



outthink
limits

What's New in Spectrum Scale

User Group Meeting NYC, Sep 2017
Steve Duersch, IBM


What's New by Scrum Team

Agile organization into 14 scrums

- Changed from waterfall in 2015
- One month iterations
- Some scrums are all co-located
- Combination of dev, test, docs
- Two cross-scrum Guilds



File System Core



- Information lifecycle management for snapshots – allows snapshot data to be migrated to another storage pool
- *mmchqos* improvements – assign IOPS to individual nodes, to a node class, to a list of nodes in a text file, or to a remote cluster. Display detailed stats on IOPS for QoS programs.
- *mmfsck* displays a summary of errors that were found with the severity of each error
- GPFS log time stamp with time zone information (ISO 8601)
- *mmkeyserv* improvements– get a fresh certificate from an Remote Key Management server without rebuilding the connection

Coming soon

- More than 32-sub-blocks per block (4M default blocksize for new file systems)
- Substantial metadata ops rate improvements for shared directory (file create, etc)
- Ability to change nsd servers with the file system mounted

Cloud and Object Scrum

- Openstack Swift and Keystone packages upgraded to Mitaka release
- Open source project to auto-tag objects with appropriate metadata using cognitive services (available on GitHub)
- Object migration tool—migrate data from compliant cloud storage to object storage (currently with LBS and presales as a service tool)
- Juju charms for automated deployment
- Transparent Cloud Tiering improvements
- Amazon AWS Quick Starts Program (90 day trial)
 - <https://aws.amazon.com/quickstart/architecture/ibm-spectrum-scale/>

Coming soon

- Object support on Ubuntu
- Improved upgrade support for Object
 - Removes intermediate levels in upgrade paths
 - Simplifies upgrades on clusters running the object protocols



Data Protection and Efficiency

- White papers
 - Varonis file audit logging
 - Bulk antivirus scanning with Symantec AV

Coming soon

- Optimized compression and support for LZ4
- Lightweight events based file audit logging
 - Events that can be captured are: Open, Close, Destroy (Delete), Rename, Unlink, Remove Directory, Extended Attributed Change, ACL change
 - Allows auditing all protocols and native GPFS file access
 - Events are logged in a parsable, JSON formatted string



Active File Management

- Monitoring for AFM with mmdiag, system health and GUI
- Recovery throttling on Gateway node with --afmMaxParallelRecoveries
- Distribution of filesets across Gateway Nodes with --afmHashVersion=3
- Optimized AFM fileset recovery to avoid repeated recovery failure

Coming soon

- Better distribution of filesets across gateway node with --afmHashVersion=4
- Scale testing of AFM DR 100 filesets having 1B files across filesets
- Compression support
- Support for snapshot ILM



Deployment



- Extended platform architecture support to Power LE
- Heterogeneous OS clusters
- Enhanced problem determination: pre-checks including checks for base OS repository setup, base software requirements, authentication prerequisites and error reporting to console for faster problem resolution
- Config populate: Install toolkit config retrieval from existing clusters (allowing for seamless transition to further operations with the install toolkit usage on clusters created without it)
- Cluster upgrade enhancements: extended upgrade pre-checks, detection of GPFS kernel module unload errors and guidance on host reboot and Spectrum scale upgrade when in an LTFS environment

Coming Soon

- Extended operating system support (Ubuntu 16.04 x86, SLES12 x86, RH Power LE)
- Improved deployment integration with ESS: the ability to detect ESS nodes (EMS and IO) and ensures validation of permitted operations when adding protocol/client nodes to an ESS cluster.
- File Audit Logging installation and configuration
- Callhome configuration enabled by default
- Cumulative object upgrade support through install toolkit
- Further network connectivity pre-checks: passwordless SSH validation from admin node and validate /etc/hosts file formatting

File (NFS and Samba)

- Netgroup caching – performance improvements by resolving deeply nested net groups
- Failover (write verifier) – avoids data loss on NFS failover
- Option to use file system blocksize instead of fixed 1M in Ganesha config
- Improvements to debugging tools: *mmadquery* (Active Directory) and *mmprotocoltrace*

Coming soon

- Ganesha 2.5 and Samba 4.6
- Ubuntu 16.04 support as a CES protocol node
- Dynamic modification of NFS exports
 - Add or remove a host from an NFS export without impacting other hosts



System Health



- More monitored services: GUI, GNR, Transparent Cloud Tiering
- Performance monitoring bridge for Grafana
- *mmhealth* changes: improved performance and reduced impact on system load, show date and time of last state change, detail info for events, notification thresholds, AFM, RDMA
- *mmnetverify* – can be used before or after a cluster is created for improved debugging. Tests include address checks, ping tests, remote shell and file copy tests, time-date checks, TCP connection checks, message size tests, bandwidth tests, and flooding tests.
- New TIPS “event” for *mmhealth* (ie `gpfs_pagepool_small`)
- Ability to control the health monitoring frequency (tradeoff between failure detection time vs resource consumption)

Coming soon

- Monitoring of critical threads
- Time sync checks
- *mmcallhome* improvements: upload snap data to existing PMRs, semi-automatic PMR creation

GUI

- New features for ReSTful API v2:
 - Query performance data
 - Manage thresholds
 - Manage nodeclasses
 - Add and remove nodes
 - Create and delete peer snapshots
 - Cancel jobs
- New GUI Panels and Function
 - AFM Monitoring – Cross cluster
 - Network Monitoring (IP and RDMA)
 - Threshold Management
 - File System Creation on existing NSDs
 - Detailed display for Filesets
 - Detailed display for Pools
 - Mount management
 - Transparent Cloud Tiering monitoring
 - Health Tips
 - Configuration of call home

File System Creation

Create File System

Name

Define Storage Pools

Define Replication Policy

Configure Pools

- system NSDs
- system Failure Groups
- system NSD Order
- data NSDs**
- data Failure Groups
- data NSD Order

Configure Properties

Configure Mounting

Summary

Select NSDs for data Pool

Choose the NSDs that provide **data** storage to the **data** pool.

