# IBM Spectrum Archive


Khanh Ngo
Senior Technical Staff Member and Master Inventor
IBM Spectrum Archive Development

IBM

# Store everywhere. Run anywhere.
## *Remove data-related bottlenecks*

- **Challenge**
  - Managing data growth
    - Lowering data costs
    - Managing data retrieval & app support
    - Protecting business data

- **Unified Scale-out Data Lake**
  - File In/Out, Object In/Out; Analytics on demand.
  - High-performance native protocols
  - Single Management Plane
  - Cluster replication & global namespace
  - Enterprise storage features across file, object & HDFS

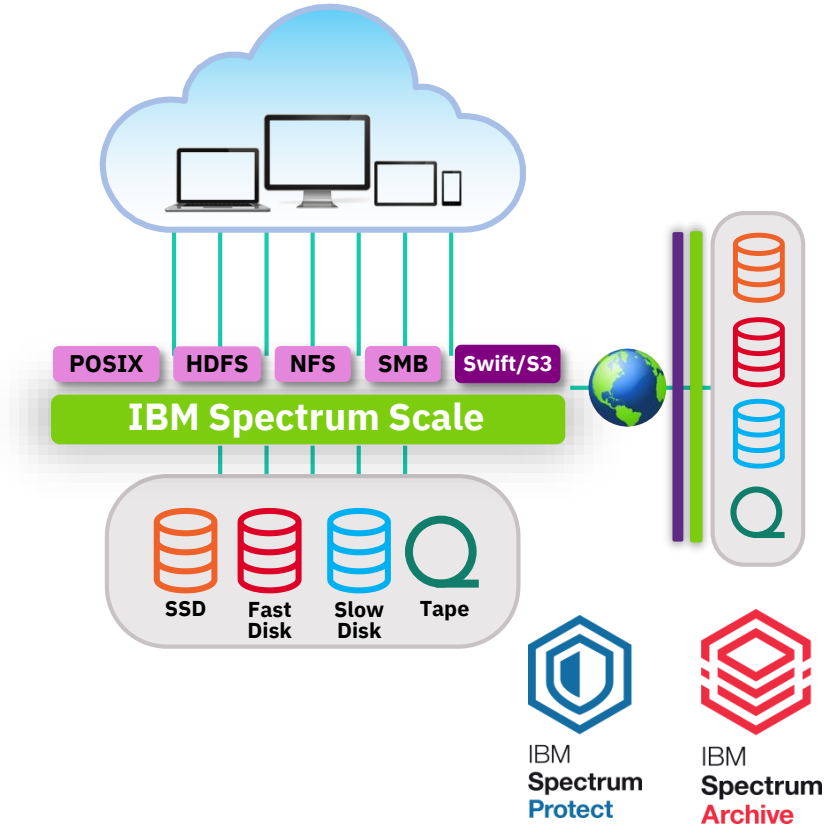


© Copyright IBM Corporation 2018

# Agenda

- What is Linear Tape File System (LTFS)?
- What is IBM Spectrum Archive?
- Introduction into IBM Spectrum Archive Enterprise Edition (EE) and its features
- Common Use Cases
- Specific customer implementations

© Copyright IBM Corporation 2018

# What is Linear Tape File System (LTFS)?

© Copyright IBM Corporation 2018

# LTFS is the Data Format Standard


many more...

- File System designed for Long-Term Retention and Media Portability
- Award-winning technology, invented and maintained by IBM
  - Reference implementation available as open source
  - Hosted at GitHub (https://github.com/LinearTapeFileSystem/ltfs)
- Open International Standard
  - ISO/IEC 20919:2016
  - Data structure on tape
    - » Two Partitions – Index Partition and Data Partition
  - Industry Collaboration - SNIA Technical Working Group
    - – Version 2.4 approved in 2017
    - – Now discussing Version 2.5
  - Logo Program (LTFS Compatibility Testing) by LTO Consortium

© Copyright IBM Corporation 2018

# Why the Data Format Matters?

- 3 typical use of tape storage

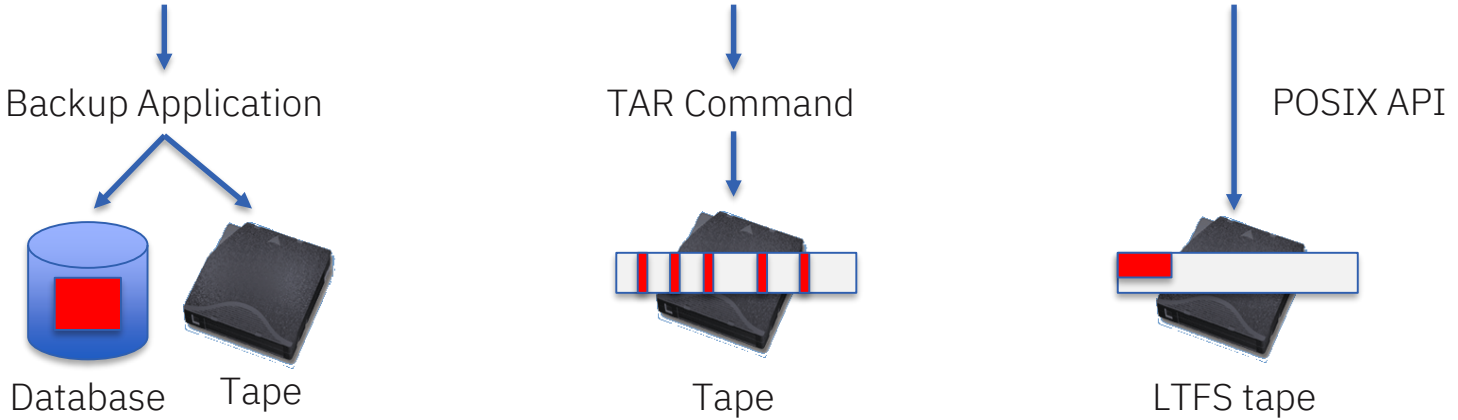


- What are requirements of Archival Storage?
  - Where/how the metadata (information of tape contents) are stored?
  - Is the tape portable across different locations or different applications?
  - Is the metadata centralized or scattered?
  - Can the files be accessible directly from end user application, or indirect?

© Copyright IBM Corporation 2018

# What is IBM Spectrum Archive?

© Copyright IBM Corporation 2018

# IBM Spectrum Archive: LTFS-based SDS software for data archive

- Member of IBM Spectrum Storage family
- Three Editions: Enterprise, Library, Single Drive
- Available as the standalone software or a part of IBM Spectrum Storage Suite (EE) only

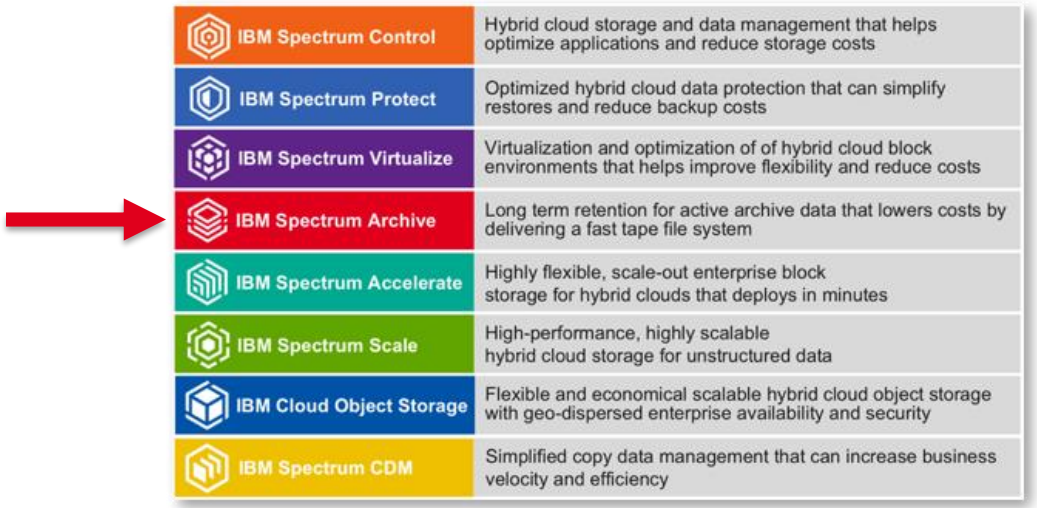


© Copyright IBM Corporation 2018

# IBM Spectrum Archive Editions

LTFS Storage Manager

IBM Archive & Essence Manager (GBS asset)

SwiftHLM

Research

LTFS DM

Licensed Software

## Enterprise Edition (EE) 1.3

- Integration of GPFS and LTFS
- Multi-node scale-out capability
- Policy based Data Management

