



IBM Spectrum Scale

Computing Insight UK 2015

Patrick Byrne
Cameron McAllister

Dec 2015





4.2.0 Summary

4.2.0 Summary

- GUI *
- QoS *
- Compression *
- Improved Object capabilities
- Ssh and scp wrappers
- Native Hadoop HDFS support *



GUI Demo

No GUI Demo?

Never work with children, animals or live demos

QoS

- Configurable limit for number of IOPS that maintenance tasks can use
- Can be configured at a per-storage pool level
- The IOPS being used can be viewed with `mmlsqos`
- Limited to just “maintenance” and “other” classes in 4.2.0

Compression

- Designed for cold data or write-once objects and files
- Compression is at a per-file level
 - Mmchattr for a single file
 - mmappypolicy to compress many files
- Can be done synchronously or deferred until a restripe
- Uses zlib, favouring compression over speed



Big Data & Analytics on Scale

Big Data and Analytics on Scale

Spectrum Scale is a natural fit in this space

Data locality (FPO & Hadoop Connector)

POSIX data access (Spark / Tachyon)

Massive scale out (Big Data)

Big Data and Analytics on Scale

- Current support:
 - FPO (data locality)
 - Hadoop (HDFS connector)
 - HDFS Transparency (support for HDFS:// protocol)
 - “How to” Guides
 - Watson Explorer
 - InfoSphere Streams
 - BigInsights & Ambari

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

Big Data and Analytics on Scale

- Future:
 - Big Insights
 - Cloudera
 - Hortonworks
 - Cognos
 - SPSS
- Performance tuning guidelines for:
 - Map Reduce
 - Hive
 - HBase
 - BigSQL
 - Spark

Anything missing? Speak now, or corner Cam later!



Priorities for 2016

Priorities for 2016

- *Every year we define a set of imperatives / goals / hills*
 - *Based mainly on client feedback*
 - *Address their needs*

- *Target is to achieve them within the year*
 - *Aim to build on last year's successes*

Priorities for 2016

- Documentation
- Defect backlog
- Security
- Problem Determination

Documentation

- We are looking for feedback on existing Documentation
 - What is important to Improve / add
 - Please talk to me (Cam) and let me know areas that are particularly important to you
 - And ways to improve it

Reduce the Defect Backlog

- It's not sexy..