Hide NSD filter ^

Display NSDs that meet any of these criteria:

- - - +

<input checked="" type="checkbox"/>	NSD	Capacity	Disk Class	Storage System	Servers
<input checked="" type="checkbox"/>	disk3	10.00 GiB		system2	home-11.localnet.com (+2 mo
<input checked="" type="checkbox"/>	disk4	10.00 GiB	SATA	system2	home-11.localnet.com (+2 mo

A wizard guides users through Pool configuration, NSD selection, Failure group assignment, NSD ordering, Mount rules, Tuning parameters

◀ Back Next ▶ Cancel

Linux on z Systems


- Ubuntu 16.04 support
- 9620 nodes in a cluster
- Geographically Dispersed Parallel Sysplex (GDPS)
- Remote cluster mount
- Heterogeneous cluster
- Spectrum Scale encryption

Coming Soon

- *mmcallhome* support



Big Data and Analytics



- Certification with HortonWorks Hadoop Data Platform 2.6
 - Certified on both Power and x86 platforms
 - Certified with Ambari 2.5 for rapid deployment
 - For further details, see the Redpaper “Hortonworks Data Platform with IBM Spectrum Scale: Reference Guide for Building an Integrated Solution”
 - <https://www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/redp5448.html?Open>
 - Solution Brief “Hortonworks Data Platform on IBM Power Systems for Financial Service”
 - <https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=POS03163USEN>

Coming Soon

- Migration support from IBM BigInsights to Hortonworks HDP
- Performance improvements to HDFS Transparency Connector

Real Fast

Releases 4.2.2 & 4.2.3

- Improved *mmrestripefs -b* performance; 20% improvement (depending on configuration, as measured in lab)
- *mmfsck* performance improvements; 50% improvement during inode scan phase of *mmfsck* (depending on configuration, as measured in lab)
- 4.2.3 performance improvements as high as 35% with NFS protocol and 13% with SMB protocol (depending on configuration, as measured in lab)

Coming Soon

- Minimize the performance jitter impact of Spectrum Scale monitoring services
- Improved file metadata performance
- Improved file data encryption performance
- Further *mmrestripefs -b* performance improvements
- `fileSizeMutex` enhancements enabling PIT operations like *mmrestripefs* to be executed in parallel with creation/deletion of snapshots as well as inode expansion.
- Performance improvements with NFS and SMB protocols.



Client Adoption

Charged with making the client experience better. Actively creating blogs and other collateral to allow for more informed customers.

- Trial VM for demo'ing new function
- New videos, red papers, blogs and presentations based on field inputs and experience
 - [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)

Real World

- Conducts customer-like, end to end testing: pulling in the functions from all scrums and making sure they all work together
- Use of Load Dynamix driver to simulate industry workloads: bio-genetics, medical imaging, archive, video surveillance, etc
- Patented home grown test tools with a focus on data integrity
- Genomics real world workloads
- Utilizes field feedback to address field escapes and unique client requirements
- Boundary testing: cinder volumes, max connections, max exports, etc
- Long running workloads
- Integrated solution testing: Spectrum Archive and Protect, TCT

Other teams in Spectrum Scale



- **GNR Scrum** – Puneet will present “What’s New” for this scrum
- **Build Guild** – Cross scrum team that works on concerns of packaging
- **Test Guild** – Cross scrum team that works to improve the quality of testing across the product

Thank you!



Legal notices

Copyright © 2017 by International Business Machines Corporation. All rights reserved.

No part of this document may be reproduced or transmitted in any form without written permission from IBM Corporation.

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or program(s) described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectually property rights, may be used instead.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER OR IMPLIED. IBM LY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. IBM makes no representations or warranties, ed or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 1 0504- 785
U.S.A.

Information and trademarks

IBM, the IBM logo, ibm.com, IBM System Storage, IBM Spectrum Storage, IBM Spectrum Control, IBM Spectrum Protect, IBM Spectrum Archive, IBM Spectrum Virtualize, IBM Spectrum Scale, IBM Spectrum Accelerate, Softlayer, and XIV are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

IT Infrastructure Library is a Registered Trade Mark of AXELOS Limited.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

ITIL is a Registered Trade Mark of AXELOS Limited.

UNIX is a registered trademark of The Open Group in the United States and other countries.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

This presentation and the claims outlined in it were reviewed for compliance with US law. Adaptations of these claims for use in other geographies must be reviewed by the local country counsel for compliance with local laws.

Special notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Transparent Cloud Tiering Coming Soon

- 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

OpenStack Mitaka Release

- Build Mitaka packages for Swift, SwiftonFile and Swift3 for RHEL platforms
- Used Version: Mitaka
 - Swift 2.7.3
 - Keystone 9.3.1
 - Swift3 v1.10.2
 - Swiftonfile v.2.7.2

Object Migration Tool

- Object Migration Tool helps migrating objects from different storage cloud to Spectrum Scale object store based on OpenStack Swift
- Currently it supports migration of:
 - Buckets & Objects from a Cloud Storage which supports S3 API
 - Containers & Objects from a Cloud Storage which support OpenStack Swift API
 - Files from local/remote Filesystem. Mainly we have tested this by migrating files from xfs/ext4 filesystem to Spectrum Scale Object storage, but it can be used with any network filesystem also
- The tool can be extended to support migration from other provides like Azure, using libcloud APIs
- The tool is not being released with Spectrum Scale but will be delivered to LBS. Note we don't ship the tool so it's a best effort support model.

