

# What's new in Spectrum Scale and the Elastic Storage System (ESS)?

ISC 2022 - May 30<sup>th</sup>, 2022

Chris Maestas, Chief Architect, Storage  
for Data and AI Solutions

[cdmaestas@us.ibm.com](mailto:cdmaestas@us.ibm.com)



# Disclaimer

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code, or functionality. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

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.

# Featured Updates

GPU Direct Storage (GDS) on **RoCE** environments

Continued enhancement to integration and deployment features with containers that now allows containerized storage deployment within Spectrum Fusion technology.

High Performance Object (HPO) as part of next generation protocol services stack.

Enhanced reliability, monitoring and security with Active File Management with expanded support for **Azure**.



## Survey – tell me about upgrades?

### Why?


- To address requests for quarterly updates to bring new features out more rapidly

### Maintain Extended Update Support concept

EUS with PTFs every 18 months

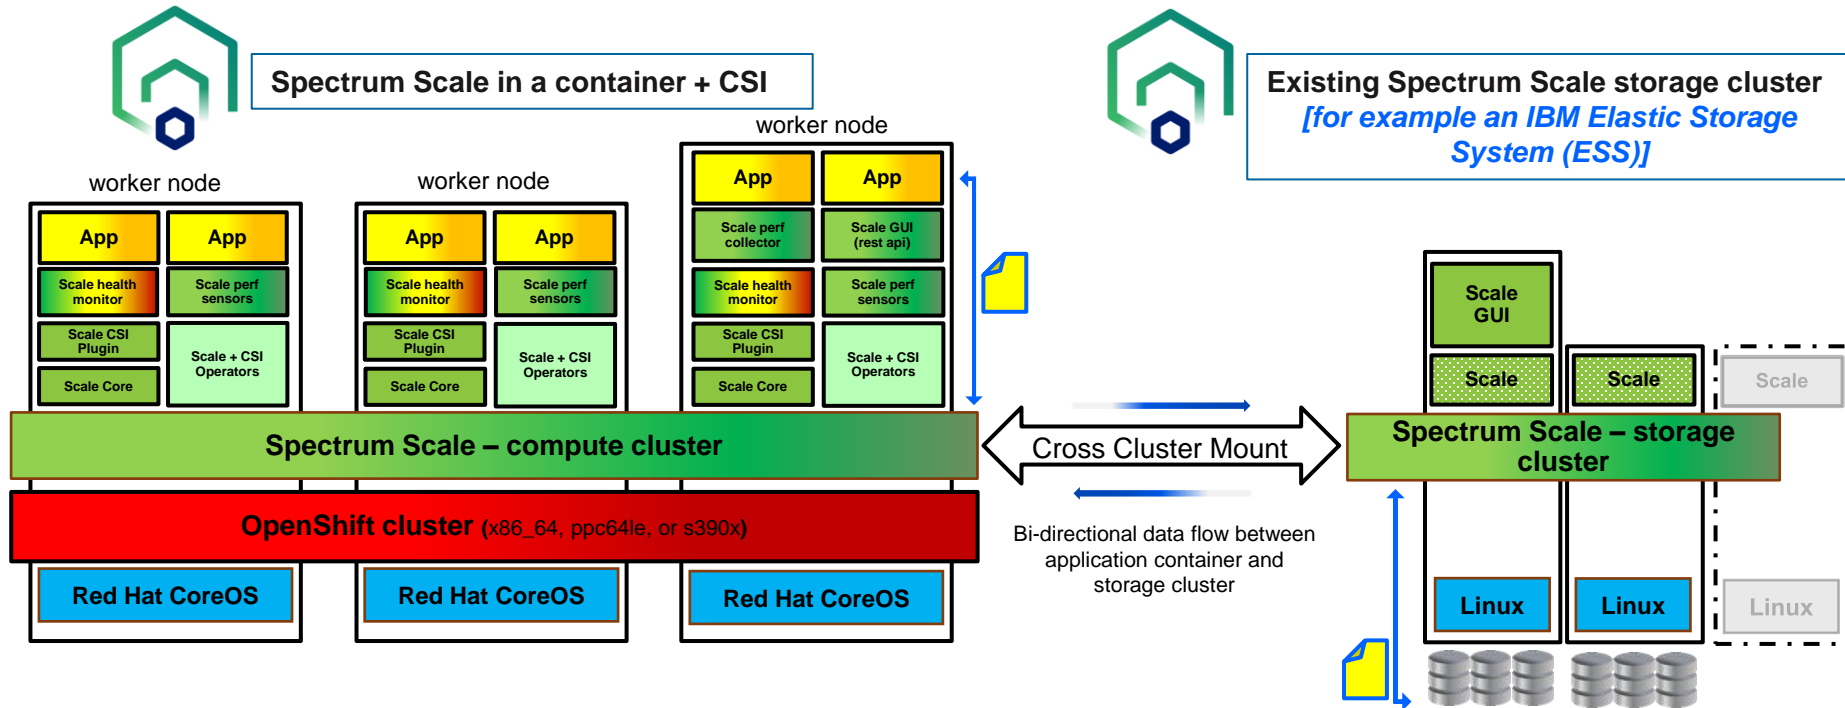
Extended support on last EUS within a release ( example: V.R.x, 4.2.3, 5.1.2, 5.1.last)

