



IBM Spectrum Protect & IBM Spectrum Scale

# mmbackup and TSM Integration

GPFS User Group UK / SPXXL Meeting  
York 05/20/2015

## Presenter Information



Stefan Bender

IBM Spectrum Protect Development for HSM, B/A and VE clients  
IBM Germany Research & Development

[stefan.bender@de.ibm.com](mailto:stefan.bender@de.ibm.com)

# Agenda

## IBM Spectrum Protect and Spectrum Scale in a Nutshell

(Terminology And High Level Overview)

## Large File System Backup

(State Of The Art And Recent Improvements)

## Hierarchical Storage Management

(Components, Responsibilities And Workflow)

## Fast Disaster Recovery

(SOBAR – Scale Out Backup And Restore)

# Agenda

## IBM Spectrum Protect and Spectrum Scale in a Nutshell

(Terminology And High Level Overview)

## Large File System Backup

(State Of The Art And Recent Improvements)

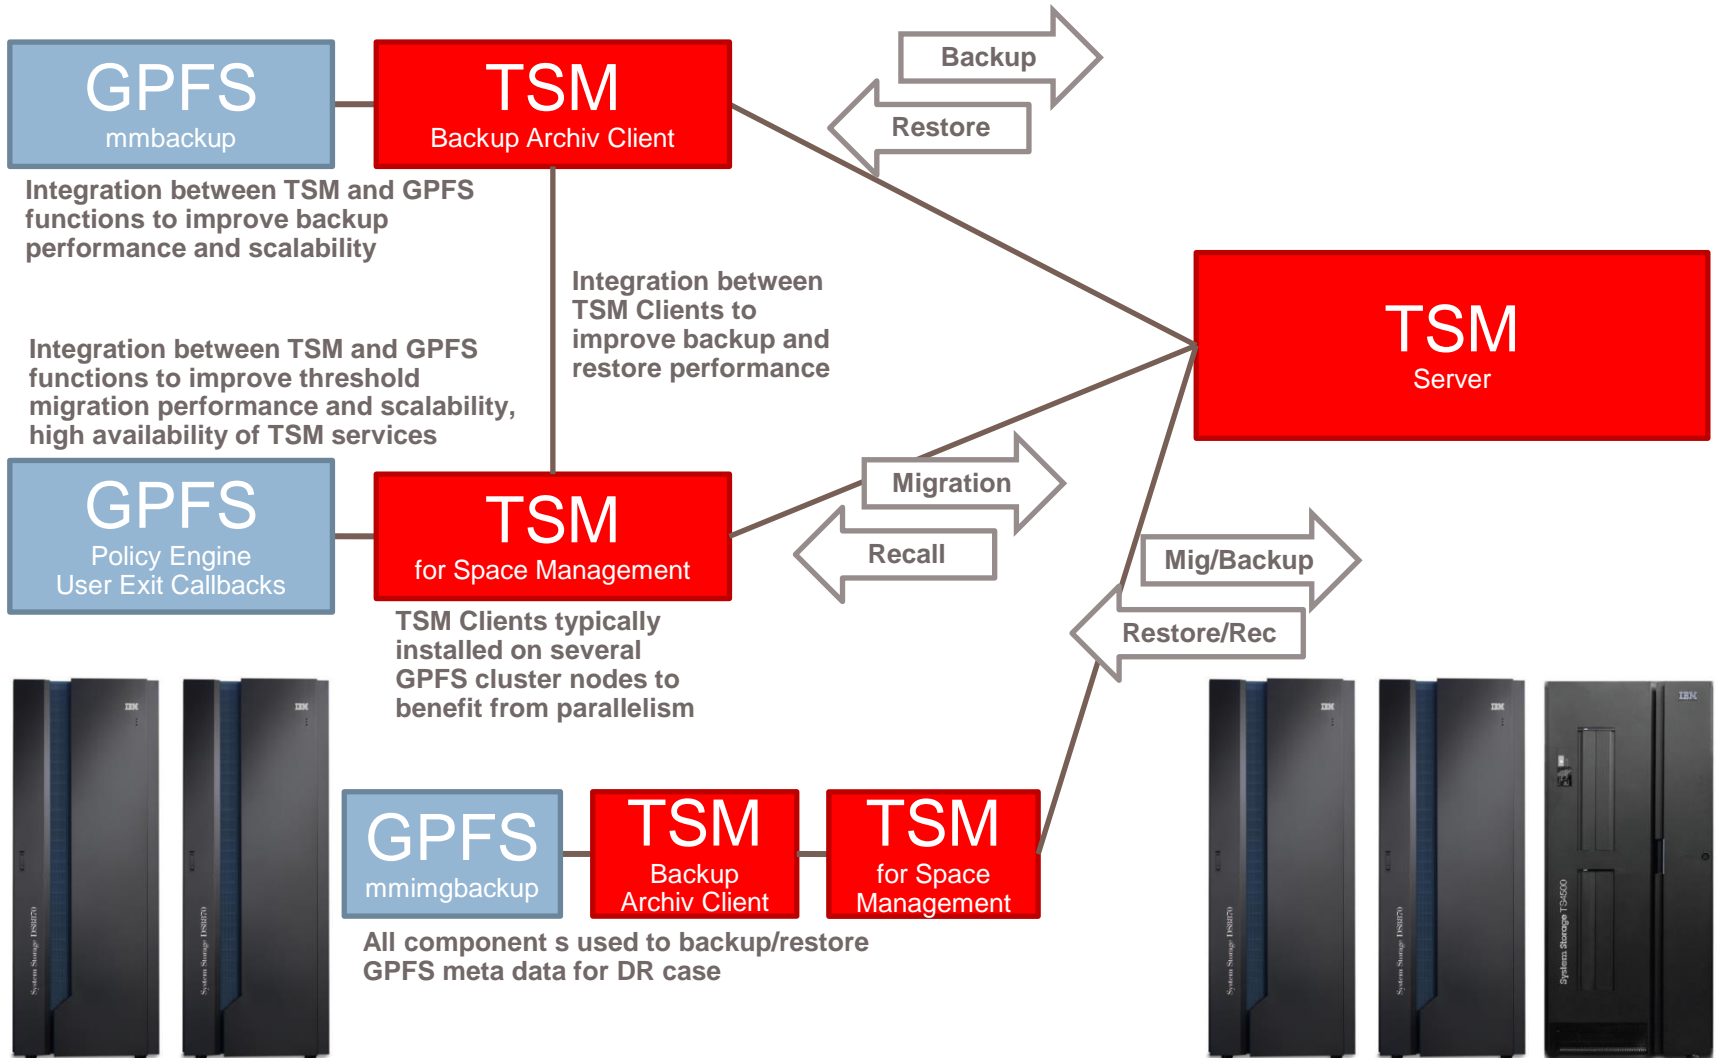
## Hierarchical Storage Management

(Components, Responsibilities And Workflow)

## Fast Disaster Recovery

(SOBAR – Scale Out Backup And Restore)