Object Migration Tool – Advantages and Limitations

▪ Advantages

- Currently Tool can migrate objects from Source like S3, Swift and Filesystem. We can extend the tool to support Cloud Storage like Azure, Atmos etc using libcloud API.
- Support for data as well as meta-data migration.
- Tool Supports migrating Swift and S3 ACL.

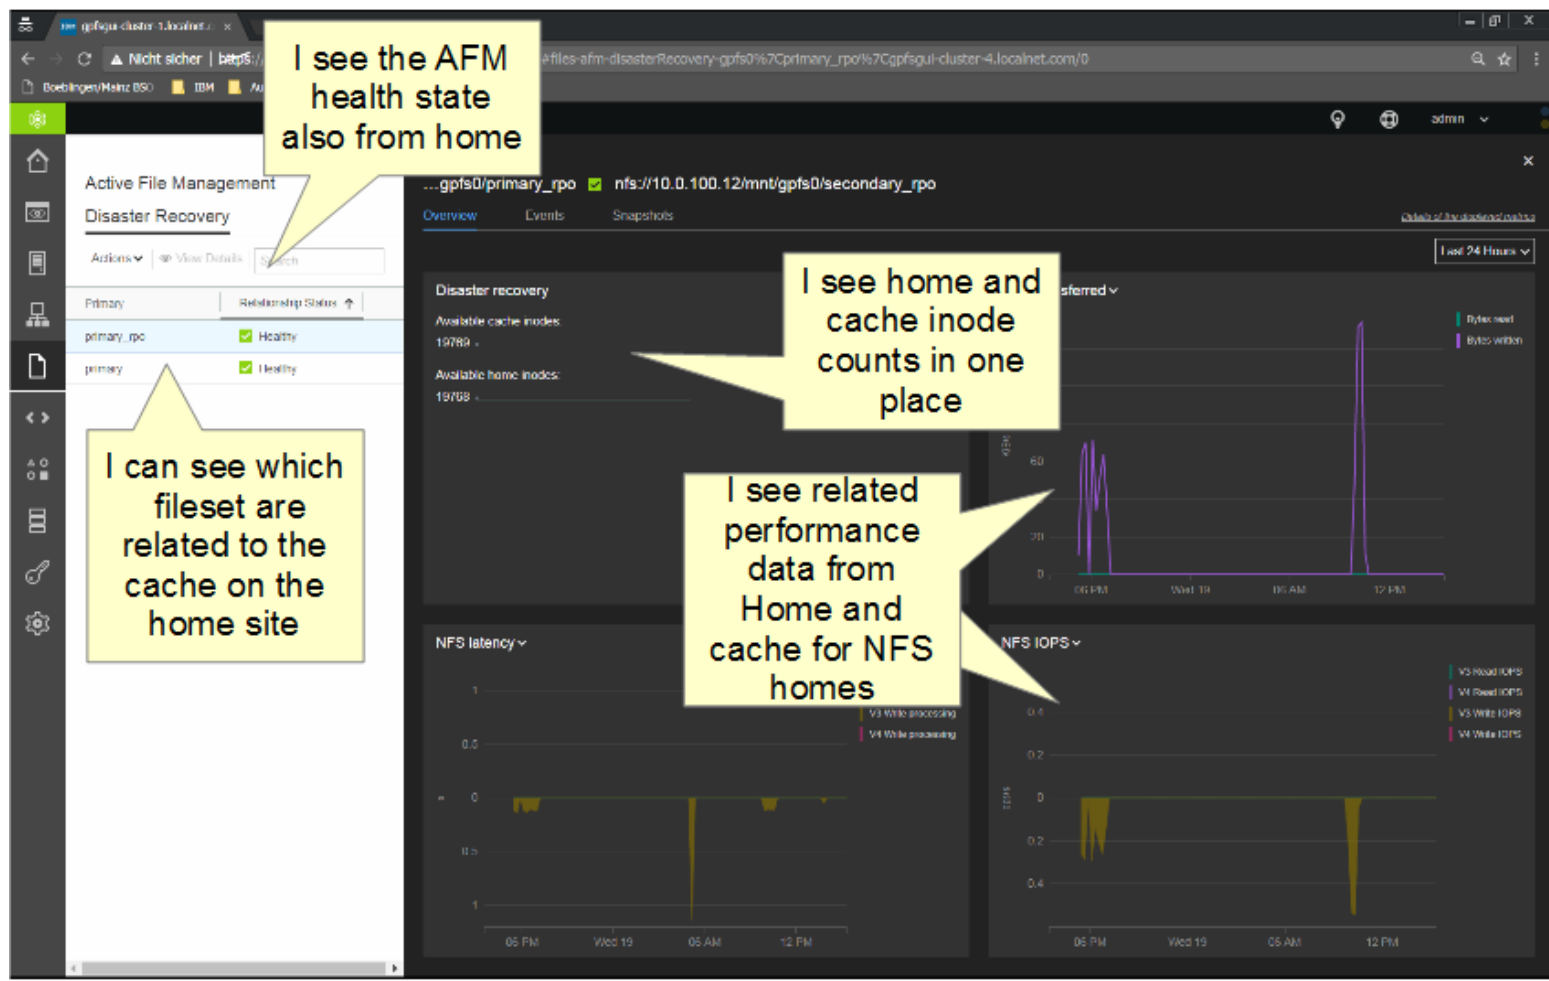
▪ Limitations

- The tool can only migrate to Spectrum Scale and not to IBM COS
- Currently username and passwords to access remote cloud are stored in clear-text format in container meta-data. The tool cleans up this meta-data after migration.
- The tool migrates object data, user meta-data and ACL. Which means after migration certain other features which S3 supports like versioning etc will not work.
- S3 tags are not migrated

