Redbooks
ibm.com/redbooks

Monitoring and Managing IBM Spectrum Scale Using the GUI

Markus Rohwedder

Alexander Wolf-Reber

Stefan Roth

Liju Jose

Przemyslaw Podfigurny

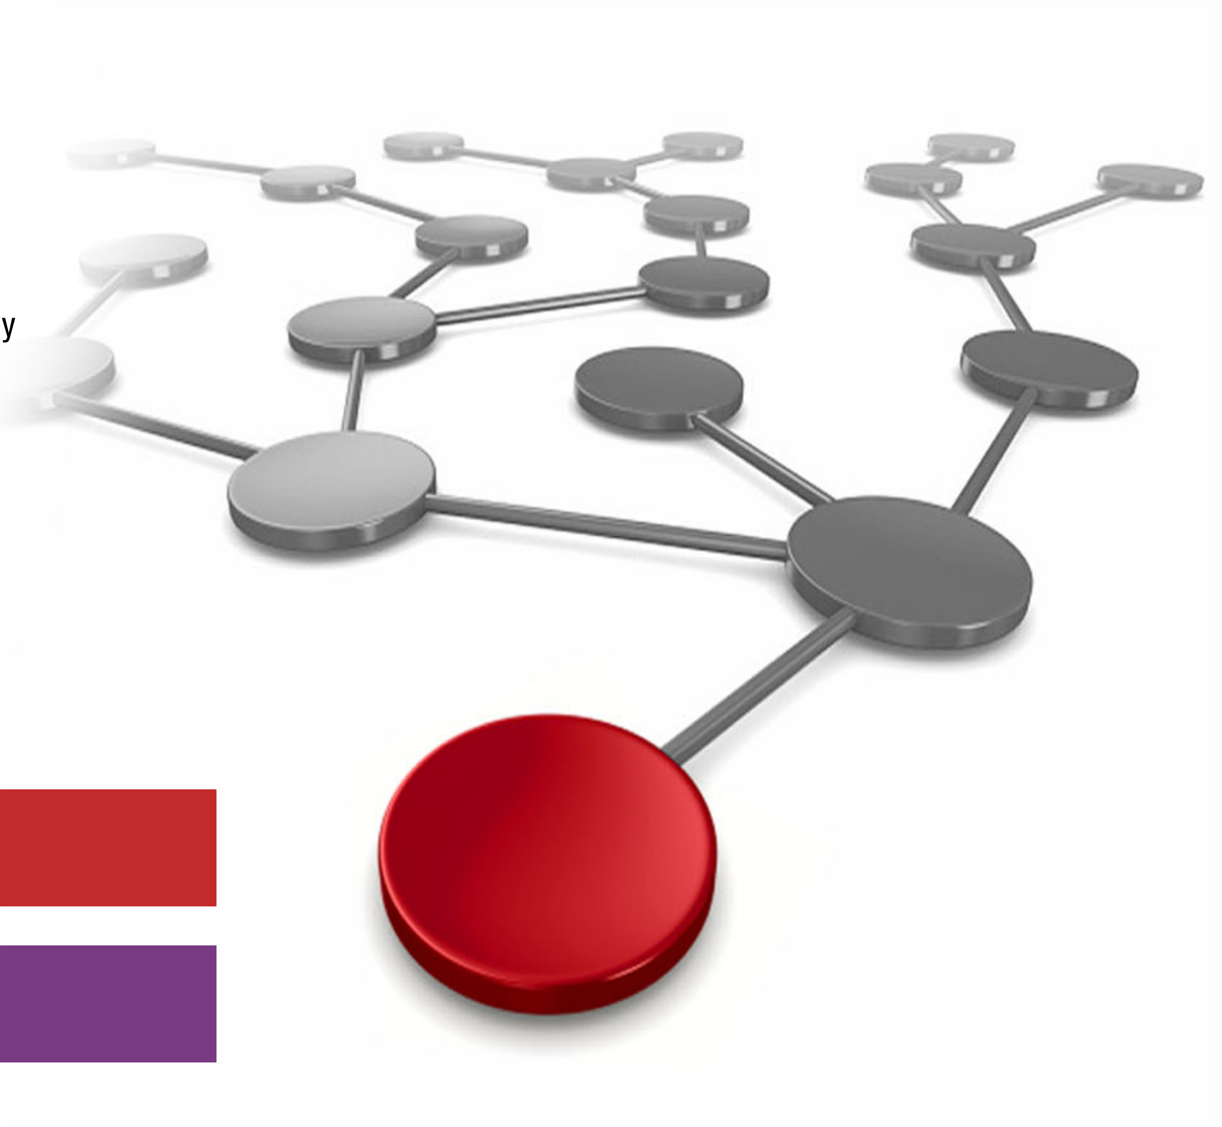

Cloud

Storage

IBM

Redpaper

International Technical Support Organization

Monitoring and Managing IBM Spectrum Scale Using the GUI

October 2019

REDP-5458-01

Note: Before using this information and the product it supports, read the information in “Notices” on page xi.

Second Edition (October 2019)

This edition applies to IBM Spectrum Scale version 5.0.3.

This document was created or updated on October 18, 2019.