Integrated Solution

Licensed / Free

## Library Edition (LE) 2.4

- Integrates the support of tape automation
- Scalable storage space by 1U TS2900 to TS4500 or Dell

Open Source SDE On GitHub

## Single Drive Edition (SDE) 2.4

- Free download
- Support IBM LTO and Enterprise Tape TS11xx
- Supports Linux, Mac, and Windows
- Bundled with OEM tape drive

Free

Tape Appliance/Device from IHV

Storage Toolkit

© Copyright IBM Corporation 2018

# Introduction into IBM Spectrum Archive Enterprise Edition (EE) and its features

© Copyright IBM Corporation 2018

# IBM Spectrum Archive Enterprise Edition (EE)

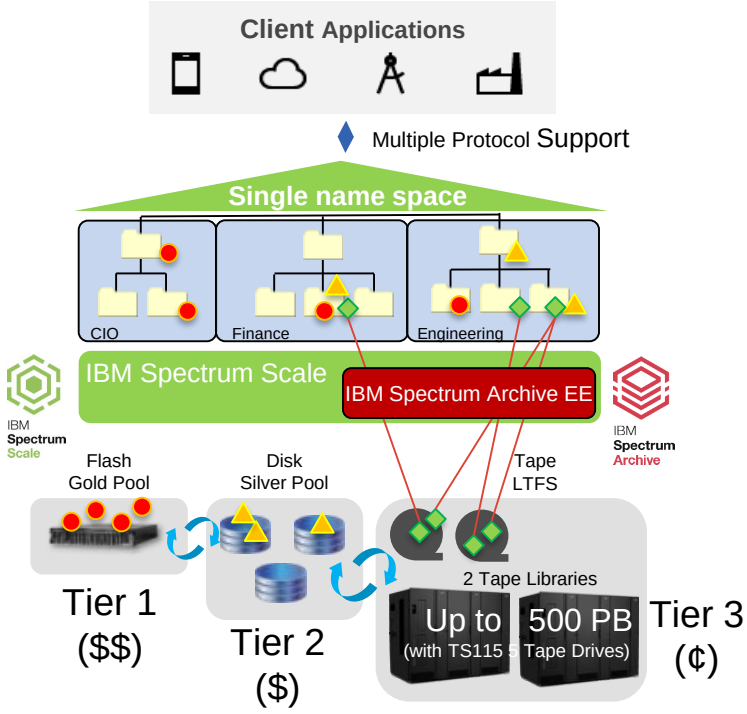


- Persistent view of the data - tape storage under the single namespace
  - Policy-based data placement for cold/idle data
  - Recall data from tape on demand
- Integrated Tape Tier
  - Up to 3 data replicas
  - Data Encryption with IBM SKLM server (LME)
  - WORM tape for anti-tampering
  - Offline tapes to store the media in an isolated environment – “air gap” for greater protection of sensitive corporate data, or extend the storage capacity beyond the library limit
  - Automated Tape Validation available with TS4500
- Export the LTFS tapes for data exchange
  - Remove data from Scale namespace, and export tapes for the use in other application

Linux:
Orderable from AAS or PPA
Trial Version available from IBM Web site

# ESS with IBM Spectrum Archive

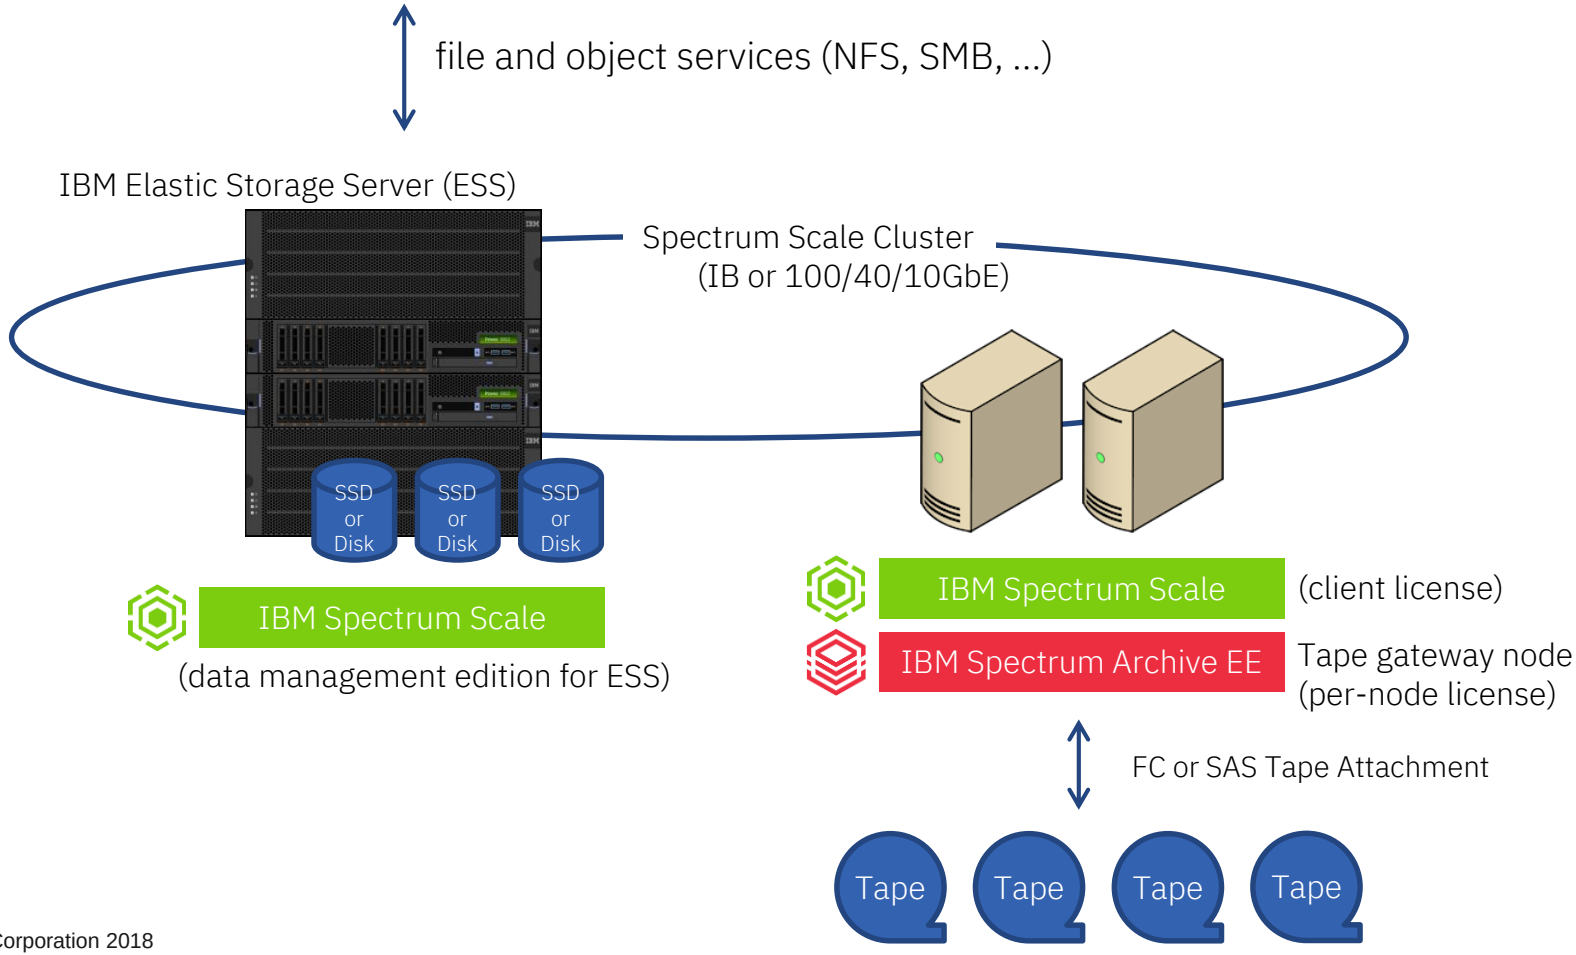


© Copyright IBM Corporation 2018

# EE: Building Block Options for Scale Out

1. Tape Gateway Servers
   - CPU - x86 or POWER Little Endian
   - Per-node license
2. Disk Storage
3. Tape Drive and Tape Media
4. Tape Library