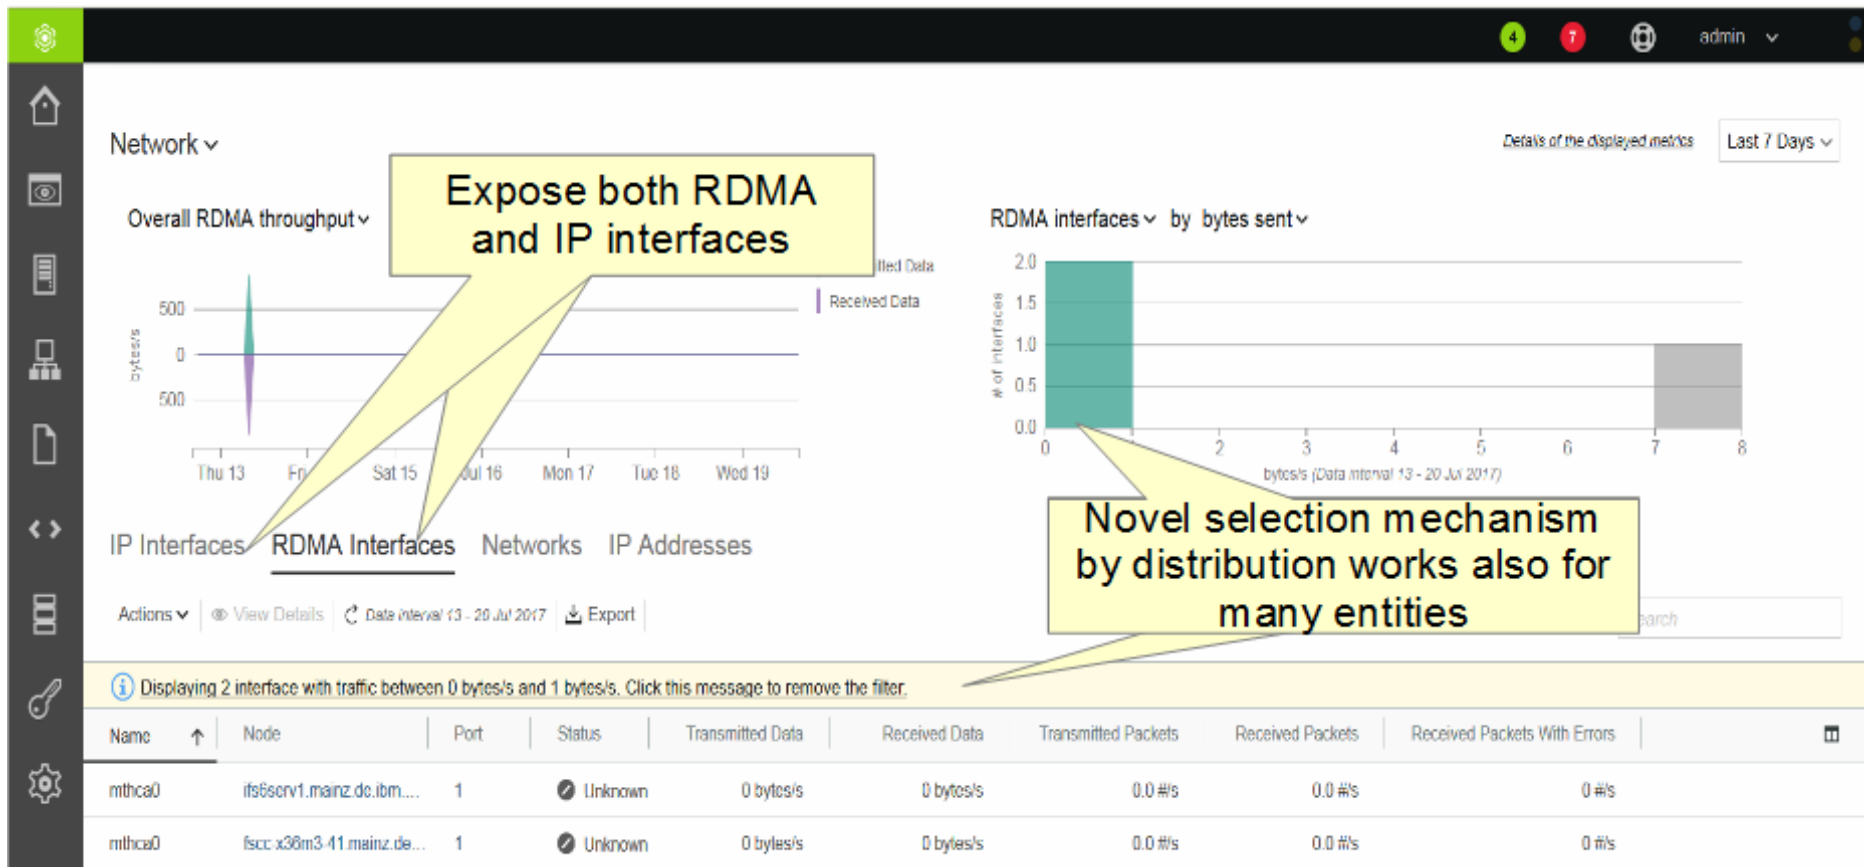
Juju Charms for Spectrum Scale Deployment

