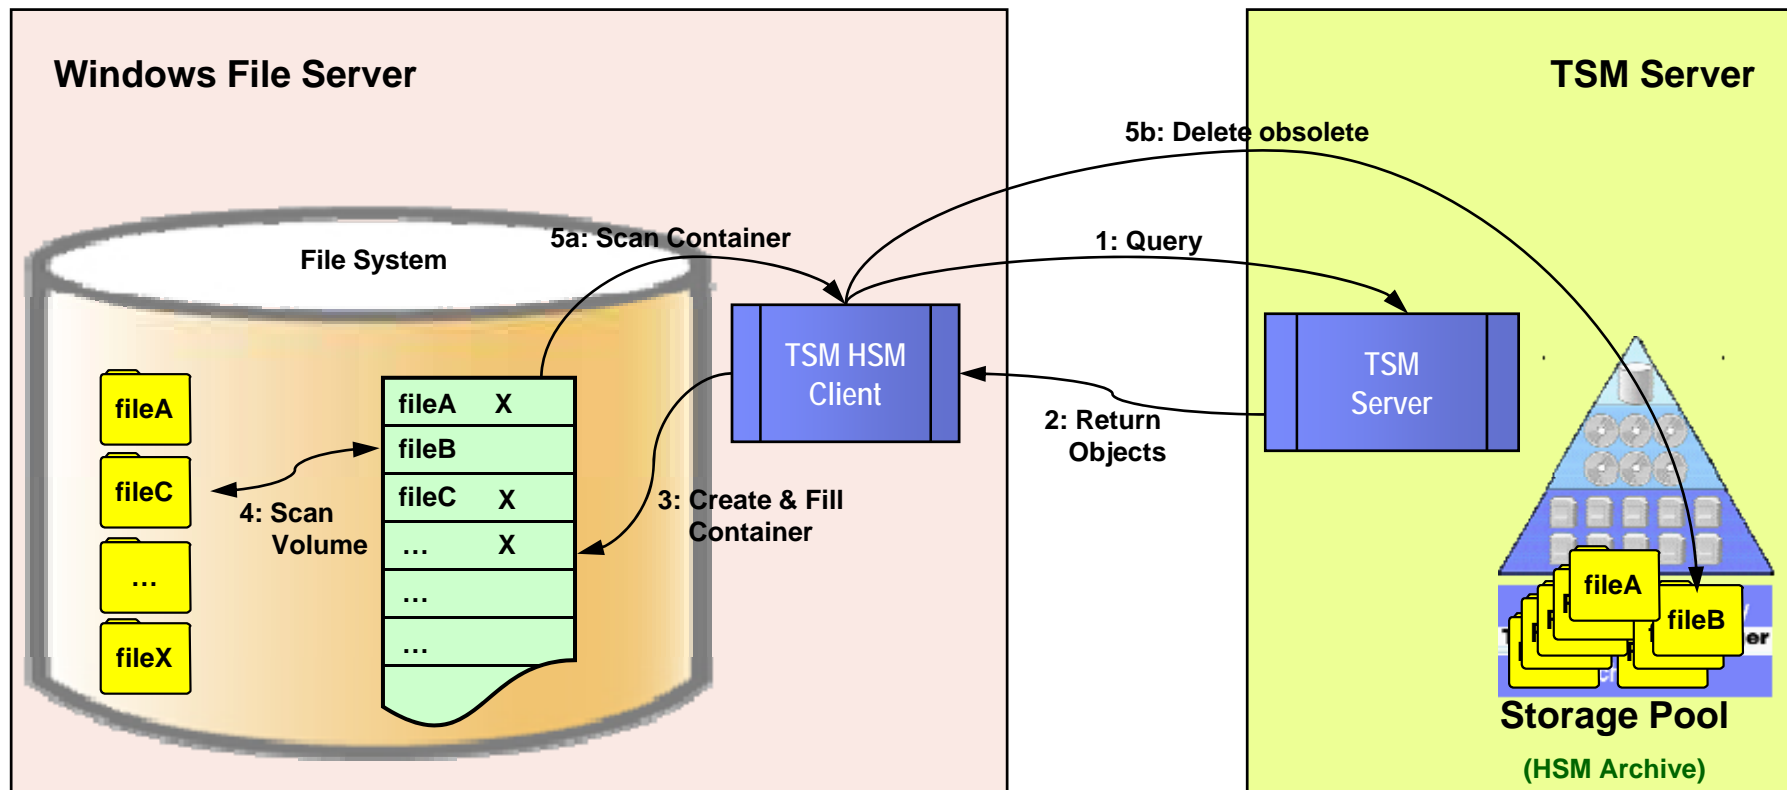
TSM HSM for Windows – Closing the Gap

Reconciliation (Planned for 5.5)

