Spectrum Scale Expert Talks

Episode 14:

What is new in Spectrum Scale 5.1.1?

Show notes:
www.spectrumscaleug.org/experttalks

Join our conversation:
www.spectrumscaleug.org/join

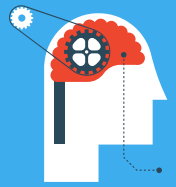

SSUG::Digital

Welcome to digital events!

Show notes:
www.spectrumscaleug.org/experttalks

Join our conversation:
www.spectrumscaleug.org/join

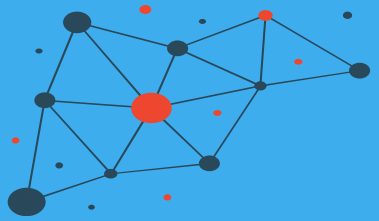

About the user group

- Independent, work with IBM to develop events
- Not a replacement for PMR!
- Email and Slack community
- https://www.spectrumscaleug.org/join

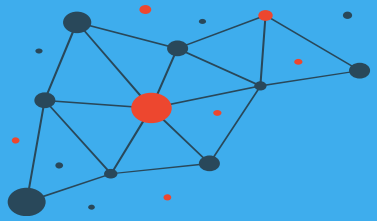

We are …

- Simon Thompson (UK)
- Kristy Kallback-Rose (USA)
- Bob Oesterlin (USA)
- Bill Anderson (USA)
- Chris Schipalius (Australia)

IBM CHAMPION

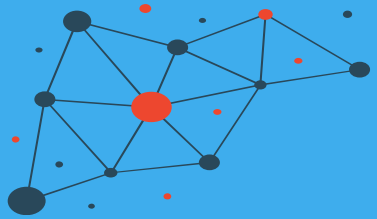

Check https://www.spectrumscaleug.org/experttalks for charts, show notes and upcoming talks

- 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
- Today:
  - May 19: What is new in Spectrum Scale 5.1.1?

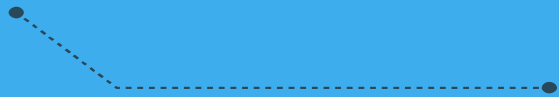

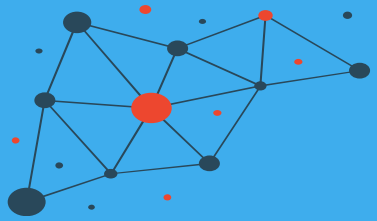

Speakers

- Chris Maestas (IBM)
- Ismael Solis Moreno (IBM)

Disclaimer

Spectrum Scale

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

Spectrum Scale

Enhanced integration and deployment features with containers.

Improved network performance for RoCE and TCP/IP configurations to ensure better bandwidth efficiencies

Enhanced DevOps installation and upgrade experience with ansible playbooks.

GPU Direct Storage (GDS) technology preview to test and validate NVIDIA pre-released code with Spectrum Scale.

Ensure immutability compliance in disaster recovery scenarios.

