



IBM Spectrum **Archive**

**Introduction to
IBM Spectrum Archive™**

Agenda

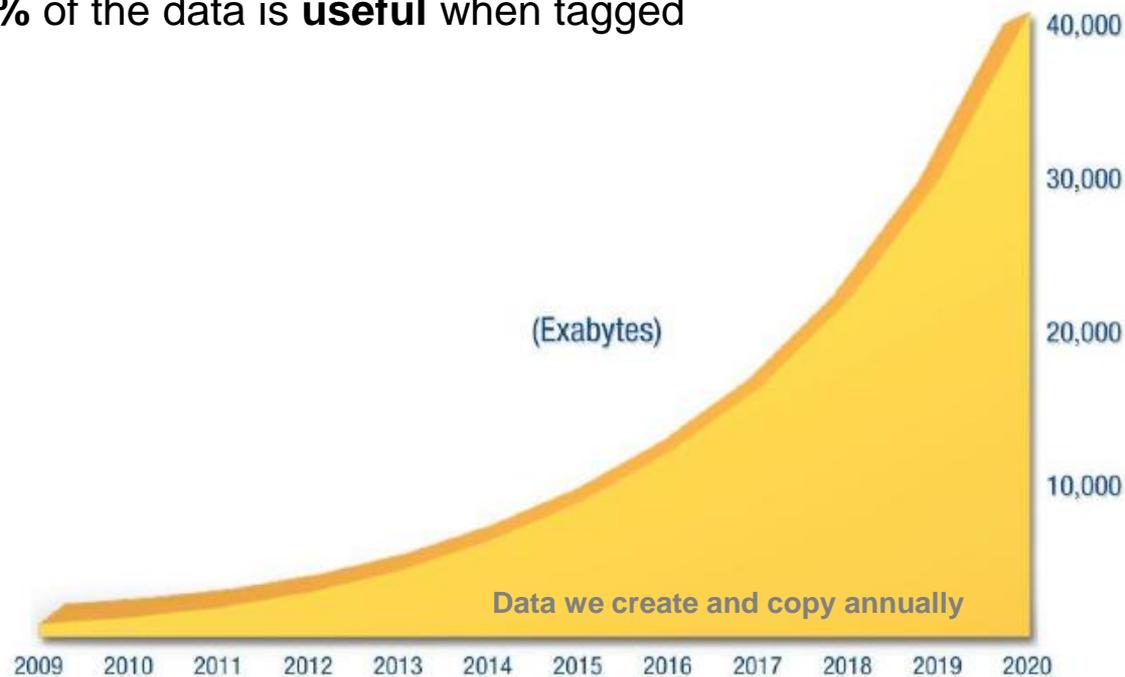
→ Motivation

Spectrum Archive Overview

Positioning & Use cases

The digital universe

- By 2020 the digital universe will reach 44 zettabytes, or 44 trillion gigabytes
 - It is **doubling in size every two years**
 - **85 % of the data** is created by enterprises
 - **22 %** of the data is **useful** when tagged



Source: IDC's Digital Universe Study, sponsored by EMC, December 2012

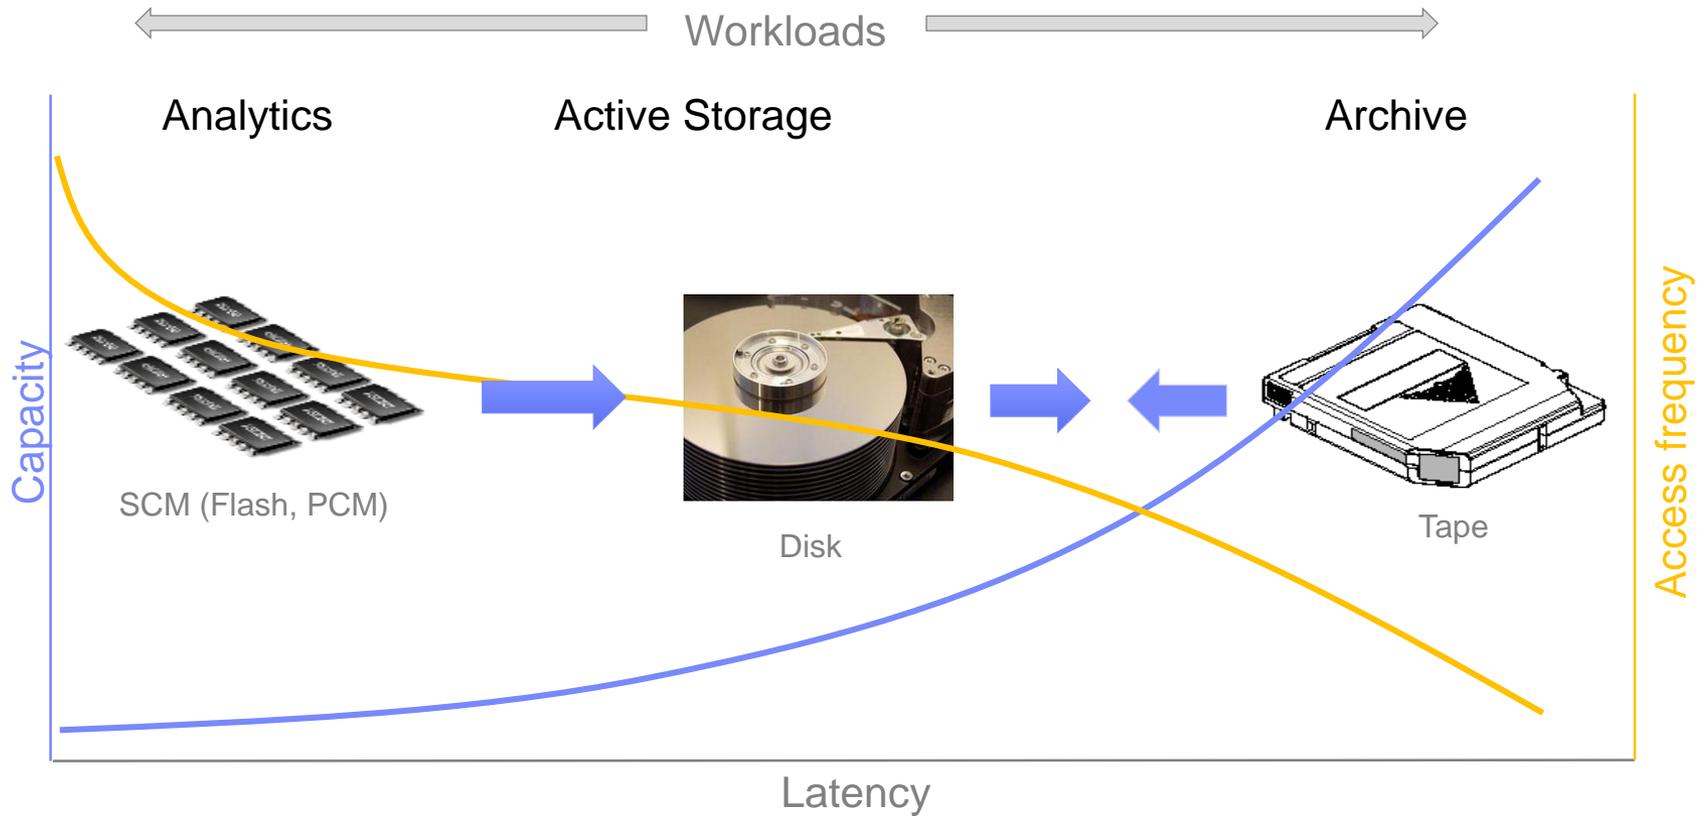
Observation for large storage demands

- Characteristic of large capacity storage
 - High ingest volumes – requires high ingest performance
 - High processing performance
 - Low retrieve rates once data has aged
 - Data is static after processing
 - Data needs to be kept for longer period of times (+5 years)