- What is Juju Charm
 - *Juju* is an open source, application and service modelling tool from Ubuntu that helps you deploy, manage and scale your applications on any cloud.
 - *Charms* are software components that contain all the instructions necessary for deploying and configuring cloud-based applications.
 - Juju solution for OpenStack is a fully integrated and optimized combination of the latest release of Ubuntu Server and the latest release of OpenStack
- Developed Juju Charms for Spectrum Scale deployment & Storage drivers' configuration and integrated to work with base OpenStack charms for a highly automated deployment of OpenStack configured to use Spectrum Scale as the storage backend for Cinder, Glance and Nova
- Enhanced GPFS Cinder driver for Spectrum Scale/GPFS to work in an Ubuntu OpenStack environment where Cinder/Glance etc services are deployed in LXD
- Charms are not tied to a particular Spectrum Scale release but are publicly made available in the online “*Charm store*”
- Detailed deployment guide with instructions is coming up from Canonical

Monitoring AFM Cross Cluster



Network Monitoring



GUI Threshold Management

Metric, Filter and grouping selection powered by zimon metadata Scan

The screenshot displays the 'Create Threshold' dialog box with the following configuration:

- Metric category:** Network
- Metric name:** Bytes received
- Name:** netdev_bytes_r_node_sum_custom
- Filter by:** Adapter (with an 'Add Filter' button)
- Group by:** Node
- Warning level:** 20 MiB
- Warning message:** Warning message to be displayed if the threshold were triggered
- Error level:** 100 MiB
- Error message:** Warning message to be displayed if the threshold were triggered
- Aggregator:** Sum
- Sensitivity:** 15 Minutes
- Hysteresis:** 20 %
- Direction:** High

The graph shows three data series: home-11.localnet.com (green), home-12.localnet.com (purple), and home-13.localnet.com (orange). A red dashed line indicates the warning threshold at 20 MiB. A yellow callout points to the graph with the text: 'Visualize metrics, groupings, filters, hysteresis in an intuitive way'. The 'OK' and 'Cancel' buttons are at the bottom right.

GUI Threshold Monitoring

The screenshot displays a monitoring interface with a sidebar on the left containing navigation icons. The main content area is titled "Thresholds" and features a table of threshold configurations. The "cpu_system_custom123" threshold is selected and highlighted in blue. To the right, the "Overview" tab for this threshold is active, showing a line graph of performance data over time. The graph has a y-axis labeled "iF" ranging from 0 to 40 and an x-axis showing dates from Jul 11, 2017, to Jul 19, 2017. A vertical orange line marks a specific time point on Jul 19, 2017, at 2:05 PM. Below the graph, the "Occurrences" section shows a total count of 2, with the first occurrence at 7/19/17, 2:05 PM and the latest at 7/19/17, 2:15 PM. The "Related Events" section shows a warning event at 7/19/17, 2:05 PM, stating: "The value of cpu_system for the component(s) cpu_system_custom123/gpfsq-13.novalocal exceeded threshold warning level 20.0 defined in cpu_system_custom123." A yellow callout box with a pointer to the event time on the graph contains the text: "Relate threshold health events to performance data".

Thresholds

Name	Target Type
a_new_threshold	Generic
cpu_system_custom123	Generic
cpu_system_node11_spatial_r...	Generic
DataCapUtil_Rule	Pool
DataPool_capUtil_80465	Pool
Fileset_Inode_49743	Fileset
InodeCapUtil_Rule	Fileset
MemFree_Rule	Generic
MetaDataCapUtil_Rule	Pool

cpu_system_custom123

Overview Events Properties

Help Topic: Thresholds
Knowledge Center
About

gpfsgui-11.novalocal
gpfsgui-12.novalocal
gpfsgui-13.novalocal

40
30
20
10
0

01:15 01:30 01:15 02:PM 02:15 02:30 02:45 03:PM 03:15 03:30

Jul 11, 2017, 4:13:30 PM Jul 19, 2017, 3:43:46 PM

Occurrences

Total Count: 2

First: 7/19/17, 2:05 PM

Latest: 7/19/17, 2:15 PM

Related Events

7/19/17, 2:05 PM

The value of cpu_system for the component(s) cpu_system_custom123/gpfsq-13.novalocal exceeded threshold warning level 20.0 defined in cpu_system_custom123.

Relate threshold health events to performance data

File System Creation

Create File System

Name

Define Storage Pools

Define Replication Policy

Configure Pools

- system NSDs
- system Failure Groups
- system NSD Order
- data NSDs**
- data Failure Groups
- data NSD Order

Configure Properties

Configure Mounting

Summary

Select NSDs for data Pool

Choose the NSDs that provide **data** storage to the **data** pool.