Support latest technology integrations with Cloudera and OpenShift

Spectrum Scale: A quick note on releases …

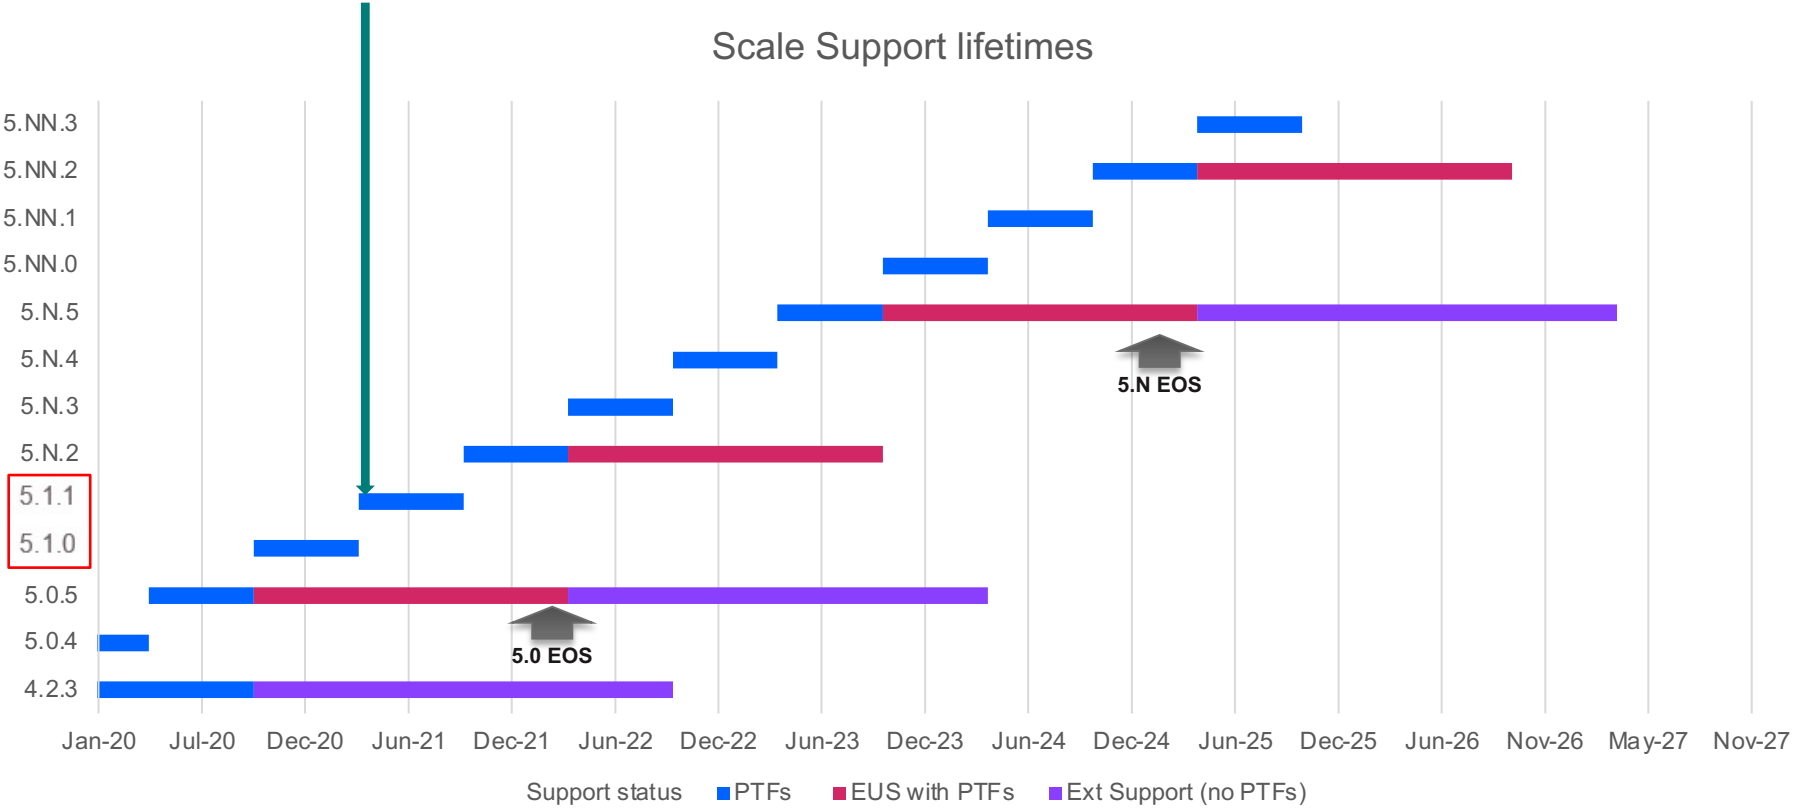

Container Native Storage Access

Spectrum Scale

Deploy Spectrum Scale on containers.

IBM Spectrum Scale container native storage access (CNSA) supports the Red Hat OpenShift Container Platform with a fully containerized deployment.

This allows deployment of IBM Spectrum Scale on Red Hat CoreOS worker nodes where classic packages cannot be installed.

Container Native Storage Access

Spectrum Scale

Improvements introduced in CNSA 5.1.0-3

Wider support to use the latest CNSA functionality.

- Support for upgrading IBM Spectrum Scale Container Native Storage Access (CNSA) from v5.1.0.1 to v5.1.0.3.
- Support for RedHat OpenShift Container Platform 4.6*
- Operator now maintains status and creates events on ScaleCluster CR.
- CA Certificate support for GUI communication.
- Support for X86, Power and Z.

