




IBM Spectrum Scale

What's new in 5.0.0

Mathias Dietz

Spectrum Scale RAS Architect

IBM Research and Development
in Kelsterbach/Frankfurt, Germany
mdietz@de.ibm.com



Before we start

Which Spectrum Scale version is used ?

- Version 5.0
- Version 4.2.3
- Version 4.2.x
- Version \leq 4.2

How many use CES/Protocols ?

- SMB
- NFS
- Object

How many use the Spectrum Scale GUI ?

- REST API

Highlights in 5.0

Scaling and Performance enhancements

- Improved space efficiency and performance for small files
- Improved performance over RDMA transports
- Improved metadata performance (directory locking) and reduced CPU jitter
- Improved encryption performance (GSKit V8)

Currency & OS Support

- Ubuntu 16.04.03 (incl. Protocol support and zLinux)
- RHEL 7.4 & SLES 12 SP3
- Dropped support for RHEL6, SLES11, Ubuntu 14.04 and Debian

New License model

- Data Management License

File Audit Logging

Scaling and Performance enhancements

Improved space efficiency and performance for small files

- More than 32 subblocks per block
- Requires 5.0 Cluster release level
- Only for new filesystems !
- Default block size is now 4M

Blocksize	Subblocks	Subblock size
64K/128K/256K	32	2K/4K/8K
512K/1M/2M/4M	64/128/256/512	8KB
8M/16M/32M	512/1024/2048	16KB

RDMA enhancements – improved performance over RDMA transports

- Improve performance of VERBS RDMA send and recv and scale to 1000's of nodes
- Make RDMA primary code paths mutex free and improve cache line efficiency
- NUMA aware allocation of send and recv buffers

File Audit Logging

Logs filesystem operations to a retention-enabled fileset to track user access

- Events are monitored at a filesystem level
- Events that can be captured are: **Open, Close, Destroy (Delete), Rename, Unlink, Remove Directory, Extended Attributed Change, ACL change**
- Local events from each node are placed on a multi-node message queue (Kafka) for consumption, eliminating bottlenecks and providing redundancy
- Events are logged in a parsable, JSON formatted string

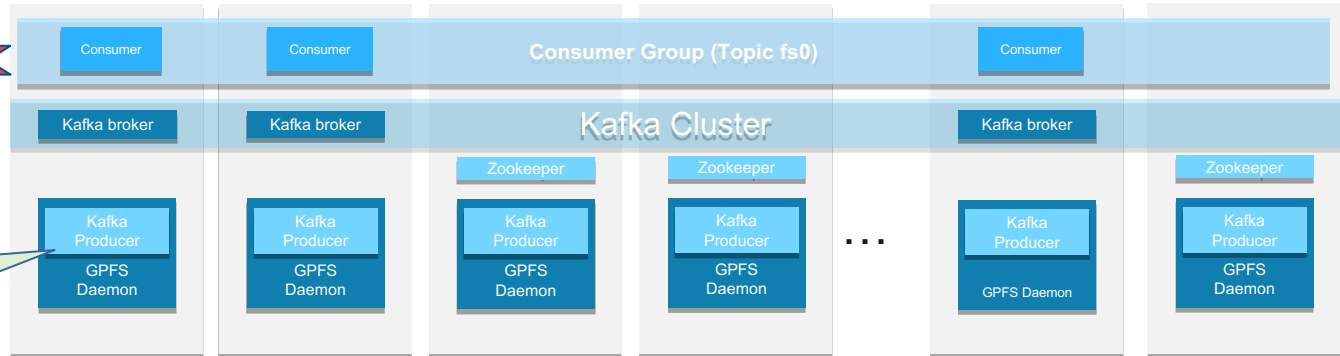
→ Join the Audit Logging Talk tomorrow 10 AM

JSON audit log

```
{
  "operation": "rm",
  "source": "fs0",
  "path": "/opt/ibm/audit-log",
  "timestamp": "2019-03-28 22:28:07.010",
  "user": "root",
  "group": "root",
  "size": 0,
  "mode": "0755",
  "type": "file",
  "attributes": {
    "acl": "user::rxwx,group::rxwx,world::rxwx",
    "owner": "root",
    "group": "root"
  },
  "error": "EEXIST",
  "message": "rm: cannot remove 'fs0': File exists"
}
```



LWE
+
librdkafka



Filesystem Updates

Usability

- Concurrent modification of preferred NSD servers
- Status/progress of a running filesystem check (`mmfsck --status-report`)
- Status of a running trace (`mmtracectl --status`)

Encryption and Compression

- Optimized compression and support for LZ4 (result in better read performance)
- LROC for encrypted files (`IrocEnableStoringClearText`)
- Improved encryption configuration (simplified IBM SKLM setup)

More enhancements

- Faster file system rebalancing (lenient vs strict round robin)
- InfiniBand addresses can be specified for RDMA transfers
- `mmnetverify` support for GPFS subnets

Protocol Support

SMB

- Update to Samba 4.6
- Detect and warn about version mismatch
- Option to use “unix primary group” in AD

NFS

- Update to Ganesha 2.5
- Modify NFS exports without impacting connected clients (Dynamic exports)

