

IBM Spectrum Scale

Spectrum Scale Release

Update 5.0.4 / 5.0.5

Mathias Dietz

STSM, Spectrum Scale RAS

Architect



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.

Before we begin....



What Spectrum Scale version is used ?

Version 5.0.4

Version 5.0.3 / 5.0.2

Version \leq 5.0.1 or \leq 4.2.3

How many use CES/Protocols ?

SMB/NFS/Object

How many use:

Spectrum Scale GUI

REST API

Performance Monitoring (mmperfmon/Zimon) / Grafana Bridge

5.0.5 Release is planned for 2Q2020*
5.0.4 Release available since 3Q2019
(PTF2 is available)

Scale Developer Edition is out!



- Scale Developer Edition is full function but size limited
 - Derived from **DME**, all functionality available
 - Limited to 12 TBs: enough for a small test cluster
 - Available from the Scale "try and buy" page on ibm.com
- Refreshed after each Scale Mod level (first PTF)
- Free for non-production use, e.g. test, learning, upgrade prep...
 - If you have to ask, it's probably not permitted
- Not formally supported

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

5.0.4 Whats New



Redhat Enterprise Linux 8 support added (RHEL8)

TRIM support for NVMe devices

- NSD configuration and the new `mmreclaimspace` command to enable TRIM support, which reduces write amplification on solid-state devices under certain workloads

Thin provisioning support **(RPQ only)**

- Add the ability to use thinly provisioned and compressed volumes for both file system data and metadata
- Contact your sales or account team (RPQ / SCORE process) for assessment on the use of thin provisioning

Improved recovery of CCR enabled clusters

- Support `mmsdrrestore --ccr-repair` option with sudo wrapper and on Windows environments
- `mmsdrrestore -N` in CCR clusters with `adminMode=central`

5.0.4 Whats New – Network & more

Attempt to **reconnect socket** before expel

- *mmchconfig proactiveReconnect=yes*
- *Default in 5.0.4 is now “yes”*

nsdperf for “stress” testing now opensource!

- https://github.com/IBM/SpectrumScale_NETWORK_READINESS/blob/master/nsdperf.C

Improvements in the usability of **mmprotocoltrace** and **gpfs.snap**

- Snap checks if the output directory (-d) is in a file system managed by the same cluster you are running against

5.0.4 Whats New – RAS and mmhealth updates



New monitoring and **health events** for

- Monitor SSD wear level
- Firmware level of NICs (ECE only)
- Nameserver issues related to AD authentication
- ESS3000 HW Monitoring

New performance measurements for thresholds

- DiskIoLatency_read, DiskIoLatency_write, MemoryAvailable_percent

Daemon startup service waits for active RDMA port (configurable)

- Adjust the length of the timeout period. See *verbsPortsWaitTimeout verbsRdmaFailBackTCPIfNotAvailable*

Monitoring critical threads in mmfsd for stuck or overloaded critical threads

5.0.4 Whats New - Protocols - SMB updates



vfs_fruit module: enhances the support of Mac OS SMB2 clients.

Enabling results in changes how Apple particular metadata is handled that improves file browsing.

Enhancements for using **immutable files** from SMB clients: Files in an immutable fileset can now be set immutable from SMB clients by setting the READONLY attribute. The retention time can be set by modifying the LastAccessTime from a SMB client. After the retention time expires, the READONLY attribute can be cleared from an SMB client and the file can be deleted.

Spectrum Scale Release	General Availability	Samba Version	Platform Support (accum.)
4.1.1	2Q15	4.2	x86_64/RHEL7
4.2.0	4Q15	4.3	ppc64/RHEL7
4.2.1	2Q16	4.3	x86_64/SLES12
4.2.2	4Q16	4.4	ppc64le, ppc64, x86_64 / RHEL7.2
4.2.3.0 - 4.2.3.8	2Q17	4.5	x86_64, ppc64, ppc64le / RHEL 7.3, 7.4
5.0.0	4Q17	4.6	x86_64/Ubuntu 16.04.2
5.0.1	1Q18	4.6	RHEL 7.5 (5.0.1.1)
5.0.2 >= 4.2.3.9	3Q18	4.6	+ Ubuntu 18.04
5.0.3	2Q19	4.9	RHEL 7.6 (bringing mutex fixes)
5.0.4	4Q2019	4.9	RHEL8

5.0.4 Whats New - Protocols – NFS updates



Ganesha 2.7.5

Ganesha grace period is changed from 60 seconds to 90 seconds

Simplified tuning: New configuration keyword: `RPC_IOQ_THRDMAX`

- Deprecates several existing configuration parameters (`NB_WORKER`, `Dispatch_Max_Reqs`, `Dispatch_max_Reqs_Xprt`)