- Disk is used primarily to address large storage capacity demands
 - But what about cost (power, cooling operations) and environment

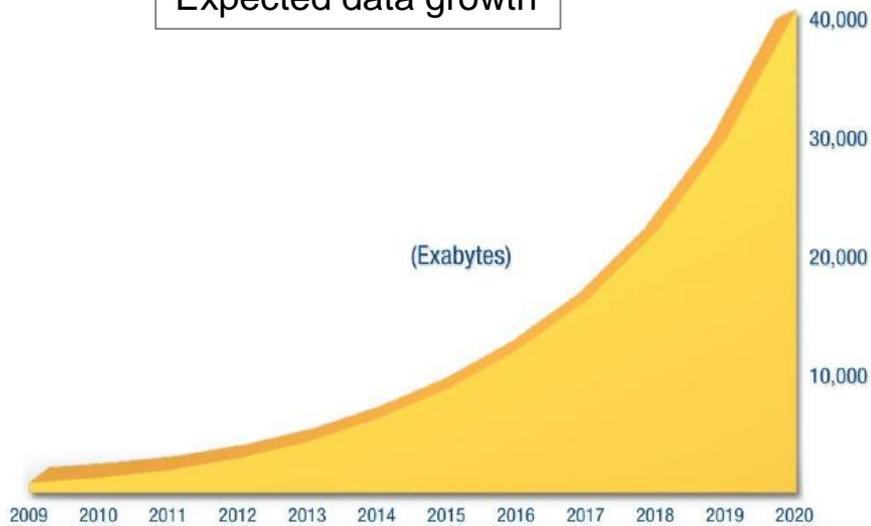
- What other storage options do we have?