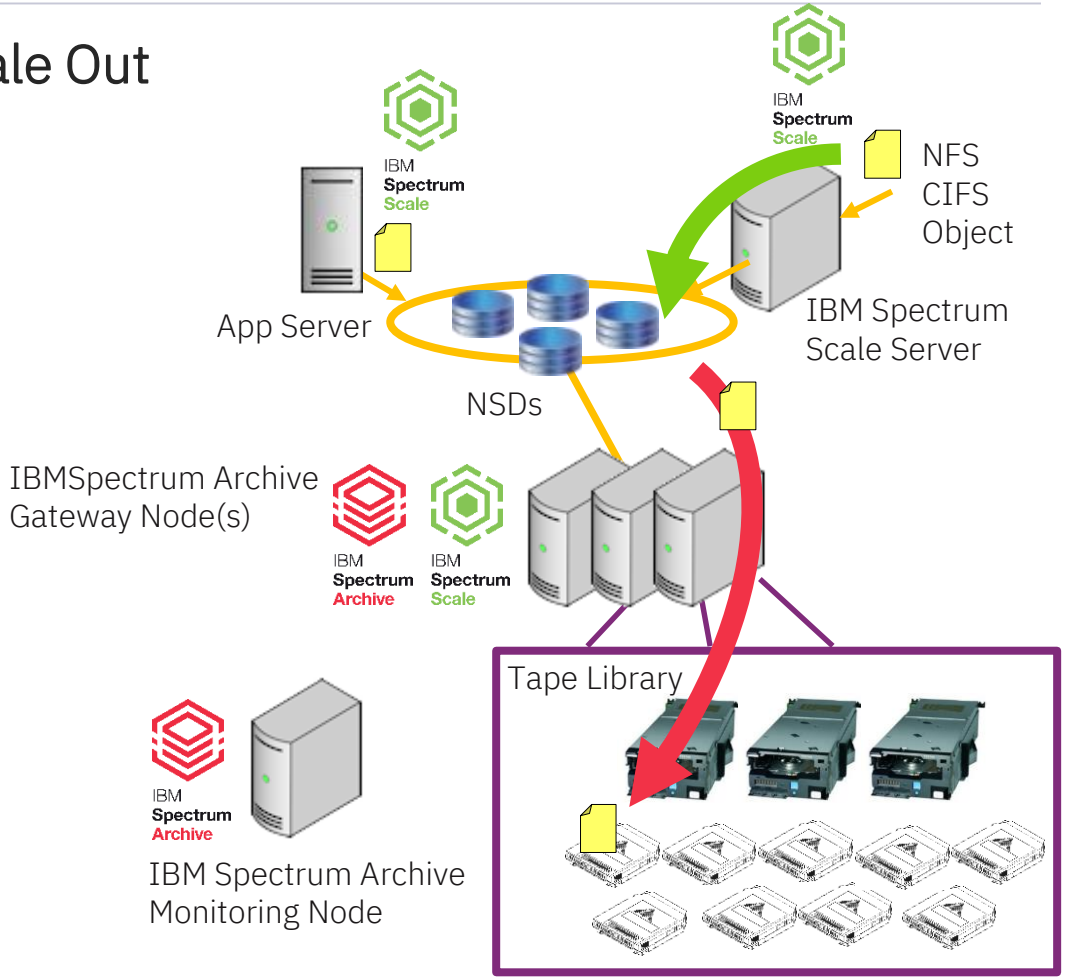


© Copyright IBM Corporation 2018

# Functional Overview

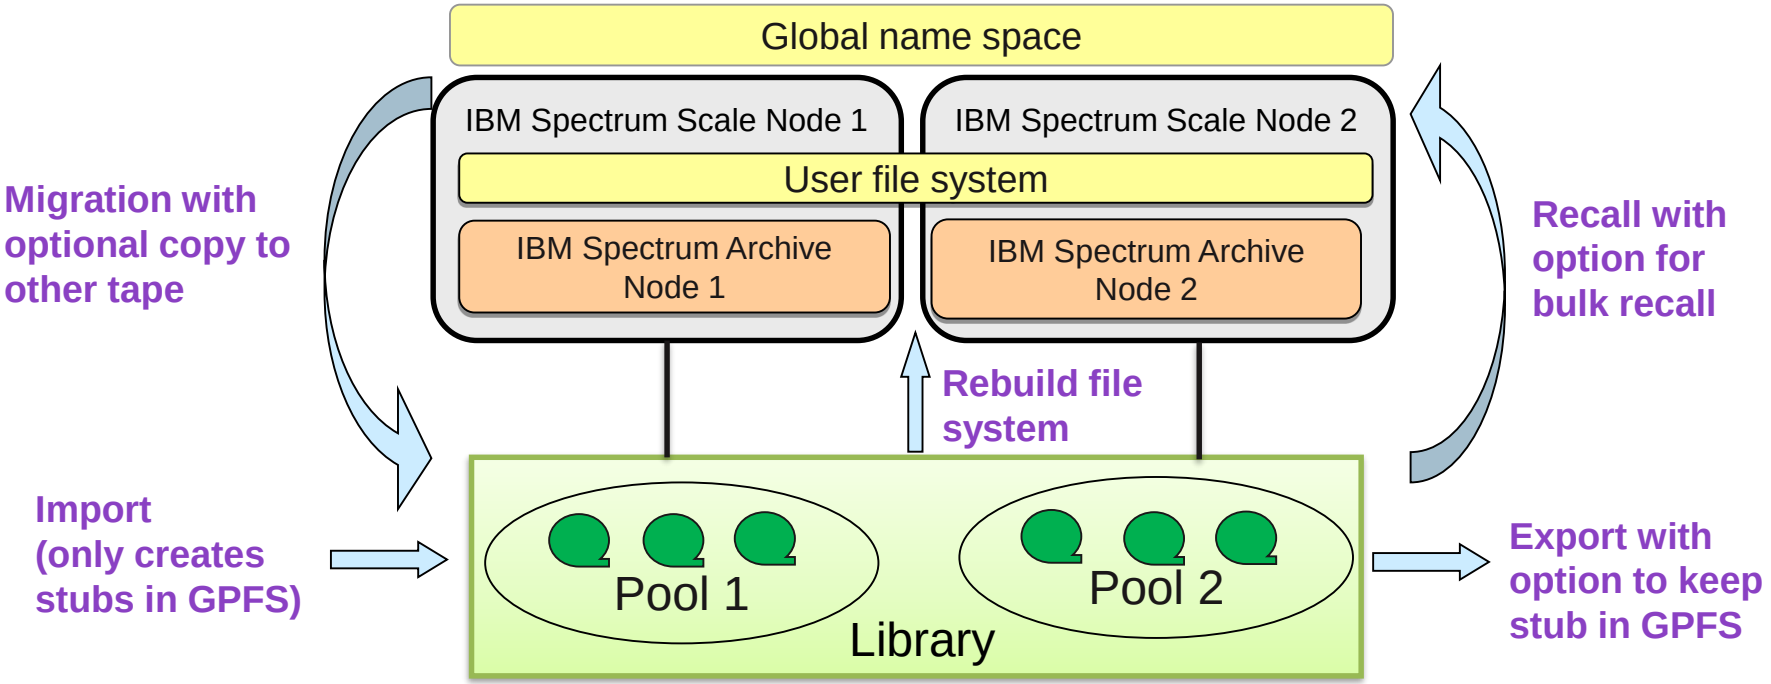


© Copyright IBM Corporation 2018

# IBM Spectrum Archive Update Sequence

## v1.2

- Multiple tape library attachment (up to 2) support to a single IBM Spectrum Scale cluster
- Data recording on WORM tape cartridges - TS1100 only
- Expand storage capacity with LTO7 support
- Performance improvement for large-scale systems
- Flexibility in pool-based data management including transparent recall retries

## v1.2.2

- New –E option to removing tapes with no file references
- Improved performance of administrative commands for reconcile, import/export
- Automated the recover process of write failures tapes
- Improved method for recovering read failure tapes

## v1.2.4

- Support of IBM Spectrum Scale Active File Management (AFM) Independent Writer (IW) mode
- RESTful API
- Control node failover
- Monitoring dashboard
- TS1155 support
- IBM SwiftHLM support

## v1.2.5/v1.2.5.1

- LTO8/M8 Support
- Library Replacement Procedure phase one (conversion method)

## 1.2.6 Updates

- Library replacement procedure phase two (translation method)
- Assisted tape technology upgrade for in-pool data migration and pool-to-pool data migration
- POWER Little Endian with Linux (RHEL) version 7.4, or later
- New datamigrate command for technology upgrade