Enhancements to Ganesha stats (`ganesha_stats`) and Ganesha Mgr (`ganesha_mgr`) command

- Reset & Duration field for statistics and new authentication related statistics
- Manager: Memory trim options, File system cache display

Enhanced Memory Management methods

5.0.4 Active File Management (AFM) & Watch Folders IBM Spectrum Scale

afmParallelMounts - for parallel data transfers by using multiple remote mounts

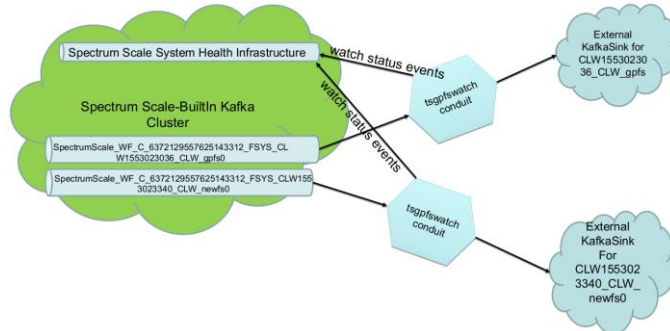
Support **dependent filesets** inside an AFM or AFM-DR fileset.

AFM and AFM DR is supported when SELinux is enabled.

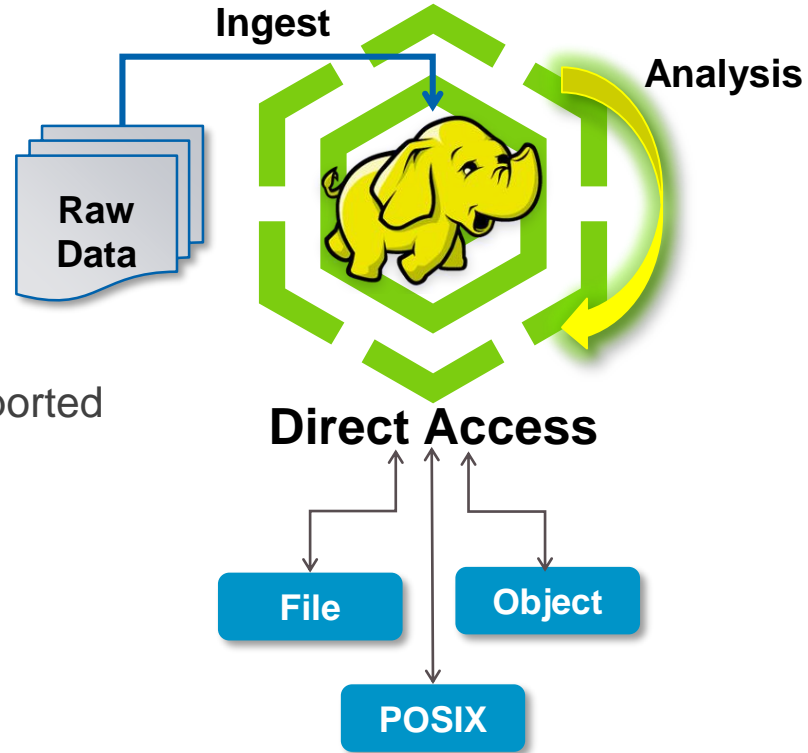


Watch Folders (mmwatch)

Support for sending events to a secondary sink when a watch is suspended



5.0.4 Whats New - Big Data & Analytics (BDA)



Support for Hortonworks Data Platform (HDP) 3.1.4

Open Source Apache Hadoop version 3.1.1 is now supported

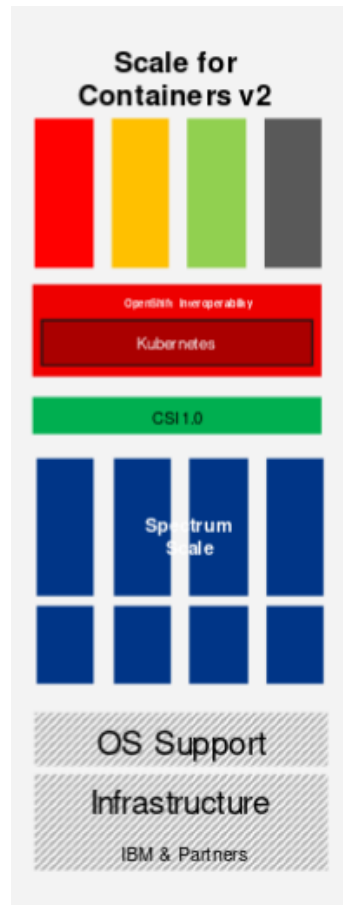


5.0.4/5.0.5 – Container support

Container Storage Interface (CSI) version 1.0

Allow a Spectrum Scale file system to surface into a pod/container running inside OpenShift 4.2