Increase the number of Modification levels with new function



# IBM Spectrum Scale Container Native Storage Access (CNSA)

## Cluster Overview



# Container Native Storage Access

Improvements introduced in CNSA 5.1.3

<https://www.ibm.com/docs/en/scalecontainernative?topic=overview-supported-features>

***Wider support to use the latest CNSA functionality.***

- Support for upgrading IBM Spectrum Scale Container Native Storage Access (CNSA) from v5.1.2.1 to 5.1.3
- Support for RedHat OpenShift Container Platform 4.10
- CNSA images now hosted on the entitled IBM Cloud Container Registry.
- Automated deployment of the CSI driver
- **Support for storage cluster encryption**
- **Rolling upgrade of IBM Spectrum Scale is supported**
- Support for a limited set of IBM Spectrum Scale configuration settings to be set directly
- Grafana support
- Support for X86, Power and Z.
- Direct storage attachment on x86 and power servers.
- Automatic quorum selection is Kubernetes topology aware.





# AFM, Compression, ILM for Containerized environments

- AFM for Caching or DR
  - Setup AFM gateway nodes **in the external storage cluster** for Caching or DR use case
  - No configuration required on the Scale compute cluster
- Compression
  - Run compression policies **in the external storage cluster**
  - Scale compute cluster will read compressed data by uncompressing it on the clients
- Storage Tiering
  - Define Storage pools, apply placement, migration policies **in external cluster**
  - No way to define placement policies from Scale compute cluster

# Container Storage Interface

## Improvements introduced in CSI 2.5

***Upgrades for OpenShift, Kubernetes and Ansible as well as improved functionality that support simpler administration and configuration.***

- Support for Red Hat [OpenShift 4.10](#) and [Kubernetes 1.23](#).
- Upgraded CSI specification from 1.3.0 to 1.5.0
- Added support for Consistency Group (**version=2**)
- Support to enable the compression for persistent volumes
- Support to enable the tiering for persistent volumes
- Increased attacher statefulset's replica count to two for high availability of attached volumes
- Upgraded Kubernetes CSI sidecar containers
- Migrated from CSI Ansible® operator to CSI Go operator





# Data Access Services – S3 object access

*Containerized S3 object access integrated within Spectrum Scale delivering high performance object for AI and analytics workloads*

