



IBM Spectrum **Archive**

**Introduction to
IBM Spectrum Scale™ ILM
and IBM Spectrum Archive™**

Agenda


→ Spectrum Scale ILM

Spectrum Archive

Integration of Spectrum Archive with Spectrum Scale

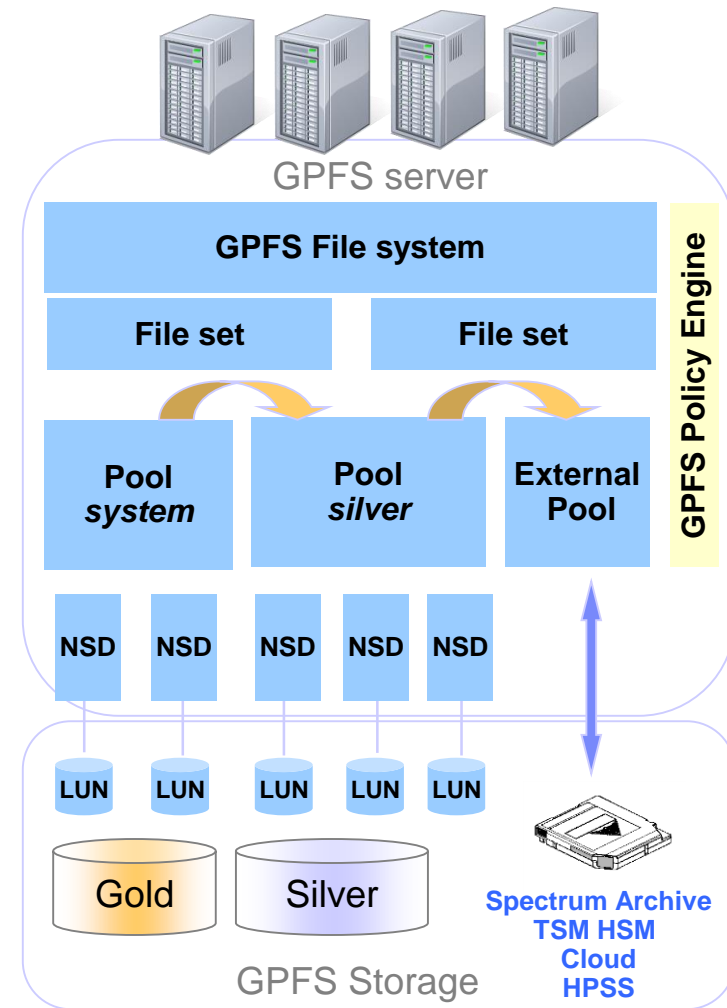
Spectrum Scale and Information Lifecycle Management

- Spectrum Scale supports Information Lifecycle Management
 - ILM means storing data on the most appropriate medium over the lifecycle
- Key techniques to manage storage cost comprise:
 - **Initial placement** of files on the most appropriate storage medium
 - **Automated based migration** during the lifetime of the files based on age, size, etc.
 - Transparent file access in in the original name space
- ILM provides cost efficiency, especially for longer data lifecycles



Spectrum Scale ILM tools

- Files stored in file system and underlying pools
 - Pool is a collection of disks of the same type
- Placement policy controls where files are placed
 - Files are placed in pools
 - Placement policy is applied during file creation
- Migration policies control transparent migration of files from one pool to another
 - Applied during life cycle
 - Migrated files can be transparently accessed
 - Files can also be migrated to tape
- Policy engine applies and executes policies



GPFS storage pools

- Two types of Storage Pools: Internal and External
- **Internal pool** – allows placement and migration
 - A collection of disks or arrays with similar properties that are managed together
 - Every file system has at least a “system” pool
 - GPFS file system metadata is only stored in “system” pool
 - Maximum of 8 storage pools are possible per file system
- **External pool** – allows migration
 - External pool is an interface (script) to an external application
 - Scripts supported for Spectrum Archive, TSM HSM, Cloud Object Storage or HPSS
 - External application is invoked by policy engine or by command
- Pools are designated by names
 - Pool name “system” is reserved for the standard pool

GPFS policies and rules

- A **policy** is a set of rules that describes the life cycle of files
- Each **rule** defines selection criteria for files and operation to be applied
 - File selection based on file attributes (name, size, age, ...)
 - Operation can be placement, migration, deletion, encryption, list, ...
- A policy can include a mix of different rules