Container Storage Interface

Spectrum Scale

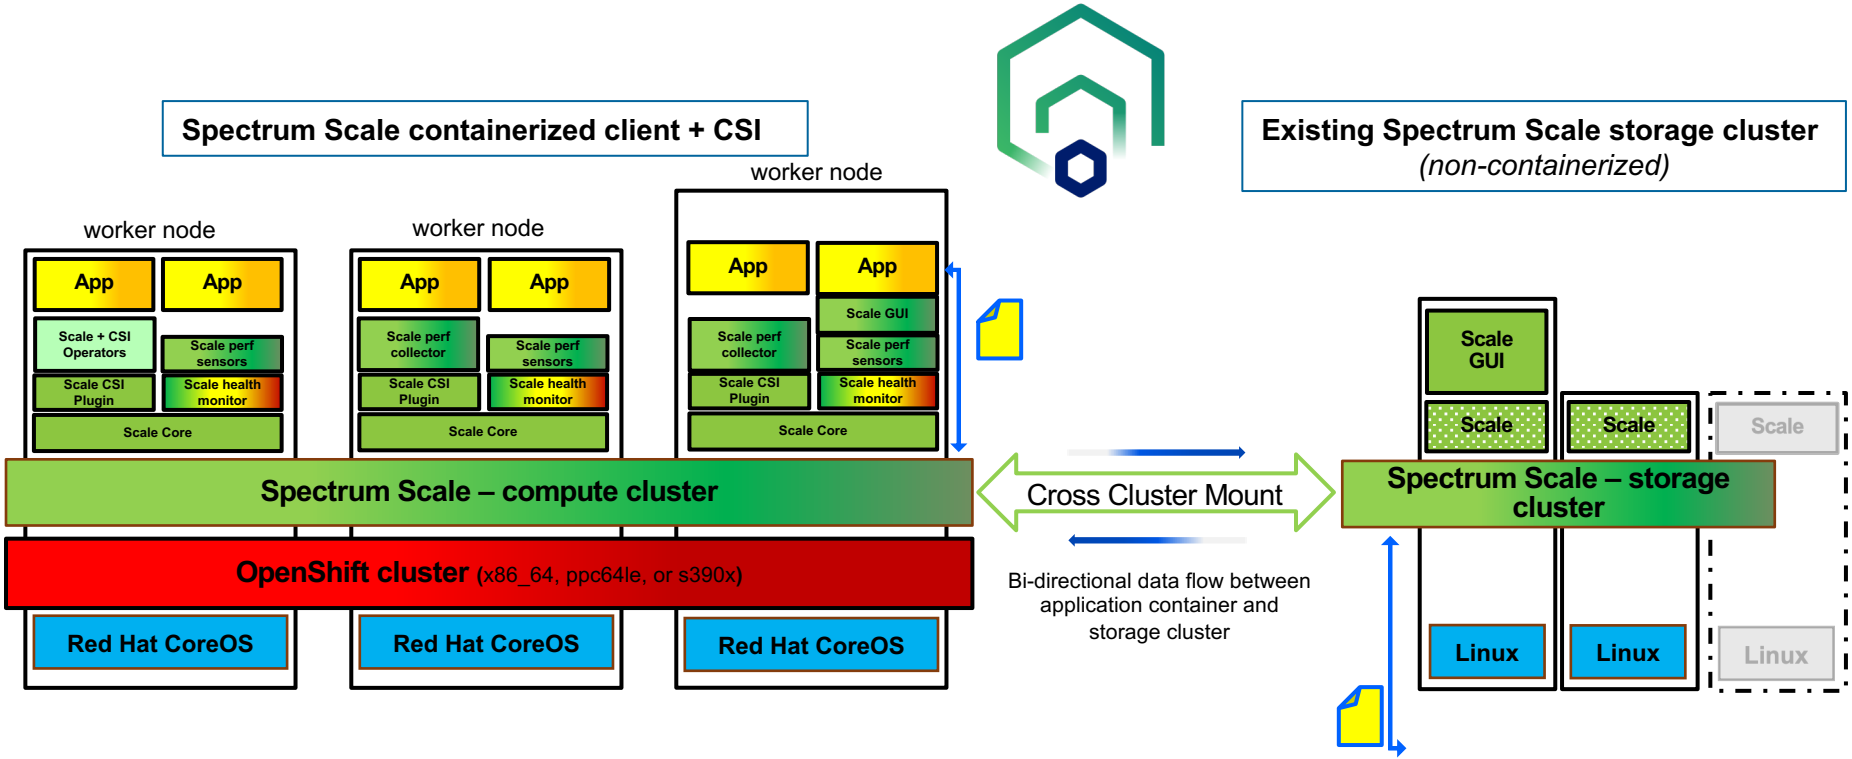

Container Storage Interface

Spectrum Scale

Improvements introduced in CSI 2.2.0

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

- Support for Red Hat OpenShift 4.7 and Kubernetes 1.20, 1.21.
- Upgraded the ansible-operator version from 0.19.4. to 1.5.1. Keep current with the latest secury, stability and functional enhancements.
- Support for IBM Spectrum Scale Container Storage Interface driver Volume Snapshots. Now, you can create volumen snapshots for independent filesets based PVCs.
- Removed scaleHostpath parameter from CR and gpfs-classicvolume mount from CSI driver daemonsetpods. These are not longer needed help to enhance your security.
- Customizable kubeletroot directory path –kubeletRootDirPath. Thus, flexible configuration

Installation Toolkit Enhancements

Spectrum Scale

Simpler installation through enhanced automation. Scale 5.1.1 installation toolkit automates the steps that are required to install GPFS, deploy protocols, and install updates and patches.

•Migration to the Ansible automation platform:

- Enables scaling up to a larger number of nodes.
- Avoids issues that arise due to using an agent-based tooling infrastructure such as Chef.
- Enables parity with widely-adopted, modern tooling infrastructure.

•Support for IPv6 addresses.

•Support for the CES interface mode.

Ansible

Spectrum Scale

A new and enhanced way to deploy Spectrum Scale

- Move to ansible installation for ALL use cases
  - Common code base
  - Open-source project https://github.com/IBM/ibm-spectrum-scale-install-infra
  - Dropping chef-based toolkit also avoids issues that arise with agent-based tooling infrastructure.
  - Enables scaling up to a greater number of nodes

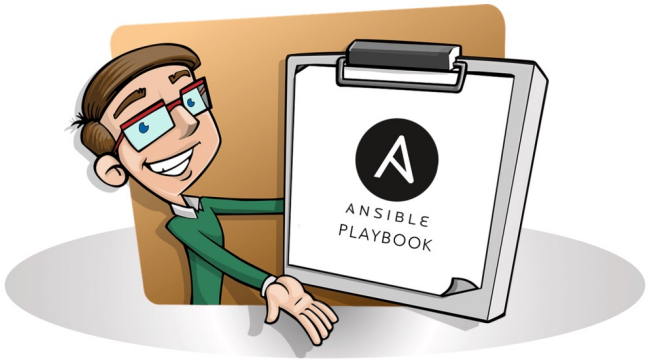

- Configuration management code language changes from Ruby to Ansible YAML.
  Cluster configuration file changes file from txt to JSON format.
- Installer to cluster node communication is completely happening through SSH, previously it was happening through both SSH and chef knife bootstrap.
- Installer cluster configuration file location change from installer/configuration/clusterdefinition.txt to ansible-toolkit/ansible/ibm-spectrum-scale-install-infra/vars/scale_clusterdefinition.json
- Ansible toolkit will deprecate ntp and auth configuration.