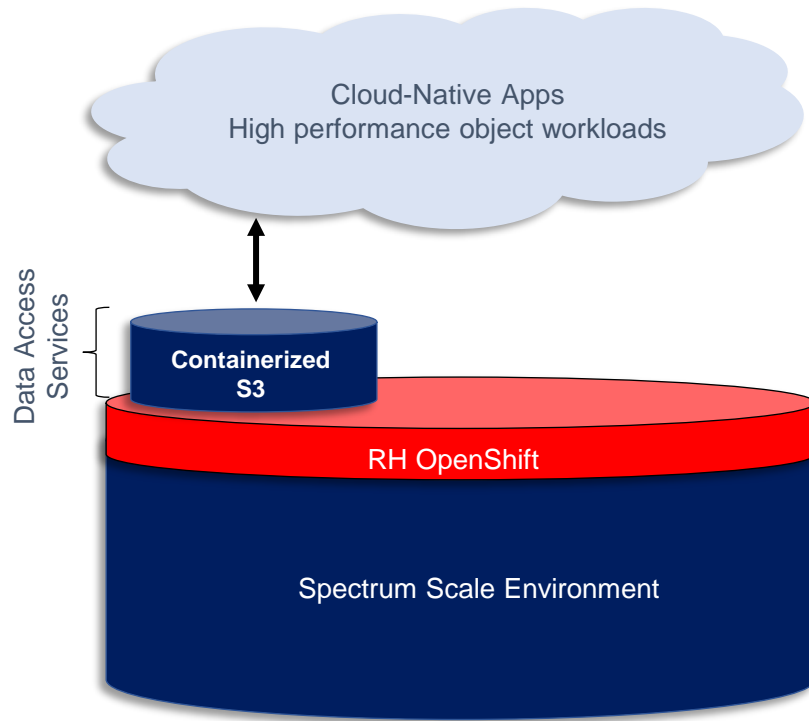
## Customer Requirements & DAS Dependencies:

- Spectrum Scale 5.1.2.1 and newer: DAE, DME, ESS for DAE, ESS for DME, ECE (future)
- OpenShift 4.9.8
- CNSA 5.1.2.1 / CSI 2.4.0
- ESS models at GA followed by any storage (supported by CNSA)

**Performance:** MVP baseline 40 GB/s w/ 3 DAN (Data Access) nodes on vanilla ethernet, scales linearly, increased performance with each release

## Advanced S3 Capabilities on roadmap

- S3 Encryption Headers
- S3 Select
- S3 Versioning
- S3 Object Lock



# GPU Direct Storage (GDS)

Scale with NVIDIA

*Understand how to get GDS and the requirements.*

## **Spectrum Scale Knowledge Center:**

<https://www.ibm.com/docs/en/spectrum-scale/5.1.3?topic=summary-changes>

<https://www.ibm.com/docs/en/spectrum-scale/5.1.3?topic=architecture-gpudirect-storage-support-spectrum-scale>

## **Nvidia GDS Documentation:**

<https://docs.nvidia.com/gpudirect-storage/index.html>

<https://developer.nvidia.com/gpudirect-storage>

For help getting started: [scale@us.ibm.com](mailto:scale@us.ibm.com)

\* For details on supported versions, refer to the Spectrum Scale FAQ

### Which GDS Release\*?

- CUDA 11.4 or later
- CUDA 11.5 for RoCE

### Supported Storage\*

- Spectrum Scale 5.1.3 and newer
- ESS or any NSD client-server storage model

### Supported Network\*

- Infiniband (RDMA)
- Ethernet (RoCE)

### GPUs\*

- NVIDIA Ampere (e.g. NVIDIA A100)

# Big Data & Analytics

## ***Support:***


- Cloudera Data Platform (CDP) Private Cloud Base is certified with IBM Spectrum Scale on x86\_64 and ppc64le since December 2020.
- Cloudera Hortonworks Data Platform (HDP) 3 and HDFS Transparency 3.1.0 end of service on December 31<sup>st</sup>, 2021.