Storage options



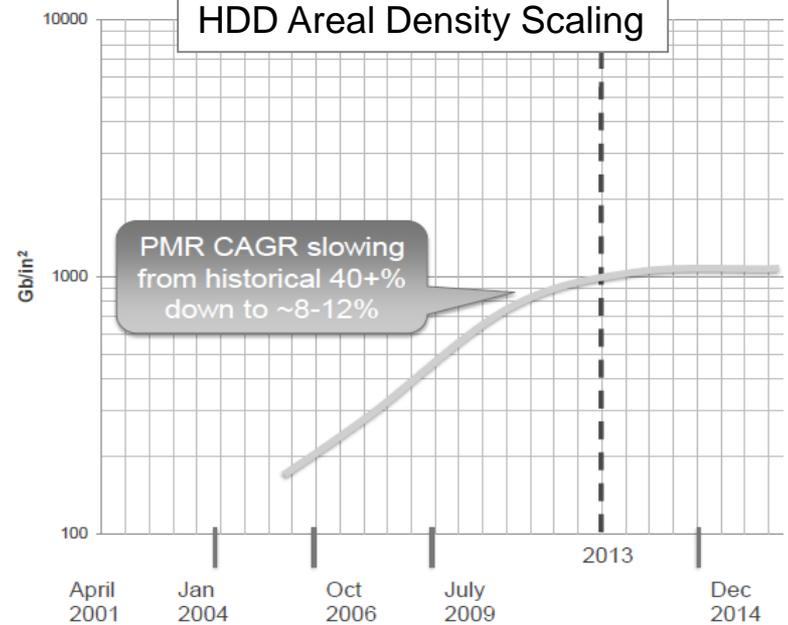
Scalability of disk

Expected data growth



Source: IDC's Digital Universe Study, sponsored by EMC, December 2012

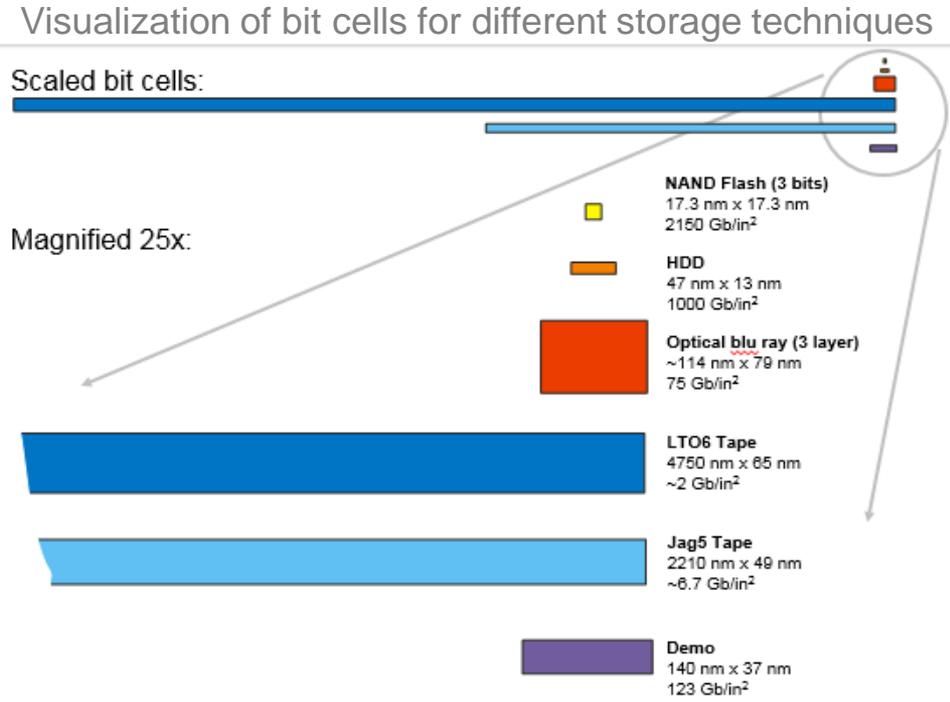
HDD Areal Density Scaling



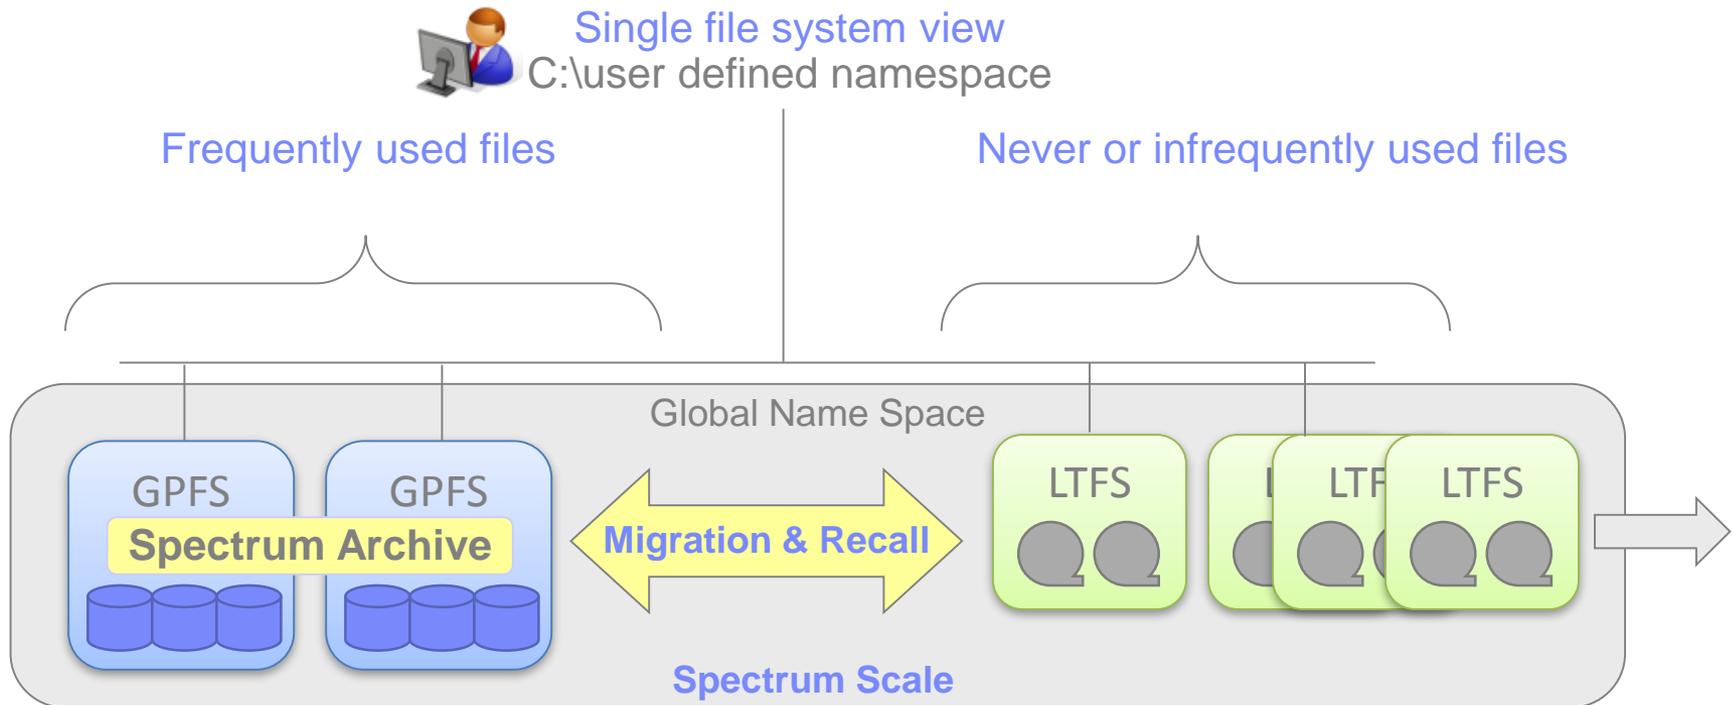
Source: D. Anderson, 2013 IEEE Conf. on Massive Data Storage

Scalability of tape

- April 2015: IBM Research demonstrated a new record of:
 - 123 Gbits/In²**
 in areal data density on magnetic particulate tape
 - LTO-6 has 1.38 Gbits/in²
- At this areal density, a standard LTO size cartridge could store up to
 - 220 terabytes**
 of uncompressed data*
 - 88 times improvement over an LTO-6



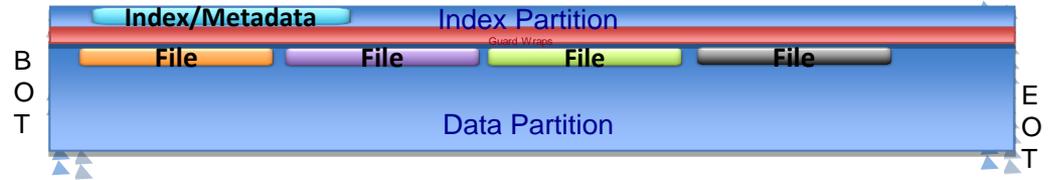
Combine disk and tape optimizing storage cost in scaling environments



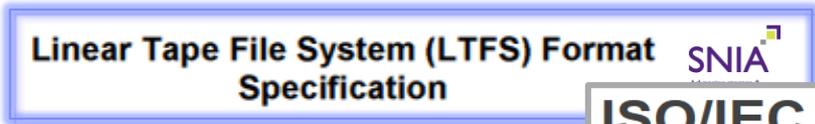
Move files which are no longer accessed to LTFS tape, leveraging automation, transparent access and standardized format.

LTFS is a file system on tape comprised of the following components

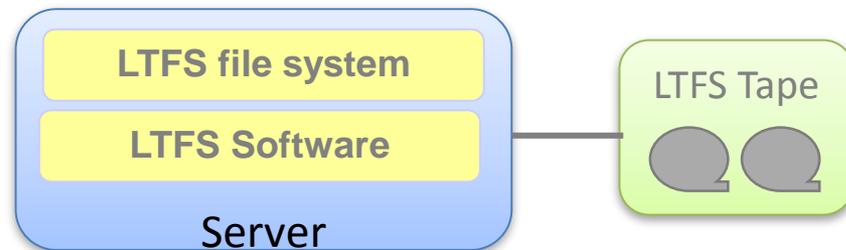
- **Tape partitioning** in index and data partition



- **Standardized LTFS format specification** describes the index and data layout



- **LTFS software** presents the file system and manages data on tapes based on LTFS format specification

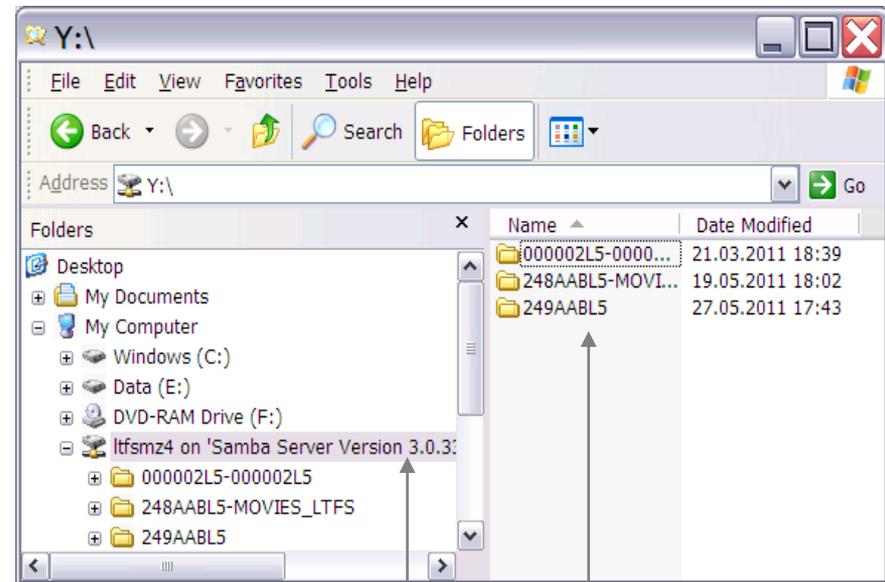


Challenges with Standard LTFS

- LTFS LE presents each tape as sub-directory under the LTFS root file system
 - Files are dropped in sub-directories

- Challenges
 - All file I/O is directed to tape
 - No caching, not user friendly
 - User does not see if a tape is full
 - File system does not show this
 - No reclamation
 - Must be scripted