Contents

Preface vii
Authors. vii
Now you can become a published author, too! viii
Comments welcome. viii
Stay connected to IBM Redbooks ix

Notices xi
Trademarks xii

Chapter 1. Overview 1
1.1 IBM Spectrum Scale GUI architecture 2
1.1.1 GUI node requirements. 2
1.2 Role-based access control with GUI users 3
1.3 High-availability with multiple GUI nodes in a cluster 3
1.4 Integration hub for configuration, health, and performance data 3
1.5 Support matrix. 4
1.5.1 Operating system levels 4
1.5.2 IBM Spectrum Scale GUI software requirements 5

Chapter 2. Installing and configuring the IBM Spectrum Scale GUI 7
2.1 Installing the IBM Spectrum Scale GUI by using the installation toolkit 8
2.2 Manually installing the IBM Spectrum Scale GUI 8
2.2.1 Prerequisites. 8
2.2.2 Installation steps 10
2.3 Enabling performance tools in the IBM Spectrum Scale GUI 11
2.3.1 Checking the GUI and performance tool status 13
2.4 Manually upgrading the IBM Spectrum Scale GUI 14
2.4.1 Prerequisites. 14
2.5 Securing the IBM Spectrum Scale GUI. 16
2.5.1 Firewall recommendations and supported ports. 16
2.5.2 Creating and using an HTTPS certificate to secure communications between GUI web server and web browsers 17
2.5.3 Root privilege considerations for the IBM Spectrum Scale GUI 18
2.6 Configuring the IBM Spectrum Scale GUI to use sudo wrappers 18
2.7 Ensuring high availability of the GUI service. 19
2.8 Node classes used for the IBM Spectrum Scale GUI 21
2.9 Modifying the IBM Spectrum Scale GUI property file 21
2.10 Distributed GUI preferences 22

Chapter 3. Understanding the GUI options. 23
3.1 Login 24
3.2 Navigation 25
3.3 Header area 26
3.4 Assistance for understanding the features associated with a GUI page. 28

Chapter 4. Monitoring options available in the IBM Spectrum Scale GUI 29
4.1 Monitoring performance 30
4.1.1 Display options in the Statistics page of the GUI 32
4.1.2 Display options in performance charts 32

4.1.3 Select performance and capacity metrics 33
4.1.4 Create customized performance charts 34
4.2 Using the dashboard to view performance charts 35
4.2.1 Custom Dashboards 36
4.2.2 Display options 36
4.2.3 Widget options 36
4.3 Monitoring waiters 36
4.4 Monitoring capacity 36
4.4.1 Capacity data obtained from the GPFS quota database 37
4.4.2 Capacity data collected through the performance monitoring tool 38
4.4.3 Capacity information for file systems, pools, NSDs, and filesets 39
4.4.4 Viewing the capacity for licensing 39
4.5 Monitoring system health 39
4.6 Monitoring nodes 41
4.6.1 Nodes tables 41
4.6.2 Performance monitoring of nodes 42
4.6.3 Creating and managing user-defined node classes 43
4.7 Monitoring Transparent Cloud Tiering 44
4.8 Monitoring Active File Management 45
4.9 Monitoring file systems 48
4.10 Monitoring filesets 50
4.11 Monitoring pools 51
4.12 Monitoring NSDs 52
4.13 Monitoring networks by using GUI 53
4.14 Monitoring remote cluster through GUI 55
4.15 Monitoring thresholds 58
4.16 Monitoring command audit log 63

Chapter 5. Configuring and managing tasks by using the IBM Spectrum Scale GUI . 65
5.1 Creating file systems 66
5.2 Mounting a file system through the GUI 66
5.3 Unmounting a file system by using the GUI 67
5.4 Creating filesets 68
5.5 Creating and managing snapshots 68
5.5.1 Scheduling snapshot creation by using snapshot rules 69
5.5.2 Example scenario for retention policy 69
5.6 Deleting snapshots 71
5.7 Configuring quota 71
5.8 Information lifecycle management 72
5.8.1 Creating and applying policy 73
5.8.2 Editing a policy by using the text editor 73
5.8.3 Defining the policy run settings 73
5.8.4 Log files 75
5.9 Managing storage 75
5.10 Managing access control lists 75
5.11 Managing Object Storage, SMB shares, and NFS exports 76
5.12 Managing IBM services 77
5.12.1 GPFS daemon 77
5.12.2 CES 77
5.12.3 CES network and CES IPs 78
5.12.4 NFS 78
5.12.5 SMB 78
5.12.6 Object 78

5.12.7 File authentication . . . 79
5.12.8 Object Authentication section . . . 79
5.12.9 Hadoop connector . . . 79
5.12.10 GUI . . . 79
5.12.11 Performance monitoring . . . 80
5.12.12 File auditing . . . 80
5.12.13 Message queue . . . 80
5.13 Configuring role-based access for GUI users . . . 81
5.13.1 Managing GUI users locally in the IBM Spectrum Scale system . . . 81
5.13.2 Assigning roles to user groups . . . 82
5.13.3 Managing GUI administrators in an external authentication server . . . 83