# IBM Spectrum Protect (formerly Tivoli Storage Manager or TSM) and IBM Spectrum Scale (formerly GPFS) in a Nutshell



# Agenda

**IBM Spectrum Protect and Spectrum  
Scale in a Nutshell**  
(Terminology And High Level Overview)

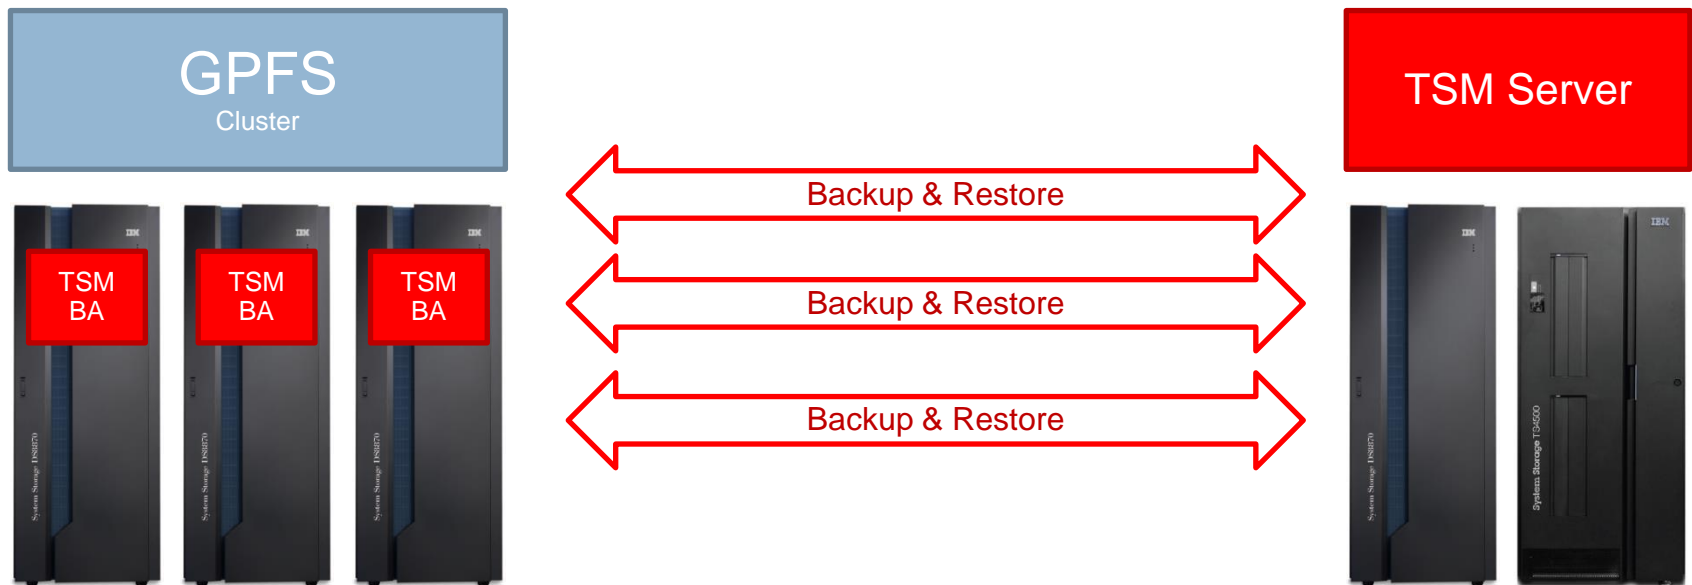
**Large File System Backup**  
(State Of The Art And Recent Improvements)

**Hierarchical Storage Management**  
(Components, Responsibilities And Workflow)

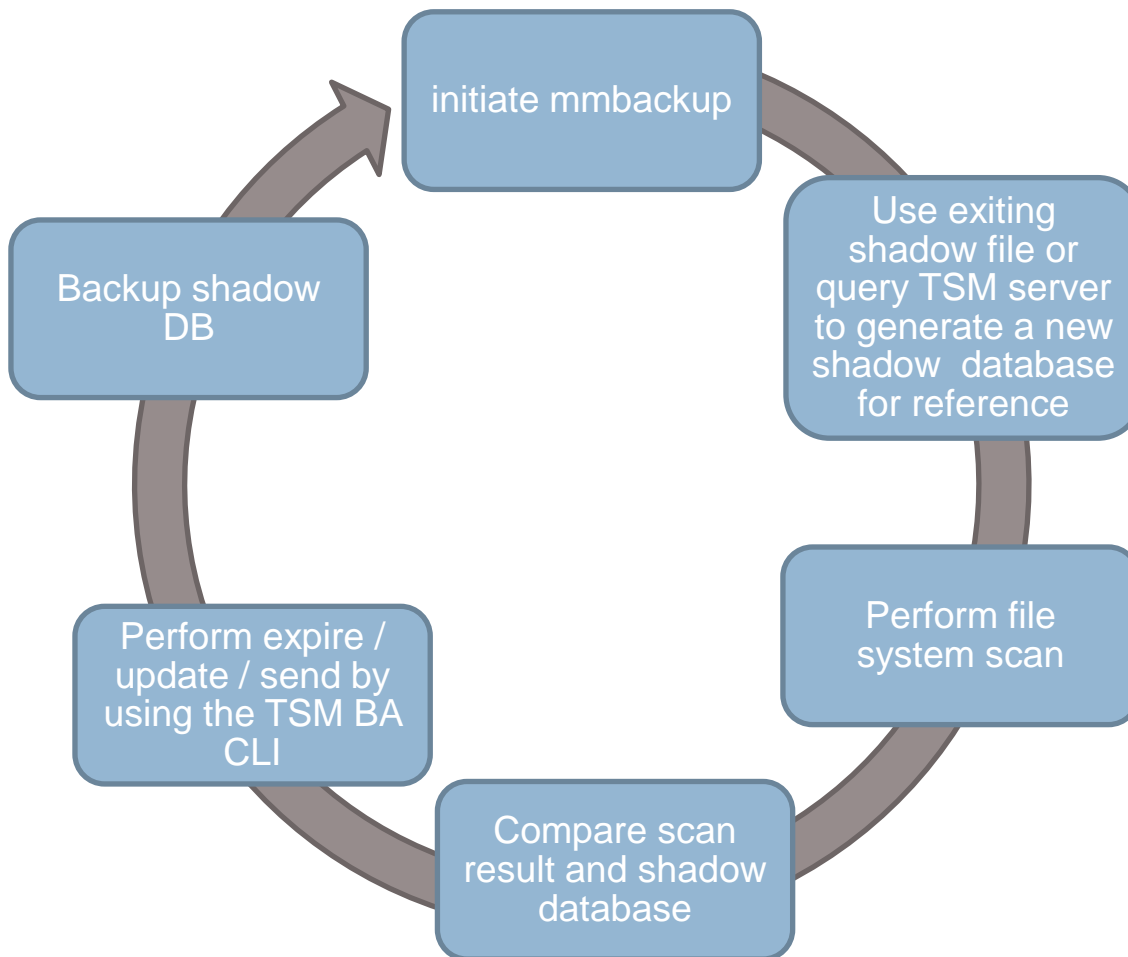
**Fast Disaster Recovery**  
(SOBAR – Scale Out Backup And Restore)

# Large File System Backup – IBM Spectrum Protect & Elastic Storage

- Much faster file system scanning times allows TSM backups to scale to many more objects compared to TSM progressive incremental
- mmbackup can utilize multiple GPFS nodes to scan the file system and take backups
- TSM GUI or CLI can be used to traverse the protected data for individual file restore



# Large File System Backup



Backup cycle:

- The initial backup run requires a query against the TSM DB2 repository
- In a regular run a new shadow file is created and compared against the previous one
- The shadow file is backed up to TSM at the end of the mmbackup run



## mmbackup Help (Version 4.1.0.0)

```
root > mmbackup --help
```

Usage:

```
mmbackup {Device|Directory} [-t {full|incremental}]  
[-N {Node[,Node...] | NodeFile | NodeClass}]  
[-g GlobalWorkDirectory] [-s LocalWorkDirectory]  
[-S SnapshotName] [-f] [-q] [-v] [-d]  
[-a IscanThreads] [-n DirThreadLevel]  
[-m ExecThreads | [--expire-threads ExpireThreads]  
  [--backup-threads BackupThreads]]  
[-B MaxFiles | [--max-backup-count MaxBackupCount]  
  [--max-expire-count MaxExpireCount]]  
[--max-backup-size MaxBackupSize] [--quote | --noquote]  
[--rebuild] [--tsm-servers TSMServer[,TSMServer...]]  
[--tsm-errorlog TSMErrorLogFile] [-L n] [-P PolicyFile]
```

# Large File System Backup – IBM Spectrum Scale 3.5 TL 3 Updates

## ▪ New shadow data base design

- Reduce the number of sort() iterations to increase the processing performance.
- Allow parallel updates (multiple processes) on the shadow DB.
- Restart able backups - shadow DB shows current progress and work remaining.
- Elimination of the post processing compensation phase
- Improved failure detection for TSM failures.