Ansible

Spectrum Scale

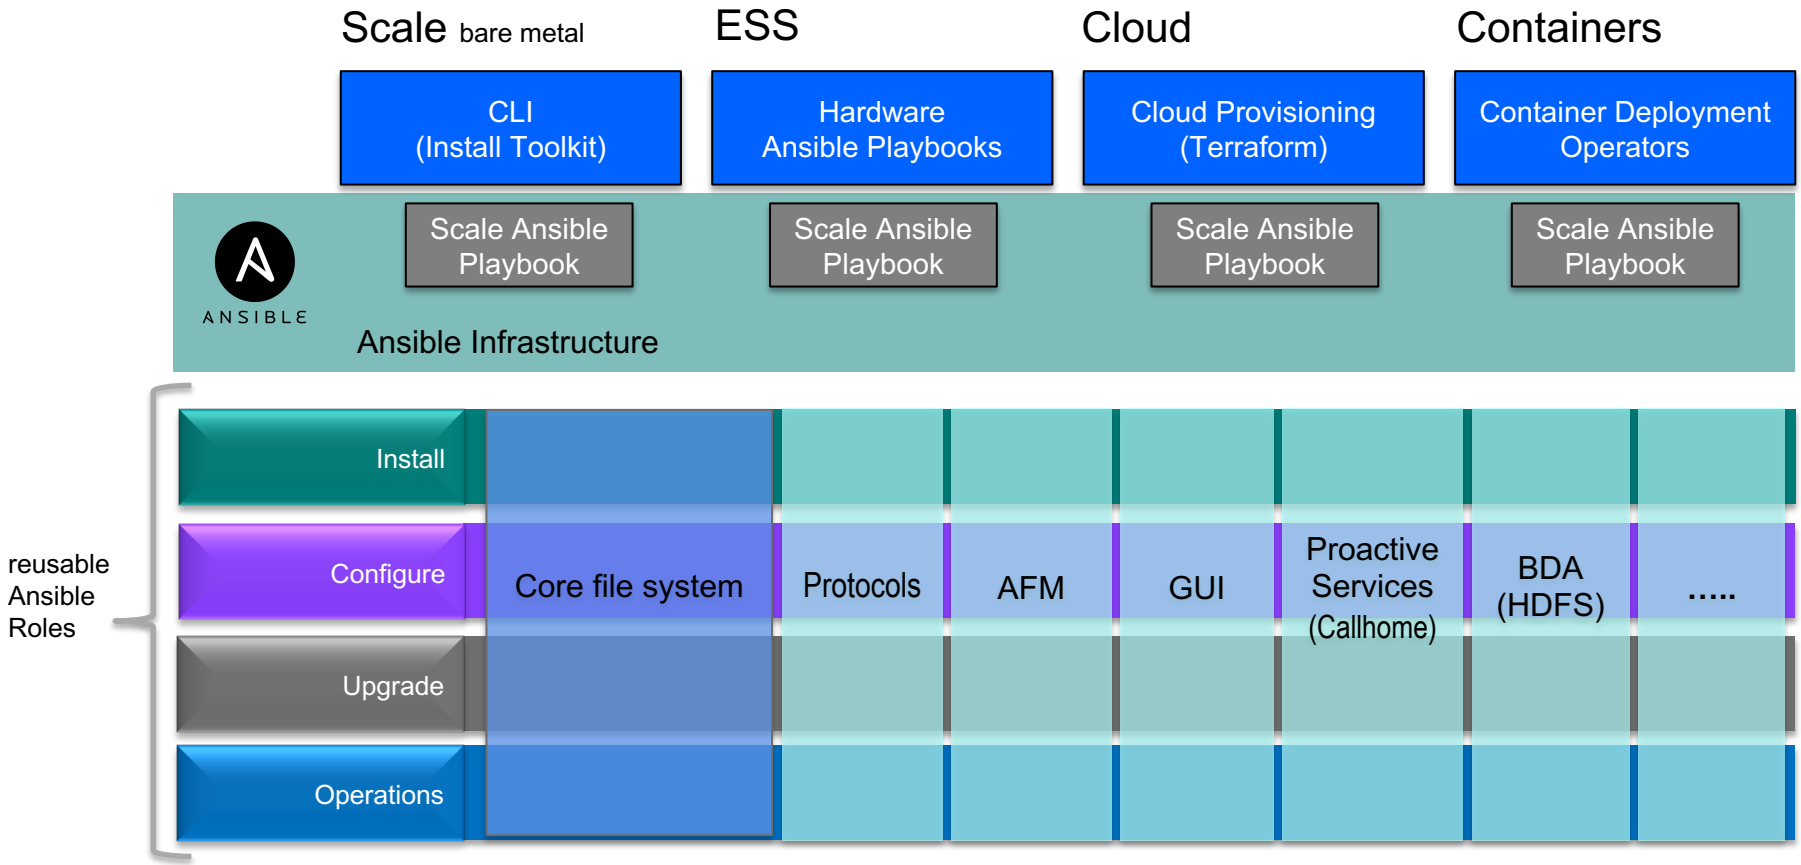

GPU Direct Storage (GDS)

Spectrum Scale

Scale with NVIDIA

Greatly reduce CPU overhead when talking to IBM Spectrum Scale from the GPU

IBM Spectrum Scale 5.1.1. Tech Preview

Requires NVIDIA's GDS SDK,
which is currently in Beta

GPUDirect Storage (GDS) enables an NVIDIA developer to

- Make a direct memory access (DMA) between GPU memory and storage
- Bypass the CPU and system memory
- Reduce latency, Increase bandwidth, and lower CPU utilization for a specific read from storage

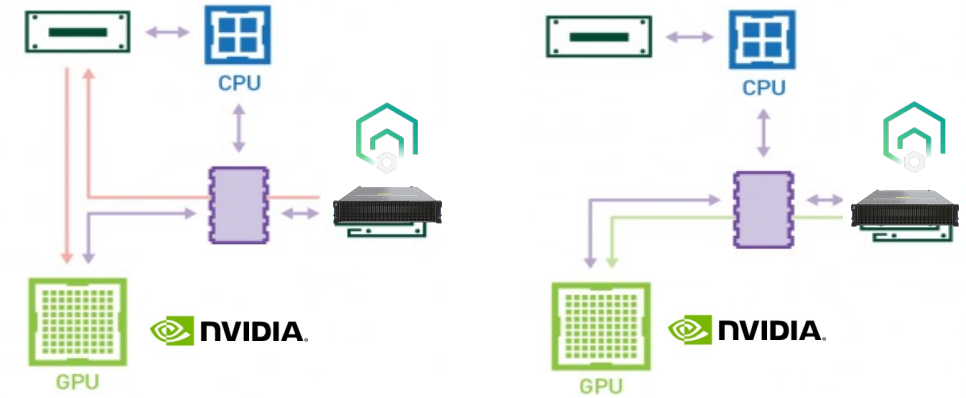

GPU Direct Storage (GDS)

Spectrum Scale

Scale with NVIDIA

Understand how to get GDS and the requirements.

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

Start here: https://developer.nvidia.com/gpudirect-storage

For help getting started: scale@us.ibm.com

Which GDS Release?

- Closed Beta v. 0.9x
- Expected to be part of a future CUDA release

Supported Storage

- Tech Preview in Spectrum Scale 5.1.1
- Any Spectrum Scale deployment

Supported Network

- Infiniband (RDMA)
- MOFED 5.2-1.0.4.0

GPUs and OS

- GPUs:
  NVIDIA Ampere
  (e.g. NVIDIA A100)
- OS:
  RHEL 8.3,
  Ubuntu 20.04 and 18.04

Spectrum Scale Core Improvements

Spectrum Scale

Multi-Connections Over TCP/IP

Enhanced TCP/IP network utilization and stronger failover by using multiple connections.

- Enhances Spectrum Scale communication subsystem to establish multiple TCP connections between each pair of daemons.
- This enables client to fully use Ethernet bandwidth, for ethernet networks or bonding configurations

Improves performance
by allowing networks to be fully utilized

Improves resilience
by providing failover from one connection to another

Spectrum Scale Core Improvements

Spectrum Scale

Multi-Connections Over TCP/IP

Simple to configure and administer. Using a simple command you can designate how many connections to establish between nodes.

- If nodes are running 5.1.1 the new feature is exploited without administrative action and works with remote cluster mounts.
- Both nodes in each pair of nodes need to be running at 5.1.1 for multiple connections to be established.
- Configuration variable maxTcpConnsPerNodeConn can be used to configure the number of TCP connections between each pair of nodes.
- Large clusters might not benefit from the improvement if the links to the NSD servers are already saturated (fully utilized).

Big Data & Analytics

Spectrum Scale

Extended support:

- Cloudera Data Platform (CDP) Private Cloud Base is certified with IBM Spectrum Scale on x86_64 and ppc64le
- CDP Private Cloud Base was 1st certified on IBM Spectrum Scale 5.1.0.1 with HDFS Transparency 3.1.1-3 in December 2020

Simplified automated deployment:

- Ansible toolkit deployment for CES HDFS in Scale 5.1.1

Improved performance:

- HDFS Transparency 3.1.1-3 implements performance enhancement using fine-grained file system locking mechanism when Ranger is enabled.
- The gpfs.ranger.enabled field is changed starting in HDFS Transparency 3.1.0-6 and 3.1.1-3 to set to “scale” for fine-grained file system enhancement.