...but is important

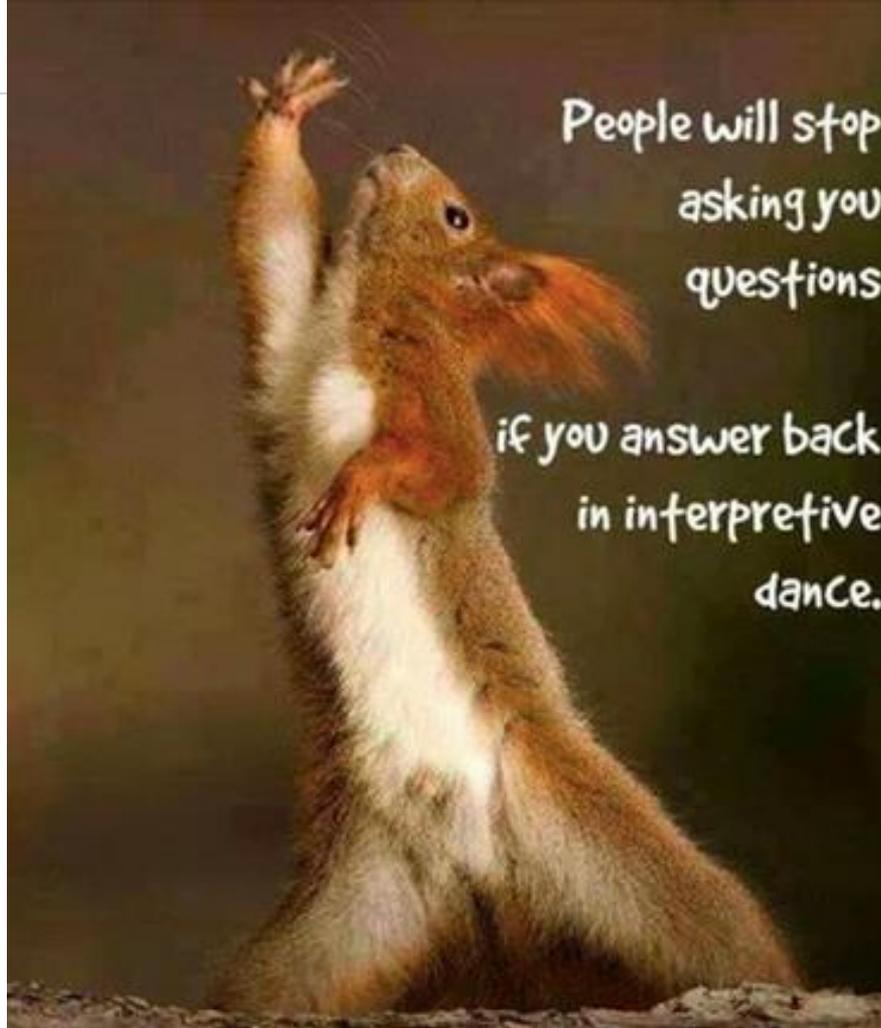
Security

- NIST-compliance
- FIPS-certified encryption
- Vulnerability scanning and security hardening
- Improvements in file encryption consumability (configuration of key server)
- Execution in an environment with restricted password-less ssh capability
- White papers: Spectrum Scale Security

Problem Determination

- The main area that needs to be improved based on feedback from users
- Three main focus areas for Problem Determination:
 - Active monitoring of the system
 - Improved documentation and guidance for fixing problems
 - Pro-active configuration checking
- Make information more easily accessible

Questions?