© Copyright IBM Corporation 2018

# December 2018 release

- Enterprise Edition (EE) 1.3.0.0
  - User Task Control and Reporting: Usability enhancements with new command-line interface (CLI) with additional support for monitoring the progress and results of user operations, and for tape maintenance
    - Active/Completed task listing including detailed information and output of command
    - Task results including file state transition results
    - Ability to run the command in background, with –async option
  - Supports the Storage Networking Industry Association's LTFS format specification 2.4.
  - Expanded storage capacity with the TS1160 tape drive.
  - Supports the IBM Spectrum Scale backup function (mmbackup) for the same file system managed by IBM Spectrum Archive.
  - Bundles the open source package for the external monitoring of Spectrum Archive through a GUI/dashboard
  - Use of /dev/sgX device

© Copyright IBM Corporation 2018

# REST API

- 7 **GET** endpoints returning json-formatted output
  - http://localhost:7100/ibmsa/v1/pools
  - http://localhost:7100/ibmsa/v1/tapes
  - http://localhost:7100/ibmsa/v1/drives
  - http://localhost:7100/ibmsa/v1/nodes
  - http://localhost:7100/ibmsa/v1/libraries
  - http://localhost:7100/ibmsa/v1/nodegroups
  - http://localhost:7100/ibmsa/v1/tasks
- Common GET parameters
  - Pretty
  - fields
  - sort


```
[{
  "id": "d9dcb712-2cc3-4a10-b6ac-bb54c520cb5d",
  "model": "03584L22",
  "name": "TS4500",
  "serial": "0000078AA0040405"
}
```


© Copyright IBM Corporation 2018

# Dashboard/GUI

- IBM Spectrum Archive supports a dashboard to monitor system performance, statistics, and configuration, based on
  - Logstash, to collect data
  - Elasticsearch, to store the data
  - Grafana, to visualize data

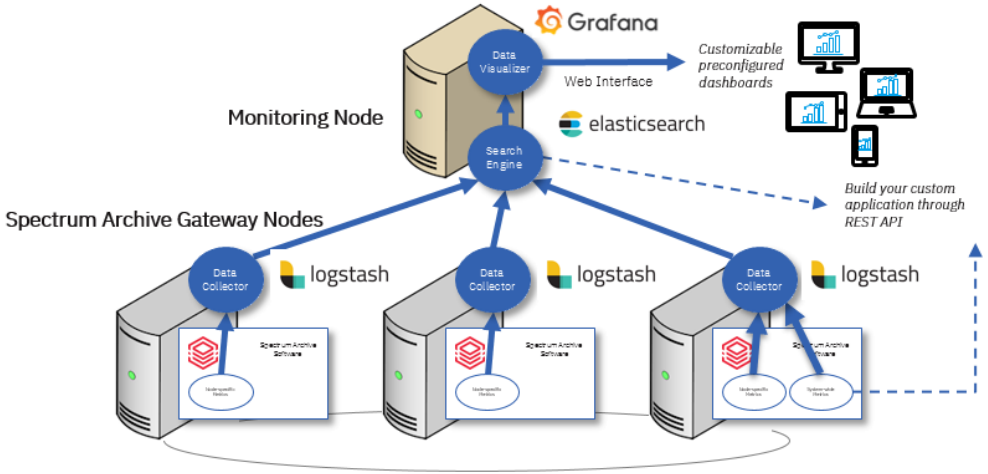


https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/General%20Parallel%20File%20System%20(GPFS)/page/Monitoring%20the%20statistics%20of%20Spectrum%20Archive

© Copyright IBM Corporation 2018

# IBM Spectrum Archive EE Dashboard

- System Health View
- Storage View
- Resource View
- Performance View
- Task View