Chapter 6. Troubleshooting options in the IBM Spectrum Scale GUI . . . 87
6.1 Monitoring events . . . 88
6.1.1 Event deduplication . . . 88
6.1.2 Event Filtering . . . 88
6.1.3 Marking events as read . . . 88
6.1.4 Running fix procedure . . . 89
6.2 Monitoring tip events . . . 89
6.3 Configuring event notifications . . . 89
6.4 Configuring email notifications . . . 90
6.5 Configuring SNMP manager . . . 91
6.5.1 Understanding the SNMP OID ranges . . . 91
6.5.2 Example for SNMP traps . . . 92
6.5.3 SNMP MIBs . . . 92
6.6 Collecting diagnostic data by using the GUI . . . 92
6.6.1 GUI logs . . . 94
6.7 Configuring IBM Call Home by using GUI . . . 94
6.8 Directed maintenance procedures . . . 95
6.9 Troubleshooting issues with capacity data displayed in the GUI . . . 96

Appendix A. Current limitations for the IBM Spectrum Scale GUI . . . 99

Related publications . . . 101
IBM Redbooks . . . 101
Other publications . . . 101
Online resources . . . 101
Help from IBM . . . 102

Preface

The IBM® Spectrum Scale GUI provides an easy way to configure and monitor various features that are available with the IBM Spectrum® Scale system. It is a web application that runs on common web browsers, such as Chrome, Firefox, and Edge. The IBM Spectrum Scale GUI uses Java Script and Ajax technologies to enable smooth and desktop-like interfacing.

This IBM Redpaper publication provides a broad understanding of the architecture and features of the IBM Spectrum Scale GUI. It includes information about how to install and configure the GUI and in-depth information about the use of the GUI options. The primary audience for this paper includes experienced and new users of IBM Spectrum Scale.

Authors

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization, Tucson Center.

Markus Rohwedder is an IT Architect at IBM Research™ and Development in Kelsterbach, Germany. He joined IBM in 1999 after acquiring a PhD in Physics. At IBM, Markus has focused on data-centric projects, such as design and administration of a large, continuously available data warehouse for production support, and implementation of a searchable tape archive for structured data or benchmarking clustered databases. Since 2008, he has worked on creating graphical user interfaces for storage systems, such as Information Archive, IBM Storwize® V7000 Unified, and IBM Spectrum Scale.

Alexander Wolf-Reber is an IT Architect in IBM Research and Development Kelsterbach, Germany. His current role is technical lead for the GUI and REST-API components of IBM Spectrum Scale. He joined IBM in 1999 and worked on various storage products, such as IBM SAN Volume Controller, IBM Enterprise Storage Server®, and tape libraries. Since 2007, his focus is on clustered file systems. During his career, he also contributed to standardization bodies, such as the Storage Networking Association and the Java Community Process. He holds a PhD degree in Physics from the Johann Wolfgang Goethe University in Frankfurt, Germany.

Stefan Roth is a Software Engineer in IBM Research and Development in Kelsterbach, Germany. He works with the IBM Spectrum Scale development team on the graphical user interface. He joined IBM in 1996, and in the first years he developed software for IBM disk drives and semiconductor factories. Since 2008, he worked on graphical user interfaces for various IBM storage products, such as Scale Out Network Attached Storage, V7000 Unified, IBM Spectrum Scale, and Elastic Storage Server. He holds a technical college degree in Electrical Engineering from University of Applied Sciences, Darmstadt.

Liju Jose is an Information Developer with the IBM ISDL ID team. He is responsible for writing and editing the customer-facing documentation for various storage products, such as IBM Spectrum Scale, IBM Elastic Storage® Server, and IBM Storwize V7000 Unified. He has been with IBM for the last five years and holds a Bachelor’s degree in Physics and a Master’s degree in Electronics Science from the Mahatma Gandhi University.

Przemyslaw Podfigurny is a Software Engineer in IBM Research and Development in Kelsterbach, Germany. He joined IBM in 2015 and worked as a developer on IBM Spectrum Scale GUI Backend, REST API components and build setup. His primary interests include Java technologies, databases, distributed Linux environment, and big data. He holds a Master's degree in Software Engineering from the Wrocław University of Science and Technology.

Thanks to the following people for their contributions to this project:

Larry Coyne
International Technical Support Organization

Sandeep Ramesh
Dietmar Fischer
Andreas Koeninger
Alifiya A Lohawalla
Dharmendra Rai
IBM Systems

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this paper or other IBM Redbooks® publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks
- Follow us on Twitter:

  http://twitter.com/ibmredbooks
- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806
- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm
- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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 and cannot confirm the accuracy of 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.

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

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. 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 terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

AIX®
Enterprise Storage Server®
IBM®
IBM Elastic Storage®
IBM Research™
IBM Spectrum®
Redbooks®
Redbooks (logo) ®
Storwize®
Tivoli®
WebSphere®

The following terms are trademarks of other companies:

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

Microsoft, Windows, 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.

Other company, product, or service names may be trademarks or service marks of others.

1

Overview

The IBM Spectrum Scale GUI provides an easy way to configure and monitor various features that are available with the IBM Spectrum Scale system. The IBM Spectrum Scale GUI is a web application that runs on common web browsers, such as Chrome, Firefox, and Edge. It uses Java Script and Ajax technologies to enable smooth and desktop-like interfacing. Therefore, no client-side installation is required to use the GUI.

This chapter provides a basic overview of the GUI and includes the following topics:

- 1.1, “IBM Spectrum Scale GUI architecture” on page 2
- 1.2, “Role-based access control with GUI users” on page 3
- 1.3, “High-availability with multiple GUI nodes in a cluster” on page 3
- 1.4, “Integration hub for configuration, health, and performance data” on page 3
- 1.5, “Support matrix” on page 4

1.1 IBM Spectrum Scale GUI architecture

The architecture for the IBM Spectrum Scale GUI is shown in Figure 1-1.