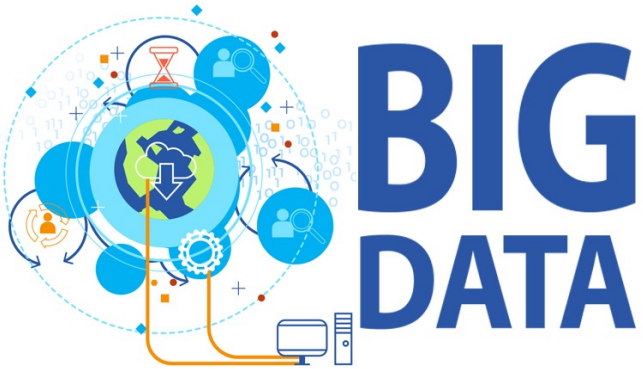

Big Data & Analytics

Spectrum Scale

Enhanced security and reliability for HDFS Transparency:

- Version 3.1.1-4 in IBM Spectrum Scale 5.1.1 contains security fixes and logging enhancements.
- Version 3.1.0-7 in Jan 2021 contains fixes to NullPointerException in NameNode logs and correct the reporting of JMX open operations.

Enhanced support and management with Mpack 2.7.0.9

- Support Ambari maintenance mode for clusters with shared and remote IBM Spectrum Scale storage types
- Customized hdfs-site.xml configuration will not be overwritten when upgrading Mpack
- The Mpack version/build information can be obtained from the tar.gz package without installation

AFM-DR Immutability

Spectrum Scale

Immutability & appendOnly for AFM-DR filesets

Enhanced recoverability and security by adding immutability & appendOnly capabilities to AFM-DR filesets. Typically, used in heavily regulated environments such as health care and financial services.

- After setting these flags on a file or a directory in an AFM-DR primary fileset, AFM automatically synchronizes these metadata flags on the secondary fileset.
- Immutable files cannot be changed or renamed.
- On a file or directory with the appendOnly flag enabled, one can do append operations, but no delete, modify, or rename operations are allowed.
- Flags can be set independently – if both Immutability and appendOnly flags are set on a file, the Immutability flag takes effect.

By enabling an integrated archive manager (IAM) mode on an AFM-DR primary and secondary fileset, user can set the Immutability and appendOnly flags on a file or a directory in an AFM-DR fileset.

File Audit Logging (FAL) Improvements

Spectrum Scale

Improved data search and management. Now it is simpler to consult and filter FAL Logs.

FAL adds a new way to parse/query FAL logs for specific information or to generate a JSON file with all logs so they can be used outside of this tool.

```
[root@dodgers-vm1 ~]# /usr/lpp/mmfs/samples/util/audit_parser --filepath
/ibm/gpfs0/.audit_log/SpectrumScale_151_7775138800282916176_1617966376_FS
YS_gpfs0_audit/auditLogFile.latest_dodgers-vm1
Items checked: 425/425
Results: 425/425
FAL log query results located in query_results.json
```

```
[root@dodgers-vm1 ~]# /usr/lpp/mmfs/samples/util/audit_parser --filepath
/ibm/gpfs0/.audit_log/SpectrumScale_151_7775138800282916176_1617966376_FS
YS_gpfs0_audit/auditLogFile.latest_dodgers-vm1 --event OPEN,CLOSE,UNLINK
Items checked: 425/425
Results: 406/425
FAL log query results located in query_results.json
```

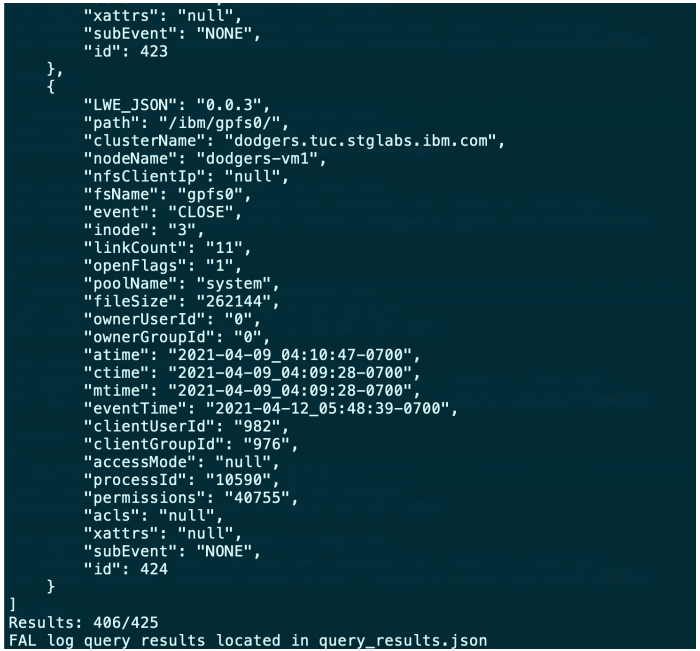

```
        "xattrs": "null",
        "subEvent": "NONE",
        "id": 423
    },
    {
        "LWE_JSON": "0.0.3",
        "path": "/ibm/gpfs0/",
        "clusterName": "dodgers.tuc.stglabs.ibm.com",
        "nodeName": "dodgers-vm1",
        "nfsClientIp": "null",
        "fsName": "gpfs0",
        "event": "CLOSE",
        "inode": "3",
        "linkCount": "11",
        "openFlags": "1",
        "poolName": "system",
        "fileSize": "262144",
        "ownerUserId": "0",
        "ownerGroupId": "0",
        "atime": "2021-04-09_04:10:47-0700",
        "ctime": "2021-04-09_04:09:28-0700",
        "mtime": "2021-04-09_04:09:28-0700",
        "eventTime": "2021-04-12_05:48:39-0700",
        "clientUserId": "982",
        "clientGroupId": "976",
        "accessMode": "null",
        "processId": "10590",
        "permissions": "40755",
        "acls": "null",
        "xattrs": "null",
        "subEvent": "NONE",
        "id": 424
    }
]
Results: 406/425
FAL log query results located in query_results.json
```

Spectrum Scale Core Improvements

Spectrum Scale

Administration and reliability

Simpler and more flexible administration that allows you better control.