## ***Simplified automated deployment:***

- Ansible toolkit deployment for CES HDFS was included in Scale 5.1.2
- Updated support for parallel offline upgrades for CES HDFS available in Scale 5.1.3

## ***Improved performance:***

- Improved memory efficiency for HDFS Transparency NameNode.
- Optimized parallelism for DataNode request processing via [delete, du and list configuration options](#).



# AFM and AFM DR Updates

## ***AFM updates for Object Storage***

- New **Manual Update (MU)** mode to support manual replication of files using a file list or ILM policy
- Added support for expiring and automatically refreshing cloud storage keys
- Added support for Azure Blob object storage by using MinIO as an S3 gateway
- Security Token Service (STS) Token support
- Support eviction of partially cached files (see `mmafmctl` documentation for details)

## ***Configuration simplification***

- The `afmDIO` parameter is now set by default based on the backend protocol

## ***New queries and statistics available in `mmafmcscctl`***

- Enhanced cloud object upload and download statistics
- See `mmafmcscctl` man page for more details

# Spectrum Scale Core Improvements

## Administration and reliability

***Simpler and more flexible administration that allows better control.***

- Support for 64-bit inode count variables in `gpfs_quotactl()`
- `--filesetdf` is enhanced to report inode capacity and usage at the independent fileset level
- Removed a restriction for `mmfsctl`, allowing it to be run on disks configured as CCR tiebreakers
- New `tsCmdAllowRemoteConnections` attribute to specify whether `ts*` commands can use a remote TCP/IP connection
  - Avoids operation not permitted errors on remote clusters in a multi-cluster environment (see `mmchconfig` man page)




# Spectrum Scale Core Improvements

## Features

*Features that allow you to improve your resource utilization and problem determination.*

- New options to automatically flush buffers to disk on file close (see [mmcrfs](#), [mmchfs](#))
- New permission inheritance mode for filesets (see [mmcrfileset](#) and [mmchfileset](#) for details on [-allow-permission-inherit](#))




# Spectrum Scale Core Improvements

## Performance Improvements

***File system utilities and network enhancements that allow you to maximize your system utilization.***

- New **appendShipEnabled** attribute improves performance of workloads appending to shared files
- Quiesce can now be done at the fileset level, so independent fileset operations, such as snapshot and unlink, only quiesce the affected fileset and not the entire file system.



# Spectrum Scale Core Improvements

## Scale on Z Systems

### ***Better resource allocation and support for NFS clients***

- IBM z16 support
- RDMA/RoCE support for Linux on Z
- z/OS NFS client support





# Security Improvements

- Support for [IBM Security Guardium Key Lifecycle Manager](#) (GKLM, formerly SKLM) [4.1.1](#)
- Additional JSON attributes available in file audit logging (see knowledge center for updates)
- File audit logging now supports updates to filesets under audit. See mmaudit for details
- SMB ACL security features (host allow/deny) for shares

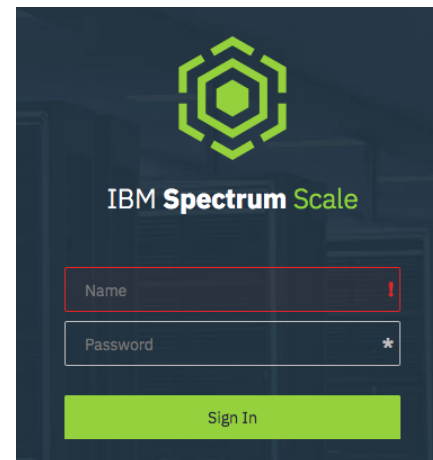


# Management GUI/API Changes

## Administration and reliability

### ***Simpler management.***

- Updates to cache tables on AFM management pages
- Enhanced [Node info](#) and [gpfs trace](#) endpoints.