- Bottom Line: LTFS LE requires additional management



LTFS-2 file system

LTFS Tapes in Library

Agenda

Motivation

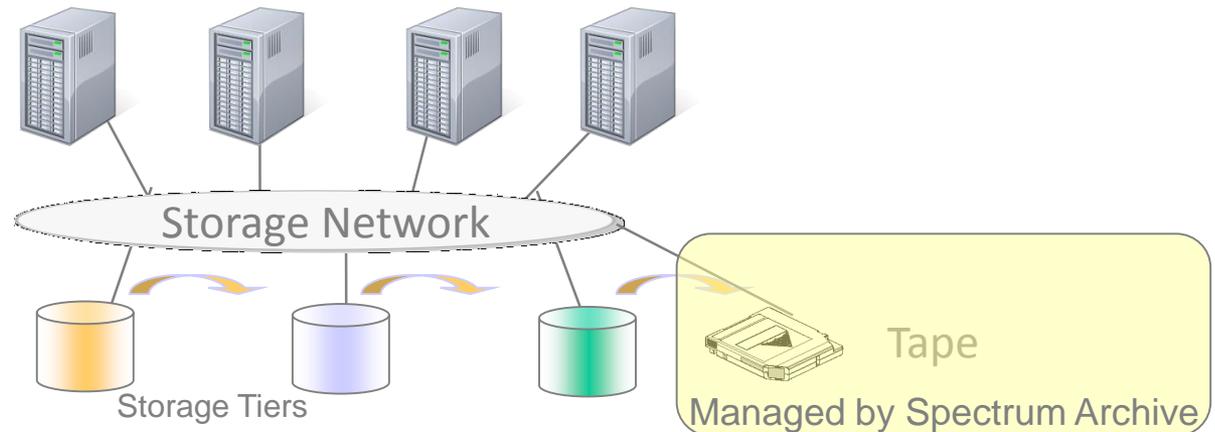
→ **Spectrum Archive Overview**

Positioning & Use cases

Spectrum Scale Information Lifecycle Management

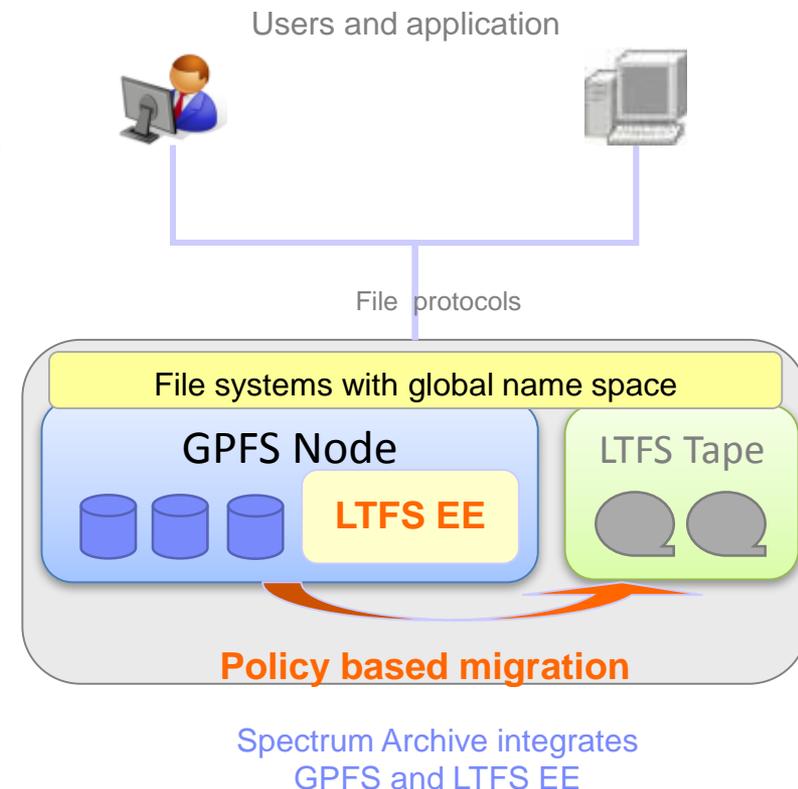
- Spectrum Scale implements Information Lifecycle Management
- Key ILM techniques to manage storage cost include:
 - Initial placement of files on the most appropriate storage medium
 - Policy based migration during the lifetime of the files based on age, size, etc.
 - Transparent file access in in the original name space
- **Spectrum Archive** integrates with GPFS and provides tape storage tier
- ILM provides cost efficiency, especially for longer data lifecycles

Spectrum Scale with LTFS EE
(Spectrum Archive)



Spectrum Archive Architecture

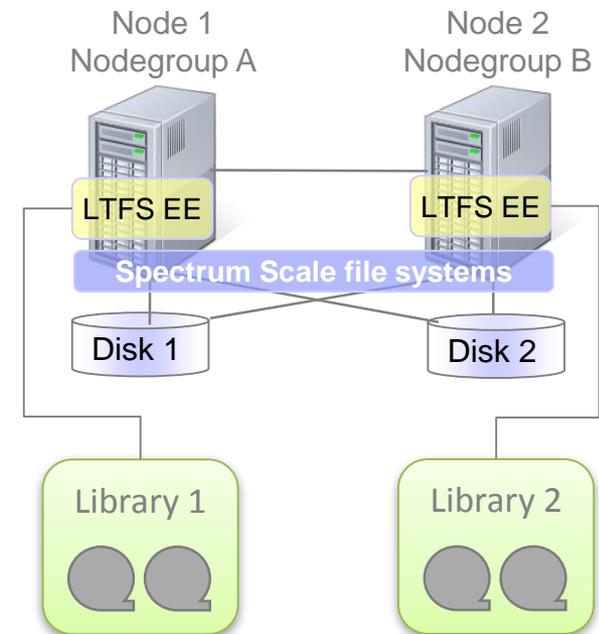
- Spectrum Archive integrates with Spectrum Scale as tape tier
 - Spectrum Scale provides global name space
 - Spectrum Archive migrates data to tape
- Each Spectrum Archive node has tape drives
 - Supports up to two libraries, one per node
- Files are (pre-) migrated from disk to tape
 - Based on policies 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