<https://github.com/IBM/ibm-spectrum-scale-csi-driver>



Spectrum Scale 5.0.5

- ***Extended Update Support Release (EUS)***
- ***Performance Improvements***
- ***Operational Efficiencies***
- ***Security & Compliance***



An Extended Update release simplifies ownership of Scale

An extended update support release allows customers to receive a fix stream without taking new features

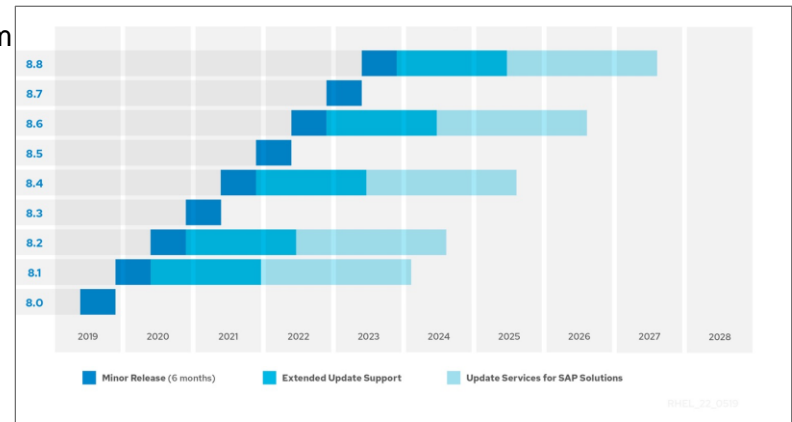
- Security and other fixes continue even after the next release
- On Linux distros, known as EUS (Extended Update Support) or LTS (Long Term Support)
- Eliminates need for extensive re-testing in order to stay secure
- Allows better dependency management both upstream (OS) and downstream (apps such as SAP or Oracle)

This time, we are doing EUS officially with V5.0.5

- You can remain current on security and other fixes without the expense of functional upgrades
- You can choose to take advantage of Versions and Releases while skipping Mod packs
- You will know even before GA how long to expect support and can plan accordingly



Compare the predictable RHEL plan



5.0.5 Outlook – Performance Improvements



mmcheckquota improved performance

- Fix slow down in mmcheckquota while traversing very very large sparse inode zero file
- Improvement for mmcheckquota for certain FSs with several independent filesets

Improve mmrestripefs performance

- Better self-tuning on the number of threads for ESSs



Speed up "ls -l" on a large directory with lots of fileset junctions

- Less time spent on lookups

Eliminate zero link count inodes in snapshots

- E.g. avoid moving unnecessary blocks when snapshot is deleted

Add a configuration option that will establish all-to-all daemon connections after mmfsd starts

5.0.5 Outlook – General Improvements

RHEL 8.1 support

Support `cp --preserve=xattr` (Will be moved to 5.0.5 PTF)

- E.g. for migrating into a 5.x "variable subblock" file system

Install / Upgrade toolkit

- Better handling short vs FQDN hostnames during install / upgrade
- ECE Rolling Upgrade from 5.0.4.3 to 5.0.5

Spectrum Scale on Z

- NFS Client support
- Active Directory Authentication for NFS
- KVM Support

5.0.5 Outlook – Health Improvements

Detect OOM killer runs

- Detect if Linux Out of memory killer has run and warn about it
- New command “*mmhealth event resolve*” to manually clear OOM event

NumaMemoryInterleave enabled but dependencies missing

- Numactl command required to control NUMA settings

NFS Monitoring improvements

- NFS Monitoring RPC Nullcheck behavior improved
- Prevent NFS from starting up when remote mount is not available



5.0.5 Outlook – Performance Monitoring

Top-k Performance Sensor


- Collect the Top 10 processes (by CPU Usage) and provide a reporting with CPU/Memory
- New command “mmperfmon report”

Improve Network sensor for Docker

- allow customer filters to filter out virtual temporary docker interfaces
- Use “mmperfmon config” command to modify the filter

Performance Threshold Improvements

- New measurement MemoryAvailable_percent
- Default threshold rules deactivated if the corresponding sensor is not configured



5.0.5 Outlook – Call Home Improvements

Show Configuration changes between two call home uploads

- New command “mmcallhome status diff”

```
[root@g5050-11 blah_old]# mmcallhome status diff --last-days 2
```

```
Active Config (modified)
  Maxfilestocache : 100010 --> 10
```

```
Cluster Data (modified)
  Maxfilestocache : 100010 --> 10
```



Enforce config options like customer name, customer ID.