- After migrated files have been modified or deleted, obsolete objects exist on the TSM server in the HSM archive
- Reconciliation removes these objects from the TSM server and also detects orphans in the local file system
- Thus frees up storage on the TSM server and reduces license costs
- Configurable in the HSM GUI or on the Windows command line
- Automatically executed periodically on Volume basis

TSM HSM for Windows – Closing the Gap

Reconciliation (Planned for 5.5)



TSM for Space Management – Lifting the Limit

- One of the major limitation of UNIX HSM to scale for file systems with 10's of millions of files is
 - The candidates selection
 - And reconciliation
- With TSM 5.5 these two tasks are completely redesigned.

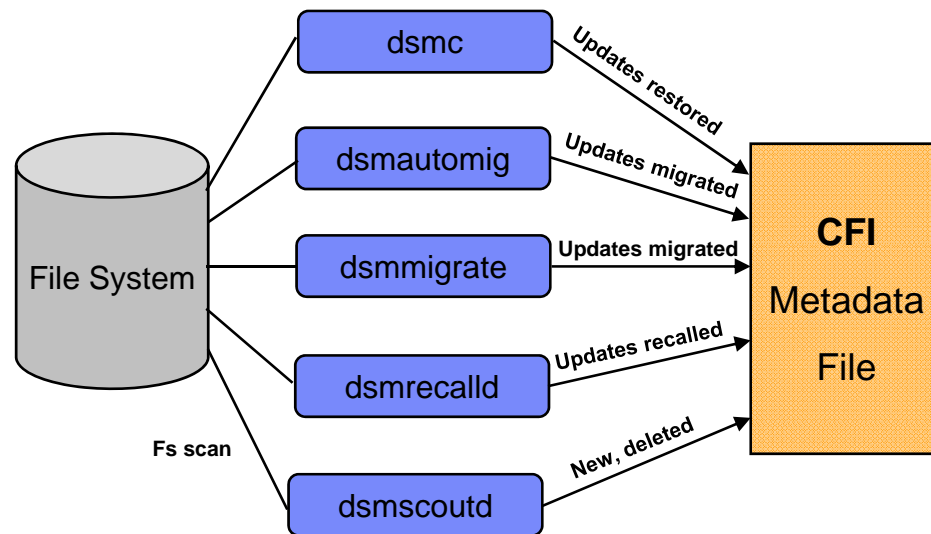
TSM for Space Management – Lifting the Limit

New Scout Daemon (planned 5.5)

- UNIX HSM uses a threshold based automigration that starts migrating files when a high-threshold of a file system is reached.
- The scout daemon is responsible to deliver candidates for the automigration.
 - Situation as of 5.4:
 - HSM uses a candidate list containing a subset of files from the file system.
 - This list is created by a scan which queries only parts of the file system.
 - Files not contained in the list cannot be identified as candidates.
 - In file systems with inhomogeneous distribution of files for criteria size and age, it might lead into situations where no eligible files are found in a reasonable amount of time.