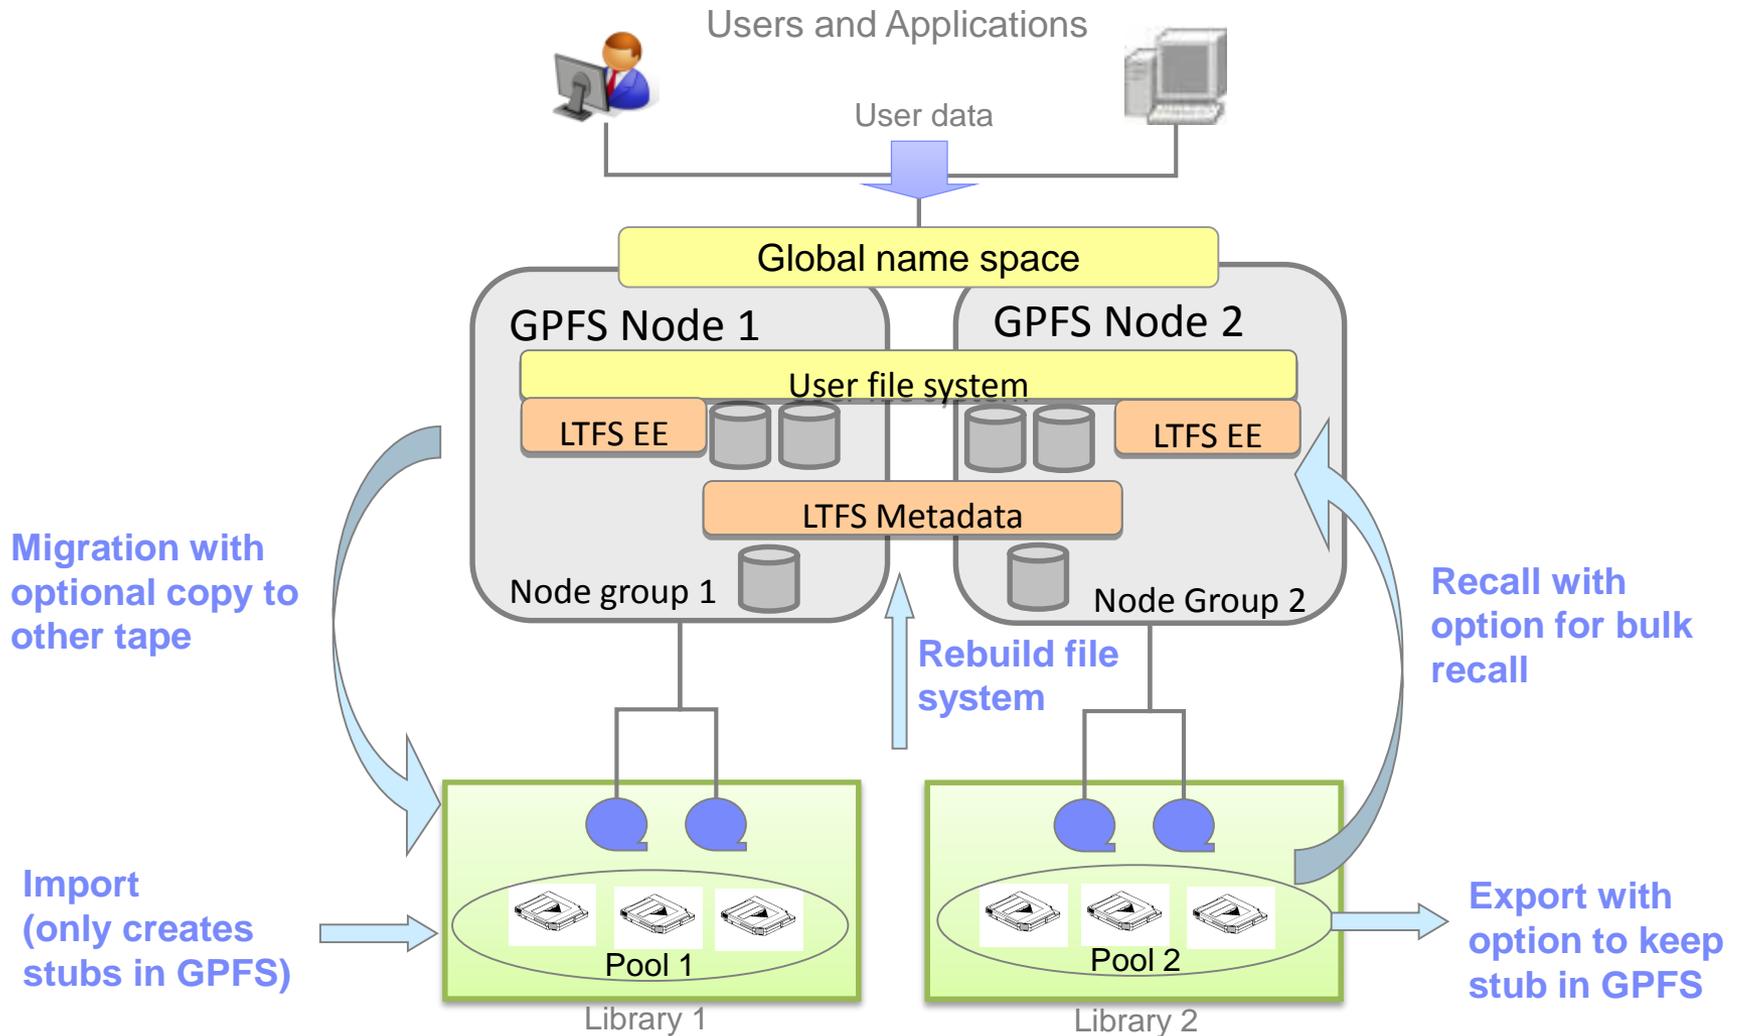
Spectrum Archive Architecture with two tape libraries

- At least two Spectrum Archive nodes required
 - Each manages one tape library
- Each node is in one node group
 - Nodes in one node group are connected to the same library and share all tape resources
 - Additional nodes can be added to one node group
- Node groups can be stretched over two locations
 - Files are replicated by Spectrum Scale on disk
 - Files are migrated by Spectrum Archive to two tapes
 - Read locality can be configured

Spectrum Archive stretched cluster



Spectrum Archive functional overview



Tape management: reclamation (free space) and reconciliation (synchronize)

Advantages of Spectrum Archive

- Lower TCO by leveraging cost effective tape storage
- Seamless data access in continuous name space
- Automated, policy based movement from disk to tape
- Tape optimized recall to accelerate retrieves
- Standardized LTFS format facilitates data exchange
- Support for transparent tape encryption
- Data protection through multiple copies on tape
- Support for immutable files on WORM tapes
- Easy administration and management



IBM
Spectrum
Archive

Agenda

Motivation

Spectrum Archive Overview

→ Positioning and use cases

Positioning Spectrum Archive in general

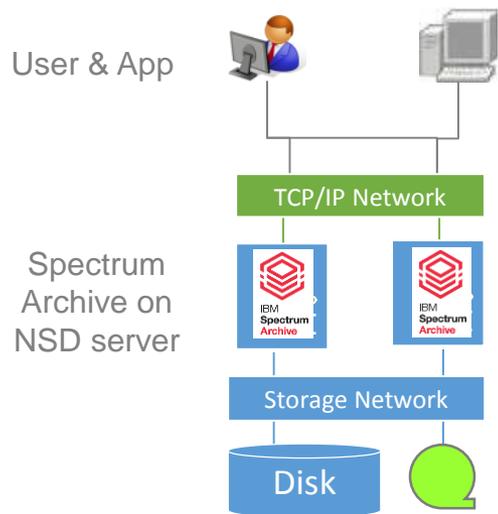
Storing large volumes of larger files which are infrequently accessed on tape

100s of TB > 10 MB Files are never or rarely accessed

- Optimize Total Cost of Ownership leveraging tape
 - Providing easy access to files stored on tape in a tiered storage system
 - Transparent user access to files via GPFS file system layer
 - Automated migration from disk to tape using GPFS policies
 - Exchanging data on LTFS tape
 - Leverage copy, export and import functionality for LTFS tapes
- Spectrum Archive does not make tape faster but much easier to use in many industries and branches**