# Monitoring, Availability & Proactive Services (MAPS) Updates

## System Health & Monitoring

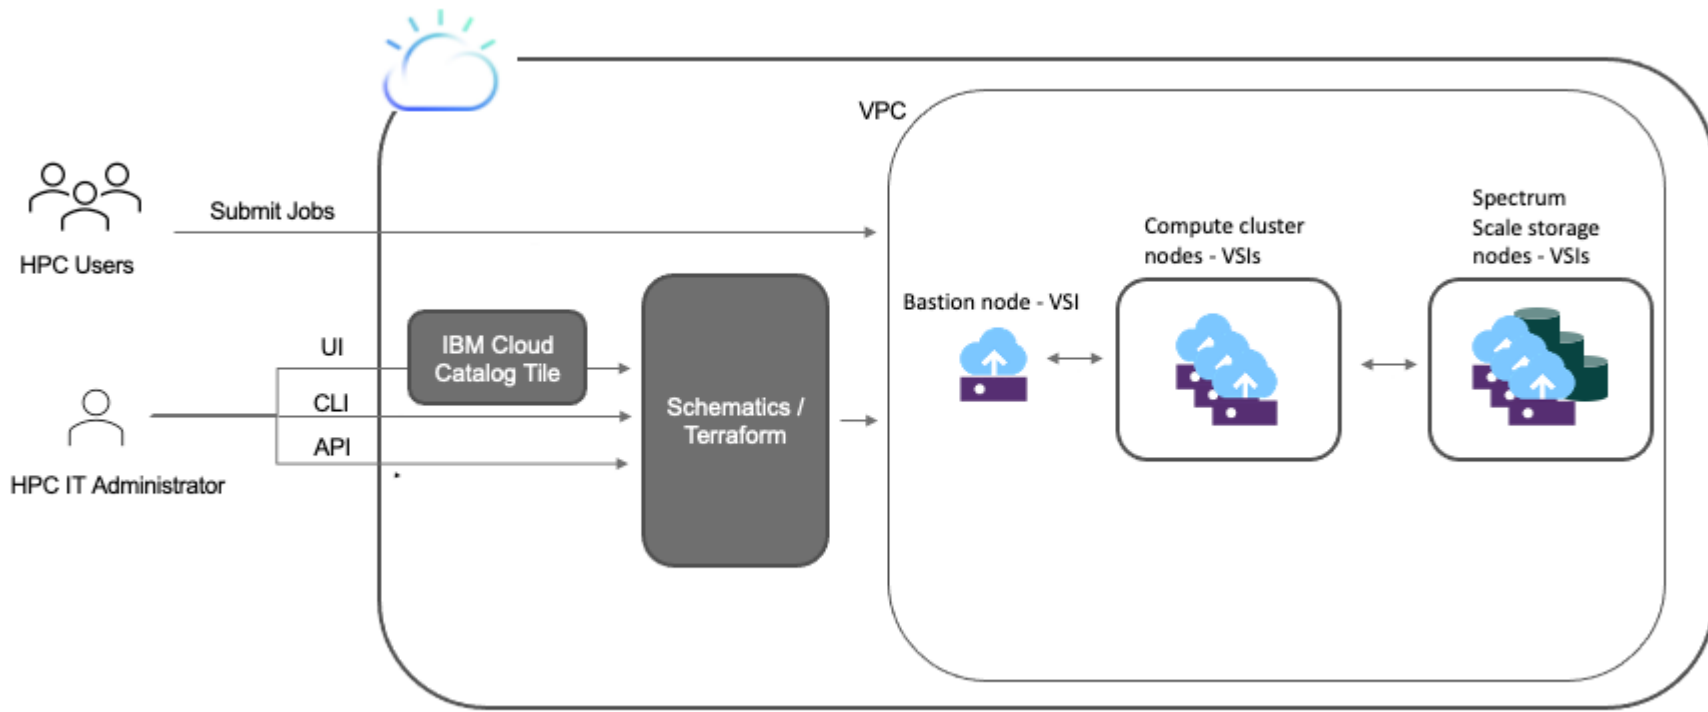
*Enhanced awareness on the status of your system components*

- New gpfs.snap options for HDFS and Hadoop information
- Enhanced stretch cluster monitoring via a new **STRETCHCLUSTER** component
- New **-server** option in the mmcallhome command allows call home servers to be specified explicitly