Active and scheduled policies

- One active policy can be configured for a file system
 - Is automatically invoked based on rule type and definitions
 - Active ILM policy for each file system can include:
 - File placement rules
 - Threshold migration rules (space triggers)
- One or more scheduled policies can be run for a file system
 - Are invoked manually or via scheduler
 - Scheduled policies include migration rules
 - Based on file attribute (size, age)
- Active migration policies must be threshold based and simple
- Scheduled migration policies do not require threshold and can be more advanced

GPFS policy engine

- Policy engine evaluates and applies policy rules by:
 - Selecting files based on criteria defined in rules
 - Initiating an action defined by the rules (placement, migration)

- Policy engine runs in three phases on multiple GPFS nodes
 1. Metadata scan: reading directory file metadata from inodes
 2. Rule evaluation: matching file metadata against rules (top down) and selecting file
 3. File operation: performing the action (placement, migration, etc.)

Starting GPFS policy engine

- Active policy for a file system is configured with command: `mmchpolicy`
 - Starts policy engine automatically for placement and threshold based migration

```
mmchpolicy filesystem policyfile [-I yes|test]
```

- When threshold is reached `mmapplypolicy` is invoked via callback

- Scheduled policy for a file system is invoked with command:

```
mmapplypolicy filesystem -P policyfile [-I yes|test| defer]
```

GUI integration of ILM in R4.2

- GUI intuitively allows to manage placement and migration policies
 - Including syntax and prerequisite check, displays rules

Information Lifecycle

Create, manage, and delete policies that manage automated tiered storage of information.

Active Policy Policy Repository

File System

Pool **system** (0% used)

Rules Add Rule

1. Placement **Default Placement**
2. Migration **Migrate 30 days (*)** ⊗


⚠ There are some issues with the policy

Apply Changes
Cancel Changes


Migration


Migrate 30 days

Enabled



Source





Target
 !

Migration thresholds: ⚙

Start when source capacity exceeds:
 80 %

Stop when source capacity reaches:
 70 %

Migration scope:

Cancel Changes
Delete Rule

Migration criteria:

Include files if any of these criteria are met:

Agenda

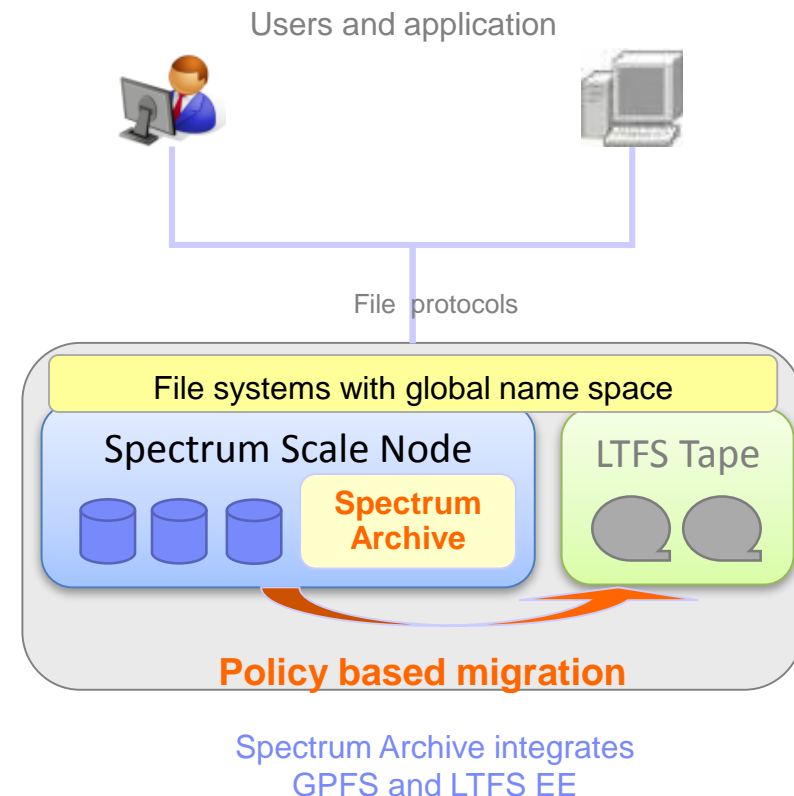
Spectrum Scale ILM

→ **Spectrum Archive**

Integration of Spectrum Archive with Spectrum Scale

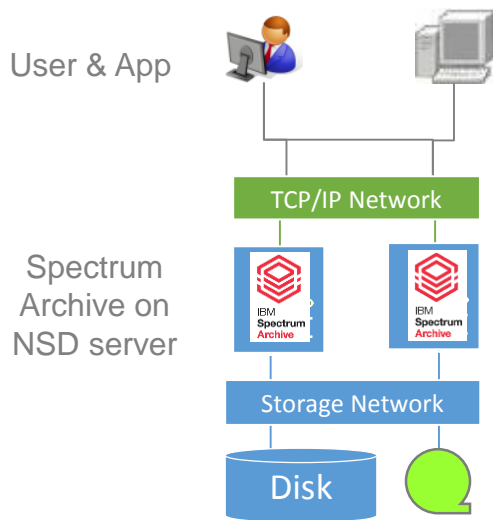
Spectrum Archive Architecture

- Spectrum Archive integrates with Spectrum Scale as tape tier
 - Spectrum Scale provides global name space
 - Spectrum Archive migrates data to tape
- Each Spectrum Archive node has tape drives
 - Supports up to two libraries, one per node
- Files are (pre-) migrated from disk to tape
 - Based on policies or file lists
 - Supports multiple copies on distinct tapes
- Files are recalled on access or by command
 - Supports tape optimized recalls
- Tapes can be exported and imported
- Workload is distributed across nodes



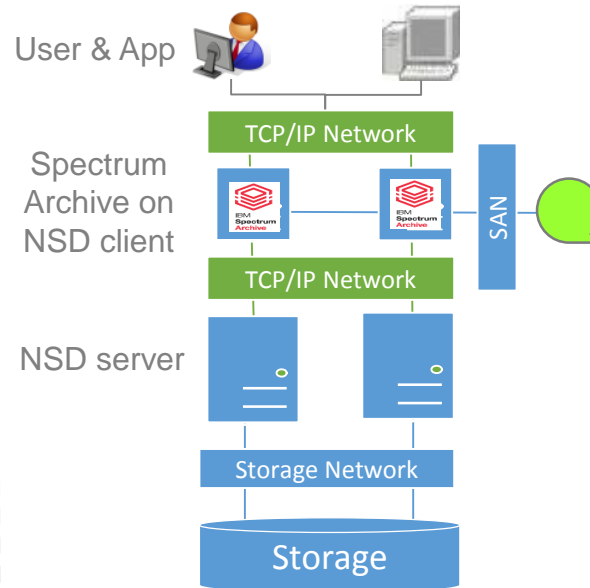
Spectrum Archive - deployment options

On NSD server



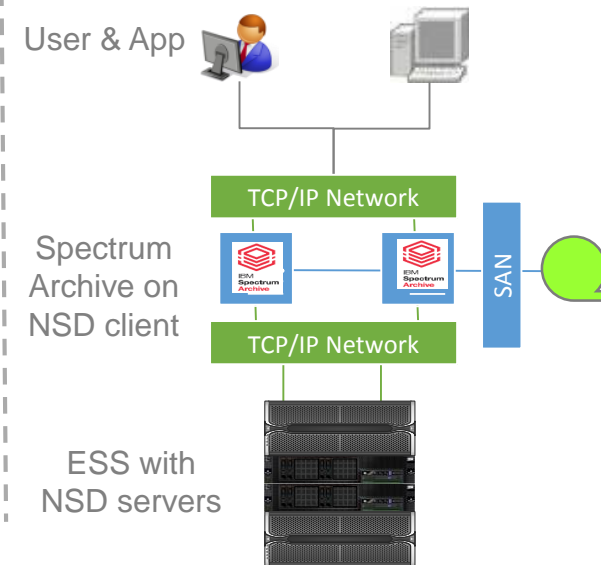
- Spectrum Archive on NSD server
- Direct SAN access disk and tape
- Requires less infrastructure