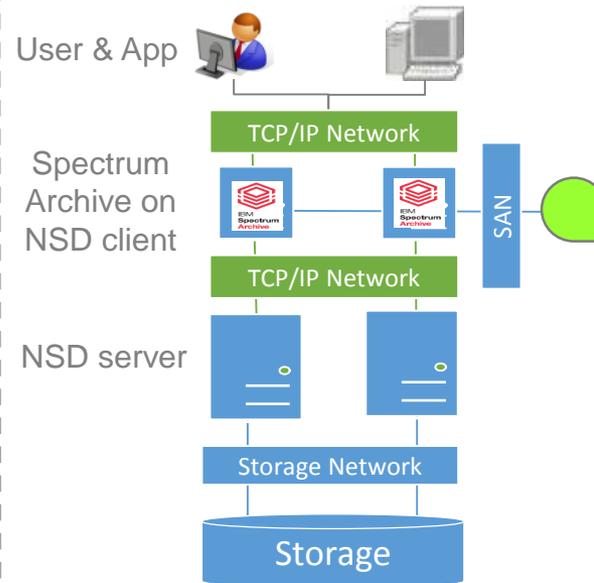
Spectrum Archive - deployment options

On NSD server



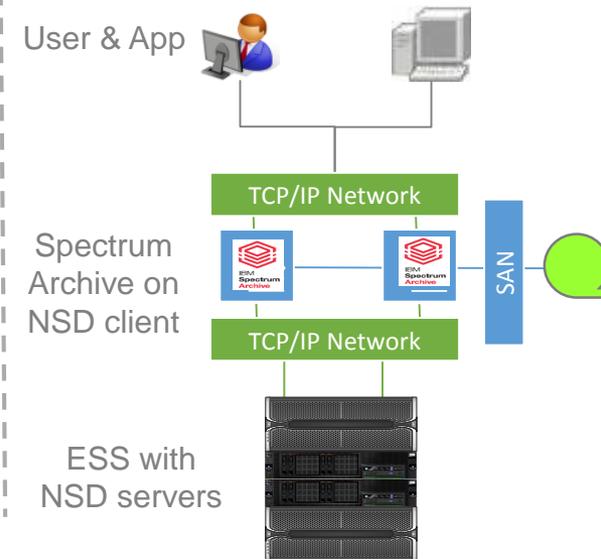
- Spectrum Archive on NSD server
- Direct SAN access disk and tape
- Requires less infrastructure

On NSD client



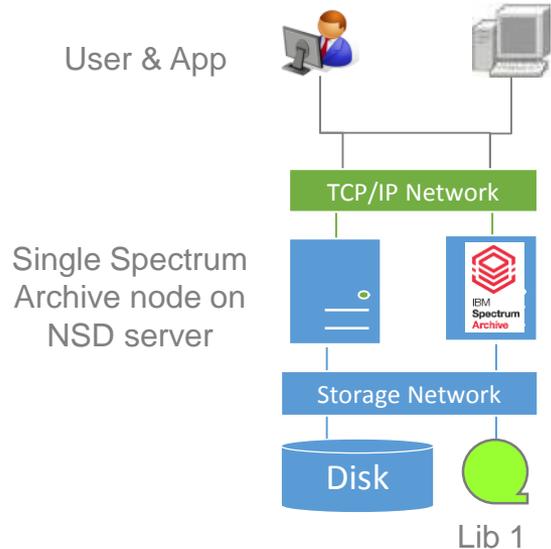
- Spectrum Archive on NSD client connected to NSD server / ESS
- LAN access to disk, SAN access to tape
- Requires more infrastructure

With Elastic Storage Server



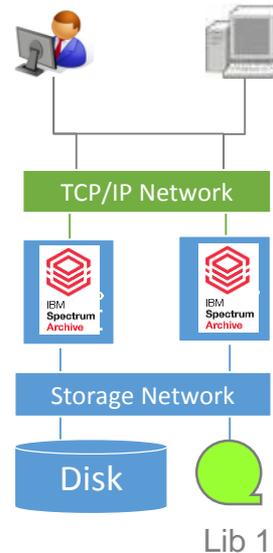
Spectrum Archive – configuration options

Single Node



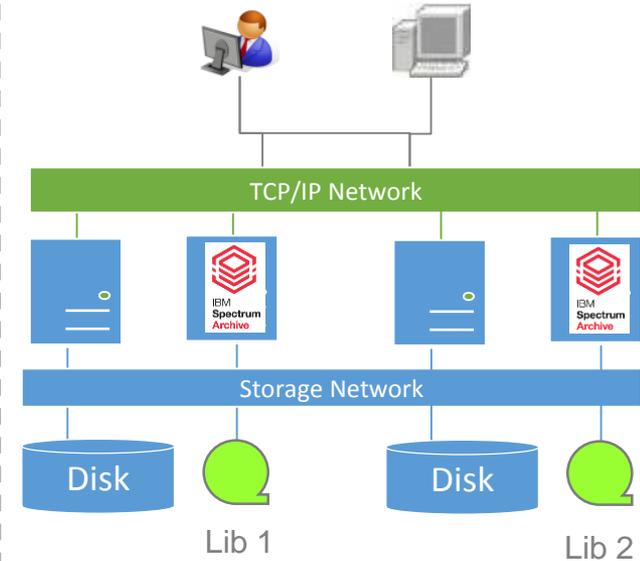
- Single Spectrum Archive node
- Less high availability
- Cost effective

Multi-node single library



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

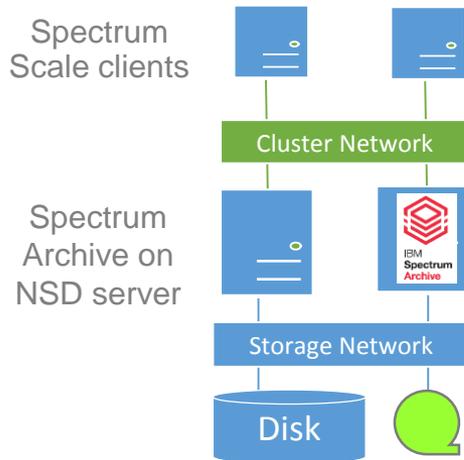
Multi-node, two libraries



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

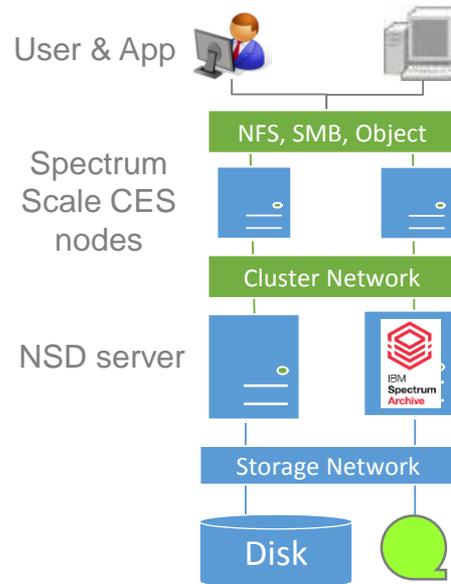
Spectrum Archive – file access options

Access via GPFS



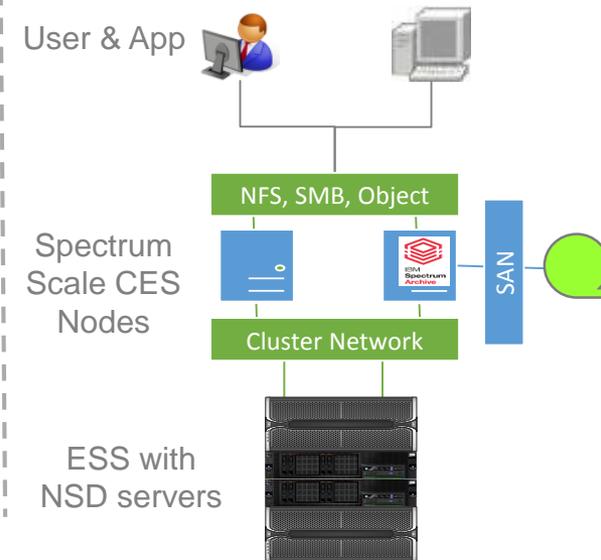
- Direct access to Spectrum Scale file system
- Spectrum Archive can run on NSD server or NSD clients
- Parallel access performance