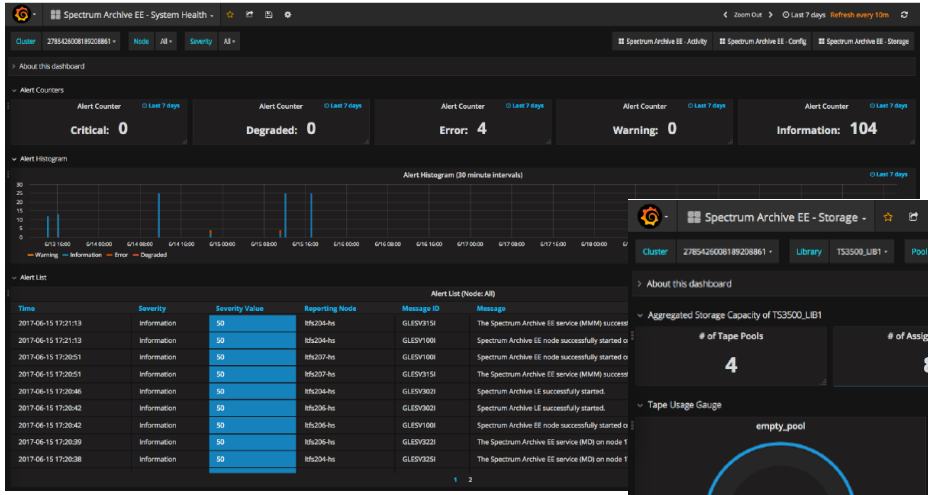


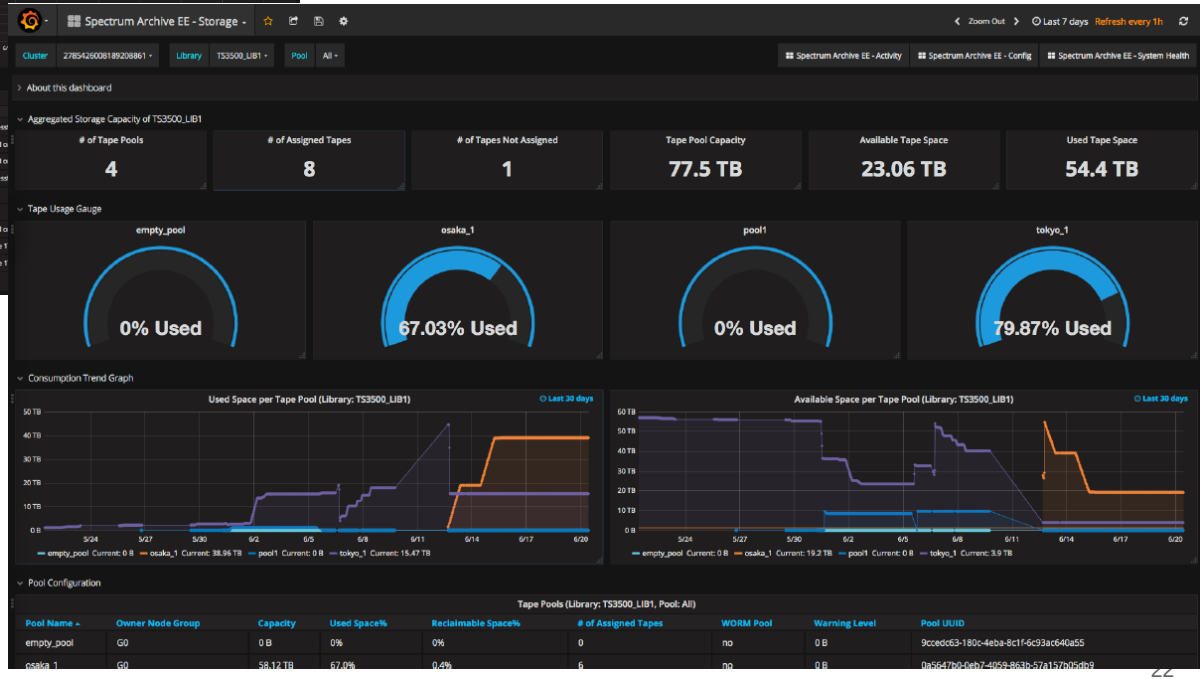


© Copyright IBM Corporation 2018

# Common Use Cases

© Copyright IBM Corporation 2018

# Active Archive

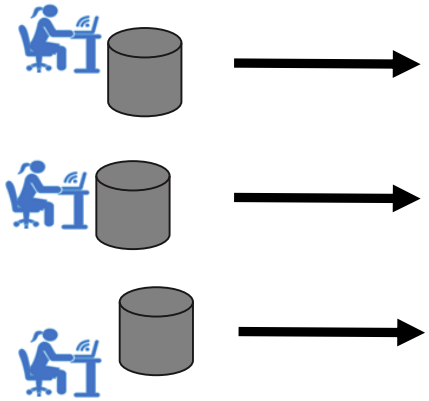

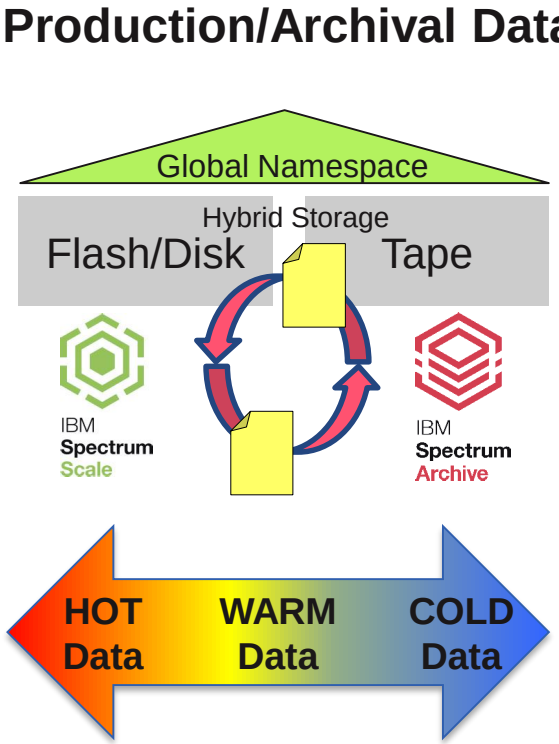


- Resiliency
  - Up to 3 copies on tape
  - Up to 2 libraries
- Data Collocation
  - Global Pool or Separated Pools per Application and/or User
- Recall Options
  - Transparent Recall
  - Bulk Recall/Prefetch by command

© Copyright IBM Corporation 2018

# IBM Spectrum Archive EE managing continuous data growth

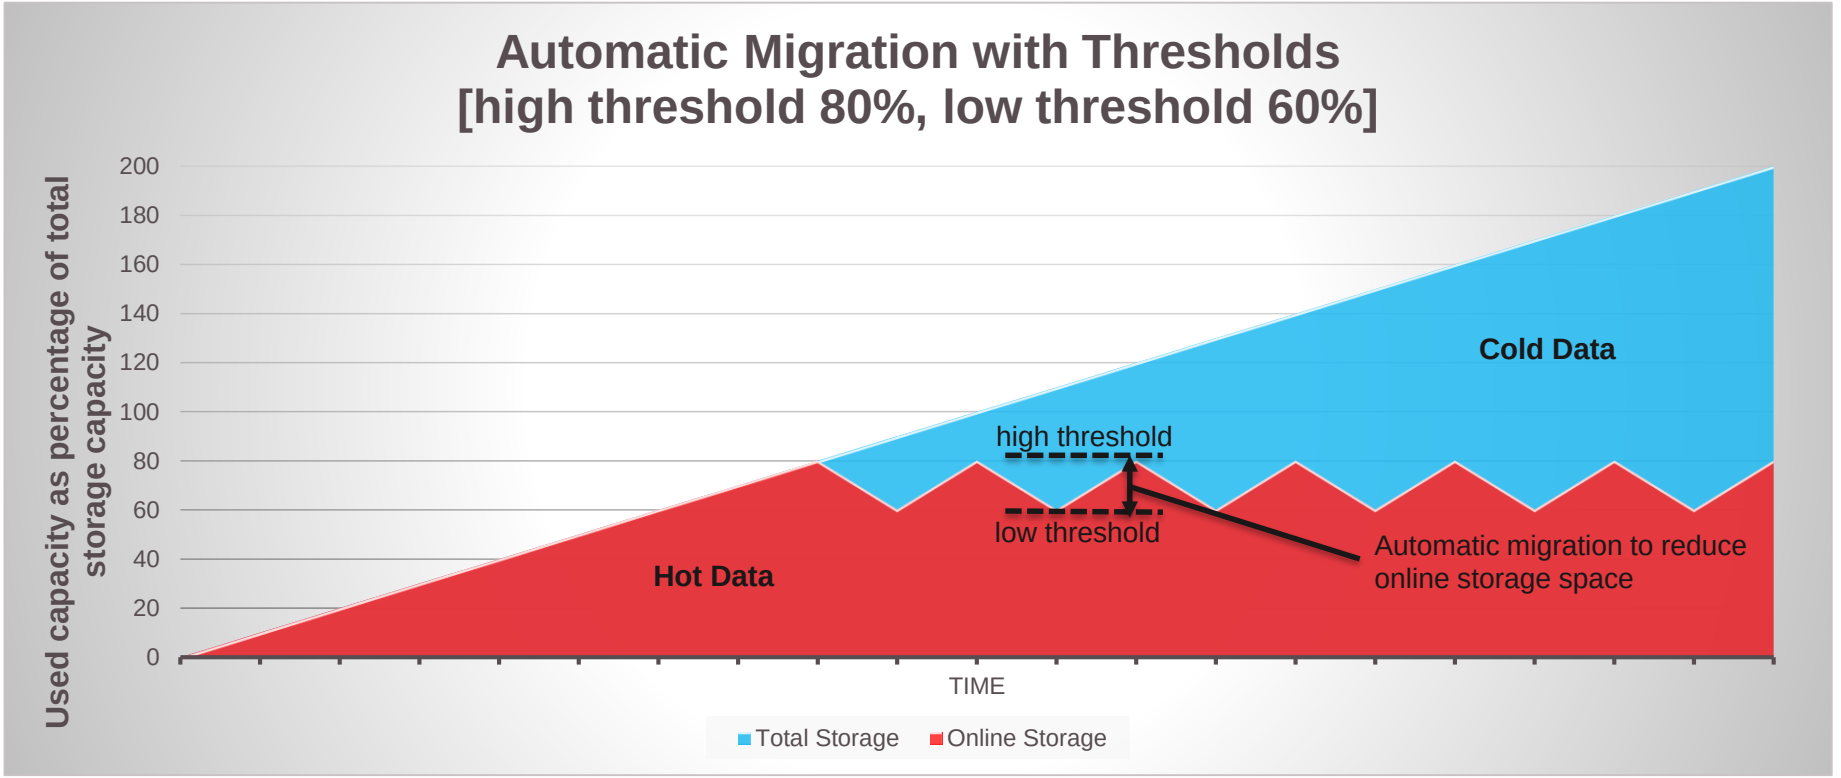


© Copyright IBM Corporation 2018

# Operational Storage

## Production Data

NFS, SMB, …

- IBM Spectrum Scale or external Filer

## Disk Cache

IBM Spectrum Scale

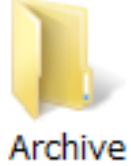


- Landing Zone for archive/retrieval
- Folder per Application/User
- With or without quota
- May use Immutability Flag

## Archival Data

IBM Spectrum Archive

Migrate all to tape

- WORM or Non-WORM Tapes
- Up to 3 copies
- Up to 2 libraries
- Global pool or Pool by App/User