## ▪ Exploit incremental backup function

- Detect CTIME changes without data or MTIME changes and run incremental:
  - CTIME, owner, group, mode change - run dsmc incremental
  - MTIME or file size change - run dsmc selective
- Detect HSM migration changes:
  - migration state change only - run dsmc incremental

## Large File System Backup – IBM Spectrum Scale 4.1 Updates

### ▪ Improved Environment Verification

- Verify whether the file system is mounted.
- Verify whether the TSM BA Client is installed.
- Verify whether TSM BA Client version is equal on all nodes
- Verify whether a session with TSM Server can be established.
- Verify whether the helper tool version is correct.
- Verify whether TSM options are set:
  - QUOTESARELITERAL (if mmbackup is used with --noquotes)
  - SKIPACL
  - SKIPACLUPDATECHECK
  - PRESERVELASTACCESSDATE (allowed in environments that have HSM management)

# Large File System Backup – IBM Spectrum Scale 4.1 Updates End To End Configuration

Define number of bytes to be used in a single TSM transaction


--max-backup-size TXNBYTELIMIT

Define number of files to be used in a single TSM transaction

--max-backup-count TXNGROUPMAX  
--max-expire-count

Define number of parallel backup processes

--expire-threads RESOURCEUTILIZATION MAXSESSIONS  
--backup-threads MAXNUMMP



## Large File System Backup – IBM Spectrum Scale 4.1 Updates

### ▪ Tuning Changes II

- Permit passing policy option "-a IscanThreads" into mmapplypolicy scan phase. Range: Enforced by mmapplypolicy
  - Permit passing policy option "-n DirThreadLevel" into mmapplypolicy during scan phase. Range: Enforced by mmapplypolicy
- TSM BA Client file list based expiration processing improved with TSM 6.4.1 (multiple expirations per transaction)

## Large File System Backup – Include and Exclude Handling

- TSM offers a rich set of include and exclude options to control which files and directories are backed up.
- mmbackup is building these options into its policy for backup
- Include and Exclude options may have significant impact on scan performance
- **Some rules to consider:**
  - Use as few EXCLUDE statements as possible
  - Avoid using INCLUDE. Use EXCLUDE instead
  - Do not use „EXCLUDE /dir/.../\*“ . Try EXCLUDE.DIR instead.
  - Do not combine EXCLUDE and INCLUDE for one subtree, like

```
exclude /home/bender*
include /home/bender/important*
```
  - If INCLUDE is only used to assign right management class in TSM

```
„INCLUDE <pattern> MGMT“
```

use mmbackup service flag is used MMBACKUP\_IGNORE\_INCLUDE

```
export MMBACKUP_IGNORE_INCLUDE=1
```
- **Watch for a new Tech Note to be published**

## Large File System Backup – Further Recommendations

- **Use latest versions**
  - GPFS 4.1 and TSM 7.1 have good improvements for mmbackup
- **Consider work load on GPFS cluster nodes**
  - mmbackup puts load on each node involved
  - Nodes need direct access to storage
  - Nodes need access to TSM server
  - Serialize backups of different file systems
  - Serialize backups with other activities on the TSM server (other node backups, Server maintenance,...)
- **Use same type of nodes**
  - Do not mix nodes. Run mmbackup on either AIX, xLinux, pLinux nodes
- **Consider TSM character limitations**
  - Files with control-X, control-Y, carriage return and the new line character in their name can't be backed up to TSM.
  - Use QUOTESARELITERAL (if mmbackup is used with --noquotes), if file names contain " or ' .
  - Use WILDCARDSARELITERAL, if file names contain \* or ?.
- **Check your log files regularly**

# Backup Methods Compared

<b><i>Backup Related File &amp; Environment Changes</i></b>	<b><i>TSM 7.1.1</i></b> <i>(progressive incremental)</i>	<b><i>GPFS 4.1</i></b> <i>(mmbackup)</i>
Detects changes in files and sends a new copy of the file to the server	Yes	Yes
Detects changes in metadata and updates the file metadata at the server or sends a new copy of the file to the server in terms of ACL/EA changes	Yes	Yes
Detects directory move, copy or rename and send a new copy of the file to the server	Yes	Yes
Detects local file deletion and expires the file at the server	Yes	Yes
Detects TSM file space deletion or node/policy change and sends a new copy of the file to the server	Yes	No*
Detects file deletion from TSM server and sends a new copy of the file to the server	Yes	No*
Detects additions of new exclude rules and expires the file at the server	Yes	Yes
Detects policy change due to new include rule and rebinds the file to the new storage pool	Yes	No**
Recognizes copy mode and copy frequency options	Yes	No*
Detects migration state changes (TSM for Space Management) and updates server object	Yes	Yes
Detects that a file wasn't processed successfully during backup operation and attempts again at the next backup	Yes	Yes

\* mmbackup queries the TSM Server only once at the time of the first backup. Changes which are performed at the TSM Server directly by using the TSM administrative client can't be detected by mmbackup. It is recommended to rebuild the mmbackup shadow data base in the case of TSM Server file space changes.

\*\* TSM include rules with associated management class bindings can't be detected from mmbackup. Therefore mmbackup doesn't rebind a file in the case of management class changes in include rules.



# Agenda

## IBM Spectrum Protect and Spectrum Scale in a Nutshell

(Terminology And High Level Overview)

## Large File System Backup

(State Of The Art And Recent Improvements)