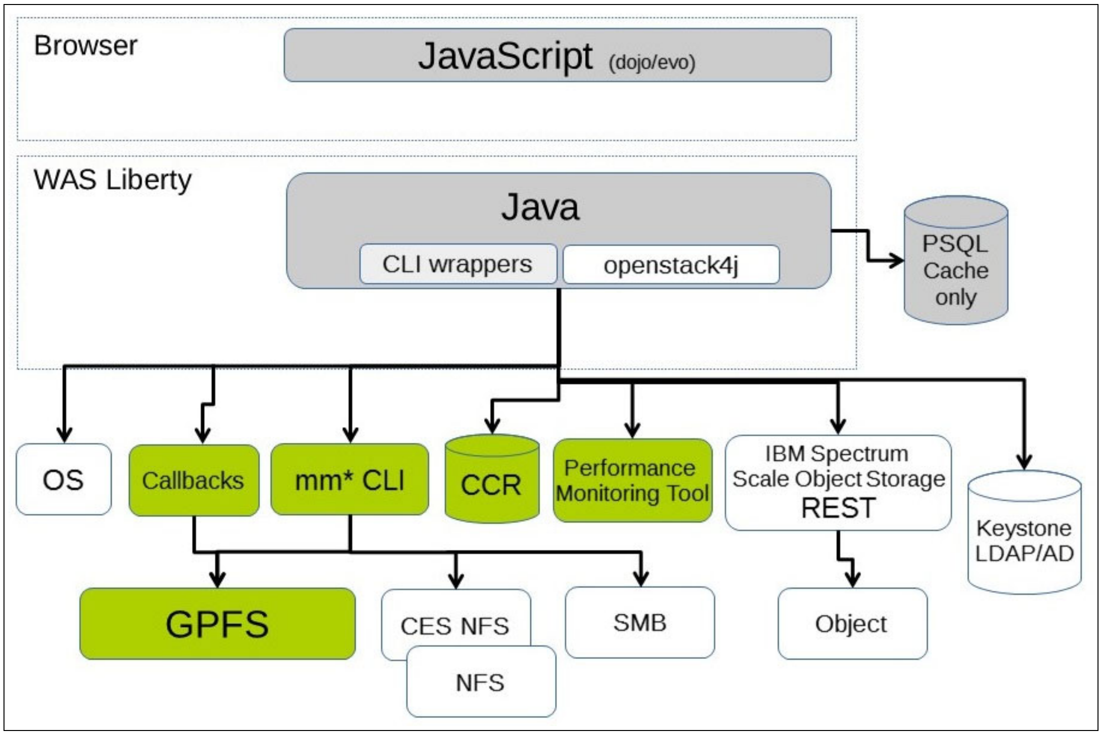

Figure 1-1 IBM Spectrum Scale GUI architecture

The IBM Spectrum Scale GUI runs on an IBM WebSphere® Liberty application server on one or more cluster nodes. The application servers provides the infrastructure for the GUI and REST API. Configuration information from the IBM Spectrum Scale cluster is cached in a database because some commands are resource intense and user interaction with the GUI should not always trigger data refresh activities from the cluster. The GUI includes interfaces to many components in the cluster.

One GUI manages a single IBM Spectrum Scale cluster, but it can also contact GUI nodes from other clusters to exchange monitoring information through REST.

1.1.1 GUI node requirements

The GUI can be installed on any cluster node, including client nodes. The GUI node does not need any specific license. It is not possible to install the GUI on a node that is not part of a cluster. The GUI can be installed on Linux operating system versions that are supported by the core IBM Spectrum Scale. Refer to the support matrix in IBM Knowledge Center for supported operating systems.

The resource consumption of the GUI service is considerably low. The memory consumption of the application server is limited to 512 MiB, and the memory limit of the GUI including all child processes is set to 2 GiB.

The GUI node also acts as the collector node for performance data. The collector process uses memory that is based on the number of elements, such as nodes or disks, and depends on collection frequency.

Network requirements

The GUI nodes must contact all nodes in the cluster. To connect to other clusters, you must establish an HTTPS connection among the local node with remote cluster nodes. The application server assumes that port 443 is available on the GUI node.

Use cases for the IBM Spectrum Scale GUI are described next.

1.2 Role-based access control with GUI users

The GUI offers a role-based access model, which is not present in the core IBM Spectrum Scale system. Different roles are available to limit access to certain functions, from a read-only monitor role to a security administrator with full access on all GUI functions.

GUI users are separate from the Operating System users. The GUI includes a built-in user database that allows you to register and manage users. It is also possible to configure the GUI to use an external authentication system, such as LDAP or AD to authenticate the GUI users. The GUI users are also used for the IBM Spectrum Scale management API (REST API) access.

For more information about GUI user management, “Configuring role-based access for GUI users” on page 81.

1.3 High-availability with multiple GUI nodes in a cluster

Up to three nodes can be running the GUI in a cluster. All of these nodes are active GUI nodes, but some services, such as GUI-based email notifications, run on a single node only. For more information, see 2.7, “Ensuring high availability of the GUI service” on page 19.