On NSD client



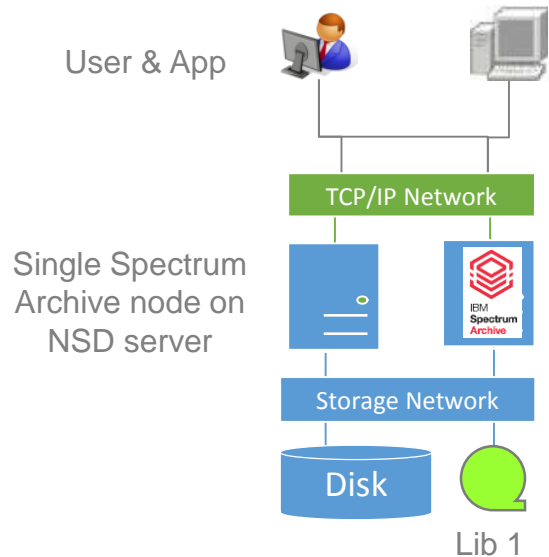
- Spectrum Archive on NSD client connected to NSD server / ESS
- LAN access to disk, SAN access to tape
- Requires more infrastructure

With Elastic Storage Server




Spectrum Archive – configuration options

Single Node



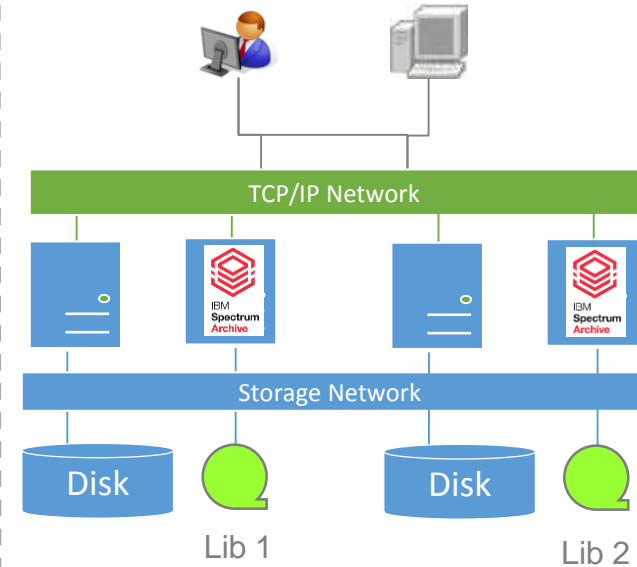
- Single Spectrum Archive node
- Less high availability
- Cost effective