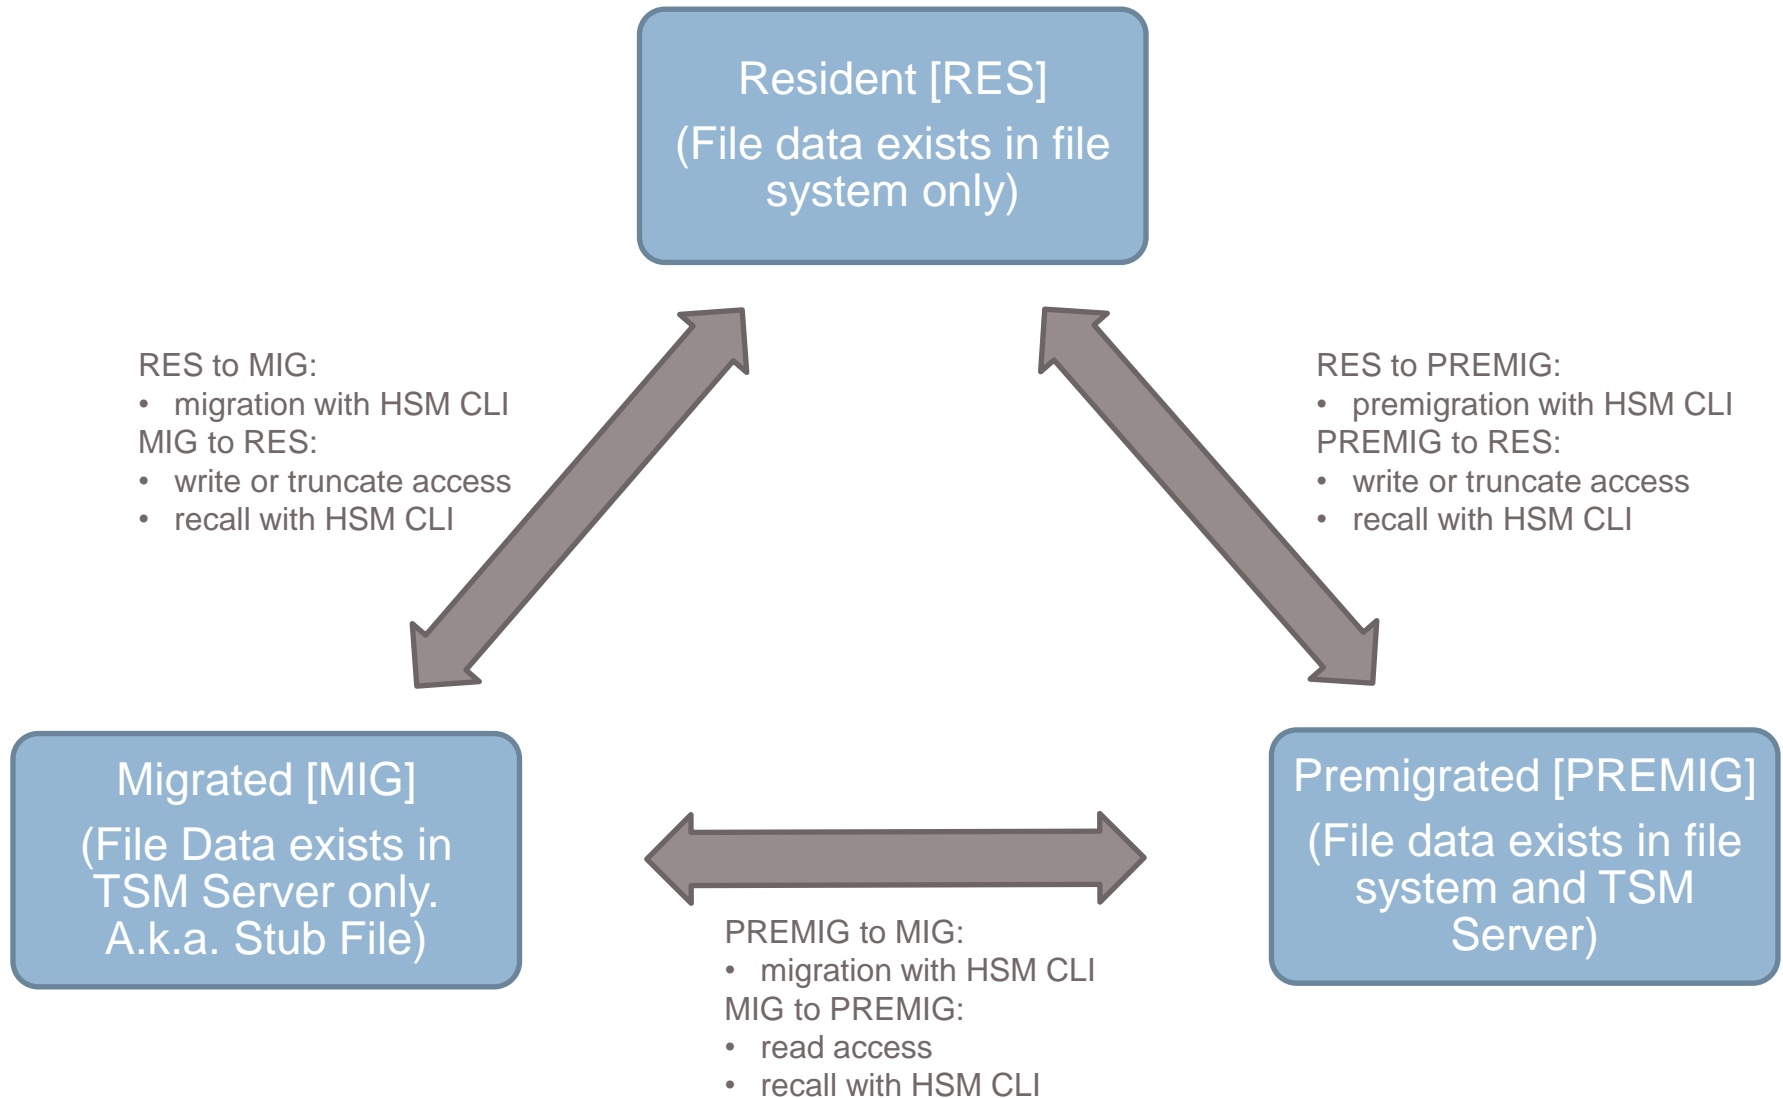
## Hierarchical Storage Management

(Components, Responsibilities And Workflow)

## Fast Disaster Recovery

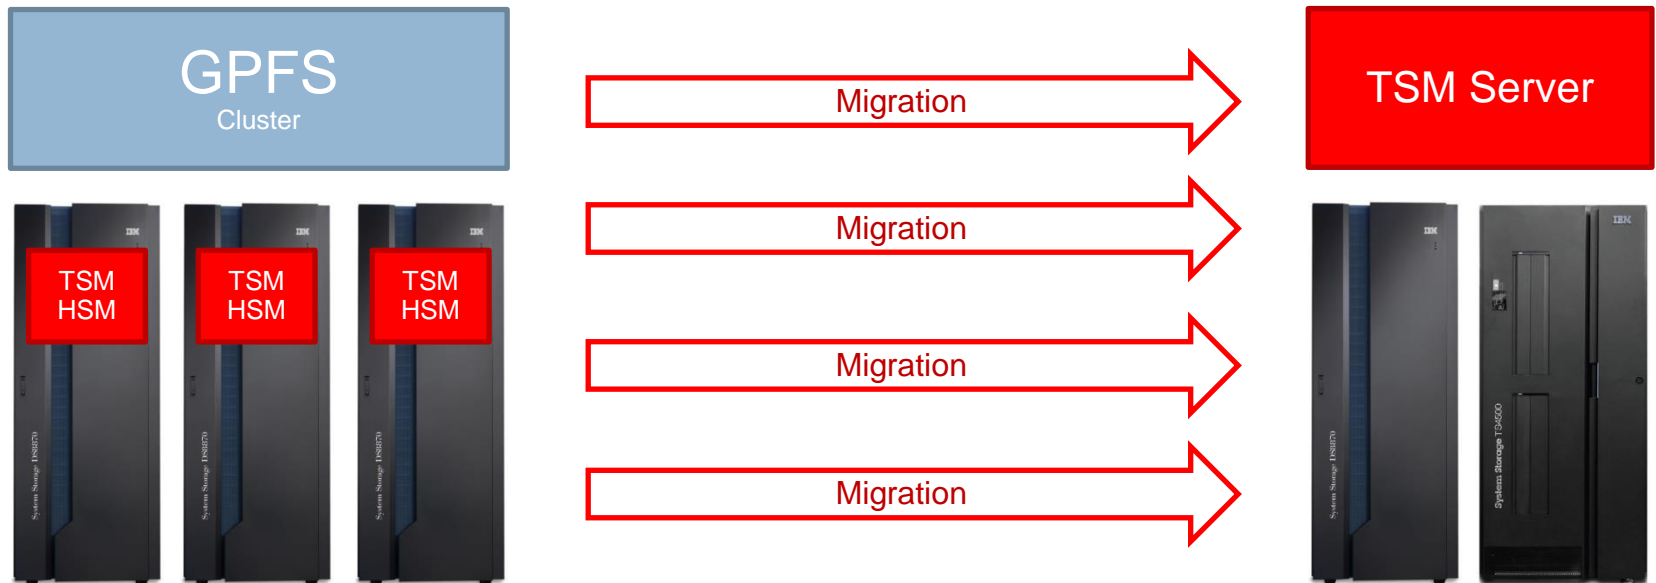
(SOBAR – Scale Out Backup And Restore)

# IBM Spectrum Protect for Space Management – File State Overview



# Functional Overview - Migration

- Supports pre/migration to address different use cases
- Parallelism node wide and cluster wide
- Close integration with TSM Backup Archive Client
- The migration of files in tape pools of the TSM server is optimized for performance



# Policy Based Threshold Migration - Components And Responsibilities

## Administrator

(Action required: Once)

- Defines GPFS policy rules
- Defines migration callback script
- Defines HSM exec script
- Enables GPFS threshold callback

## GPFS

(Action required: Continuously)