1.4 Integration hub for configuration, health, and performance data

You can perform the following important tasks through the IBM Spectrum Scale GUI:

- Monitor the performance of the system based on various aspects
- Monitor system health
- Create and manage file systems
- Create and manage filesets
- Create, manage, and schedule snapshots
- Create rules and policies for information lifecycle management
- Monitor active file management
- Monitor storage pools
- Monitor NSDs
- Monitor thresholds
- Monitor command audit logs
- Monitor file audit logs

- Monitor remote clusters
- Monitor nodes and networks in the cluster
- Manage IBM Spectrum Scale services
- Manage SMB service and SMB shares
- Manage NFS service and NFS exports
- Display and modify NFSv4 ACL for files and directories
- Create users and define roles for the GUI users
- Configure authentication method for NFS and SMB users
- Configure authentication method for GUI users
- Create and manage node classes
- Define default, user, group, and fileset quotas
- Monitor the capacity details at various levels, such as file system, pools, filesets, users, and user groups
- Configure event notifications through emails and SNMP
- Collect diagnostic data to find the root cause and troubleshoot an issue reported in the system
- Monitor events
- Perform directed maintenance procedures to fix certain issues or to optimize the system
- Enable and configure call home feature in the cluster
- Monitor Transparent Cloud Tiering service
- Manage Object Storage, and create object users and roles

1.5 Support matrix

In this section, we describe the operating systems and IBM Spectrum Scale release levels on which the GUI is supported.

1.5.1 Operating system levels

The GUI supports all Linux operating system variations, which are supported by spectrum scale, on all supported platforms. The GUI can also run on virtual nodes. The GUI does not run on Windows or IBM AIX® nodes.

The GUI nodes in the IBM Spectrum Scale cluster can be on any of the supported operating systems or hardware platforms. That is, an intermix of GUI nodes on various operating systems and hardware platforms are supported in a cluster.

For more information about the supported operating system and other software versions, see this IBM Knowledge Center web page.

1.5.2 IBM Spectrum Scale GUI software requirements

IBM Spectrum Scale software includes the following requirements for GUI:

- The GUI is supported on the cluster that runs on IBM Spectrum Scale 5.0.3 or later. Issue the **`mmlsconfig`** command to see the value that is set for the `minReleaseLevel` attribute. The cluster must have a minimum release level of 4.2.0.
- All IBM Spectrum Scale packages that are installed on a GUI node must be of the same release. For example, do not mix the IBM Spectrum Scale 5.0.3 GUI rpm with a 5.0.2 base rpm. However, GUI PTFs and fixes often can be applied without installing the corresponding PTF or fix of the base package. This feature is helpful if you want to resolve a GUI issue without changing anything on the base layer.

  Note: It is recommended to move to the latest PTF level that is available for the underlying IBM Spectrum Scale release.

- The minimum release level of the cluster must be on the latest release level to display the latest GUI features.

Installing and configuring the IBM Spectrum Scale GUI

You can install the management GUI by using the following methods:

- The installation toolkit
- Manual installation

This chapter describes both the methods and includes the following topics:

- 2.1, “Installing the IBM Spectrum Scale GUI by using the installation toolkit” on page 8
- 2.2, “Manually installing the IBM Spectrum Scale GUI” on page 8
- 2.3, “Enabling performance tools in the IBM Spectrum Scale GUI” on page 11
- 2.4, “Manually upgrading the IBM Spectrum Scale GUI” on page 14
- 2.5, “Securing the IBM Spectrum Scale GUI” on page 16
- 2.6, “Configuring the IBM Spectrum Scale GUI to use sudo wrappers” on page 18
- 2.7, “Ensuring high availability of the GUI service” on page 19
- 2.8, “Node classes used for the IBM Spectrum Scale GUI” on page 21
- 2.9, “Modifying the IBM Spectrum Scale GUI property file” on page 21
- 2.10, “Distributed GUI preferences” on page 22

2.1 Installing the IBM Spectrum Scale GUI by using the installation toolkit

The IBM Spectrum Scale GUI is installed on a node where the node is specified as a GUI server in the cluster definition with the `-g` option:

```
./spectrumscale node add gpfsnode3 –g
```

The GUI server node must be added as an administrator node by using the `-a` flag:

```
./spectrumscale node add gpfsnode3 -a
```

If no nodes are specified as GUI servers, the GUI is not installed. Generally, have at least two GUI interface servers and a maximum of three for redundancy.

The GUI is installed on specified GUI servers when you run the **`./spectrumscale`** installation command.

At the end of a successful IBM General Parallel File System (IBM GPFS) installation or protocol deployment, you can access the GUI through a web browser with the following node address:

```
https://<GUI server IP or host name>
```

> **Note:** After the installation, you must create the first GUI user to log in to the GUI and create other GUI administrative users who perform system management and monitoring tasks. When you start the GUI for the first time after the installation, the GUI welcome page provides options to create the first GUI user from the command-line prompt by using the **`/usr/lpp/mmfs/gui/cli/mkuser <user_name> -g SecurityAdmin`** command.

2.2 Manually installing the IBM Spectrum Scale GUI

You can install the IBM Spectrum Scale GUI by manually installing the required packages.

2.2.1 Prerequisites

The prerequisites for installing the IBM Spectrum Scale system are applicable for GUI installations as well. For more information about the prerequisites for installation, see IBM Spectrum Scale `Installation prerequisites`.