Hide NSD filter ^

Display NSDs that meet **any** of these criteria:

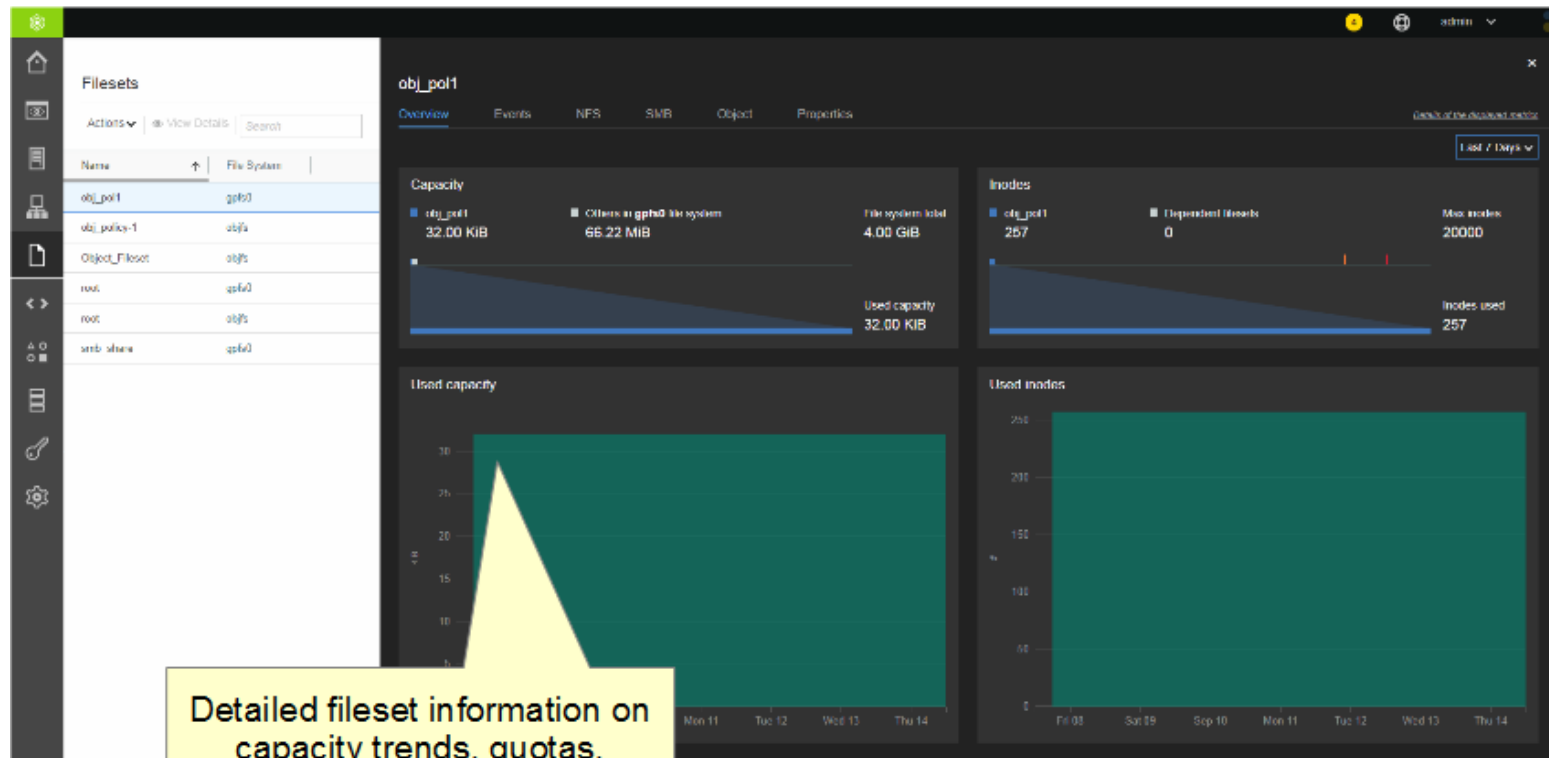
- - - +

<input checked="" type="checkbox"/>	NSD	Capacity	Disk Class	Storage System	Servers
<input checked="" type="checkbox"/>	disk3	10.00 GIB		system2	home-11.localnet.com (+2 mo
<input checked="" type="checkbox"/>	disk4	10.00 GIB	SATA	system2	home-11.localnet.com (+2 mo