Multi-node single library



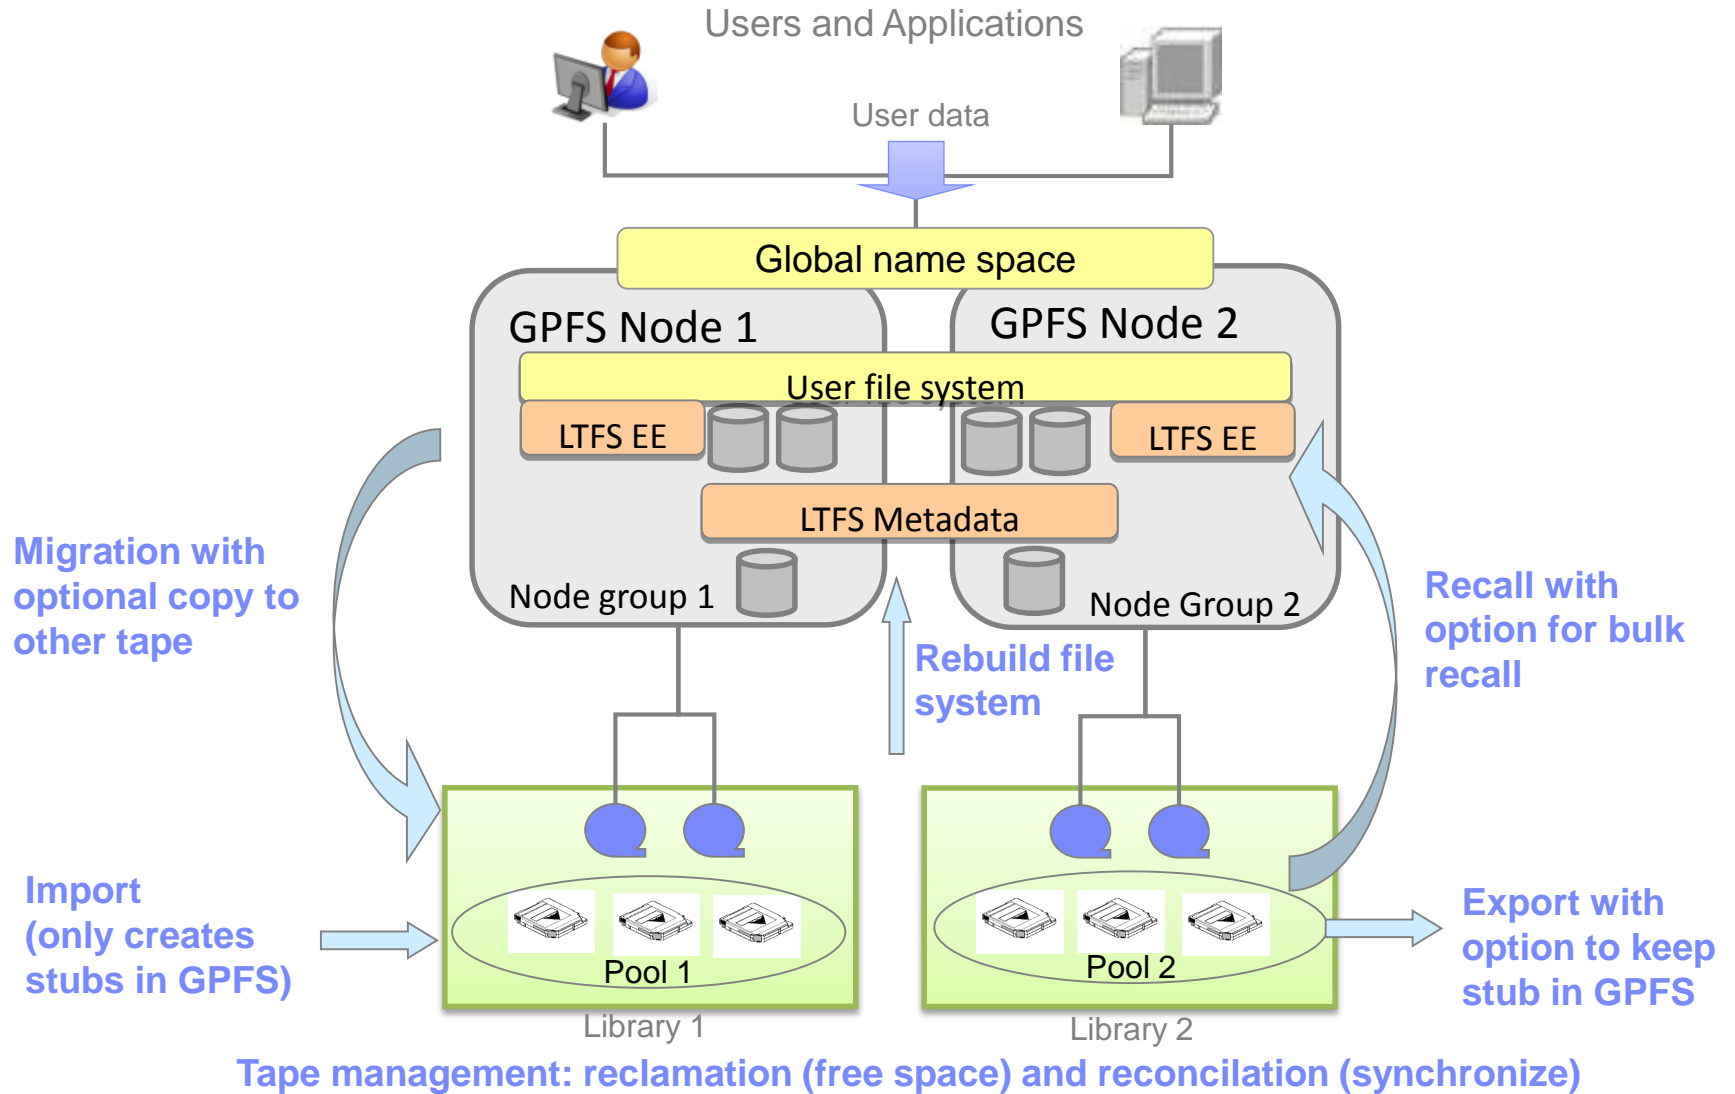
- Multiple Spectrum Archive nodes, one library
- Better high availability
- Scalable performance