TSM for Space Management – Lifting the Limit

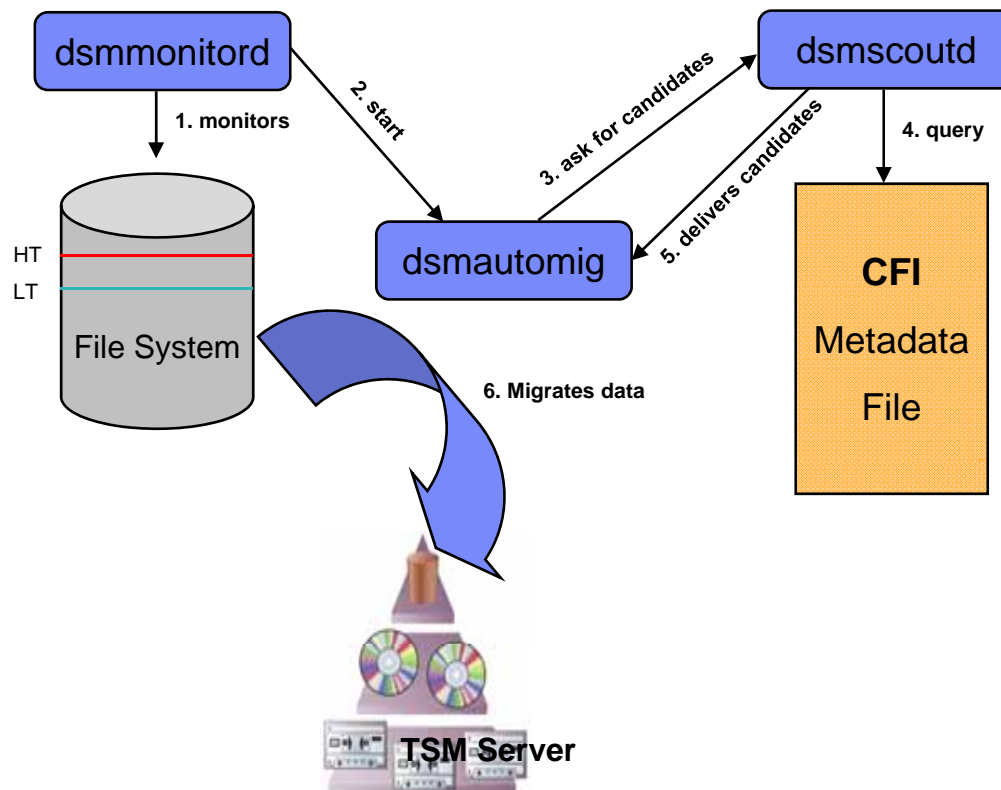
New Scout Daemon (planned 5.5)



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

TSM for Space Management – Lifting the Limit

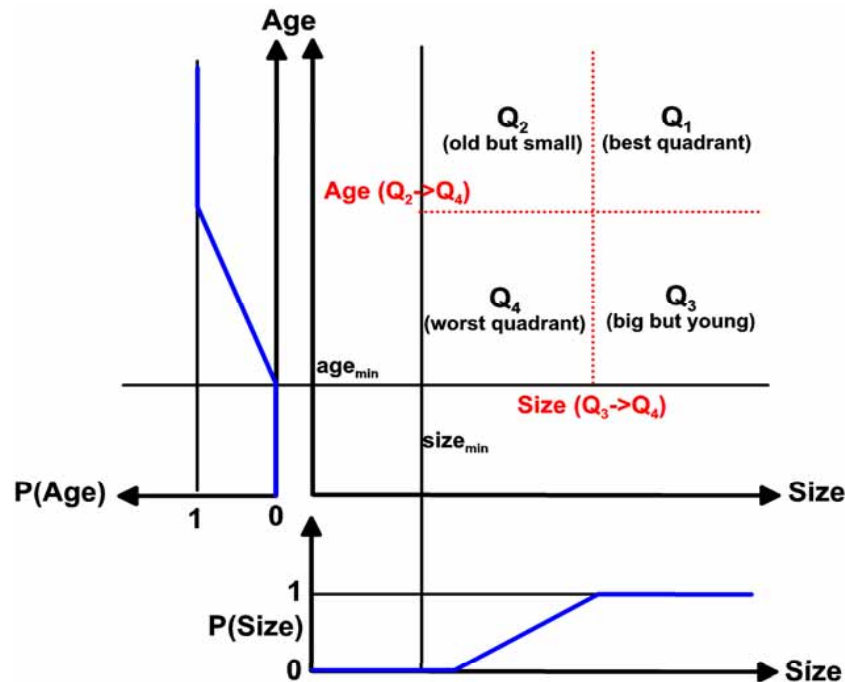
New Scout Daemon (planned 5.5)



- dsmautomig ask Scout daemon for candidates.
- The Scout daemons internal search engine queries CFI for best candidates:
 - Pre-migrated first
 - Biggest and oldest
- dsmautomig uses delivered candidates for migration to the TSM server