The IBM Spectrum Scale GUI package is also part of the installation package. You must extract this package to start the installation. The performance tool packages are also required to enable the performance monitoring tool that is integrated into the GUI. The following packages are required for performance monitoring tools in GUI:

- The performance tool collector package. This package is placed on the collector nodes only. By default, every GUI node is also used as the collector node to receive performance details and display them in the GUI.
- The performance tool sensor package. This package is applicable for the sensor nodes, if not already installed. It is recommended to install the sensor package on all nodes of the cluster to monitor their performance in the GUI.

Note: The GUI must be a homogeneous stack. That is, all packages must be of the same release. For example, do not mix the 5.0.03 GUI rpm with a 4.2.35.0.2 base rpm. However, GUI PTFs and efixes often can be applied without installing the corresponding PTF or efix of the base package. This feature is helpful if you want to remove a GUI issue without changing anything on the base layer.

The IBM Spectrum Scale GUI and performance tool packages that are required for different platforms are listed in Table 2-1.

Table 2-1 GUI packages that are required for each platform

Platform	GUI package name
RHEL 7.x	gpfs.gui-5.0.3-0.noarch.rpm gpfs.java-5.0.3-0.x86_64.rpm gpfs.java-5.0.3-0.ppc64.rpm gpfs.java-5.0.3-0.ppc64le.rpm gpfs.java-5.0.3-0.s390x.rpm
SUSE Linux Enterprise Server 12	gpfs.gui-5.0.3-0.noarch.rpm gpfs.java-5.0.3-0.x86_64.rpm gpfs.java-5.0.3-0.ppc64le.rpm gpfs.java-5.0.3-0.s390x.rpm
Ubuntu 16 and 18	gpfs.gui_5.0.3-0_all.deb gpfs.java_5.0.3-0_amd64.deb gpfs.java_5.0.3-0_s390x.deb gpfs.java_5.0.3-0_ppc64el.deb
Performance monitoring tool platform	**Performance monitoring tool rpms**
RHEL 7.x X86	gpfs.gss.pmcollector-5.0.3-0.el7.x86_64.rpm gpfs.gss.pmsensors-5.0.3-0.el7.x86_64.rpm
RHEL 7 s390x	gpfs.gss.pmsensors-5.0.3-0.el7.s390x.rpm gpfs.gss.pmcollector-5.0.3-0.el7.s390x.rpm
RHEL 7.x ppc64	gpfs.gss.pmcollector-5.0.3-0.el7.ppc64.rpm gpfs.gss.pmsensors-5.0.3-0.el7.ppc64.rpm
RHEL 7.x ppc64 LE	gpfs.gss.pmcollector-5.0.3-0.el7.ppc64le.rpm gpfs.gss.pmsensors-5.0.3-0.el7.ppc64le.rpm
RHEL6 s390x	gpfs.gss.pmsensors-5.0.3-0.el6.s390x.rpm gpfs.gss.pmcollector-5.0.3-0.el6.s390x.rpm
SUSE Linux Enterprise Server 12 X86	gpfs.gss.pmcollector-5.0.3-0.SLES12.x86_64.rpm gpfs.gss.pmsensors-5.0.3-0.SLES12.X86_64.rpm
SUSE Linux Enterprise Server 12 SP1 s390x	gpfs.gss.pmsensors-5.0.3-0.SLES12.1.s390x.rpm gpfs.gss.pmcollector-5.0.3-0.SLES12.1.s390x.rpm
SUSE Linux Enterprise Server 12 ppc64	gpfs.gss.pmcollector-5.0.3-0.SLES12.ppc64.rpm gpfs.gss.pmsensors-5.0.3-0.SLES12.ppc64.rpm
SUSE Linux Enterprise Server 12 ppc64 LE	gpfs.gss.pmcollector-5.0.3-0.SLES12.ppc64le.rpm gpfs.gss.pmsensors-5.0.3-0.SLES12.ppc64le.rpm
SUSE Linux Enterprise Server 11 ppc64 (sensor only)	gpfs.gss.pmsensors-5.0.3-0.SLES11.ppc64.rpm

Platform	GUI package name
SUSE Linux Enterprise Server 11 s390x (sensor only)	gpfs.gss.pmsensors-5.0.3-0.SLES11.s390x.rpm
Ubuntu 16.04 LTS sensor and collector packages	gpfs.gss.pmsensors_5.0.3-0.U16.04_amd64.deb gpfs.gss.pmcollector_5.0.3-0.U16.04_amd64.deb
Ubuntu 18.04 LTS sensor and collector packages	gpfs.gss.pmsensors_5.0.3-0.U18.04_amd64.deb gpfs.gss.pmcollector_5.0.3-0.U18.04_amd64.deb

Note: Ensure that the performance tool collector runs on the same node as the GUI.

Yum repository setup

You can use the yum repository to manually install the IBM Spectrum Scale GUI `.rpm` files. This method of installation is preferred because yum checks the dependencies and automatically installs the missing platform dependencies, such as the postgres module, which is required but not included in the package.

2.2.2 Installation steps

You can install the IBM Spectrum Scale GUI by using the package manager (**yum** or **zypper** commands) or by issuing the **rpm** commands individually.