Multi-node, two libraries



- Multiple Spectrum Archive nodes, two libraries, can stretch over sites
- Better high availability,
- Better disaster protection

Spectrum Archive functional overview



Spectrum Archive object: technology outlook

- File systems interfaces for tiered storage systems are blessing and cure
 - Transparent access is a blessing, however user have no awareness of data on tape
- Solution: Integrate tape awareness with OpenStack Swift API
 - Let the user / app control migration and recall and execute it efficiently

IceTier Swiftbrowser (Frankfurt) admin:admin

Containers

Name	Objects	Archived	Size	On disk	On tape	
newcont	2	1	82,6 KB	57,7 KB	24,9 KB	
testcont	2	1	130,2 KB	128,9 KB	1,3 KB	

Total size on disk: 186,6 KB

Total size on tape: 26,2 KB

- Delete container
- Migrate container
- Recall container

- “Swift associated project”: SwiftHLM middleware
 - Session 3) Object Storage und Swift HLM – 11:00 – 12:00

Agenda

Spectrum Scale ILM

Spectrum Archive

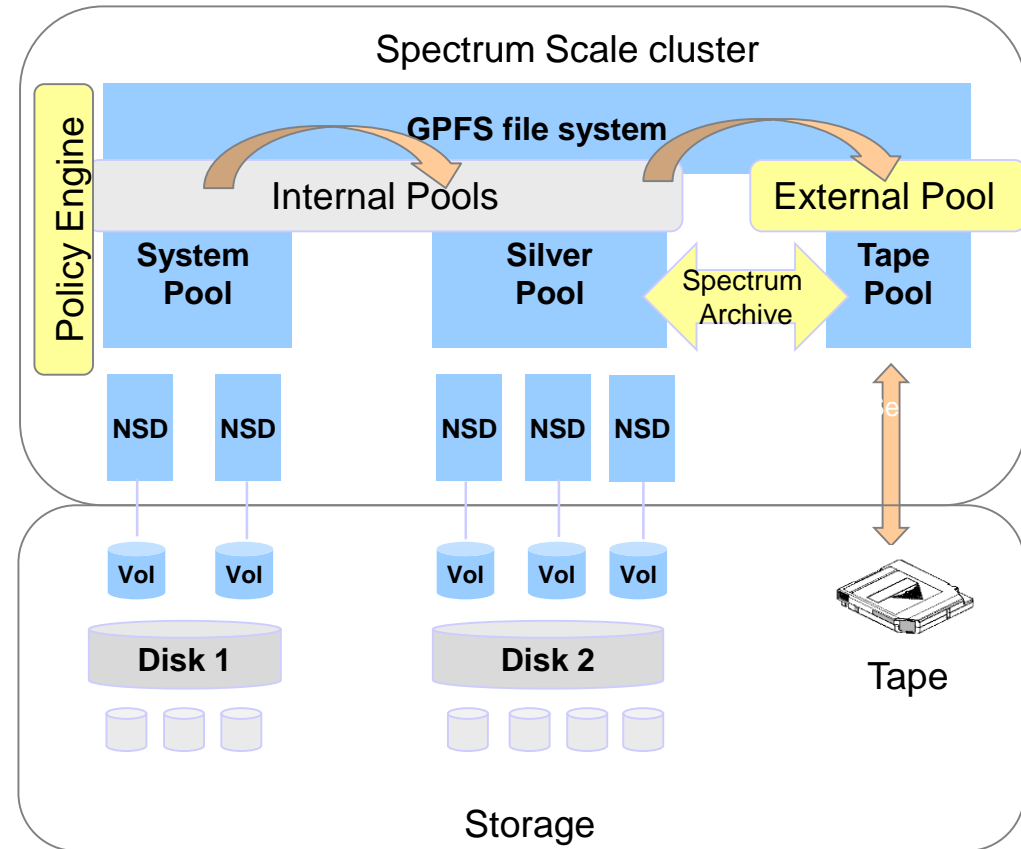
→ Integration of Spectrum Archive with Spectrum Scale

Spectrum Scale and Spectrum Archive integration

- Spectrum Archive is external pool in Spectrum Scale file system
 - Performs migration and recall on LTF5 formatted tapes

- Policy engine can be used to identify files based on attributes and invokes Spectrum Archive for migration
 - List of files can also be migrated using commands.