TSM for Space Management – Lifting the Limit

New Scout Daemon (planned 5.5) – Candidate Selection



- A criteria for good candidates are files that are big and old.
- Files are classified into four quadrants.
- Boundaries are calculated based on a sample of files.
- A random sample is representative for the distribution of size and age in the file system.
- Files of Q_1 are migrated first, then Q_2 and Q_3
- Quadrants are adapted automatically after each scan

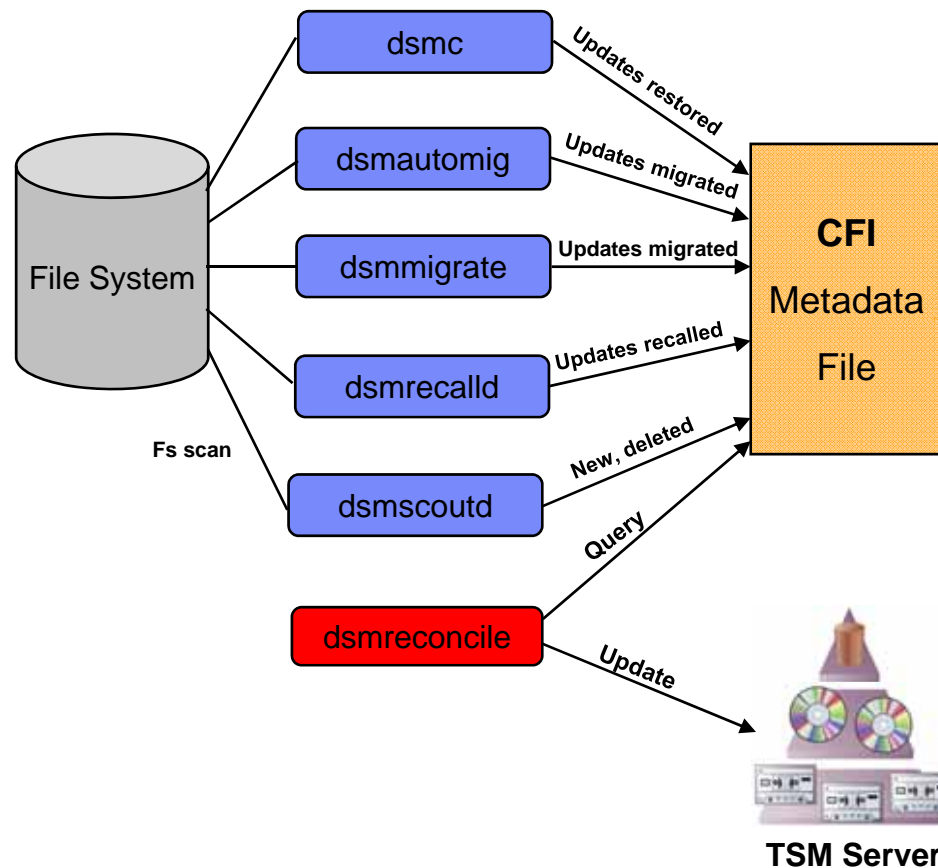
TSM for Space Management – Lifting the Limit

New Reconciliation (planned 5.5)

- The HSM reconciliation process has mainly 3 different use cases:
 - Expire migrated/pre-migrate files on the TSM server that are deleted from the managed file system.
 - Found orphan stub files that have no migrated data on the TSM server.
 - Update the metadata entries on the TSM server data base.
- Reconciliation works on a file list in RAM. It could exceed available memory, if processing more than a few million files.
- Reconciliation could be a very long running task.

TSM for Space Management – Lifting the Limit

New Reconciliation (planned 5.5)



- Reconciliation synchronizes HSM managed file system with the TSM server DB.
- The CFI is like a image of the file system in metadata format.
- The new reconciliation uses the CFI to do fast queries instead of long running file system scans.
 - Increased performance
 - No limitation in number of files
- Stub restore does no longer require full file system scan.