Access via CES



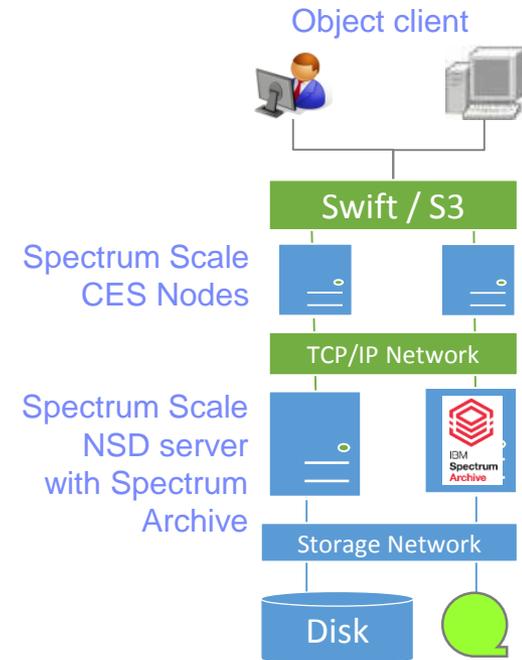
- Access to Spectrum Scale file system via CES nodes (NFS, SMB, Object)
- Spectrum Archive can run on NSD server (not with ESS) or NSD clients
- Perhaps better reliability with separated servers for CES, NSD server and Spectrum Archive

Access via CES on ESS



Spectrum Archive with Spectrum Scale Object

- Spectrum Scale support object storage (Swift, S3)
 - Object storage provided by CES nodes
 - Objects are organized in container / buckets
 - Each object container / bucket is a file set
- Spectrum Archive can migrate objects to tape
 - Either from specific containers
 - Or for all containers
- Limitations:
 - Recalls are not optimized
 - Swift background task (auditor, replicator) are inactive
- Reference: [Active Archive Implementation Guide with IBM Spectrum Scale Object and IBM Spectrum Archive](#)



Spectrum Archive object: technology outlook

- Forster tape awareness for Swift
 - Create a new Tape Data Ring with multiple Spectrum Archive instances arranged into zones and regions
 - Enrich Swift API with archive operations

- Features:
 - User/app control for migration, recall and status
 - Leverages tape optimized recalls
 - Supports replication across tape libraries
 - Scales out by adding more instances to the TDR
 - No interference with Swift background tasks

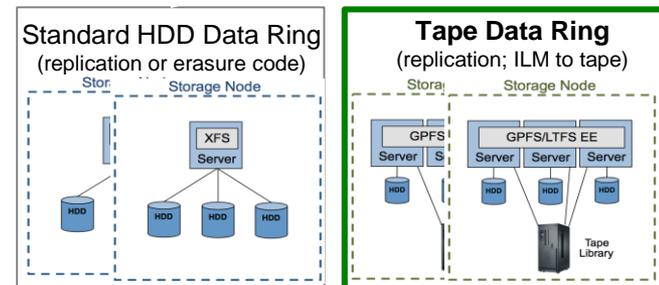
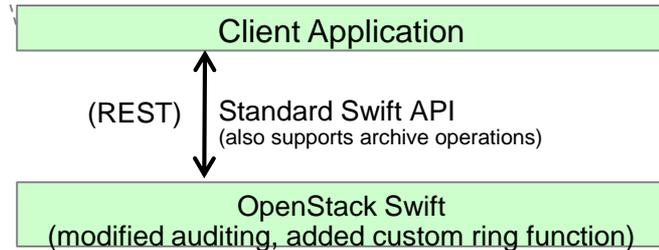
- “Swift associated project”: SwiftHLM middleware

http://docs.openstack.org/developer/swift/associated_projects.html#alternative-api

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						

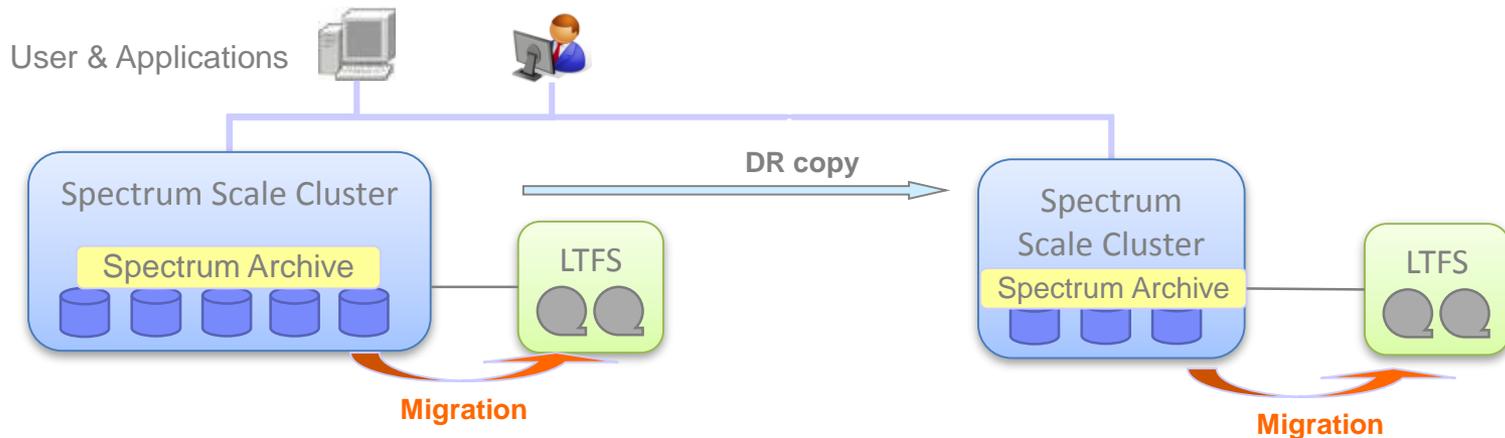
Delete container
 Migrate container
 Recall container



e.g. objects are migrated to tape after X minutes

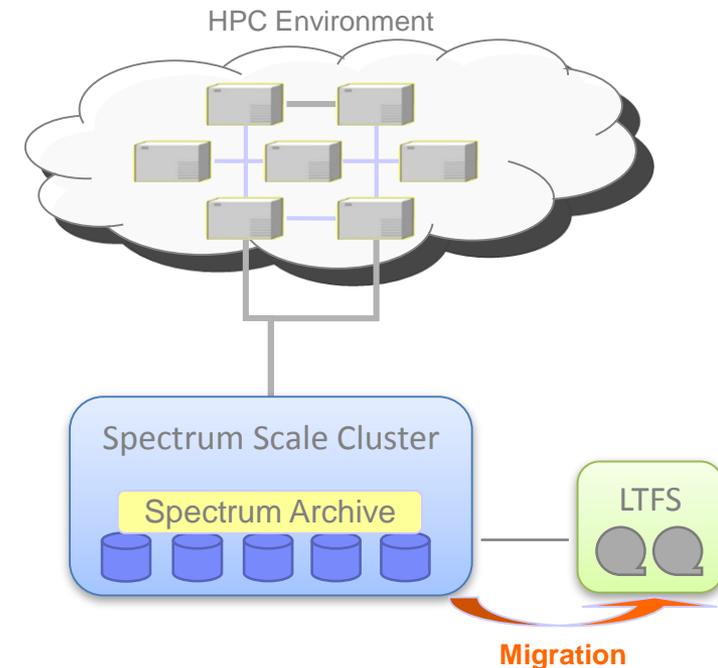
Client Example: Archiving seismic data

- IT service provider archives seismic data for long period of times
 - Cost-efficient archiving on tape for over 1 PB of data
 - Transparent access for user and applications via specialized portal
 - Disaster protection with two sites
 - LTFS tapes considered as medium for data exchange