Installing by using the package manager (yum or zypper commands)

Generally, use the package manager to install the IBM Spectrum Scale GUI because it checks the dependencies and automatically installs the missing platform dependencies. To use this method to install the IBM Spectrum Scale GUI, issue the following commands:

- Red Hat Enterprise Linux:

```
yum install gpfs.gss.pmsensors-5.0.3-0.el7.<arch>.rpm
yum install gpfs.gss.pmcollector-5.0.3-0.el7.<arch>.rpm
yum install gpfs.java-5.0.3-0.<arch>.rpm
yum install gpfs.gui-5.0.3-0.noarch.rpm
```

- SUSE Linux Enterprise Server:

```
zypper install gpfs.gss.pmsensors-5.0.3-0.SLES12.<arch>.rpm
zypper install gpfs.gss.pmcollector-5.0.3-0.SLES12.<arch>.rpm
zypper install gpfs.java-5.0.3-0.<arch>.rpm
zypper install gpfs.gui-5.0.3-0.noarch.rpm
```

Installing by using RPMs

Issue the following commands for RHEL and SUSE Linux Enterprise Server platforms:

```
rpm -ivh gpfs.java-5.0.3-0.<arch>.rpm
rpm -ivh gpfs.gss.pmsensors-5.0.3-0.el7.<arch>.rpm
rpm -ivh gpfs.gss.pmcollector-5.0.3-0.el7.<arch>.rpm
rpm -ivh gpfs.gui-5.0.3-0.noarch.rpm
```

The sensor **rpm** must be installed on any other node that you want to monitor. All sensors must point to the collector node.

Installing management GUI on Ubuntu by using dpkg and apt-g

Issue the following commands for Ubuntu platform:

```
dpkg -i gpfs.java_5.0.3-0_<arch>.deb
dpkg -i gpfs.gss.pmsensors_5.0.3-0.<os>_<arch>.deb
dpkg -i gpfs.gss.pmcollector_5.0.3-0.<os>_<arch>.deb
apt-get install postgresql
dpkg -i gpfs.gui_5.0.3-0_all.deb
```

2.3 Enabling performance tools in the IBM Spectrum Scale GUI

The performance tool consists of sensors that are installed on all nodes that must be monitored. It also consists of one or more collectors that receive data from the sensors.

The GUI expects that a collector runs on a GUI node. The GUI also queries the collector on the same node for performance and capacity data. The following process uses the automated approach to configure and maintain performance data collection by using the **mmperfmon** CLI command.

> **Note:** Manually editing the /opt/IBM/zimon/ZIMonSensors.cfg file is not compatible with this configuration mode.

To enable performance tools, complete the following steps:

1. Install the necessary software packages.

   Install the collector software package, gpfs.gss.pmcollector, on all GUI nodes. Install the sensor software packages, gpfs.gss.pmsensors, on all nodes that are supposed to send the performance data.

2. Initialize the performance collection. Use the following command to create an initial performance collection setup on the selected nodes:

   ```
   mmperfmon config generate --collectors [node list]
   ```

   The GUI nodes must be configured as collector nodes. Depending on the installation type, this configuration might be completed before. However, verify the existing configuration.

   When several collectors are defined, the system automatically configures these collectors as peers. When one of the collectors are queried, the collector contacts its peers to gather data that might be on a peer only. Automatic peer configuration was is available from IBM Spectrum Scale 5.0.2. In earlier releases, a manual peer configuration was necessary.

3. Enable the nodes for performance collection. You can enable nodes to collect performance data by issuing the following command:

   ```
   mmchnode --perfmon -N [SENSOR_NODE_LIST]
   ```

   In this command, [*SENSOR_NODE_LIST*] is a comma-separated list of sensor node host names or IP addresses. The sensor node host name or IP address must be a GPFS node or a daemon IP address, not any other node IP.

   You can also use a node class. Depending on the type of installation, nodes might be configured for performance collection. Review the perfmon designation of nodes in the **mmlscluster** command output to verify whether the node is configured for performance data collection.

4. Configure aggregation configuration for the collectors.

   The collector configuration is stored in the `/opt/IBM/zimon/ZIMonCollector.cfg` file. The performance collection tool includes predefined rules about how data is aggregated after it gets older. By default, the following aggregation domains are created:

   – A raw domain that stores the metrics uncompressed
   – A first aggregation domain that aggregates data to 30-second averages
   – A second aggregation domain that stores data in 15-minute averages
   – A third aggregation domain that stores data in 6-hour averages

   In addition to the aggregation that is done by the performance collector, the GUI might request aggregated data depending on the zoom level of the chart.

5. Configure the sensors.

   Several GUI pages display performance data that is collected with the help of performance monitoring tools. If data is not collected, the GUI shows error messages, such as “No Data Available” or “Objects not found” in the performance charts. Installation by using the IBM Spectrum Scale installation toolkit manages the default performance monitoring installation and configuration. The GUI context-sensitive help that is available on various pages shows performance metric information. The GUI context-sensitive help also lists the sensor names.

   The **Services** → **Performance Monitoring** page provides option to configure the sensor and provide hints for collection periods and restriction of sensors to specific nodes.

   You can also use the **`mmperfmon config show`** command to verify the sensor configuration. Use the **`mmperfmon config update`** command to adjust the sensor configuration to match your requirement.