TSM for Space Management – Lifting the Limit

Performance and Scalability


- The new technology of the Scout daemon and reconciliation on UNIX allows to support file systems with 10s to 100s million files.
 - No limitation through out of memory conditions.
 - We plan to certify 50 to 100 million files in 5.5 and to go beyond that in the next releases.
 - Improved automigration performance through
 - Fast migration candidate selection.
 - High ‘quality’ of selected candidates.
 - Improved reconciliation to free up TSM server space and to synchronize metadata.
 - Better performance because no file system scan necessary
 - No limitation in number of migrated files.

TSM for Space Management – Entering new Territories

- GPFS 3.1 introduces policy-based data management for automatic management of the placement and movement of the data between different storage pools.
- GPFS 3.2 integrates TSM HSM into that management.
- This allows candidates selection based on every file attribute.

TSM for Space Management – Entering new Territories

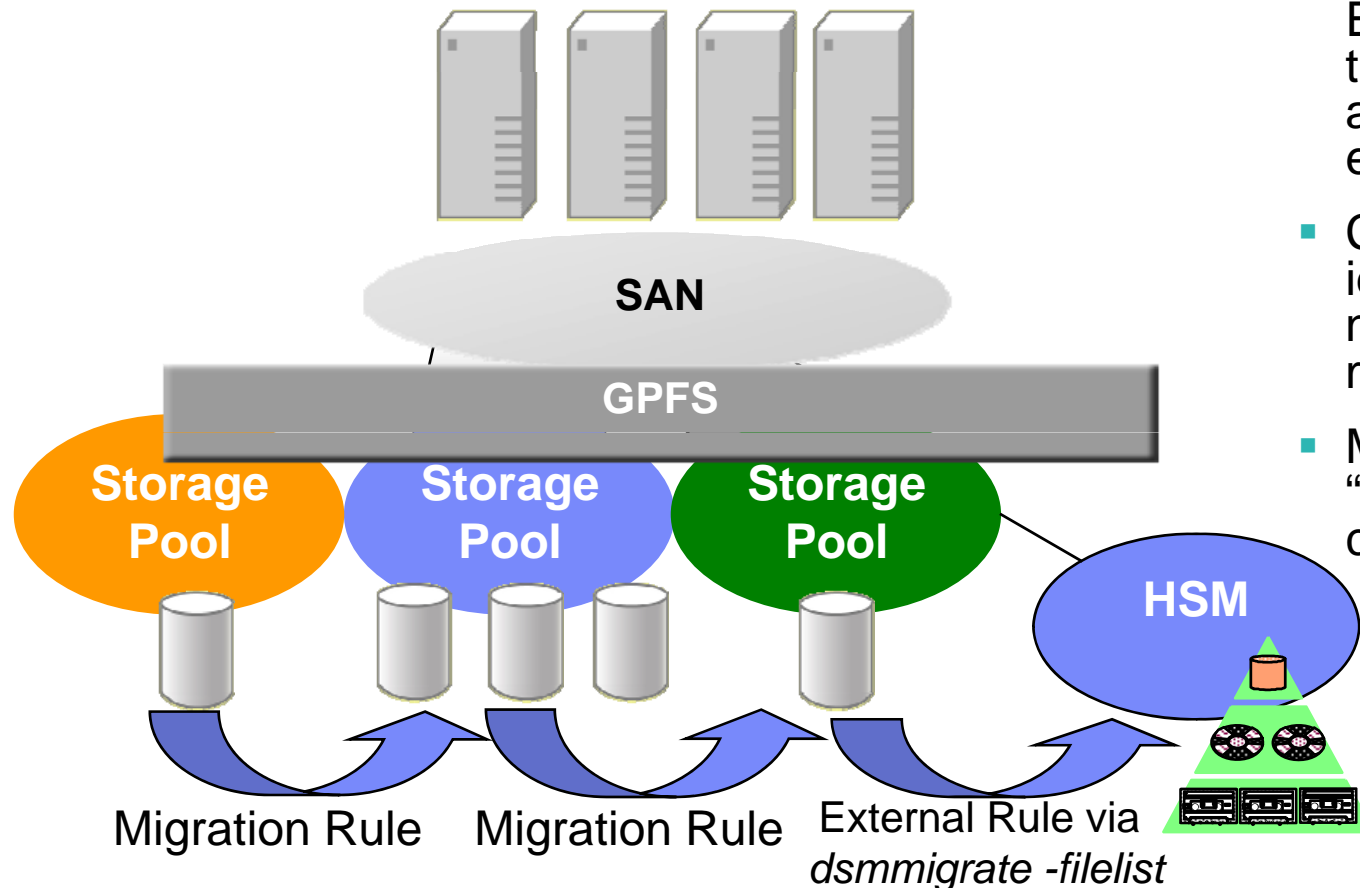
GPFS Storage Pools



- Without storage pools:
 - The entire file system is striped across all disks in one pool
- With storage pools:
 - One name space across all pools, but files in the same directory can be in different pools.
 - Files placed in storage pools at create time using policies.
 - Files can be moved between pools for policy reasons.