© Copyright IBM Corporation 2018

# Specific customer implementations

© Copyright IBM Corporation 2018

# Media and Entertainment

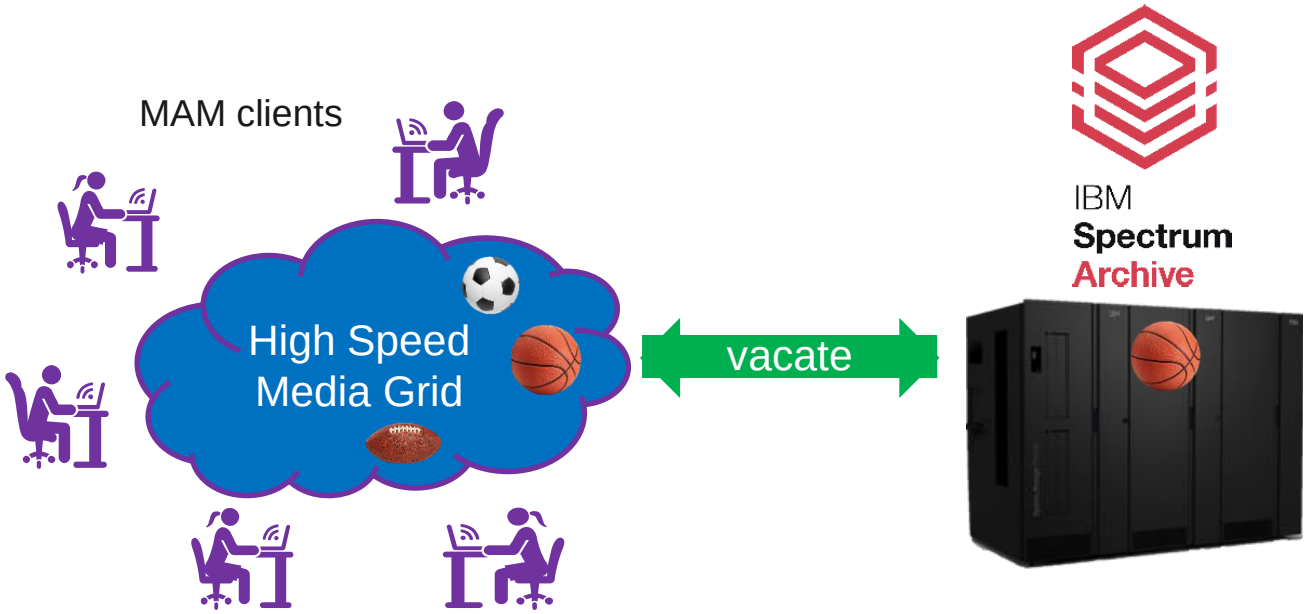


© Copyright IBM Corporation 2018

# Media and Entertainment

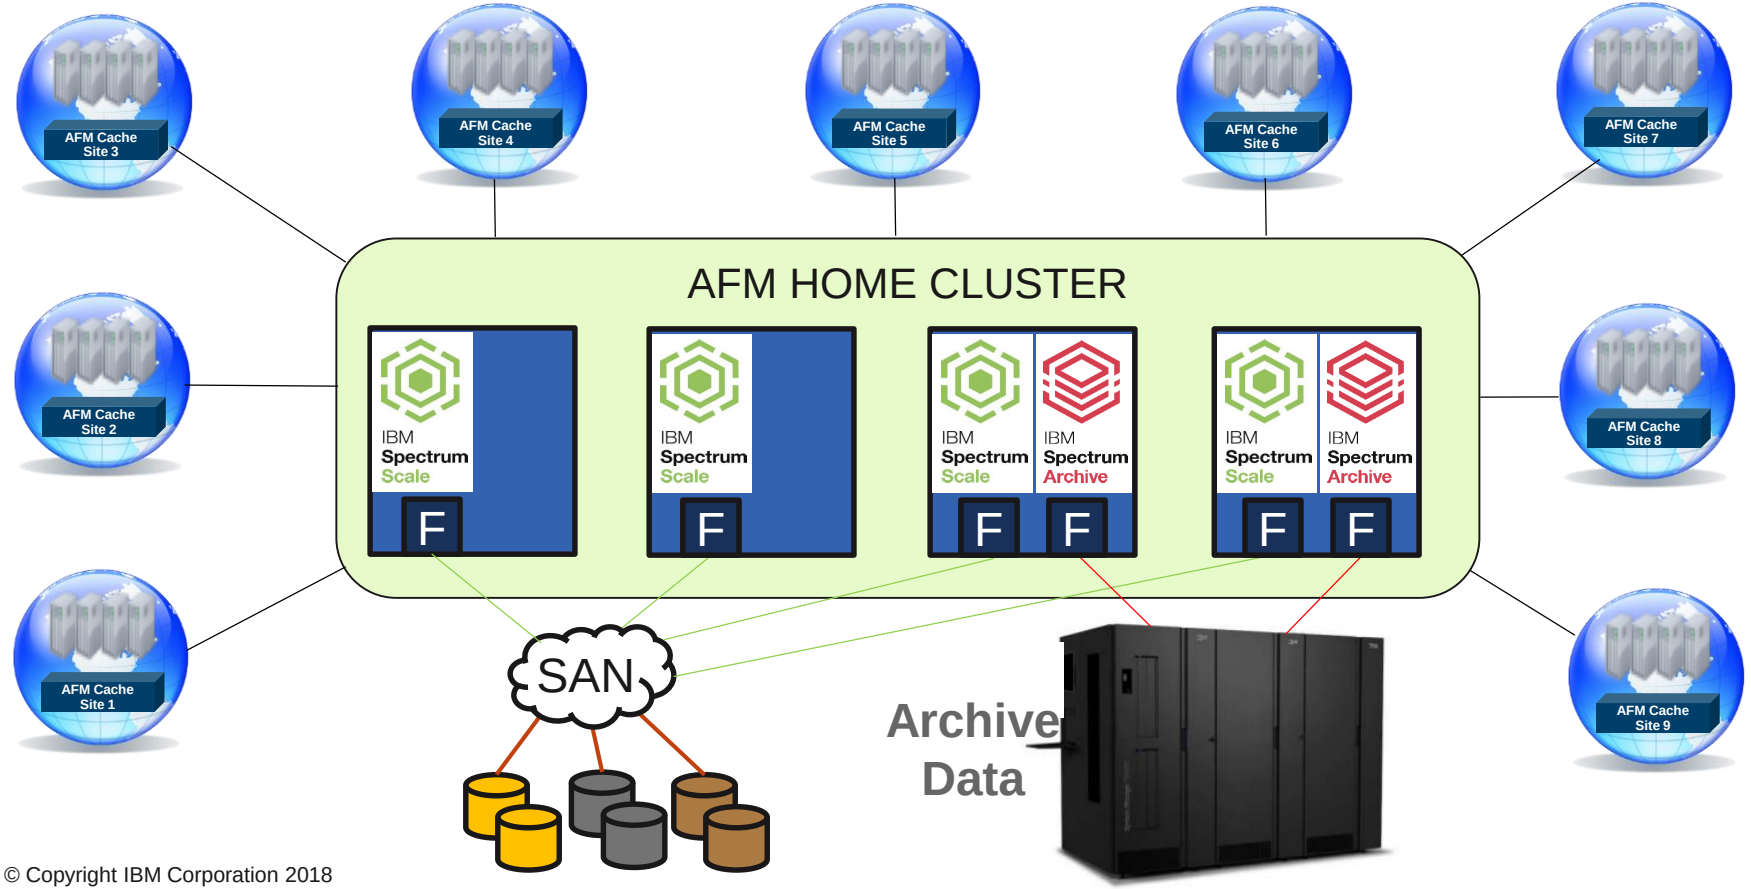


© Copyright IBM Corporation 2018

# Archive of Genomics Data

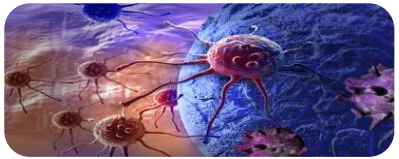

biopsy

sequence

analyze

Migration

Recall

Researchers

FastQ
~150GB
(compressed)

VCF
100 - 200 MB

BAM
~150GB

SET A

SET B

IBM Spectrum Archive

© Copyright IBM Corporation 2018

# University Scientific Data Archive

Organizational Units accessing the Scientific Data Archive

- Advanced Research
- University Libraries
- University Institute 1
- University Institute 2
- University Institute 3