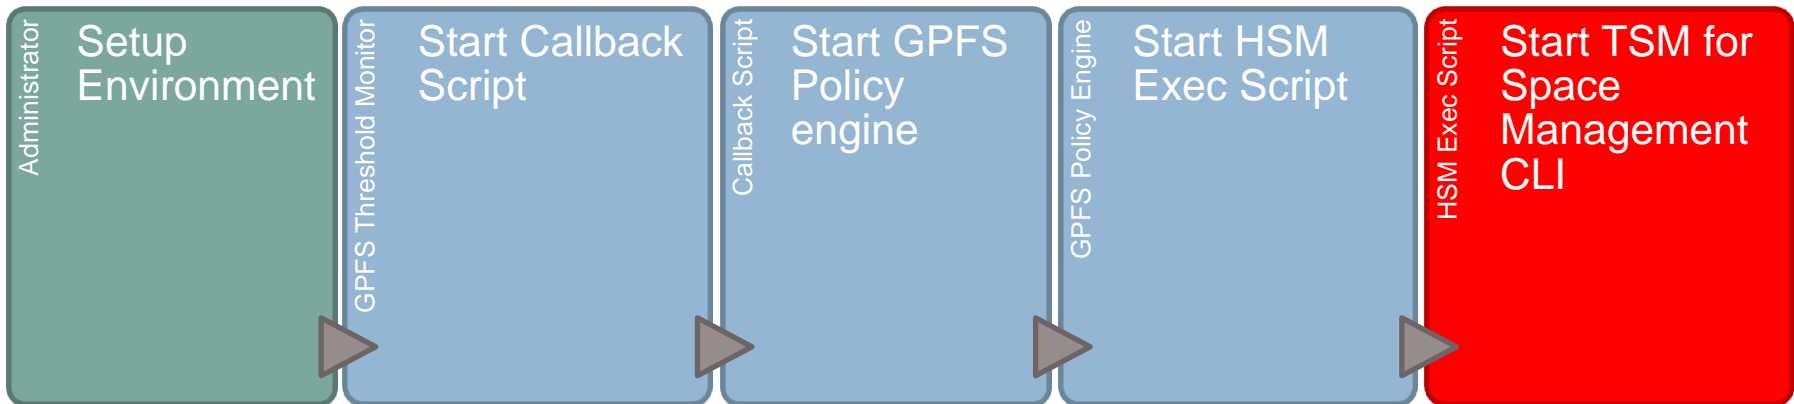
- Monitors file system thresholds
- Starts threshold callback function
- Scans file system directory and inode structure
- Generates candidates lists for migration
- Starts HSM migration commands

## TSM for Space Management

(Action required: Continuously)

- Performs premigration of files
- Performs migration of files

# Policy Based Threshold Migration - Workflow



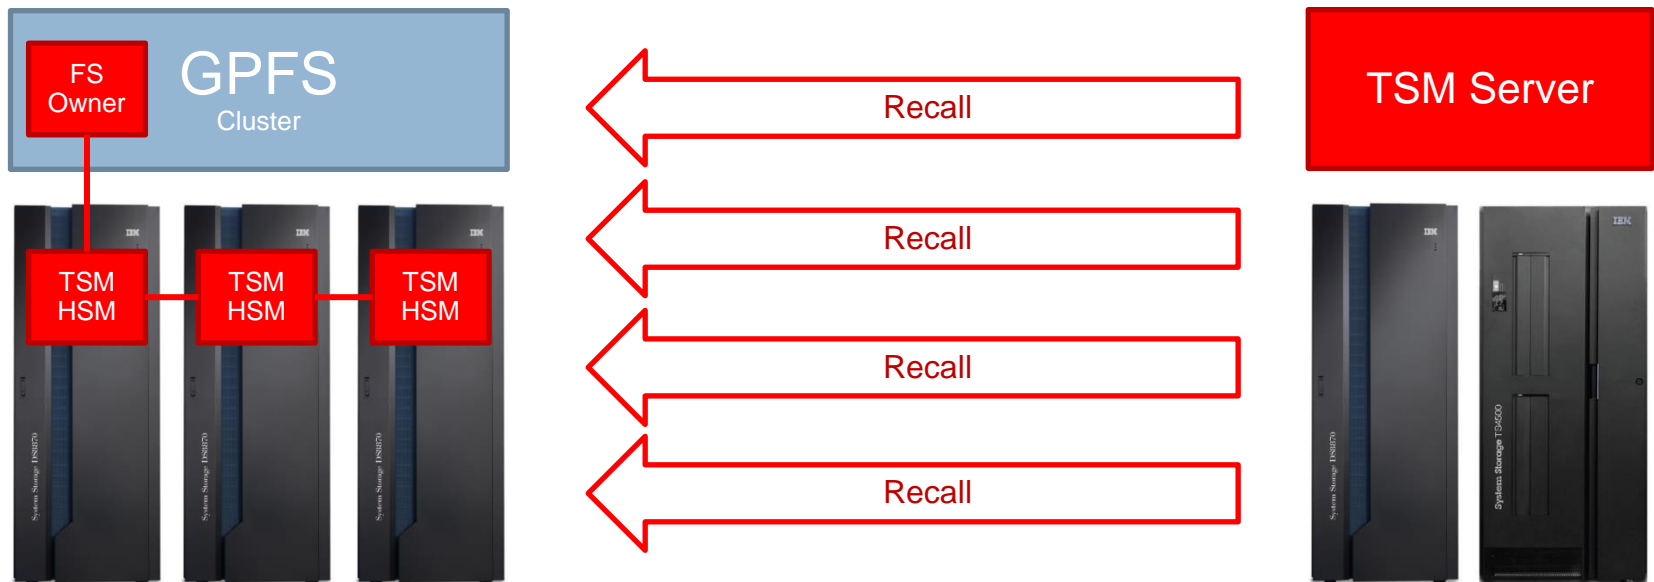
- Policy rules sample: `/usr/lpp/mmfs/samples/ilm/mmpolicyRules-lowspace.sample`
- Callback script sample: `/usr/lpp/mmfs/bin/mmstartpolicy`
  - Number of used migrate processes can be defined here (Option -m)
  - Number of objects per migration file list can be defined here (Option -B)
  - Callback runs on the GPFS file system manager node. Ensure HSM is installed on all manager nodes. See command: `mmlsmgr`
- HSM exec script sample: `/usr/lpp/mmfs/samples/ilm/mmpolicyExec-hsm.sample`
- Tivoli Field Guide “TSM for Space Management – GPFS Integration” describes base configuration
- GPFS Advanced Administration Guide describes the policy language

## Functional Overview - Recall

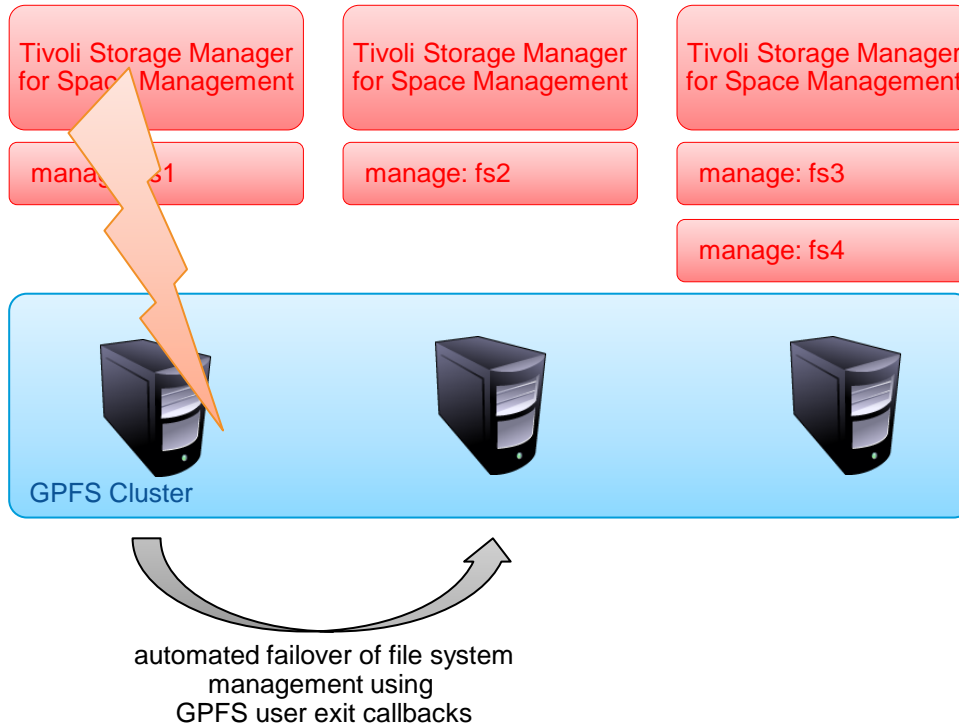
- Supports three different recall modes:
  - Normal (Application can read the file after recall has finished)
  - Streaming (Application can read the file after a administrator defined portion of the file was recalled)
  - Partial (Application can read the sub portion of the file that was requested after it was recalled)
- Recall options can be defined in detail with a common interface
- Parallelism node wide and cluster wide
- Close integration with the TSM Backup Archive Client
- A recall method is available that optimizes the recalls of files stored on tape pools in the TSM server