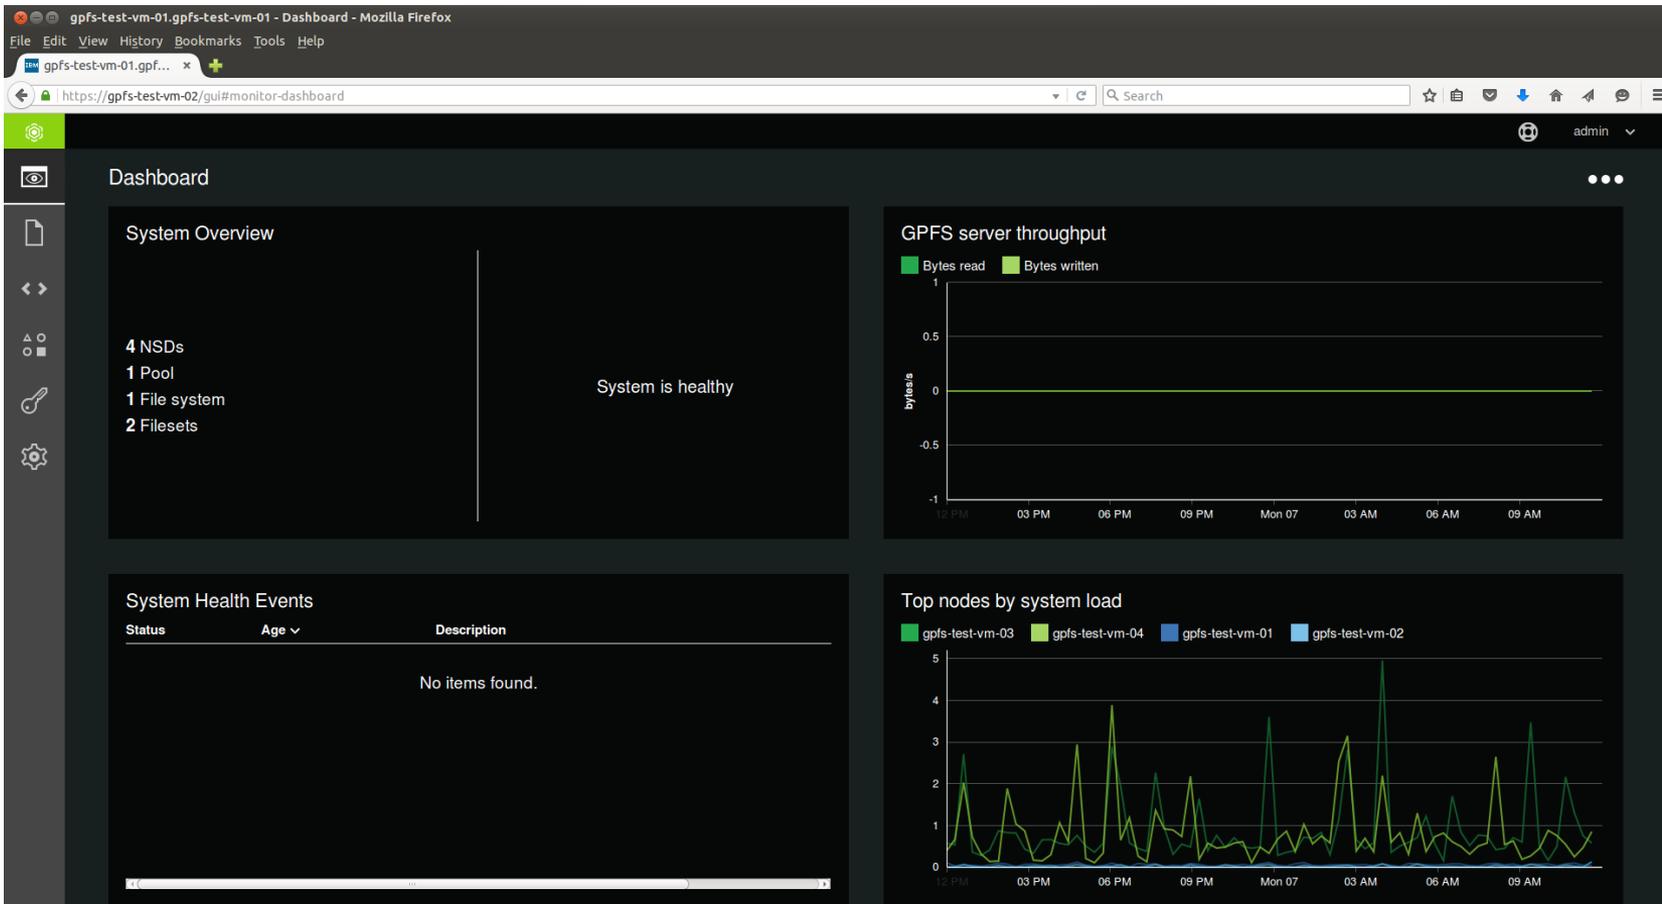


Appendix

QoS

List of tasks that are “maintenance” tasks:

- mmaddisk
- mmapplypolicy
- mmbackup
- mmcheckquota
- mmdefragfs
- mmdeldisk
- mmdelfileset
- mmdelsnapshot
- mmdf
- mmfileid
- mmfsck
- mmlssnapshot
- mmrestripefs
- mmrpldisk



gpfs-test-vm-01.gpfs-test-vm-01 - Events - Mozilla Firefox

https://gpfs-test-vm-02/gui#monitor-events

admin

Events

Next recommended action Actions All Events

Last Updated: 11:52

Search

Status	First Occurrence	Event ID	Description	Action
Information	07/12/2015 11:42:57	MS0140	The local file systems on node gpfs-test-vm-02 reached a normal level.	
Information	07/12/2015 11:42:56	MS0140	The local file systems on node gpfs-test-vm-03 reached a normal level.	
Information	07/12/2015 11:42:56	MS0140	The local file systems on node gpfs-test-vm-04 reached a normal level.	
Information	07/12/2015 11:42:56	MS0140	The local file systems on node gpfs-test-vm-01 reached a normal level.	
Information	07/12/2015 11:40:34	MS0371	The state of performance monitoring sensor service pmsensor on node "gpfs-test-vm-01" ...	
Information	07/12/2015 11:40:34	MS0371	The state of performance monitoring sensor service pmsensor on node "gpfs-test-vm-04" ...	
Information	07/12/2015 11:40:34	MS0371	The state of performance monitoring sensor service pmsensor on node "gpfs-test-vm-03" ...	
Information	07/12/2015 11:40:34	MS0371	The state of performance monitoring sensor service pmsensor on node "gpfs-test-vm-02" ...	
Information	07/12/2015 11:40:03	MS0361	The state of performance monitoring collector service pmcollector on cluster "gpfs-test-v...	
Information	07/12/2015 11:38:34	MS0520	The SMB service is healthy on node "gpfs-test-vm-04".	
Information	07/12/2015 11:38:34	MS0540	The network state is healthy on node "gpfs-test-vm-04".	
Information	07/12/2015 11:38:33	MS0530	The object service is healthy on node "gpfs-test-vm-04".	
Information	07/12/2015 11:38:33	MS0510	The NFS service is healthy on node "gpfs-test-vm-04".	
Information	07/12/2015 11:38:33	MS0500	The Hadoop DFS state is healthy on node "gpfs-test-vm-04".	