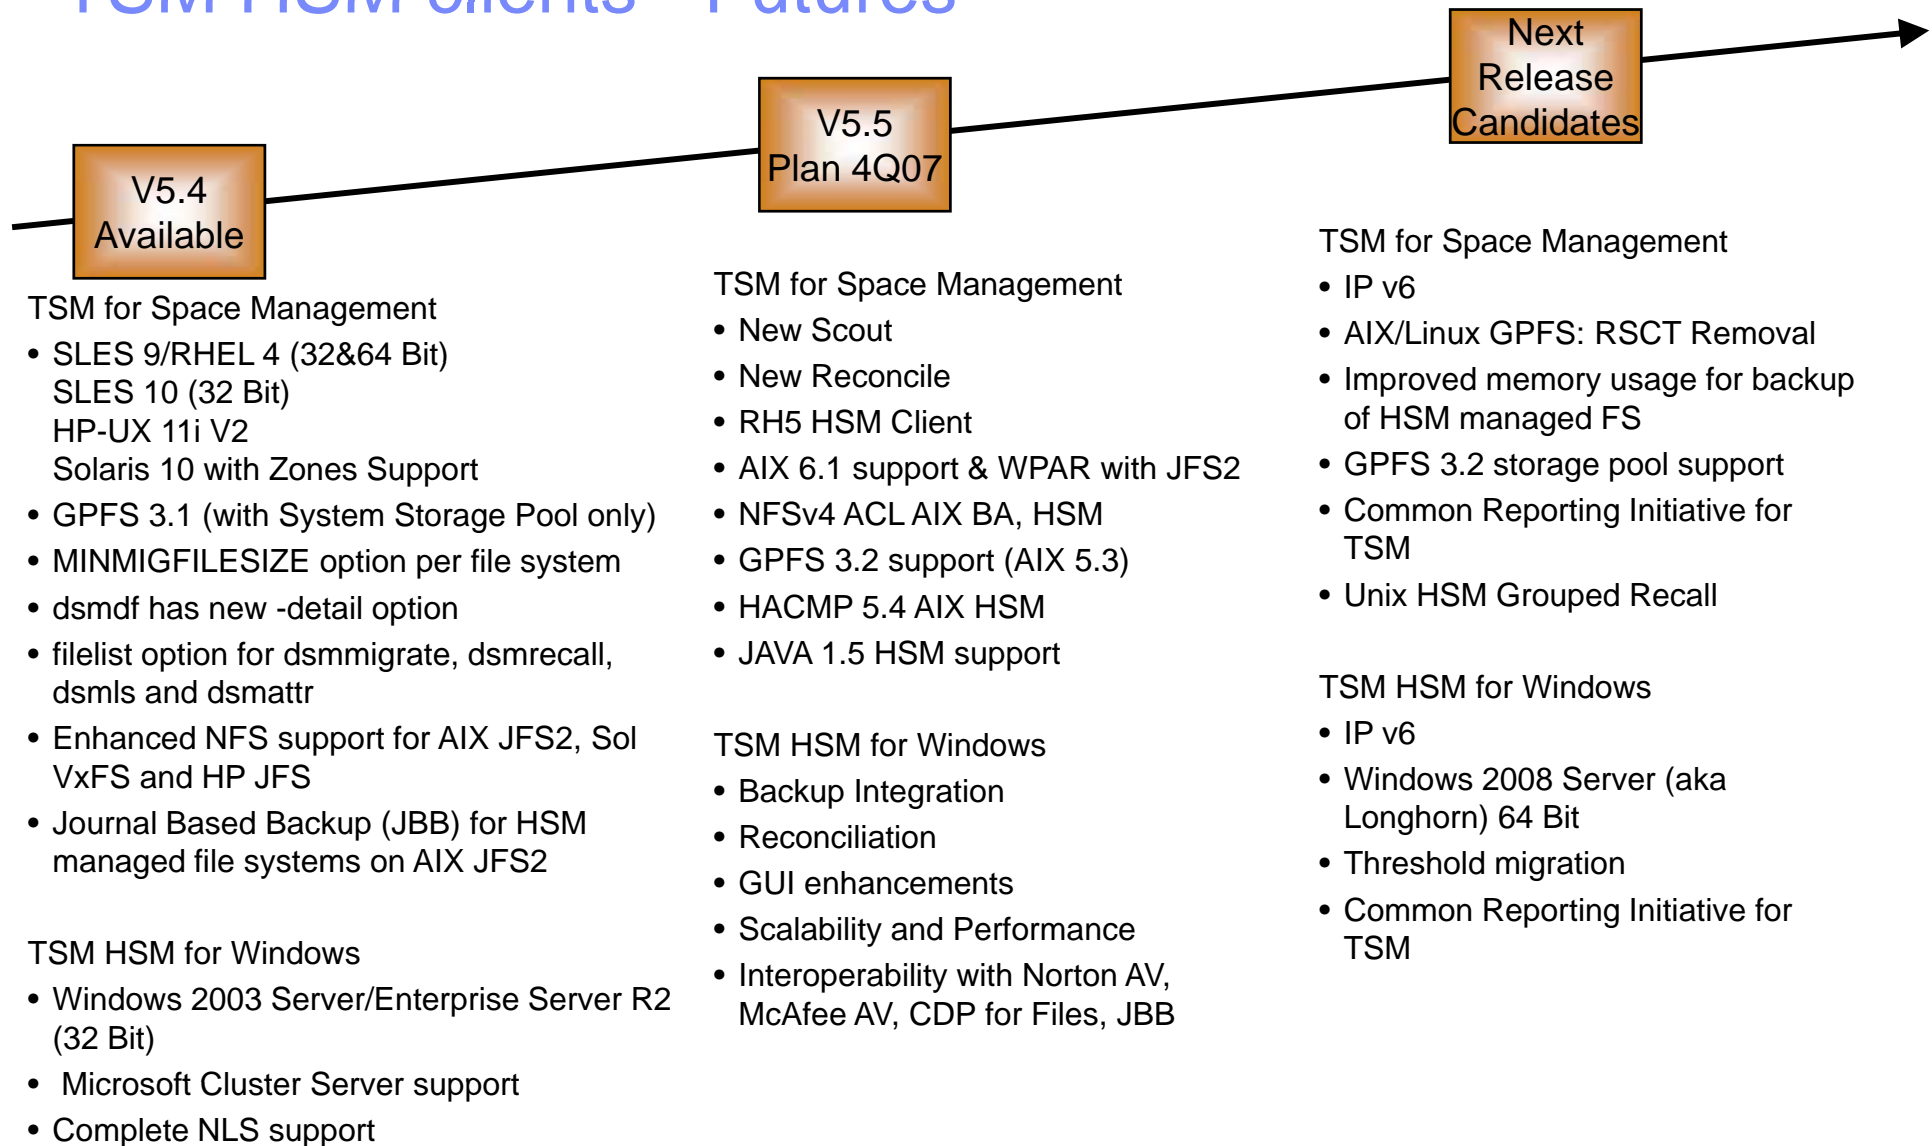
TSM for Space Management – Entering new Territories

GPFS 3.2 Storage Pools and external HSM pool / ILM to Tape




- GPFS introduce a new EXTERNAL POOL rule that extends the policy to allow migration to external pools like HSM.
- Candidate selection is identical to the GPFS native pool to pool migration rule.
- Movement based on the “*dsmmigrate -filelist*” command.

TSM HSM clients - Futures



Questions and Answers



Vielen
Dank

Obrigado!

THANK YOU

תודה



Gracias

Bedankt

Köszönettel

Teşekkürler

شكراً

Grazie