# Transparent Distributed Recall

- Recalls can be distributed to different cluster nodes in GPFS clusters automatically to improve scalability and performance
- File system owner and distribution behavior can be defined
- TSM for Space Management HA function ensures service availability



# Hierarchical Storage Management



- TSM for Space Management HA function re-implemented with TSM 7.1.1
- new implementation bases on GPFS user exit callbacks
- significant reduced service downtime
- reduced impact to GPFS cluster functions
- failover recognizes system load



# Space Management News

## TSM 7.1.0

- Consistent return codes for Backup Archive client and HSM
- Improvement of status query command allows verification of HSM service

```
root > dsmmigfs query -node=black -detail
IBM Tivoli Storage Manager
Command Line Space Management Client Interface
  Client Version 7, Release 1, Level 0
  Client date/time: 05/02/14 11:28:08
(c) Copyright by IBM Corporation and other(s) 1990, 2014. All Rights Reserved.

GPFS Node Name:          black
GPFS Node ID:           1
GPFS Status:            active
HSM Status:             active
Recall Daemon Session ID: 5357ADB700000000
Mount Disposition:      YES
Ping Recall Daemon:     YES
Watch Daemon Session ID: 5357ADB700000000

root > dsmmigfs query -node=black -detail -pars
dsmmigfs:queryNode:black:1:active:active:5357ADB700000000:YES:YES:5357ADB700000000:
```

## TSM 7.1.1

- Redesign of HA function to improve HSM service availability
- New option PREVIEWSIZE improving streaming recall service

# Agenda

**IBM Spectrum Protect and Spectrum  
Scale in a Nutshell**  
(Terminology And High Level Overview)

**Large File System Backup**  
(State Of The Art And Recent Improvements)

**Hierarchical Storage Management**  
(Components, Responsibilities And Workflow)

**Fast Disaster Recovery**  
(SOBAR – Scale Out Backup And Restore)

# Scale Out Backup And Restore

## KEY FEATURES


- High backup scalability. Only file system metadata (inode & path) has to be backed up.
- High restore performance. File data resides on the TSM Server and recall happens on demand.
- Recreates the whole directory tree with all permissions and all files in stub format in one step.

## Customer Experience (FZJ):

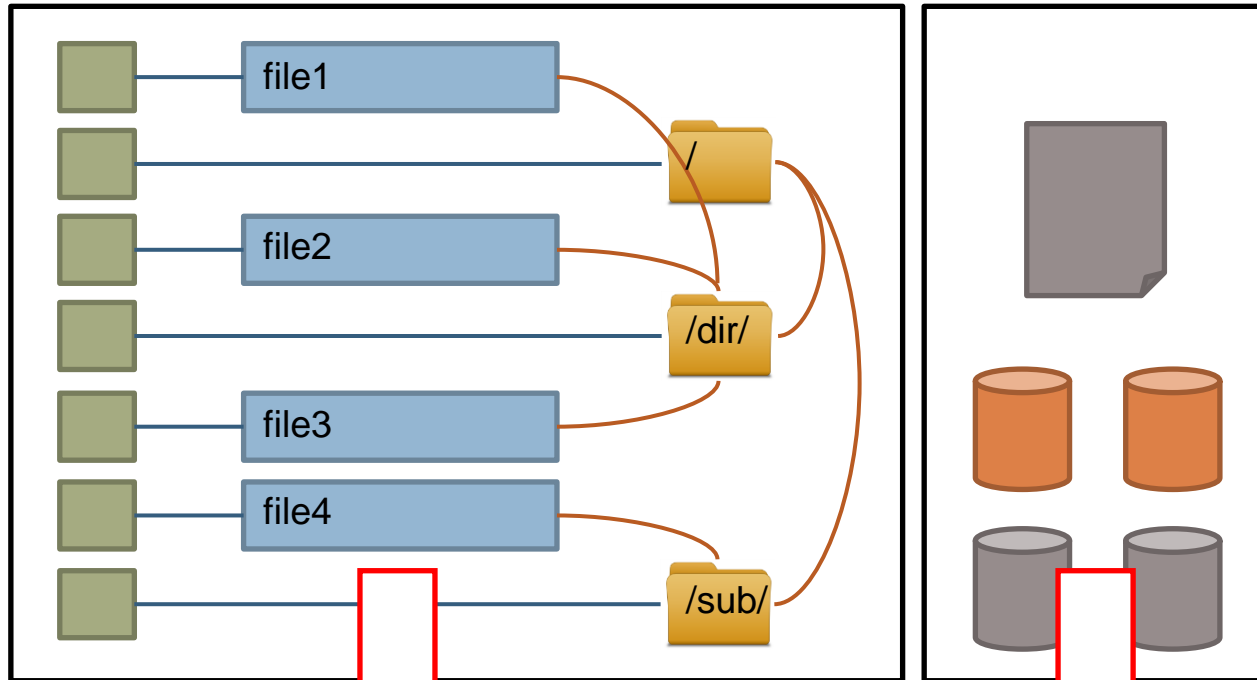
SOBAR backup (3 node GPFS, 60 million files) : 4 hours

SOBAR restore (create file system, restore and inflate image, start production) : 4 hours