# Spectrum Scale on IBM Cloud!

Similar to AWS experience - <https://www.ibm.com/cloud/hpc>




# Spectrum Scale Erasure Code Edition Changes

***Simpler installation and enhanced mechanisms for cluster management.***

Support up to 64 disk drives per sever

Updated recommendations for 4 and 6 node configurations  
(see knowledge center for details)

## Erasure coding technology



# IBM ESS 3500: Next Generation

## NEXT GENERATION

Up to 12% better  
performance vs previous  
models

## GREEN ENHANCED

Streamlined designed for  
better thermal results

## POWERED BY SPECTRUM SCALE

Supports Latest Global  
Data Platform Data  
Services

## ENHANCED HIGH AVAILABILITY

Enhanced non disruptive  
upgrades for scale-up



**ESS 3500**

**Measured 91GB/s**  
300  $\mu$ s latency  
Measured 1+ million IOPs

\* GA May 20, 2022

# ESS 3500 & edge computing

Optimized for entry configuration

Eliminate dedicated protocol node

Virtualized protocol services for 100s of clients

- NFS (1000)

- SMB (512)


1 VM per canister

- 8 cores


- 64 GB RAM

Adapters via PCIe-Passthrough

Don't forget about your EMS! 😊



ESS 3500



# Discontinued Features

[https://www.ibm.com/docs/en/STXKQY/pdf/scale\\_deprecated\\_features.pdf](https://www.ibm.com/docs/en/STXKQY/pdf/scale_deprecated_features.pdf)

Category	Discontinued functionality	Recommended Action
Security	The use of TLS 1.0 and 1.1 for authorization within and between IBM Spectrum Scale clusters.	Upgrade to TLS 1.2 or later.
GUI/REST API	The use of TLS 1.0 and 1.1 for authorization with the GUI/REST API server	Upgrade to TLS 1.2 or later.
Platforms	Encryption acceleration library for Power7 (CLIC)	If encryption performance is critical, migrate to newer generations of Power Systems .
	Big Endian Power servers	Upgrade to newer generations of Power systems or remain on IBM Spectrum Scale 5.1.2.
	Linux support for IBM Power7 systems	Plan to migrate to newer generations of Power systems.
Protocols	iSCSI as a target for remote boot	Use some other block services provider.
Containers	Storage Enabler for Containers (SEC)	Migrate to Container Storage Interface (CSI).
HDFS Transparency	HDFS Transparency 3.1.0 and 2.7	It is recommended to contact their IBM or Cloudera representative for more details about migration options.



# Deprecated Features

[https://www.ibm.com/docs/en/STXKQY/pdf/scale\\_deprecated\\_features.pdf](https://www.ibm.com/docs/en/STXKQY/pdf/scale_deprecated_features.pdf)

Category	Deprecated functionality	Recommended Action
Platforms	AIX support for IBM Power7 systems	Plan to migrate to newer generations of Power systems.
Watch folders	Kafka queue for watch folders	If you are using Kafka for other purposes, install it separately from IBM Spectrum Scale. It is no longer installed as part of the IBM Spectrum Scale installation.
Security	Support for Vormetric DSM V5	Upgrade to Vormetric DSM V6.2 or later
Protocols	mmcesdr command (Protocols cluster disaster recovery)	Use AFM and AFM DR to set up your own replication strategies between clusters.
Cluster configuration	The primary and secondary configuration server functionality. Instead of this, clusters must use CCR.	The default configuration service is CCR, and new clusters are created using CCR. If not yet operating with CCR, change to that mode with mmchcluster --ccr-enable.
Block size	The --metadata-block-size option of mmcrfs command is deprecated. This option is used for defining metadata blocks to a different size than the data blocks.	Only a single definition for the number of subblocks per block exists per file system. Selecting a smaller metadata block size has the unintended side effect of increasing the subblock size for data blocks. Although it is supported to set metadata blocks to a different size than data blocks by using the --metadatablock-size parameter, it is not recommended to use that option. This option is deprecated.