- Access to files is transparent
 - Upon access file is recalled



Example for active migration policy with threshold

- Migration based on threshold

- Register callback to trigger migration on lowdiskspace

```
mmaddcallback MIGRATION --command /usr/lpp/mmfs/bin/mmstartpolicy  
--event lowDiskSpace,noDiskSpace --parms "%eventName %fsName  
-N ltfsnodes -m 3 -B 1000 -s /workdir --single-instance"
```

- Create policy text file:

```
RULE EXTERNAL POOL 'ltfs'  
EXEC '/opt/ibm/ltfsee/bin/ltfsee' OPTS '-p Pool1@lib1 Pool2@lib2'  
  
RULE 'eesysmig' MIGRATE FROM POOL 'system' THRESHOLD(80,60) TO  
POOL 'ltfs' WHERE (KB_ALLOCATED > 0)  
RULE 'default' set pool 'system'
```

- Set the policy for the file system

```
mmchpolicy <file-system> policyfile.txt -t "system policy" -I yes
```

Example for scheduled migration without threshold

- Migration based on schedules

- Create policy text file:

```
RULE EXTERNAL POOL 'ltfs'  
EXEC '/opt/ibm/ltfsee/bin/ltfsee' OPTS '-p Pool1@lib1 Pool2@lib2'  
RULE 'eeschedmig' MIGRATE TO POOL 'ltfs' WHERE (KB_ALLOCATED > 0)  
AND (DAYS(CURRENT_TIMESTAMP) - DAYS(MODIFICATION_TIME) > 2)
```

- Run the policy:

```
mmapplypolicy <file-system> -P policyfile -m 3 -N ltfsnodes -s  
/workdir --single-instance
```

- Schedule the policy via system scheduler

General guidance for policies with Spectrum Archive

- Threshold based migration is last line of defense
 - Size file system to not run full and use schedule policies to move files periodically
 - Policies for threshold based migration should be simple
- Policies for scheduled migration can be more sophisticated
 - Use file attributes (`KB_ALLOCATED`, `ACCESS_TIME`, `MODIFICATION_TIME`, `NAME`)
- Default WEIGHT Parameter: **KB_ALLOCATED**
 - Consider changing to e.g. `ACCESS_TIME` or `NOT(FILE_HEAT)`
- Do not migrate file from certain directories (`.SpaceMan`, `.snapshot`, `.lftsee`)
- Consider limits of maximum number of files in queue
 - Migrate. < 30.000; recall < 15.000

Considerations with other Spectrum Scale functions

- Using snapshots can cause transparent recalls of files
 - When migrated files that are in snapshots are deleted
- Using AFM with Spectrum Archive is not supported (yet)
 - Certain use cases are being tested (SW and DR mode)
- Using Spectrum Archive and TSM HSM in the same cluster is not supported
- Integration with „mmbackup“ is rudimentary
 - No inline backup
 - With release 4.2 and above mmbackup skips migrated files
 - Use Spectrum Protect for Space Management (HSM) alternatively

Learn all about migration policies for Spectrum Archive



- Spectrum Scale ILM Policy Guide

<http://w3-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102642>

What to learn more about Spectrum Scale and Spectrum Archive?

- **Spectrum Scale Standard hands-on workshop**
 - Learn about architecture & concepts, installation and configuration, ILM, CES, AFM, Backup,
 - <https://academy.avnet.com/de/training/course/141611>

- **Spectrum Scale Advanced hands-on workshop**
 - Learn about architecture & concepts, monitoring, configuration parameters, tools and problem determination
 - <https://academy.avnet.com/de/training/course/136736>

- **Spectrum Archive Standard hands-on workshop**
 - Learn about ILM concepts, installation and configuration, using policies,
 - <https://academy.avnet.com/de/training/course/141610>

→ We also offer customized workshops according to your needs !