- Designate a node as performance monitoring collector while adding a node using mmaddnode. Simply add and configure at once.
- MMBACKUP_SERVER_FROM_OPT
  option to specify the IBM Spectrum Protect server to be used (dsm.opt) for backups without the --tsm-servers option. This provides flexibility on your backup process.
- mmdf command reports inode capacity and inode usage if quota is disabled and filesetdf is enabled. Get more information for independent filesets.
- Improved checking allowed characters in the keystore password mmkeyserv. Alphanumeric plus limited set of special characters. Stronger password definition to enhance your security.

Spectrum Scale Core Improvements

Spectrum Scale

Features

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

- Support for RoCE with multiple IP aliases per RDMA adapter. Get more flexibility to configure complex RoCE environments. Define multiple aliases and Scale will perform a dynamic pairing of IPs within the same network.
- New MISC_ATTRIBUTE file attribute that can be used in the SQL expressions to identify files whose data is completely contained in the inodes. It helps you to improve your storage allocation.
- Maximum size of a Local Read-Only Cache (LROC) device has been increased to 4TB. More size to improve your read performance.
- Other LROC enhancements such as integration with system health and diagnostic improvements. This supports faster problem solving.

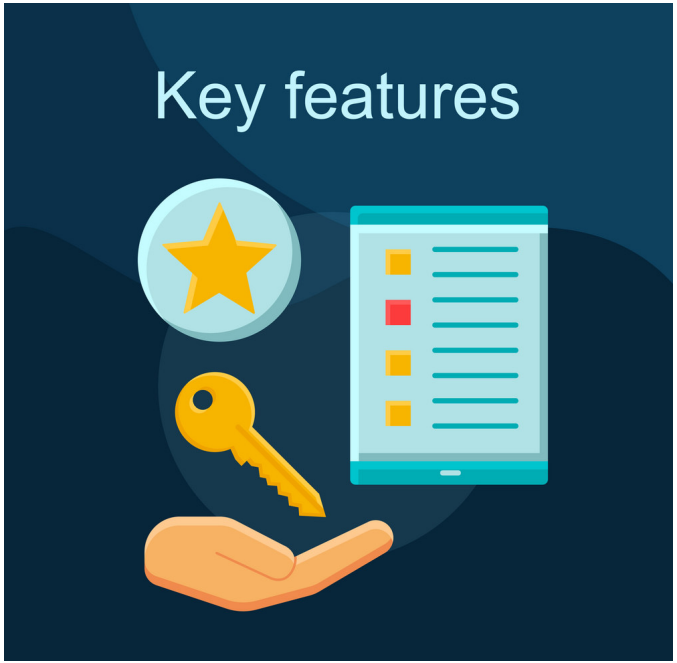

Spectrum Scale Core Improvements

Spectrum Scale

Features

Features that allow you to improve error handling and cluster awareness.

- Introduced warning for avoiding the deletion of a node that is configured as a performance monitoring collector. Increase your awareness to avoid losing your performance stats.
- Commands like mmchdisk, mmrestripefs, mmchdisk and mmdelsnapshot can continue running after a PIT phase error when using the switch --pit-continues-on-error. It will also generate a file that contains a list that failed files during the PIT process.
- New --on-error-continue option in the mmaddnode and mmdelnode commands. Enables the command to trap invalid nodes, remove those nodes from the node list, and continue to process the command.

Spectrum Scale Core Improvements

Spectrum Scale

Performance Improvements

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

- RDMA thread NUMA binding. Reduced latency and bandwidth bottlenecks in NUMA nodes with RDMA adapters by reducing cross node memory access.

- Faster way to calculate the amount of storage used by snapshots with mmlssnapshot -d –fast. Faster snapshot management and lower performance impact.

- Performance enhancement using mmfindinode command to minimize service outage. 8x faster than previous tsfindinode using multiple nodes.

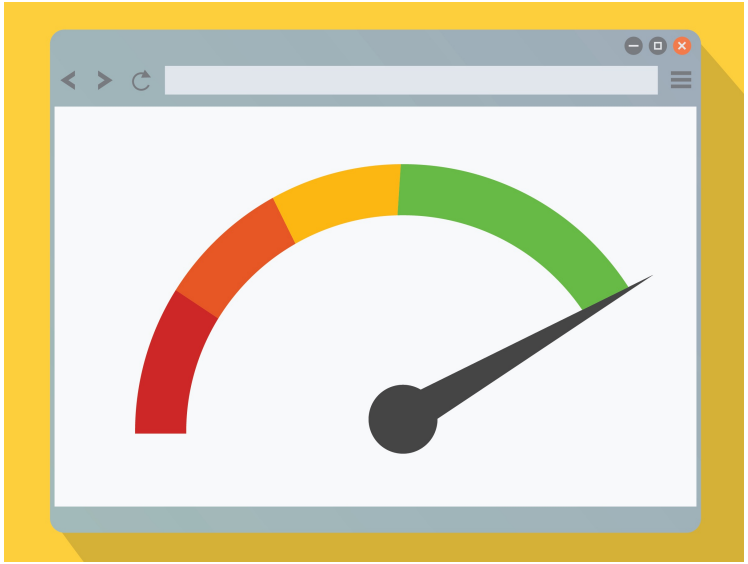

Spectrum Scale Core Improvements

Spectrum Scale

Scale on Z Systems

Better resource allocation and support for NFS clients

- Thin provisioning with Spectrum Virtualize hardware configurations (V7000). Optimize your space allocation on demand and get the most of your available resources.

- Wider support for z/OS NFS clients Including RHEL 7 & 8, SLES12 & 15 as well as z/OS V2. This means more options to deploy your Scale clusters over Z systems.

Security Improvements

Spectrum Scale

Authentication and session management for the GUI

Improved security, through an enhanced authentication mechanism. As well as more flexible way to configure your GUI for login and session timeouts

- Multi-Factor Authentication can be enabled for GUI user groups. Stroger security for GUI users.
- The Services > GUI provides options to enable the users with Security Administrator privileges to define the number of failed login attempts and the maximum duration after which the user account is locked.
- The session timeout is enabled by default and predefined at 30 minutes after which the user will be automatically logged out. The Services > GUI provides the option to configure this parameter and set a new session timeout duration as required.
- Last login details displayed for users provides for an additional security insight.