# Stabilized Features

[https://www.ibm.com/docs/en/STXKQY/pdf/scale\\_deprecated\\_features.pdf](https://www.ibm.com/docs/en/STXKQY/pdf/scale_deprecated_features.pdf)

Category	Stabilized functionality	Recommended Action
cNFS	The use of TLS 1.0 and 1.1 for authorization within and between IBM Spectrum Scale clusters.	IBM®'s strategic path is to invest in User Space solutions for NFS support of Scale workloads. Once User Space performance and function are considered to be sufficient to replace cNFS, anticipate that the support for cNFS is deprecated.

# Log your RFE!

[https://www.ibm.com/developerworks/rfe/execute?use\\_case=productsList](https://www.ibm.com/developerworks/rfe/execute?use_case=productsList)

- [Spectrum Scale \(formerly known as GPFS\) - Private RFEs](#)
- [Spectrum Scale \(formerly known as GPFS\) - Public RFEs](#)

*contact*

**Filter the page content by brand and product**

Servers and Systems So... ▼

Spectrum Scale (formerly known as GPFS) - Pu... ▼ ➤

[Hot](#)

[Top](#)

[New](#)

Search



35  
votes

## **eleminate lack of I/O on mmdelsnapshot start**

When deletion of bunch of snapshots starts we a lack of I/O for about three minutes. NFS Clients see a huge delay of I/O. Related applications hanging for this time and user connections and run into t...


**Under Consideration**

21  
votes

## **start services after gpfs (filesystems) is ready using systemd**


the current systemd units gpfs.service and <fs-mountpoint>.mount units can't be used to depend on (After/Required/.. systemd attributes) for other services. GPFS service is reporting itself as success...

**Under Consideration**



Check <https://www.spectrumscaleug.org/experttalks>  
for charts, show notes and upcoming talks

- Past talks:
  - 001: What is new in Spectrum Scale 5.0.5?
  - 002: Best practices for building a stretched cluster
  - 003: Strategy update
  - 004: Update on performance enhancements in Spectrum Scale (file create, MMAP, direct IO, ESS 5000)
  - 005: Update on functional enhancements in Spectrum Scale (inode management, vCPU scaling, NUMA considerations)
  - 006: Persistent Storage for Kubernetes and OpenShift environments
  - 007: Manage the lifecycle of your files using the policy engine
  - 008: Multi-node scaling of AI workloads using Nvidia DGX, OpenShift and Spectrum Scale
  - 009: Continental: Deep Thought – An AI Project for Autonomous Driving Development
  - 010: Data Accelerator for Analytics and AI (DAAA)
  - 011: What is new in Spectrum Scale 5.1.0?
  - 012: Lenovo - Spectrum Scale and NVMe Storage
  - 013: Event driven data management and security using Spectrum Scale Clustered Watch Folder and File Audit Logging
  - 014: What is new in Spectrum Scale 5.1.1?
  - 015: IBM Spectrum Scale Container Native Storage Access



# Thank you!

Please help us to improve Spectrum Scale with your feedback

- If you get a survey in email or a popup from the GUI, please respond

- We read every single reply



Provide Feedback



Tell IBM What You Think

Let us know what you think about IBM Spectrum Scale. It takes only a couple of minutes for you to help us improve our service. [IBM Privacy Policy](#)

Not Now

 Provide Feedback



## Spectrum Scale User Group

The Spectrum Scale (GPFS) User Group is free to join and open to all using, interested in using or integrating IBM Spectrum Scale.

The format of the group is as a web community with events held during the year, hosted by our members or by IBM.