Client Example: Archiving research data

- Research institute provides super computing and applications and services
- Must archive research datasets for long times
 - Automated migration based on policies
 - Two copies on tape for some data
 - Transparent access to data in global name space
 - Cost efficient archive on tape



Positioning Spectrum Archive with TSM HSM

- Spectrum Archive and TSM HSM are complementary products
 - TSM provides enterprise class HSM and backup functions for many environments
 - Spectrum Archive provides tape tier for Spectrum Scale on Linux

Use Spectrum Archive....

- Easy to use file storage on disk and tape
- For scaling environments (direct attached tapes)
- Data exchange on LTFS tapes
- For GPFS customers on Linux
- If TSM server is not an option

Use TSM HSM...

- For combined Backup & HSM
- Multiple tape libraries and drive sharing
- Multiple platforms
- Gain more functionality*
- When TSM (skills) is already available

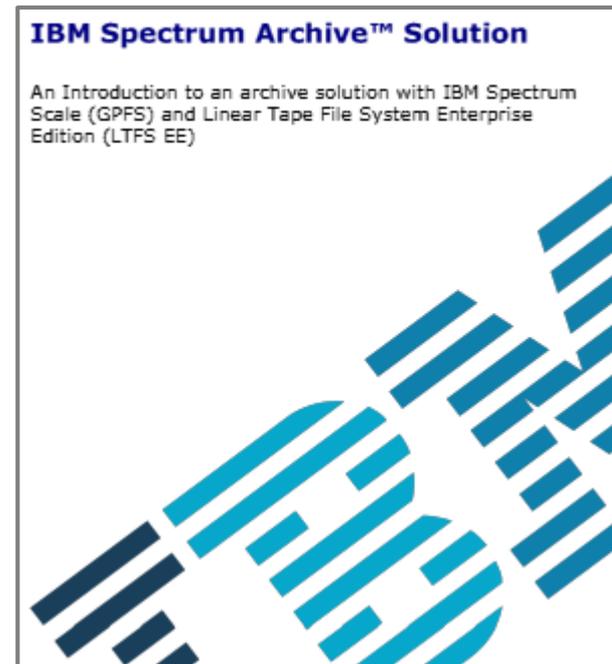
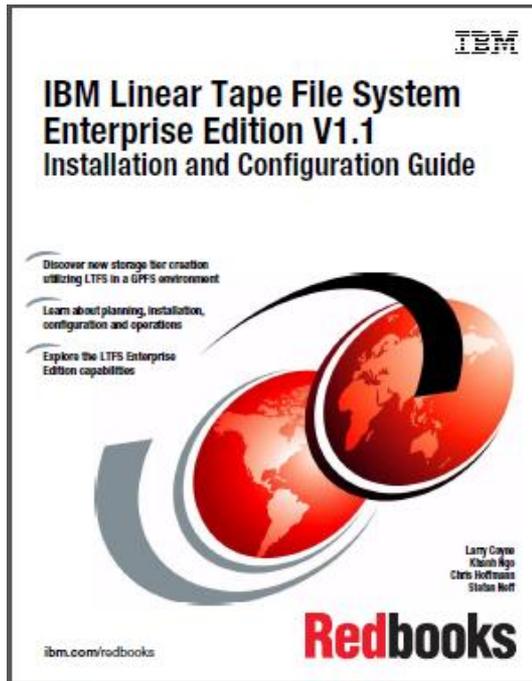
* From HSM client: streaming and partial recalls, Sobar. From TSM server: deduplication, node replication,

More information (IBM Internal): <https://w3-connections.ibm.com/files/app#/file/5a2ddc2f-a59a-4be3-bb37-663cb23c7428>

Thank You

Spectrum Archive key references

- Knowledge Center: <http://www.ibm.com/support/knowledgecenter/ST9MBR/welcome?lang=en>
- Redbook: www.redbooks.ibm.com/redpieces/abstracts/sg248333.html?Open
- Whitepaper: <http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102504>
- Sales Kit: <https://w3-03.sso.ibm.com/sales/support/ShowDoc.wss?docid=SSPG936296D29490F84&node=brands,B5000|brands,B8S00|clientset,IA>



Spectrum Archive Client references

- Research Institute in Germany

<http://w3-01.ibm.com/sales/ssi/cgi-bin/ssialias?appname=crmd&subtype=na&infotype=cr&htmlfid=0CRDD-9J6KCR>

- EVERY Seismic Data Repository:

<https://www.youtube.com/watch?v=YfmYGRWgevl>

- Kennedy Center for the Performing Arts:

<https://www.youtube.com/watch?v=jQ-mhE-319A>

Links and references

- **LTFS format specification ISO standard 20919:16**
http://www.iso.org/iso/home/store/catalogue_tc/catalogue_detail.htm?csnumber=69458
- **LTFS format specification SNIA standard**
http://www.snia.org/sites/default/files/LTFS_Format_2.2.0_Technical_Position.pdf
- **LTFS home page**
<http://www-03.ibm.com/systems/storage/tape/lvfs/index.html>
- **Spectrum Archive EE knowledge Center:**
<http://www.ibm.com/support/knowledgecenter/ST9MBR/welcome?lang=en>
- **Redbook: LTFS 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/atmsastr.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>
- **LTFS SDE Info center**
<http://publib.boulder.ibm.com/infocenter/lvfs/cust/index.jsp>
- **LTFS LE Infocenter:**
<http://pic.dhe.ibm.com/infocenter/lvflsle/cust/index.jsp>
- **LTFS Redbook: Installation and configuration**
<http://www.redbooks.ibm.com/abstracts/sg248090.html?Open>
- **Whitepaper: Using LTFS**
<http://public.dhe.ibm.com/common/ssi/ecm/en/tsl03109usen/TSL03109USEN.PDF>
- **LTFS software and IBM Device driver:**
<http://www-933.ibm.com/support/fixcentral/>
- **Almaden Research**
<http://www.almaden.ibm.com/storagesystems/projects/lvfs/>
- **The top Youtube videos:**
 - LTFS customer reference: https://www.youtube.com/watch?v=6s_hjzul9Y4
 - LTFS Intro: <http://www.youtube.com/watch?v=Qtd4CrajCYs>
 - Bamm TV: <http://www.youtube.com/watch?v=X82sC97yQeE>

Disclaimer

This information is for **IBM System Technical Events**, 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.