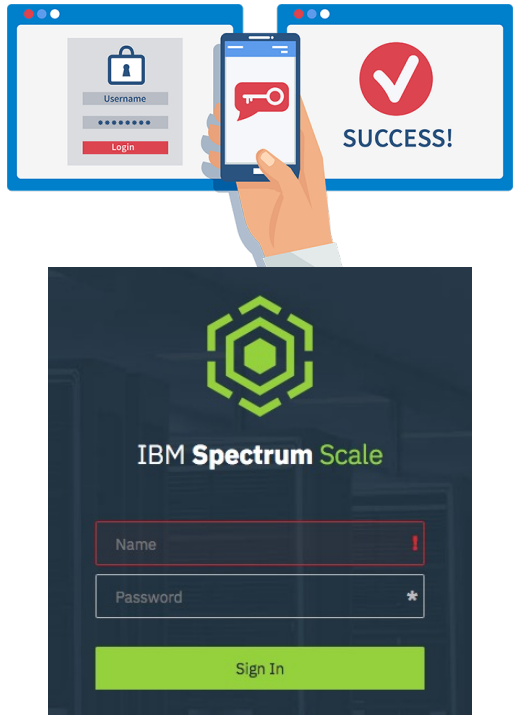

Management GUI Changes

Spectrum Scale

Administration and reliability

Simpler utilization.

- User passwords can now be configured so that they never expire. Now you decide whether changing periodically your passwords.
- Configuration parameters are provided under Services > GPFS daemon to define the way the available limit (quota) of data storage and file system sizes are displayed.
- The number of API requests that can queued in one second is defined under Services > GUI to eliminate the possibility of a Denial of Service (DOS) state owing to server resources being exhausted.
- Snapshot creation and deletion by GUI uses batch mode of mmcrsnapshot / mmdelsnapshot. This provides you a simpler way to create and delete your snapshots.

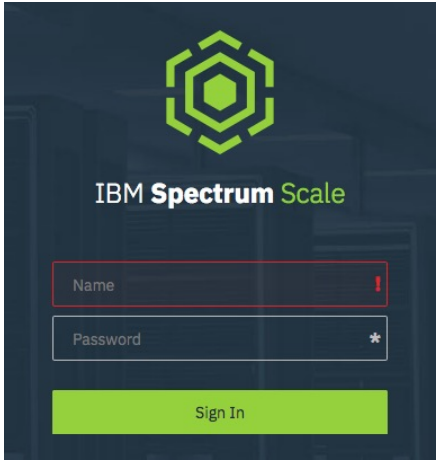

Management API Changes

Spectrum Scale

Rest API Security Enhancements

Enhanced security through the implementation of end-point based authorization.

- End-point based authorization for REST API in /usr/lpp/mmfs/gui/cli/ for chrestacl & lsrestacl

  In case there is **no** REST API ACL defined, it is presumed that the default access is ALLOW.

  In case there is **any** REST API ACL created, regardless the user group, it is presumed that the default access is DENY.

  A user is authorized to consume a certain endpoint only if there is at least one explicit allow permission which applies to his user group and no explicit deny permission.

- REST API to GET Storage Pool information. More flexible ways to get the necessary information of your system.

- REST API for snap handling. Wider possibilities to administrate your system snapshots.

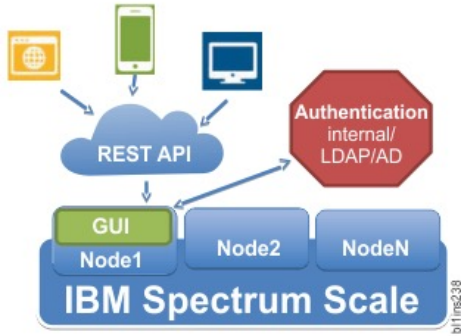

Monitoring, Availability & Proactive Services (MAPS) Updates

Spectrum Scale

System Health & Monitoring

Enhanced awareness on the status of your system components

•The mmhealth config monitor command displays the health monitoring service status of a node. A simple way to verify, pause and resume the status of the monitoring service.

•Local Cache component added to the mmhealth node show and mmhealth cluster show commands. New components, more awareness of your system health including LROC devices.

•Monitor the encryption service status using mmhealth node show command. Easily check the status of your encryption service status.

•Added security (APIKEY) to query interface mmperfmon config add/show/delete/update --apikey key_name. See how your performance stats are accessed.

Monitoring, Availability & Proactive Services (MAPS) Updates

Call Home

Help you during the diagnosis of system issues.

- Added Hadoop configuration data. Information on your Big Data environment.
- Added IB information. Information on your network.
- Improved documentation of data collected. Information on collected data.
- Collect details on call home HW (EMS Nodes only).
- Added mmnetmetrics performance tool to samples (CPU, socket, network, performance, IB). This data can be now stored and uploaded by call home to be used for diagnostic purposes.

Spectrum Scale Erasure Code Edition Changes

Spectrum Scale

Enhanced storage scaling mechanisms and network optimizations to expand the capabilities of your ECE cluster.

- Support for vertical scale up by adding new disks to ECE nodes. Start from small scale and grow to large scale step by step.
- Add storage space to file system automatically in ECE scale out. Create new vdisks and add to file system automatically without manual steps.
- Run ECE with AMD EPYC + PCIe Gen4 + Mellanox ConnectX-6. This comes along substantial system performance improvements.
- Support for multiple connections over TCP to optimize the bandwidth usage. Therefore, better utilization of your network resources.

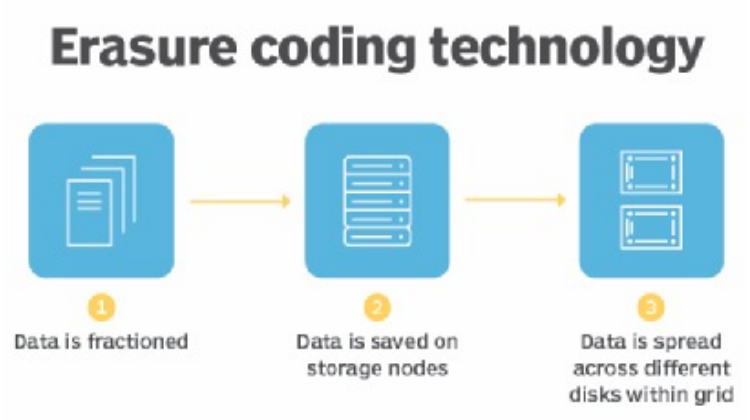

Spectrum Scale Erasure Code Edition Changes

Spectrum Scale

Simpler installation and enhanced mechanisms for cluster management.