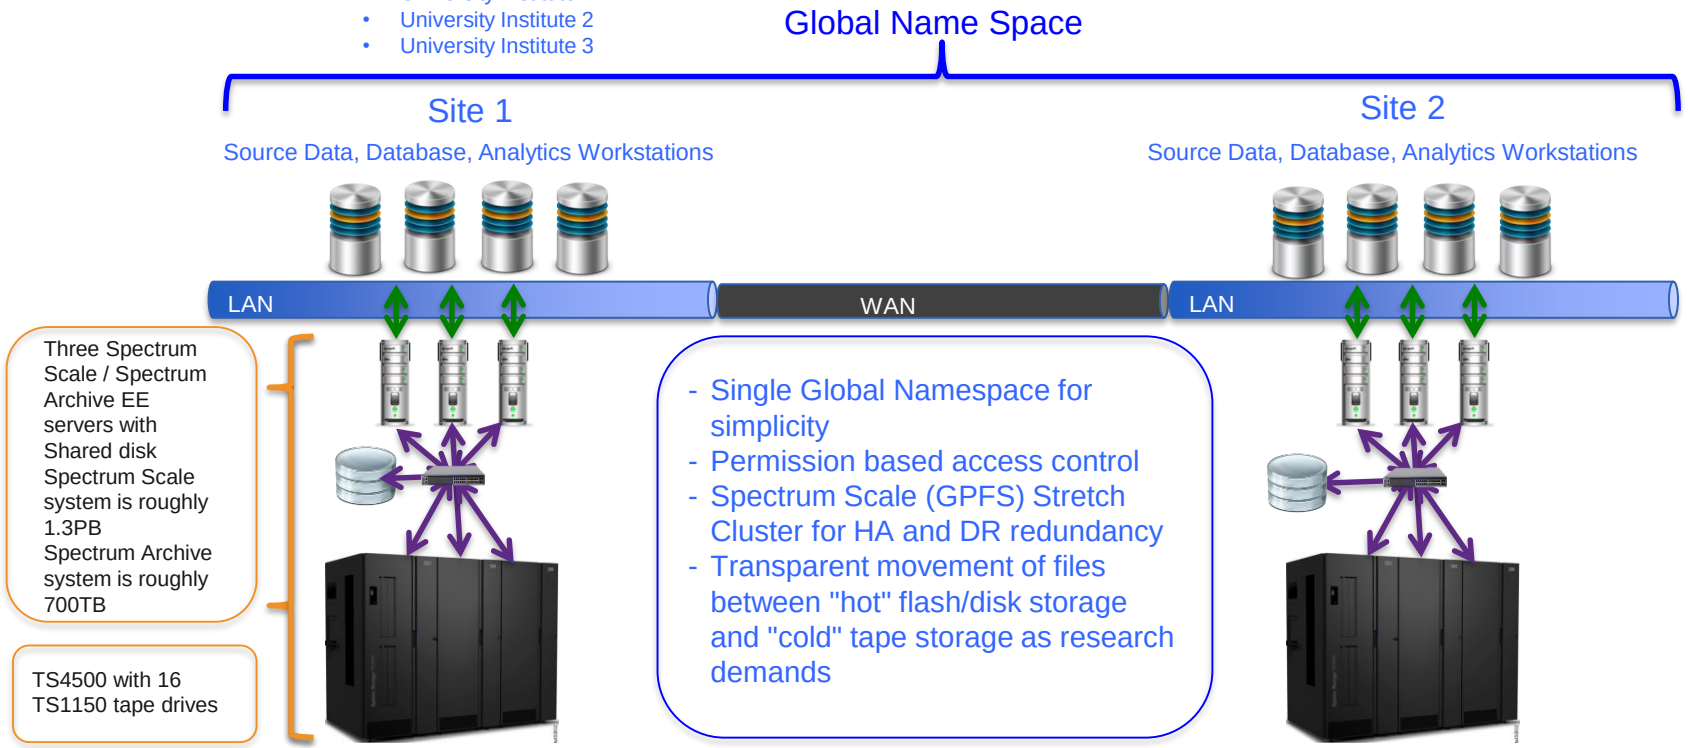


© Copyright IBM Corporation 2018

# Archive of Research/Scientific/Financial Data for Long Periods of Time

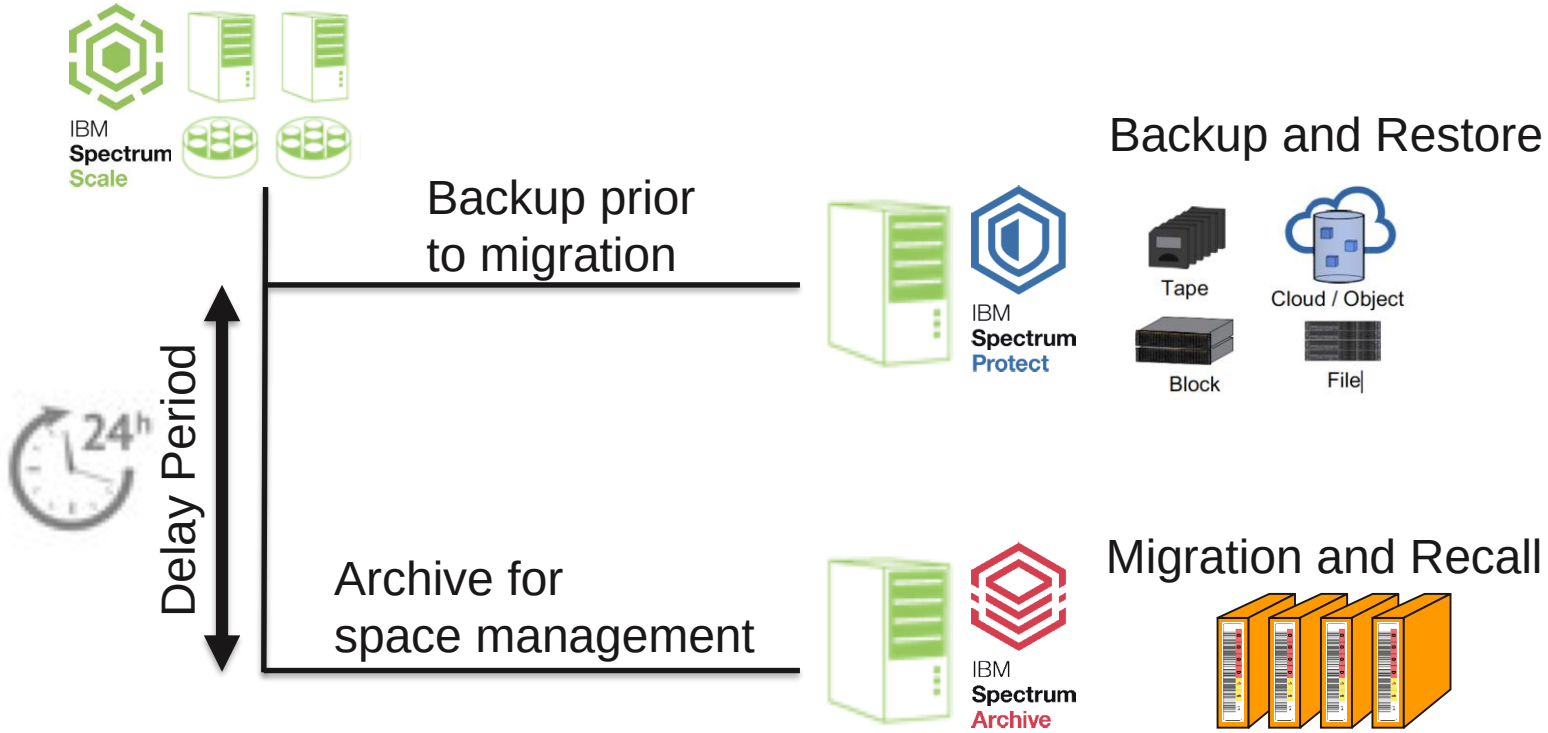


© Copyright IBM Corporation 2018

# Repository area for long-term archive of important files

- Shared storage area across all HPC systems for backup of home directories and user's data for long term archive
- Files may be initially backed up using IBM Spectrum Protect
- Then files are migrated to tape after certain conditions are met such as older than 3 months and/or larger than 50-100MB