Thank You

Links and references

- LTF5 format specification ISO standard 20919:16
http://www.iso.org/iso/home/store/catalogue_tc/catalogue_detail.htm?csnumber=69458
- LTF5 format specification SNIA standard
http://www.snia.org/sites/default/files/LTF5_Format_2.2.0_Technical_Position.pdf
- LTF5 home page
<http://www-03.ibm.com/systems/storage/tape/ltf5/index.html>
- Spectrum Archive EE knowledge Center:
<http://www.ibm.com/support/knowledgecenter/ST9MBR/welcome?lang=en>
- Redbook: LTF5 EE 1.1.1.2
<http://www.redbooks.ibm.com/redpieces/abstracts/sg248143.html>
- Redbook: Spectrum Archive 1.2
www.redbooks.ibm.com/redpieces/abstracts/sg248333.html?Open
- Redbook: Spectrum Archive with Spectrum Scale Object:
<http://www.redbooks.ibm.com/abstracts/redp5237.html?Open>
- Whitepaper: Spectrum Archive solution and use cases
<https://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102504>
- Whitepaper: Spectrum Protect HSM with Spectrum Scale AFM:
<https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Storage%20Manager/page/Configuring%20IBM%20Spectrum%20Scale%20Active%20File%20Management>
- LTF5 SDE Info center
<http://publib.boulder.ibm.com/infocenter/ltf5/cust/index.jsp>
- LTF5 LE Infocenter:
<http://pic.dhe.ibm.com/infocenter/ltf5le/cust/index.jsp>
- LTF5 Redbook: Installation and configuration
<http://www.redbooks.ibm.com/abstracts/sg248090.html?Open>
- Whitepaper: Using LTF5
<http://public.dhe.ibm.com/common/ssi/ecm/en/tsl03109usen/TSL03109USEN.PDF>
- LTF5 software and IBM Device driver:
<http://www-933.ibm.com/support/fixcentral/>
- Almaden Research
<http://www.almaden.ibm.com/storagesystems/projects/ltf5/>
- The top Youtube videos:
LTF5 customer reference: https://www.youtube.com/watch?v=6s_hjzul9Y4
LTF5 Intro: <http://www.youtube.com/watch?v=Qtd4CrajCYs>
Bamm TV: <http://www.youtube.com/watch?v=X82sC97yQeE>

Disclaimer

This information is for **IBM Spectrum Scale user day 2017 use only**, publication beyond this scope is forbidden

This information is provided on an "AS IS" basis without warranty of any kind, express or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Some jurisdictions do not allow disclaimers of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

This information is provided for information purposes only as a high level overview of possible future products. **PRODUCT SPECIFICATIONS, ANNOUNCE DATES, AND OTHER INFORMATION CONTAINED HEREIN ARE SUBJECT TO CHANGE AND WITHDRAWAL WITHOUT NOTICE.**

USE OF THIS DOCUMENT IS LIMITED TO SELECT IBM PERSONNEL THIS DOCUMENT SHOULD NOT BE GIVEN TO A CUSTOMER EITHER IN HARDCOPY OR ELECTRONIC FORMAT.

Important notes:

IBM reserves the right to change product specifications and offerings at any time without notice. This publication could include technical inaccuracies or typographical errors. References herein to IBM products and services do not imply that IBM intends to make them available in all countries.

IBM makes no warranties, express or implied, regarding non-IBM products and services, including but not limited to Year 2000 readiness and any implied warranties of merchantability and fitness for a particular purpose. IBM makes no representations or warranties with respect to non-IBM products. Warranty, service and support for non-IBM products is provided directly to you by the third party, not IBM.

All part numbers referenced in this publication are product part numbers and not service part numbers. Other part numbers in addition to those listed in this document may be required to support a specific device or function.

MHz / GHz only measures microprocessor internal clock speed; many factors may affect application performance. When referring to storage capacity, GB stands for one billion bytes; accessible capacity may be less. Maximum internal hard disk drive capacities assume the replacement of any standard hard disk drives and the population of all hard disk drive bays with the largest currently supported drives available from IBM.

IBM Information and Trademarks

The following terms are trademarks or registered trademarks of the IBM Corporation in the United States or other countries or both: the e-business logo, IBM, IBM Spectrum Scale, IBM Spectrum Archive, IBM Spectrum Protect

Intel, Pentium 4 and Xeon are trademarks or registered trademarks of Intel Corporation.

The SNIA logo is a registered trademark of Storage Networking and Industry Association.

The ISO logo is registered trademark of International Organization for Standardization

Microsoft Windows is a trademark or registered trademark of Microsoft Corporation.

Linux is a registered trademark of Linus Torvalds.

Other company, product, and service names may be trademarks or service marks of others.