gpfs-test-vm-01.gpfs-test-vm-01 - Performance - Mozilla Firefox

https://gpfs-test-vm-02/gui#monitor-performance-c0=systemresources-node/d0=1449402881808,1449489281808/m0=cpu-idle/o0=gpfs-test-vm-01,gpfs-test-vm-02,gpfs

admin

Performance

Node CPU / Usage / Idle

gpfs-test-vm-01 gpfs-test-vm-02 gpfs-test-vm-03 gpfs-test-vm-04

6 - 7 Dec 2015 5 minutes 1 hour 1 day 1 week all

Cluster: gpfs-test-vm-01.gpfs-test-vm-01 / IBM Spectrum Scale™ Client / Throughput

Bytes read Bytes written

6 - 7 Dec 2015 5 minutes 1 hour 1 day 1 week all

Resources

Resource type: System Resources

Aggregation level: Network

Name

gpfs-test-vm-01	NFS
gpfs-test-vm-02	SMB
gpfs-test-vm-03	Object
gpfs-test-vm-04	

Metrics

Load: Load

CPU

Rates: Interrupts Context switches

Usage: User Nice System I/O wait Idle

Interrupts: Hardware Software

Memory

General: Free Buffers Cached Swap cached Active Inactive Active anon Inactive anon

Apply Close

gpfs-test-vm-01.gpfs-test-vm-01 - Snapshots - Mozilla Firefox

gpfs-test-vm-01.gpf... x

https://gpfs-test-vm-02/gui#files-snapshots

admin

Snapshots

Create snapshots or snapshot rules to automate snapshot creation and retention.

Actions

Name	Remote Cluster
fs0	
root	
object_fileset	

Create Snapshot

Manual Schedule

Path: /ibm/fs0/object_fileset

Select one or more

Name	Path	Freq

Create Snapshot Rule

Name:

Frequency:

Minute:

Retention:

Hours: Days: Weeks: Months:

Prefix:

gpfs-test-vm-01.gpfs-test-vm-01 - Performance - Mozilla Firefox

https://gpfs-test-vm-02/gui#monitor-performance-c0=systemresources-node/d0=1449402881808,1449489281808/m0=cpu-idle/o0=gpfs-test-vm-01,gpfs-test-vm-02,gpfs

admin

Performance

Node CPU / Usage / Idle

gpfs-test-vm-01 gpfs-test-vm-02 gpfs-test-vm-03 gpfs-test-vm-04

6 - 7 Dec 2015 5 minutes 1 hour 1 day 1 week all

Cluster: gpfs-test-vm-01.gpfs-test-vm-01 / IBM Spectrum Scale™ Client / Throughput

Bytes read Bytes written

6 - 7 Dec 2015 5 minutes 1 hour 1 day 1 week all

Resources

Resource type: System Resources

Aggregation level: Network

1 month 1 year all

Name	Resource
gpfs-test-vm-01	NFS
gpfs-test-vm-02	SMB
gpfs-test-vm-03	Object
gpfs-test-vm-04	

Metrics

Load: Load

CPU

Rates: Interrupts Context switches

Usage: User Nice System I/O wait Idle

Interrupts: Hardware Software

Memory

General: Free Buffers Cached Swap cached Active Inactive Active anon Inactive anon

Apply Close

gpfs-test-vm-01.gpfs-test-vm-01 - Information Lifecycle - Mozilla Firefox

File Edit View History Bookmarks Tools Help

gpfs-test-vm-01.gpfs... x

https://gpfs-test-vm-02/gui#files-ilm

admin

Information Lifecycle

Create, manage, and delete policies that manage automated tiered storage of information.

Active Policy Policy Repository

File System: fs0

Pool **system** (29% used) 4.0 GiB

Rules Add Rule

1. Placement **Example** (*)

⚠ There are some issues with the policy

Apply Changes Cancel Changes

Placement **Example** Policy Text

Enabled

Pool: !

Placement scope: ? Entire File System

Placement criteria:

Place files into pool if any of these criteria is met:

- - - + -

Delete Rule

gpfs-test-vm-01.gpfs-test-vm-01 - NFS Exports - Mozilla Firefox

gpfs-test-vm-01.gpf... x

https://gpfs-test-vm-02/gui#protocols-nfs

admin

NFS Exports

Create and manage NFS exports.

+ Create Export Actions

Path File System Access Type

Create Export

Path: **!**

Owner: Owing group:

Add NFS client. Reorder the clients as needed.

+ Add NFS Client Actions

Name or IP Address	Access Type	Root Squashing	NFS Protocol
No items found.			