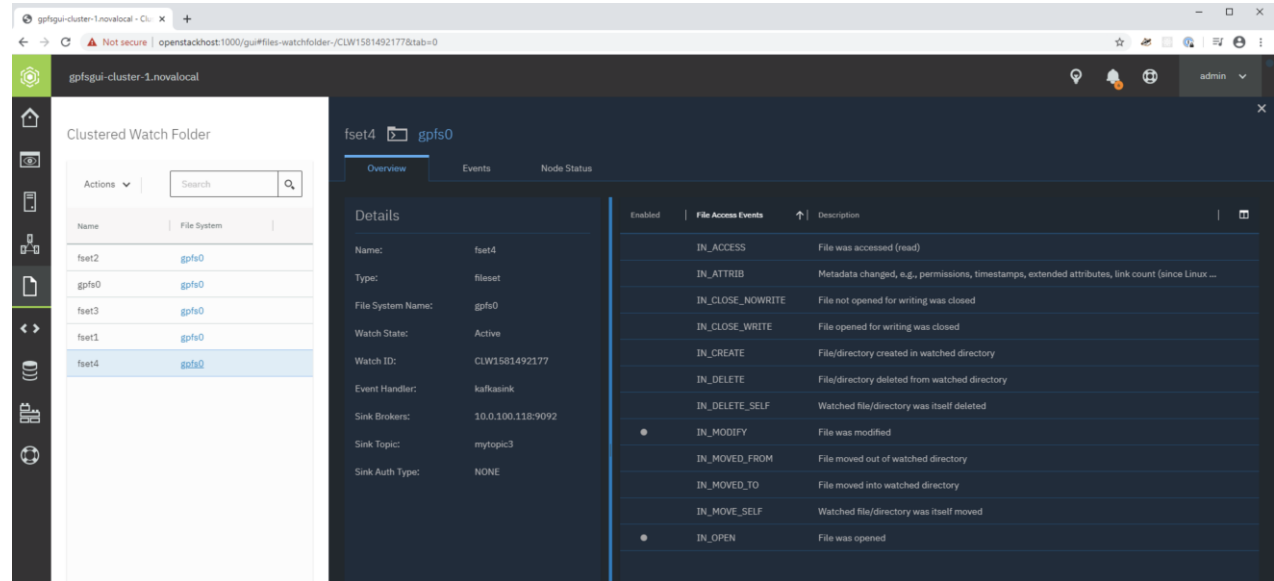
- Show health tip if missing

And more....

- Sudo wrapper support, Test connection of ungrouped nodes, add/del nodes from group

5.0.5 Outlook – GUI Improvements

Support for Watch Folders:



Display NVMe information for physical Disks that are NVMeS (for ESS and ECE)

Configure overlapping ID ranges for AD + RFC2307

5.0.5 Outlook – Protocols



SMB

- Samba 4.11
- Re-introduce recovery lock to prevent inconsistencies in case of firewall misconfiguration
- Improved MAC support (file id generation)