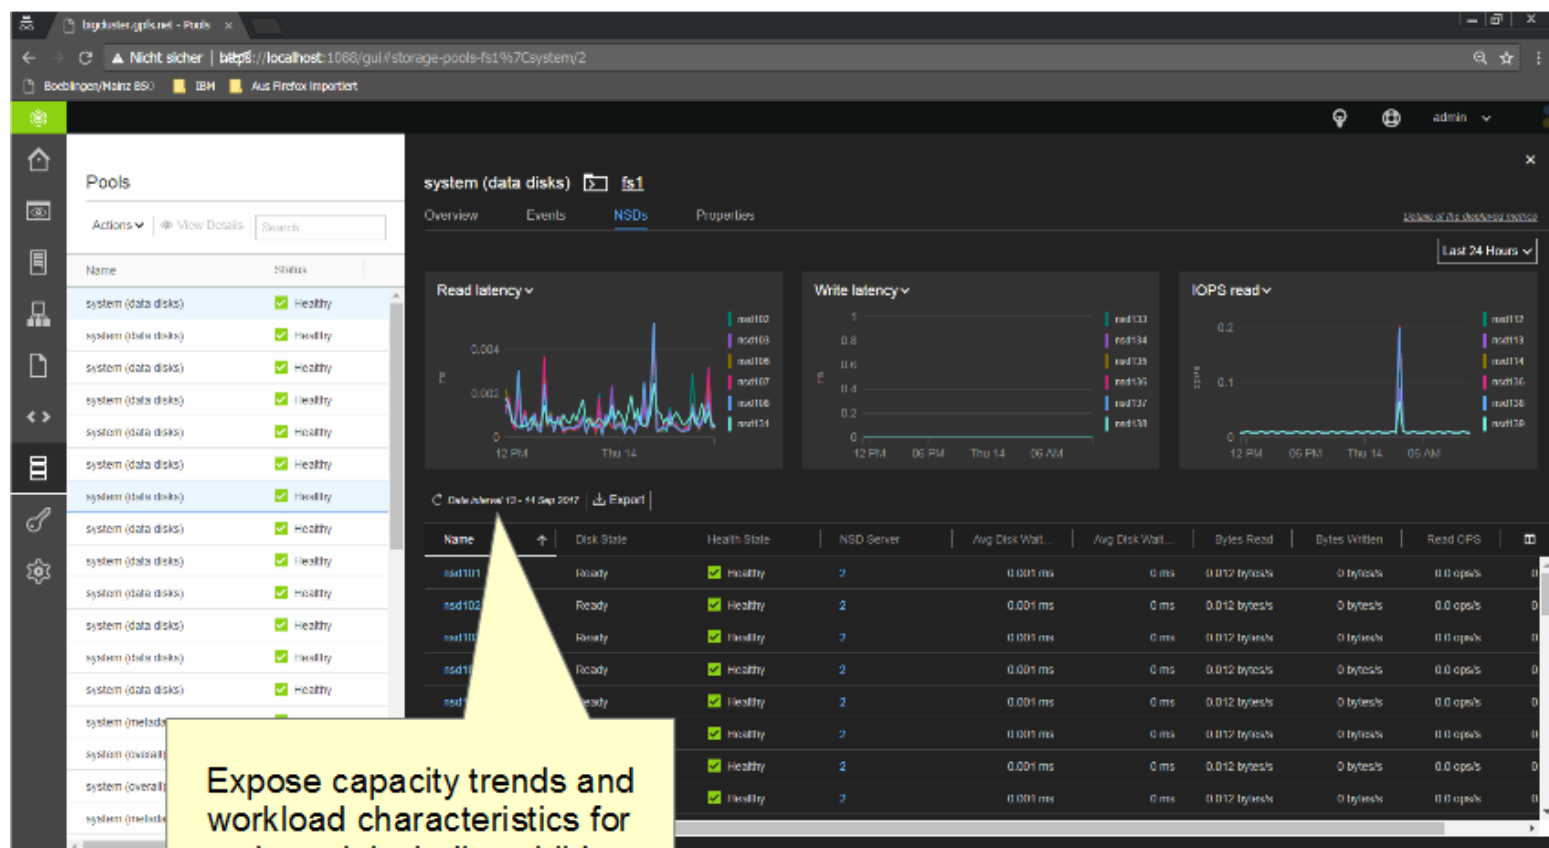
◀ Back Next ▶ Cancel

A wizard guides users through Pool configuration, NSD selection, Failure group assignment, NSD ordering, Mount rules, Tuning parameters

File Set Details



Pools Details



Expose capacity trends and workload characteristics for each pool, including children like NSDs

Mount Management

The screenshot displays a storage management interface for a file system named 'gpfs0'. A modal dialog box titled 'Prevent File System Mounts' is open, showing a progress indicator for a task that is 100% complete. The dialog contains a list of commands and their execution times:

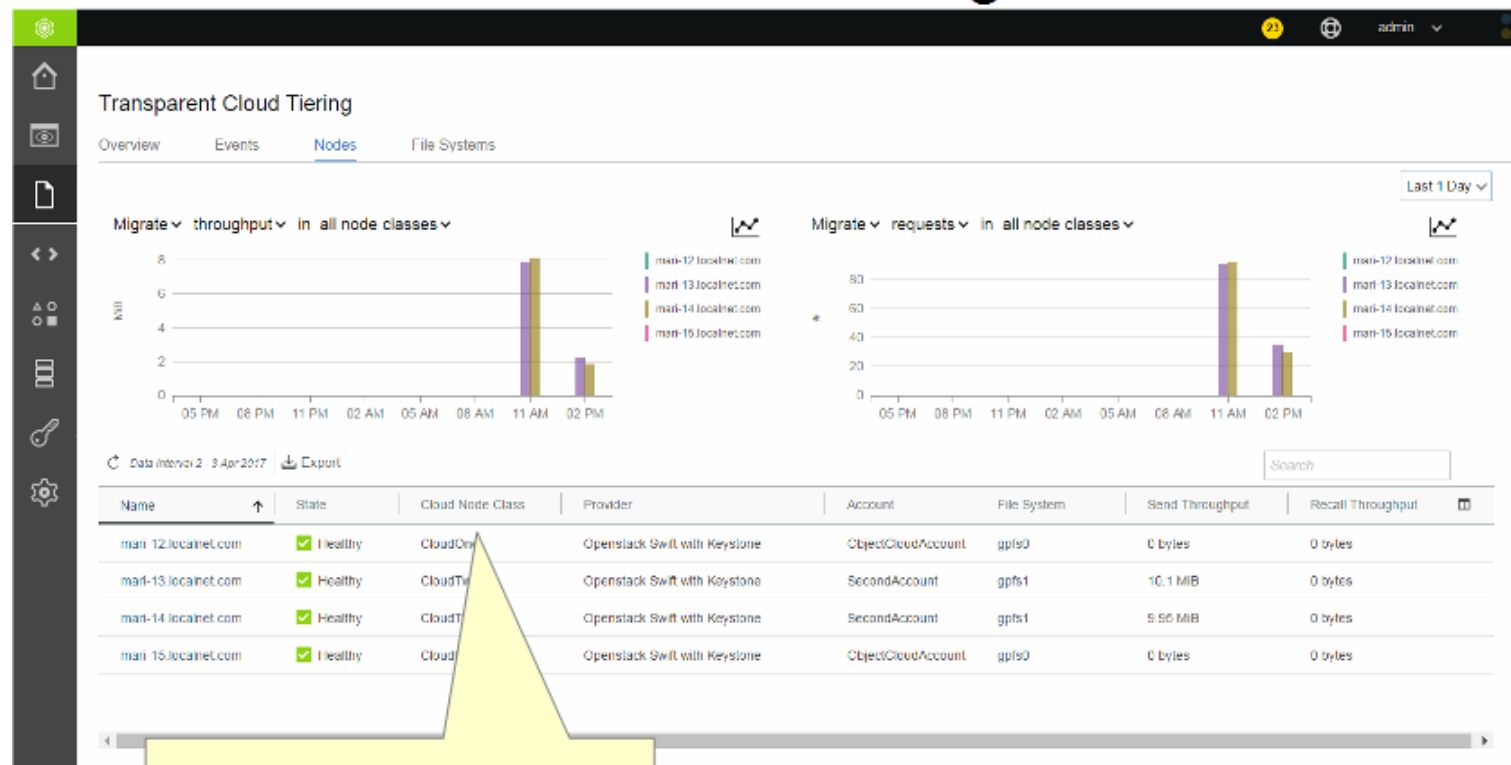
```
Task started. 10:53 AM
cache-41.localnet.com[10.0.100.41]: 10:53 AM
cd /var/mef/rel/ 10:53 AM
re IgnoreAnyMount* IgnoreStartupMount* 10:53 AM
touch IgnoreAnyMount_gpfs0 10:53 AM
The task is 100% complete. 10:53 AM
Refreshing data ... 10:53 AM
Task completed. 10:53 AM
```