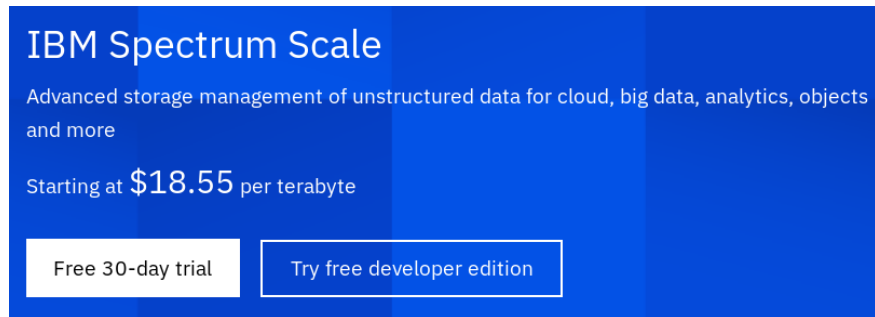
See our web page for upcoming events and presentations of past events. Join our conversation via mail and Slack.

[www.spectrumscaleug.org](http://www.spectrumscaleug.org)

# Spectrum Scale Developer Edition!

Fully functional!

- Based on first PTF of a release
- Derived from **Data Management Edition (DME)**
- Limited to 12 TBs:  
enough for a small test cluster
- Available from the Scale "try and buy" page on [ibm.com](https://ibm.com)



**IBM Spectrum Scale**

Advanced storage management of unstructured data for cloud, big data, analytics, objects and more

Starting at **\$18.55** per terabyte

Free 30-day trial    Try free developer edition

Free for non-production use, e.g. test, learning, upgrade prep...

- If you have to ask, it's probably not permitted

Not formally supported

# Spectrum Scale Early Programs

## Types of Programs:

### Alpha

Influence the development of new technology by gaining before market access to product code. Alpha programs are typically confidential and the first opportunity for you to interact with a feature or function.

### Beta

Try out a new offering with the team who owns the product and influence its usability and design. A Beta program gives you the ability to evaluate and provide feedback on IBM products before the products general availability. Beta programs are typically confidential and run prior to GA.

### Early Support Program (ESP)

Be one of the few selected participants to validate new Software or Hardware and potentially give your enterprise an edge over the competition. The IBM early support programs give you and IBM the opportunity to develop, evaluate, and gain experience with a product or a set of products in your enterprise environment.



### Customer Success

- ☐ Evaluate new IBM HW or SW in your environment.
- ☐ Validate procedures and interoperability with other products in your enterprise.
- ☐ Opportunity to Influence Product Design
- ☐ Early Enablement and education
- ☐ Strengthen Partnership with IBM

Talk to your IBM contact or Partner to be nominated!

# Spectrum Scale on GitHub!

<https://github.com/IBM/SpectrumScaleTools>

- IBM Spectrum Scale Bridge for Grafana
- IBM Spectrum Scale cloud install
- IBM Spectrum Scale Container Storage Interface driver
- IBM Spectrum Scale install infra
- IBM Spectrum Scale Security Posture
- Oracle Cloud Infrastructure IBM Spectrum Scale terraform template
- SpectrumScale\_ECE\_CAPACITY\_ESTIMATOR
- SpectrumScale\_ECE\_OS\_OVERVIEW
- SpectrumScale\_ECE\_OS\_READINESS
- SpectrumScale\_ECE\_STORAGE\_READINESS
- SpectrumScale\_ECE\_tuned\_profile
- SpectrumScale\_NETWORK\_READINESS

Find open source tools that are related with IBM Spectrum Scale.

Unless stated otherwise, the tools compiled in this list come with no warranty of any kind from IBM.



# Check out the FAQ!

<https://www.ibm.com/support/knowledgecenter/en/STXKQY/gpfsclustersfaq.html>

<https://www.ibm.com/support/knowledgecenter/STXKQY/gpfsclustersfaq.pdf?view=kc>

<https://www.ibm.com/support/knowledgecenter/SSYSP8/gnrfaq.html>

## HTML or PDF

## Spectrum Scale version compatibility with OS or kernels

## Updated regularly!