Object

- Upgrade to OpenStack Pike release
- Improved upgrade support for Object (Removes intermediate levels in upgrade paths)

General

- Ubuntu support as a CES protocol node (16.04.x)

Install Toolkit Enhancements

- Improved integration between the Installation Toolkit and ESS
 - Installation toolkit will automatically detect ESS nodes for easier deployment
- Enhanced pre-checking to find common network issues prior to deployment
- Installation and configuration of file audit logging
- Simplified callhome deployment and configuration
- Support for Ubuntu 16

Monitoring & RAS Improvements

Health & Performance Monitoring Improvements

- Warn customers about filled /var/mmfs
- Detect and display when the GPFS daemon is unresponsive
- Kafka & audit consumer monitor
- Event script callout (undocumented feature)
- User defined thresholds – priorities and additional measurements
- Significant reduction of CPU Jitter and GSKit initialization

Callhome aka Proactive Service

- Simplified setup and configuration (through Install toolkit and GUI)
- Upload of diagnostic data directly from the system (CLI/GUI)
- Added Callhome support for zLinux and Ubuntu
- Many improvements under the hood to provide the basis for proactive service

→ [Join the System Health Talk tomorrow 11AM for more details](#)

GUI & REST API Updates

Enhanced management of AFM relationships

- Ability to access remote GUIs and view AFM status
- Enhanced monitoring: View events from multiple cache sites on a home cluster

Transparent Cloud Tiering support

Health & Performance Monitoring Improvements

- Performance and capacity threshold management
- Network monitoring for both IP and RDMA transports
- Additional directed maintenance procedures to assist with failure recovery

Filesystem creation for non-ESS environments

REST API has been expanded to allow additional tasks

- Performance data collection, collect diagnostics data and manage thresholds
- Add/remove nodes from a cluster and manage node classes
- The ability to cancel jobs

→ Join the GUI & RESTAPI Talks tomorrow 9AM & 10AM for more details

AFM Updates

- File compression is now supported for files residing in AFM and AFM DR filesets
- Load balancing improvements for AFM gateways
 - Updated hash algorithm (afmHashVersion) tunable to modify how filesets are distributed across gateway nodes
- New limit on the number of parallel recoveries being run on gateway nodes
 - Can be tuned to reduce the resource requirements on a gateway node
- ILM for snapshots is now supported for AFM and AFM DR filesets
- GUI improvements to assist with AFM management

→ Join the AFM Deep Dive tomorrow at 10AM

Transparent Cloud Tiering

- Remote mounted filesystem support
 - Clients can access tiered files on a remotely mounted filesystem
- Ability to tier different filesets to different cloud containers
- Enhanced support for multiple cloud accounts and containers
- Usability improvements, updated CLI and enhanced GUI support
 - CLI has been enhanced for ease of deployment with multiple cloud accounts

→ TCT Talk is right after this slot at 11:15AM

Big Data and Analytics

- Certification with HortonWorks Data Platform 2.6
 - Certified on both Power8 and x86 platforms
 - Certified with Ambari 2.5 for rapid deployment
- Support short circuit write for better performance
- Support 2 or more file systems in one HDFS Transparency cluster
- Support local read/write bytes for Spectrum Scale nodes with internal disks
- Integration with Spectrum Conductor for Spark

→ Join the Cognitive, ML, Hortonworks Talk tomorrow 3:30PM

Disclaimer

- The information in this document is **IBM CONFIDENTIAL**.
- 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 AND TO BUSINESS PARTNERS WHO HAVE A CURRENT SIGNED NONDISCLOSURE AGREEMENT ON FILE WITH IBM. THIS INFORMATION CAN ALSO BE SHARED WITH CUSTOMERS WHO HAVE A CURRENT SIGNED NONDISCLOSURE AGREEMENT ON FILE WITH IBM, BUT THIS DOCUMENT SHOULD NOT BE GIVEN TO A CUSTOMER EITHER IN HARDCOPY OR ELECTRONIC FORMAT.**
- 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.

5.0.1 Outlook

Filesystem Enhancements

Support for end-to-end data checksums in non-ESS environments with unreliable network
Improvement in the area of Systemd integration, fsck, adding disks to full file systems

Health Monitoring Updates

Simplified performance monitoring configuration
Several improvements around CES, network monitoring and Call Home

NFS Enhancements

Automatic restart of nfs-ganesha for better error recovery
Automatic backtrace generation for better problem determination

GUI Updates

New Services Panel to manage and monitor services like SMB, NFS, GPFS, Zimon
Display performance data from remote clusters

...and more....

QUESTIONS ?

THANK YOU

LEARN MORE

Client Enablement Material

New videos, blogs and presentations:

- [https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/General%20Parallel%20File%20System%20\(GPFS\)/page/White%20Papers%20%26%20Media](https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/General%20Parallel%20File%20System%20(GPFS)/page/White%20Papers%20%26%20Media)

Watch for new blogs at:

- <https://developer.ibm.com/storage/blog/>

Turn-key Spectrum Scale VM available for download

- Try the latest Spectrum Scale enhancements
- Full functionality on laptop, desktop or server
- Incorporate external storage

Spectrum Scale Blueprints for Genomic Medicine Workloads

- Enablement and Redpapers available for these workloads