NFS

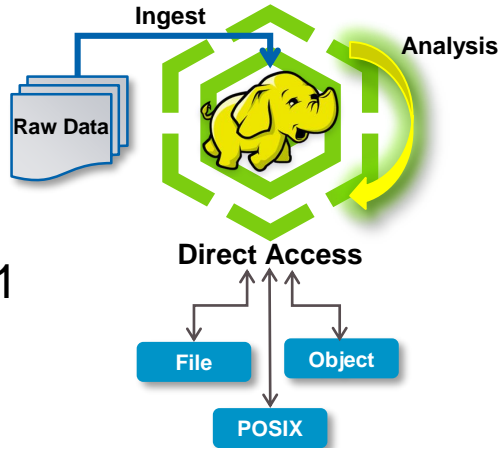
- Ganesha 2.7.5-ibm054.06

Heads-up: CES IPv6 Support under development

- Explicit Network Adapter configuration vs static marker IP address

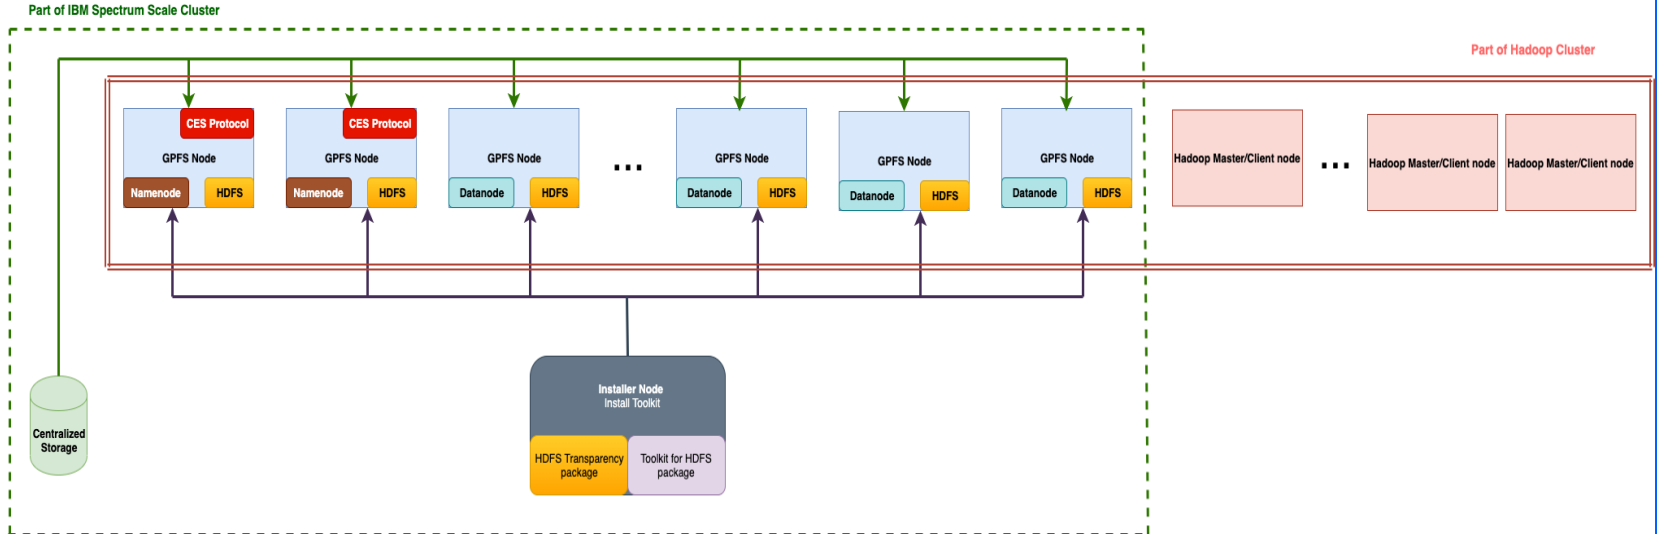
HDFS

Changes in HDFS Transparency 3.1.1



Integrates with CES protocol and IBM Spectrum Scale installation toolkit for NameNode service

- Supports Open Source Apache Hadoop distribution and RedHat Enterprise Linux operating systems.



5.0.5 Outlook – Security

TLS 1.3 Support

- Spectrum Scale GUI & REST API

Encryption

- File encryption now supports SKLM 4.0 as a Remote Key Management (RKM) server.
- Qualification of Thales DSM 6.3
- Renew expired client certificates without recreating a new client



Audit logging on fileset level

- With fileset file audit logging, you can specify a list of filesets to enable or skip from a file system.

RPM Package signing support for installation & upgrade

Healthcare Workloads

The image shows the cover of a Redbooks document. At the top right is the Redbooks logo. Below it, the title 'HIPAA Compliance for Healthcare Workloads on IBM Spectrum Scale' is displayed. The authors 'Sandeep R. Patel' and 'Sandeep Zende' are listed on the left. The central graphic is a network diagram with a prominent red node at the bottom. Below the diagram are two colored boxes: a green one with a shield icon and the word 'Security', and a purple one with the word 'Storage'. At the bottom, the IBM logo and 'Redpaper' logo are visible. A small text at the bottom right of the diagram area reads 'In partnership with IBM Academy of Technology'.

Redbooks
Research & Education

DrP Document for Review February 27, 2020 4:42 pm TRDP-0021-16

HIPAA Compliance for Healthcare Workloads on IBM Spectrum Scale

Sandeep R. Patel
Sandeep Zende

Security

Storage

In partnership with
IBM Academy of Technology


IBM Redpaper

New HIPAA Rebook has been published last week !!

Link: <http://www.redbooks.ibm.com/redpieces/abstracts/redp5591.html>

Thank you!

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](#)

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 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

IBM Early Programs website:

<https://ibm.biz/NTIPrograms>

NTI Program Interest Form:

www-355.ibm.com/technologyconnect/cna/eplInterestForm.xhtml