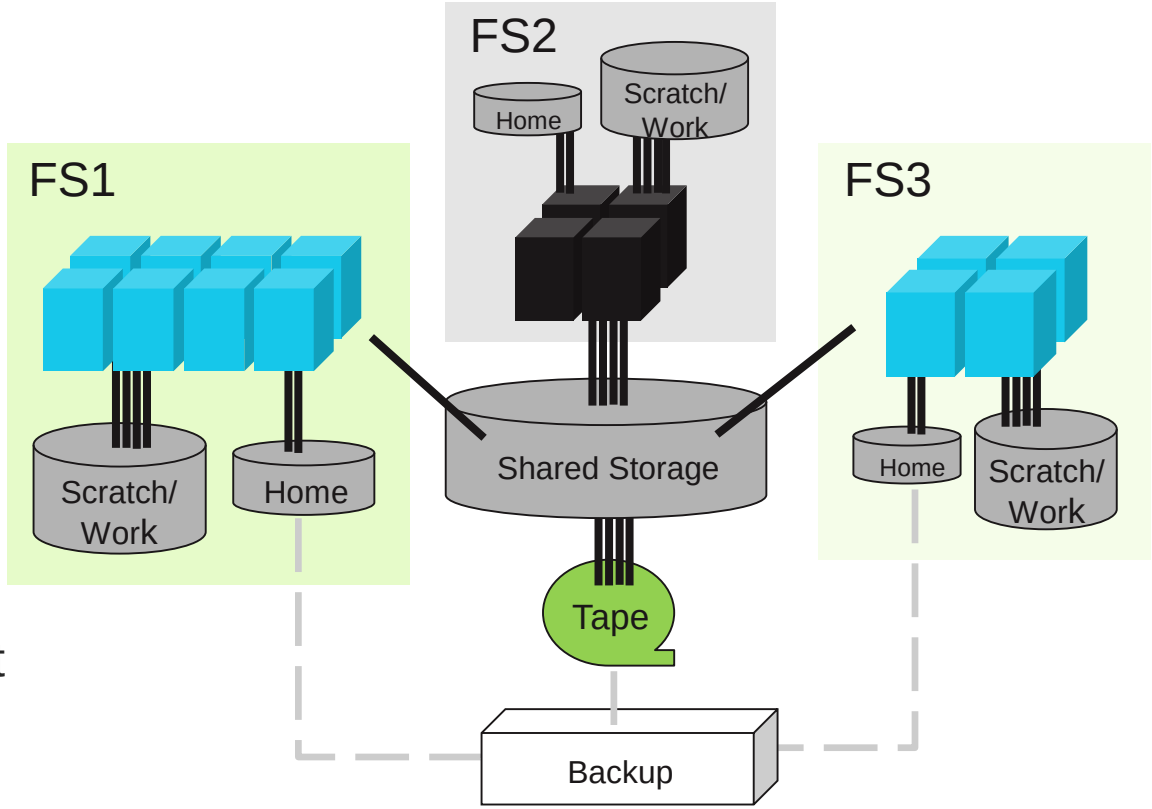


© Copyright IBM Corporation 2018

# Archive of Seismic Data

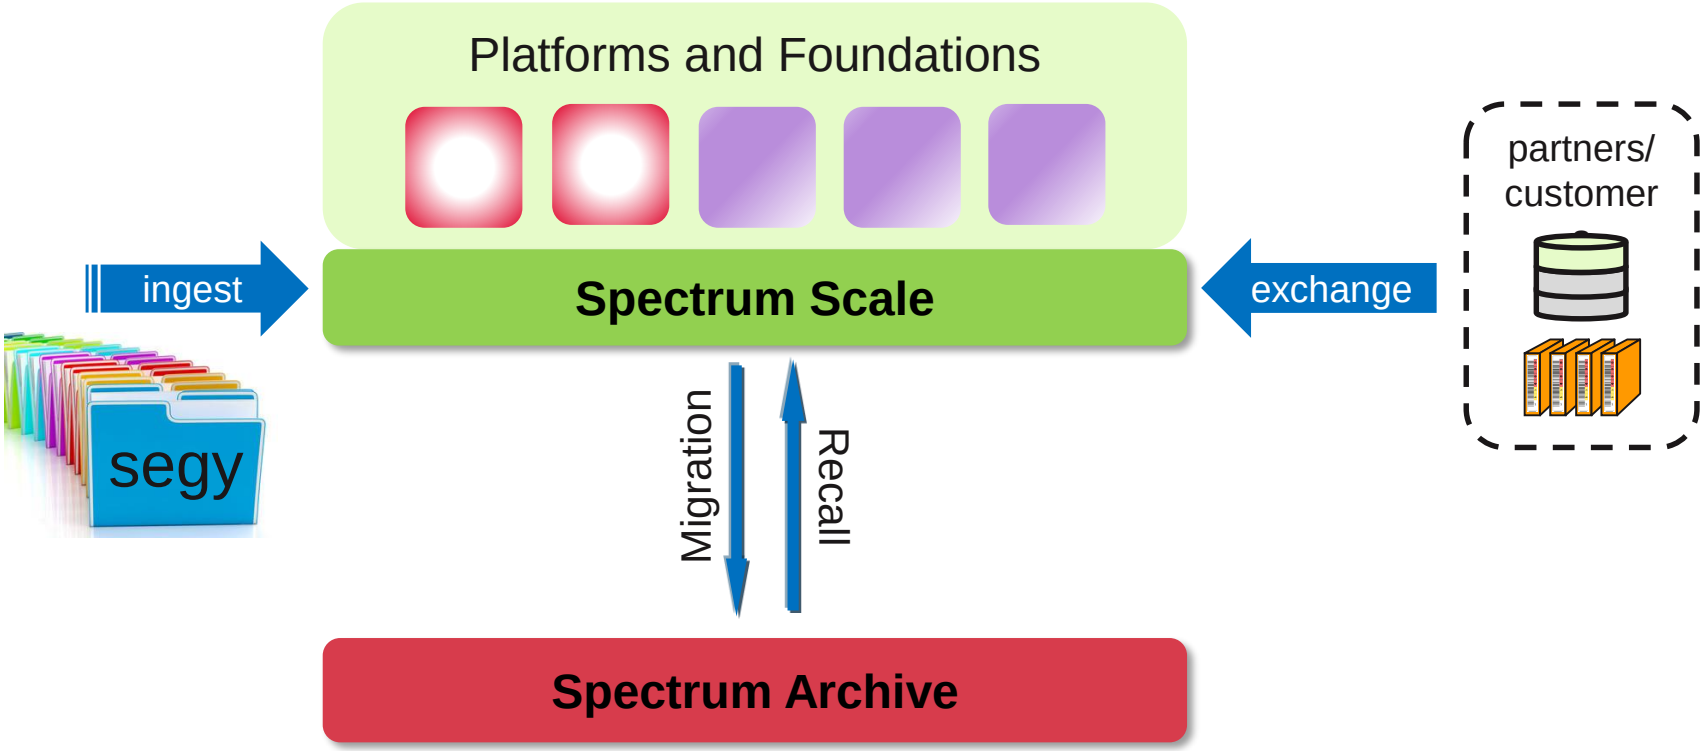


© Copyright IBM Corporation 2018

# IBM Spectrum Archive Features

- 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
- Two site replication by stretch cluster or AFM IW
- Media export/import for data sharing and/or offsite storage
- Media health check with TS4500
- Easy administration and management

IBM Spectrum Archive

© Copyright IBM Corporation 2018

# Convincing arguments for IBM Spectrum Archive EE

- Scalability (from an architectural perspective), can grow cartridges, tape drives, Archive nodes, etc
- Cost (no capacity license)
  - $60k list per node with no capacity limitations
- Standardized and open format

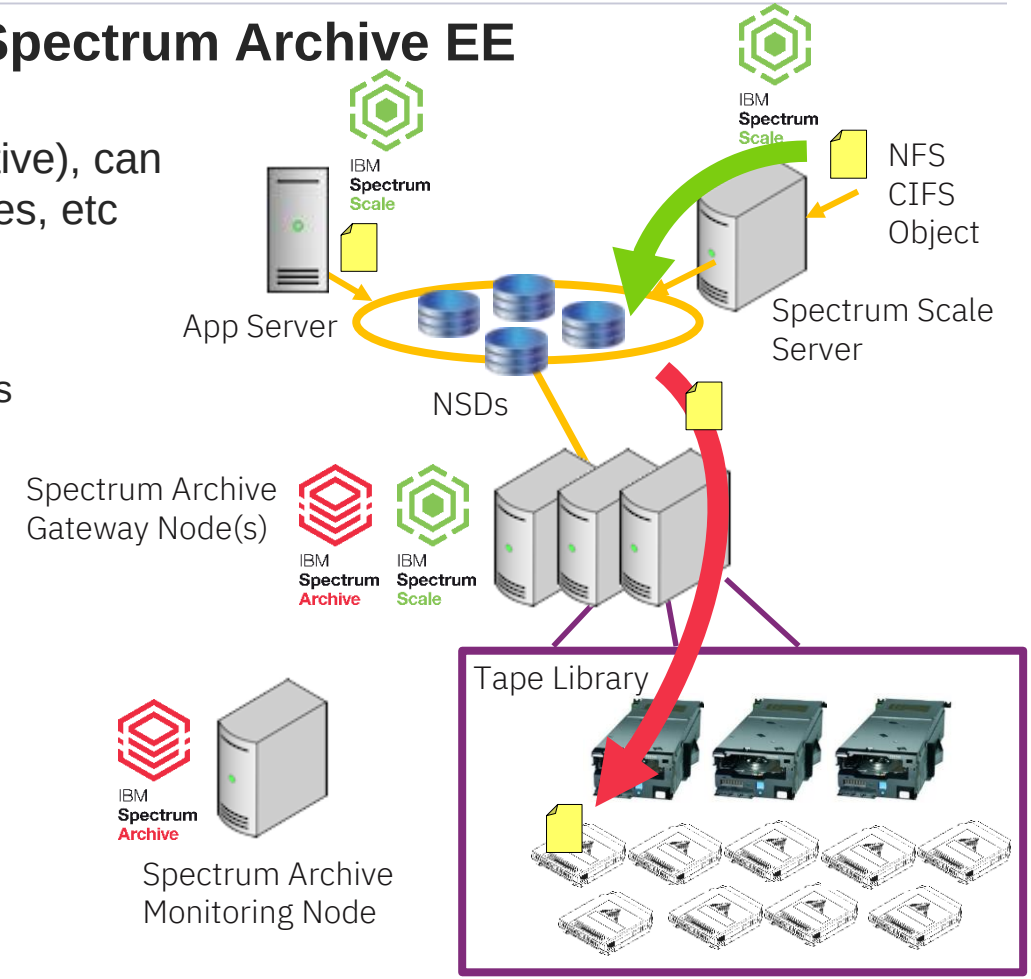


© Copyright IBM Corporation 2018


# Questions?

© Copyright IBM Corporation 2018