Improved Installation toolkit. Simplification of vdisk and file system creation, automation with Ansible, and enhanced support for multiple recover groups.

Support to add multiple nodes at the same time in ECE scale out **(ptf1).** Shorten the time-consuming scale out maintenance operations. Less data movement and performance impact.

Minimal 4 disks per small DA (ptf1). Optimized for small scale RG, especially the minimal 4 nodes configuration.

RHEL 8.4 currency (ptf1).

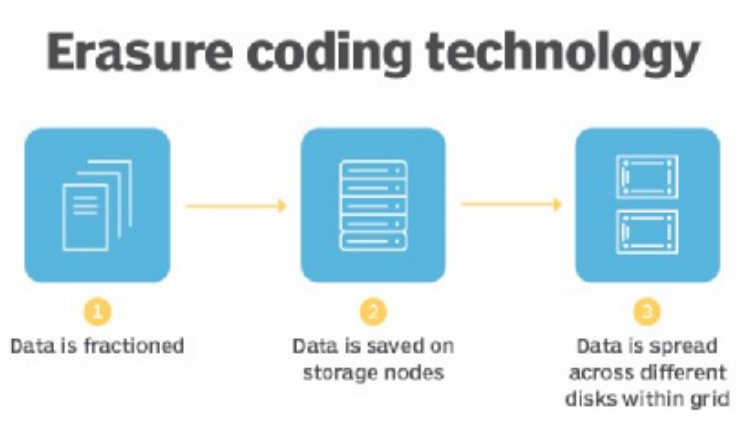

SMB Protocol Improvements

Spectrum Scale

SMB 4.13

New upgrades and security enhancements that support reliable and flexible SMB deployments.

- Updated to use Python 3.
- Improved wide links functionality. Now you decide whether to follow symlinks outside of a share path by setting wide links parameter to true.
- Deprecates Samba's original domain controller mode. Samba now uses Active directory for authentication purpose. Windows or MacOS, can join an Active Directory domain provided by Samba by the same mechanism as a Microsoft Windows Active Directory domain.

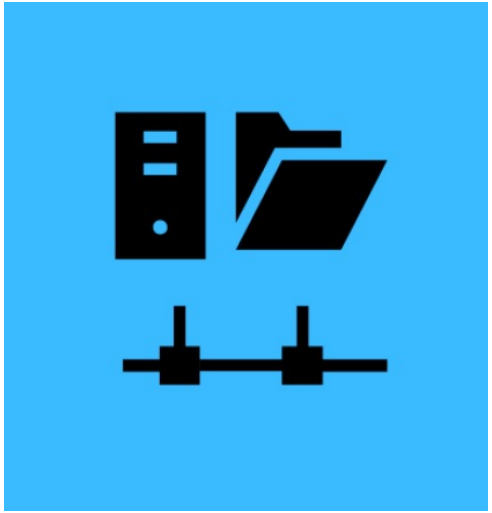

Discontinued Features

Spectrum Scale

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 5.0.5.
	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).

Deprecated Features

Spectrum Scale

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.

Log your RFE!

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

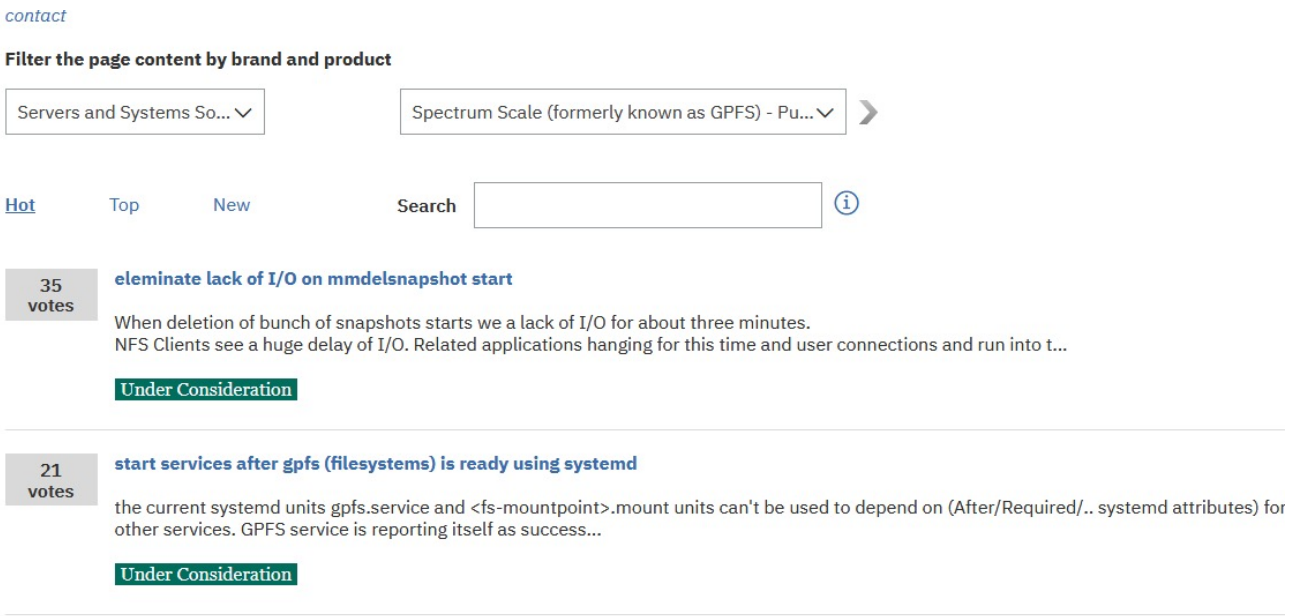

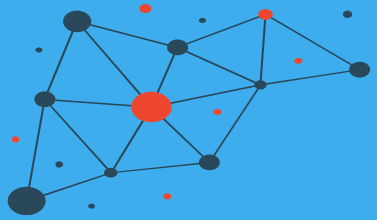

Check https://www.spectrumscaleug.org/experttalks for charts, show notes and upcoming talks

- 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
- Today:
  - May 19: What is new in Spectrum Scale 5.1.1?

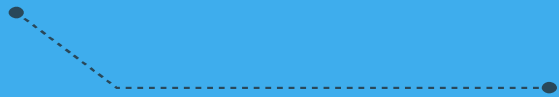

Thank you!

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

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

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

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

Spectrum Scale

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!