Below the dialog, a table lists the mounted file systems:

Name	Mount State	Bytes Read	Bytes Written	Read OPS	Write ...	Avg Disk Wait R...	Avg Disk Wait
cache-41.localnet.com	Mounted	0.002 bytes/s	0 bytes/s	0.0 ops/s	0 ops/s	0.017 ms	0 ms
cache-42.localnet.com	Mounted	0.019 bytes/s	0 bytes/s	0.0 ops/s	0 ops/s	0.011 ms	0 ms
cache-43.localnet.com	Mounted	0.037 bytes/s	0 bytes/s	0.0 ops/s	0 ops/s	0.014 ms	0 ms

Provide simple mounting actions, but also expose the partly unknown mount configuration capabilities

TCT Monitoring



TCT health and performance monitoring. Significant rework in 4.2.4 due to a configuration model change in TCT

Health Tips

The screenshot displays a management interface with a 'Tips' section. A table lists health tips, with one selected. A dialog box titled 'Fix Procedure: Configure Performance Monitoring Sensor' is open, showing a wizard to configure an NFS performance monitoring sensor. The dialog includes instructions, a data collection interval dropdown set to 10 seconds, and a terminal window showing the execution of a command and the successful completion of the task.

Current	Tip Name	Event Time
	nfs_sensors_not_configured	9/7/17 1:11:49 PM
	nfs_sensors_not_configured	9/7/17 1:11:30 PM

Fix Procedure: Configure Performance Monitoring Sensor

The NFS performance monitoring sensor (NFSIO) is not configured on one or more nodes of the cluster.

Select the frequency of the performance data collection. It is recommended to select a value greater than or equal to 10 as the data collection frequency to reduce the impact on system performance.

Data collection interval: **10** seconds

Click **Finish** to configure the NFS performance monitoring sensor and set data collection interval for the sensor.

Status: Task succeeded ✓

```
Running command: 10:56 AM
m ' /var/lib/ostx/gut/rep/sensorNFS.txt' 10:56 AM
Running command: 10:56 AM
Refreshing data ... 10:56 AM
The task is 100% complete. 10:56 AM
Task completed. 10:56 AM
```

Close

The new health event type „Health Tip“ is also exposed in the GUI. Several health tips have related fix procedures which provide an easy to use wizard to repair the issue.

Call Home

Call Home

Enable Call Home

Call home node:
mr-11.novalocal

Company Information

Company name: testcompany
Customer ID: 1234567
E-mail: rohwedder@de.ibm.com
Country: Germany

Proxy Information

Enable Proxy
Proxy host:
Proxy port:
 Enable Proxy authentication
Proxy username:
Proxy password:
Test Connection
Configure Call Home

Configure Call home via a GUI
(restricted to a single callhome group)

Standard Diagnostics Performance Diagnostics

Breakdown Diagnostics

Download or delete the diagnostic data collected in the pool

File Name	Size
/tmp/ibm.snap/MT2019/04/0000112	155.00 MiB
/tmp/ibm.snap/04/1902/04/0000112	155.00 MiB

Collected diagnostic data for the last 3 days

Functional areas to be included:
 Core IBM Spectrum Scale
 Networks
 NFS
 Object
 Protocols
 Samba
 NFS
 Crash d

Node selection criteria: All Nodes

Collected Diagnostic Data

Upload to PMR

Select PMR: 123,223,123456 |

Upload Cancel

Upload snaps and associat them to an
existing PMR