# Scale Out Backup And Restore – Protected Items



# Scale Out Backup And Restore – Backup I of II

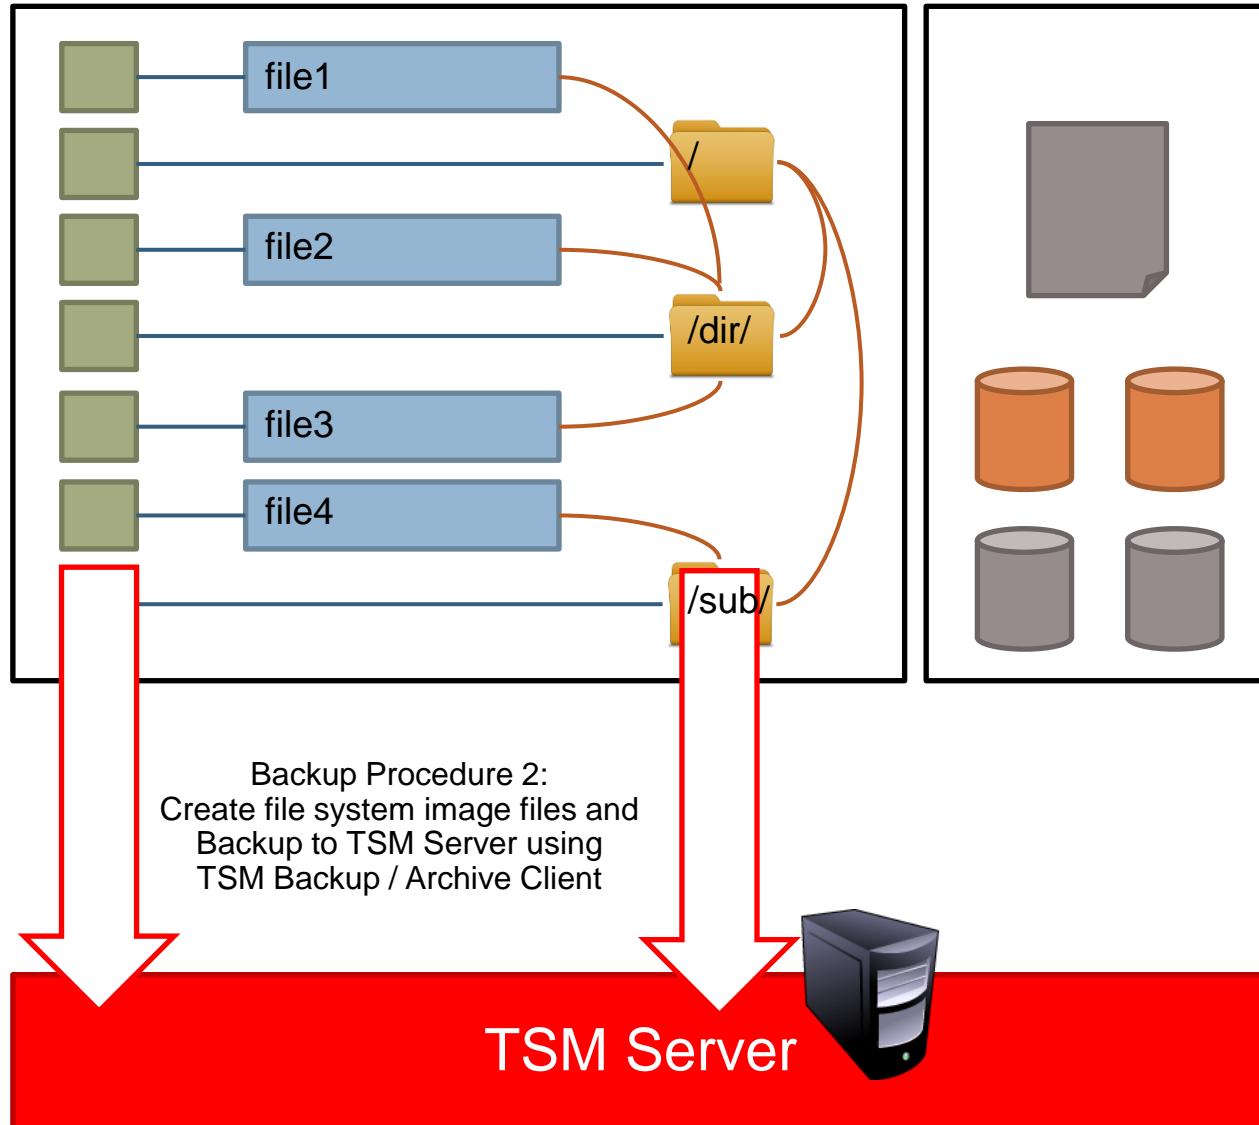


Continously:  
Premigrate all file data  
to TSM Server using  
TSM for Space Management  
and GPFS policy engine

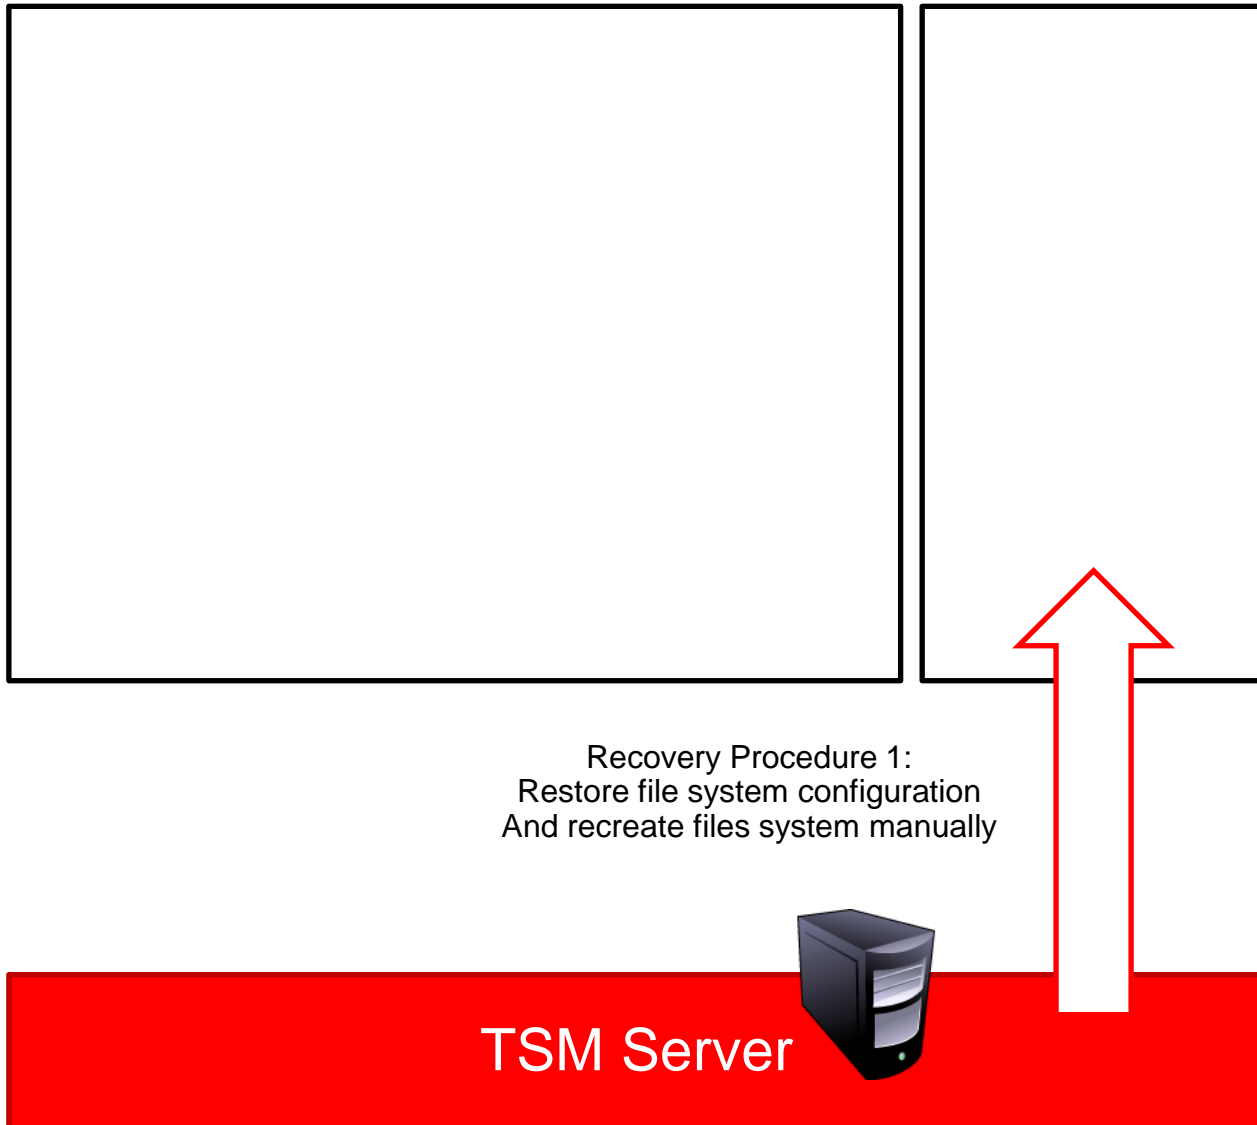
Backup Procedure 1:  
Collect and backup  
file system configuration  
to TSM Server using  
TSM Backup / Archive Client



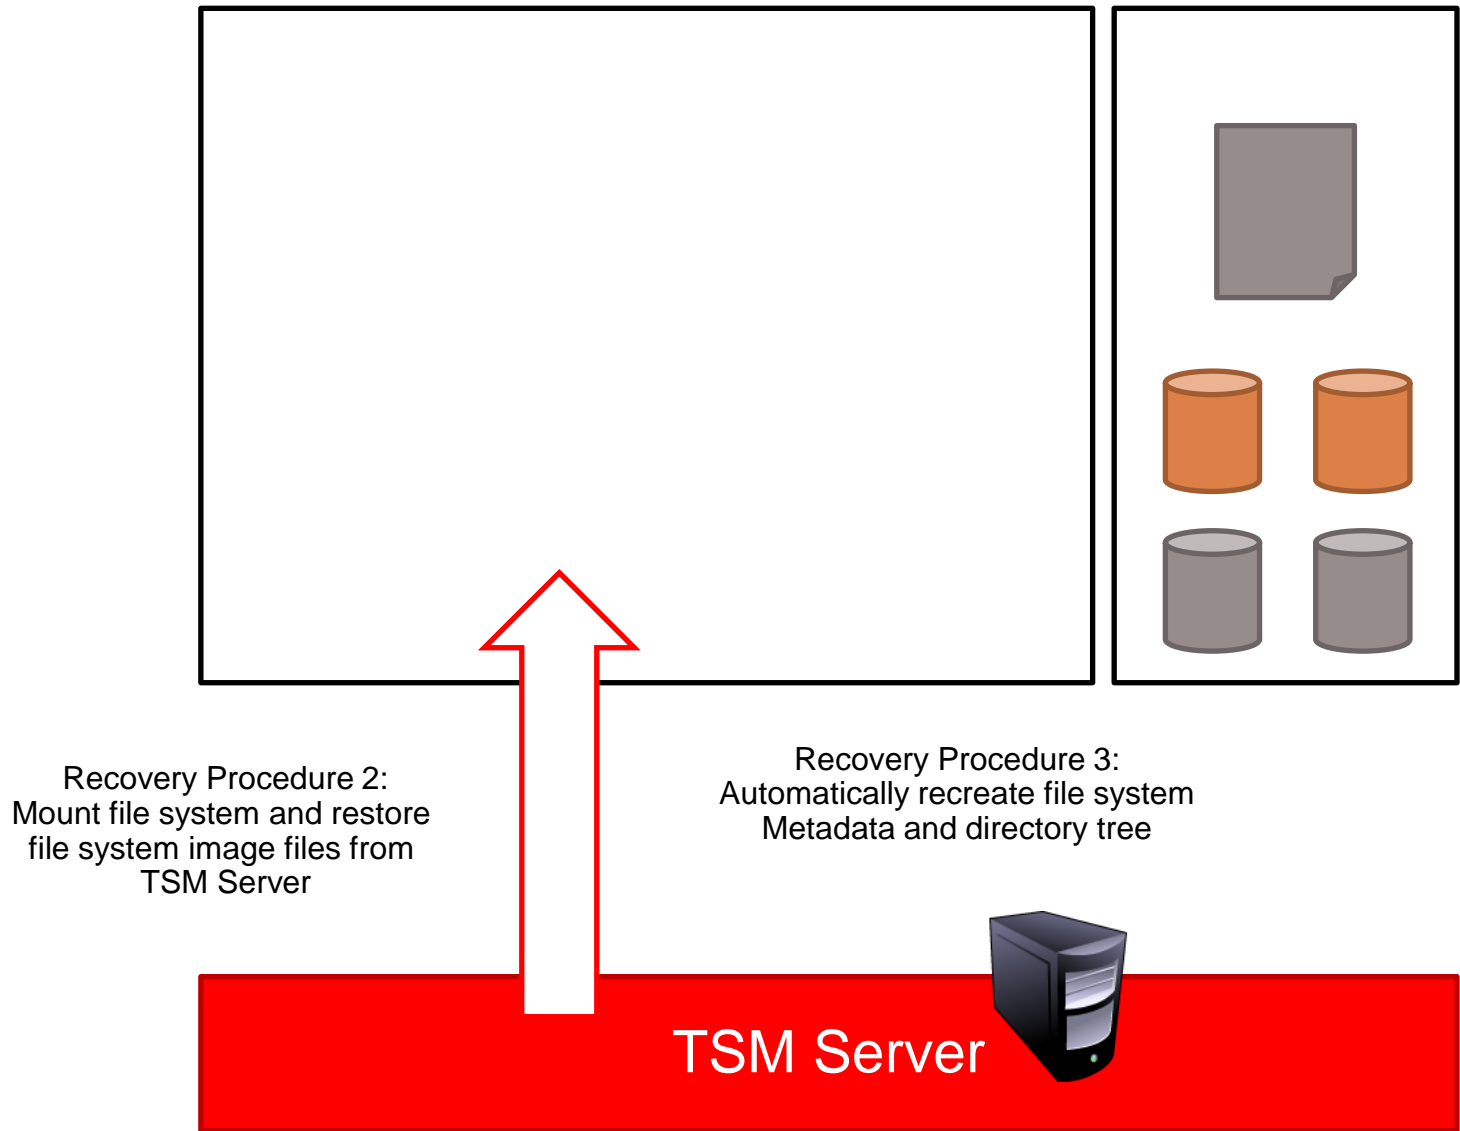
# Scale Out Backup And Restore – Backup II of II



# Scale Out Backup And Restore – Recovery I of III

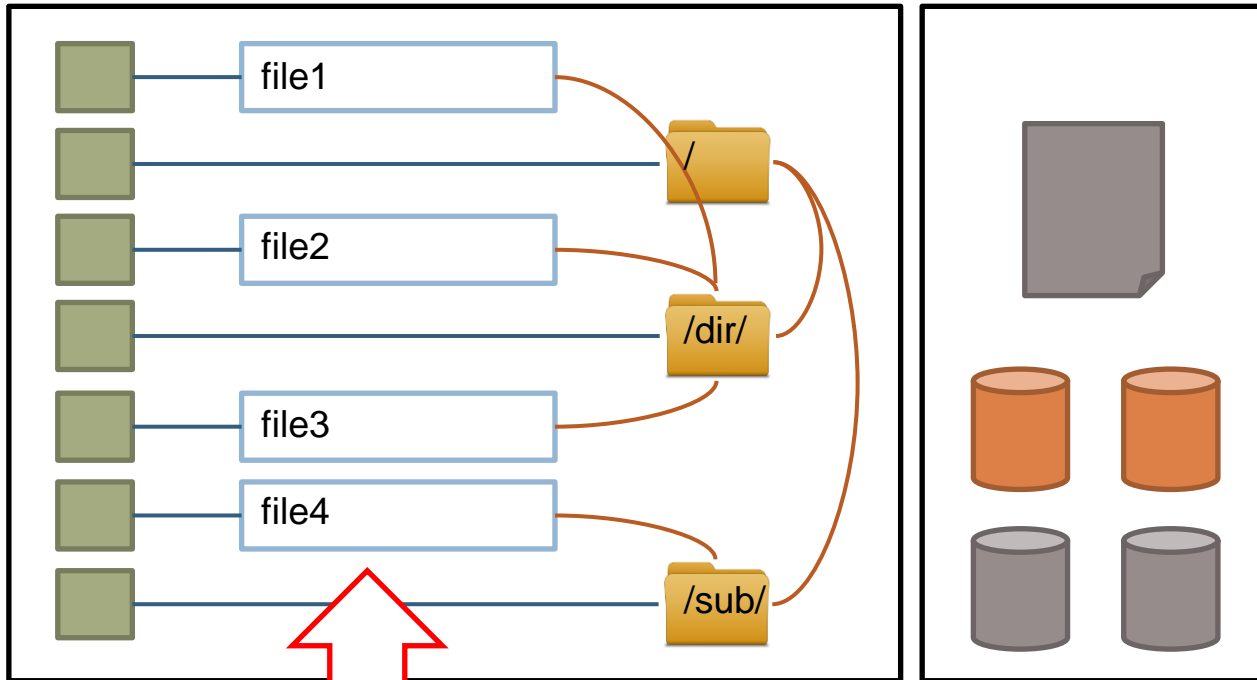


# Scale Out Backup And Restore – Recovery II of III





# Scale Out Backup And Restore – Recovery III of III



Recovery Procedure 4:  
 Enable space management  
 and start production.  
 Recall file data  
 (on demand & background  
 using GPFS policy engine)



Thank You !

# Trademarks

**The following are trademarks of the International Business Machines Corporation in the United States, other countries, or both.**

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by © are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml):

\* AS/400®, e business (logo)®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/30, VM/ESA®, VSE/ESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, BladeCenter®

**The following are trademarks or registered trademarks of other companies.**

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

\* All other products may be trademarks or registered trademarks of their respective companies.

## Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.