The `/opt/IBM/zimon/ZIMonSensors.cfg` local file can be different on every node, and the system updates this file whenever a configuration change is made. Therefore, this file must not be edited manually when using the automated configuration mode. During the distribution of the sensor configuration, the restrict clause is evaluated and the period for all sensors is set to 0 in the `/opt/IBM/zimon/ZIMonSensors.cfg` file on the nodes that did not match the restrict clause. You can check the local file to confirm that a restrict clause worked as intended.

Configuring capacity-related sensors to run on a single-node

Several capacity-related sensors must run on a single node only as they collect data for a clustered file system; for example, GPFSDiskCap, GPFSFilesetQuota, GPFSFileset, and GPFSPool.

It is possible to automatically assign restrict these sensors to a single node. In IBM Spectrum Scale 5.0.1 and later, the capacity-related sensors are configured to automatically elect a single node where the capacity collection occurs.

Use the **Services** → **Performance Monitoring** page to set appropriate periods for these sensors.

The GPFSDiskCap sensor includes a recommended period of 86400, which means once per day. As the GPFSDiskCap sensor runs **`mmdf`** command to get the capacity data, it is not recommended to use a value less than 10800 (every 3 hours).

To show fileset capacity information, quota must be enabled for all file systems where fileset capacity must be monitored. For more information about enabling quota, see the **`-q`** option in the **`mmchfs`** command and **`mmcheckquota`** command.

2.3.1 Checking the GUI and performance tool status

To check the GUI and performance tool status, complete the following steps:

1. Issue the **`systemctl status gpfsgui`** command to determine the GUI status, as shown in Example 2-1.

   Example 2-1 Determining the GUI status using the systemctl status gpfsgui command

   ```
   systemctl status gpfsgui
   gpfsgui.service - IBM_Spectrum_Scale Administration GUI
      Loaded: loaded (/usr/lib/systemd/system/gpfsgui.service; disabled; vendor
   preset: disabled)
      Active: active (running) since Tue 2017-04-11 16:09:23 CEST; 1 day 22h ago
    Main PID: 1430 (java)
      Status: "GSS/GPFS GUI started"
      CGroup: /system.slice/gpfsgui.service
              ··1430 /usr/lpp/mmfs/java/jre/bin/java
   -XX:+HeapDumpOnOutOfMemoryError -Dcom.ibm.gpfs.platf...
   ```

2. Issue the **`systemctl status pmcollecto`**r and **`systemctl status pmsensors`** commands to determine the status of the performance tool.

   You can also check whether the performance tool backend can receive data by using the GUI or alternative by using a command line performance tool that is called `zc`. This tool is available in the `/opt/IBM/zimon` folder. Example 2-2 shows some sample output.

   Example 2-2 Sample output of the zc tool

   ```
   echo "get metrics cpu_user last 1 bucket_size 60"|/opt/IBM/zimon/zc localhost
   1: node1.localnet.com|CPU|cpu_user
   2: node2.localnet.com|CPU|cpu_user
   3: node3.localnet.com|CPU|cpu_user
   Row   Timestamp           cpu_user    cpu_user  cpu_user
   1    2017-04-13 14:50:00   4.328333    null      2.343333
   2    2017-04-13 14:51:00   3.492500    null      3.305000
   ```

This example output shows that the performance data is collected from node1 and node3, whereas no data is available from node2. The sensor on node2 is not started or misconfigured.

2.4 Manually upgrading the IBM Spectrum Scale GUI

You can upgrade the IBM Spectrum Scale GUI to the latest version to get the latest features. You can upgrade one GUI node at a time without shutting down IBM Spectrum Scale on other nodes to ensure high availability.

2.4.1 Prerequisites

Ensure that you are aware of the following details before you start the upgrade process:

- If an external authentication server, such as AD or LDAP, is used to authenticate the GUI users, connections to AD and LDAP systems are stored in the `/opt/ibm/wlp/usr/servers/gpfsgui/server.xml` file. This file is overwritten during the upgrade. Ensure that you save the edited `server.xml` file to make the similar changes on the replaced file after the upgrade.

  > **Note:** This step is required only if you are upgrading the GUI from a version older than 5.0.1.

- Local users, user groups, user roles, and snapshot rules are not affected with an upgrade because they exist in the cluster configure repository (CCR).
- Data in the postgres database is retained.
- All IBM Spectrum Scale packages must be of the same release on the GUI node. For example, do not mix the 5.0.3 GUI rpm with a 5.0.2 base rpm. However, GUI PTFs and fixes can usually be applied without having to install the corresponding PTF or fix of the base package. This method is helpful if you want to resolve a GUI issue without changing anything on the base layer.
- Release levels can be different among GUI nodes and other nodes of the cluster. However, the minimum release level of the cluster must be 4.2.0.0 or later for the GUI to function.
- The `scalemgmt` user ID must not be used because the GUI requires this user ID to run the IBM Spectrum Scale GUI WebSphere Java process.
- The ports 443, 47080, and 47443 must not be used by other processes.
- You can directly upgrade the IBM Spectrum Scale GUI from 4.2.0.0 or later to the latest version.

Complete the following steps to upgrade the management GUI from 4.2.1.x or later to 4.2.2.x or later:

1. Stop the GUI services on the node by issuing the **`systemctl stop gpfsgui`** command.

   Ensure that the latest packages are available at the required location. For more information about the latest packages that are required for different platforms, see 2.2, “Manually installing the IBM Spectrum Scale GUI” on page 8.

   For more information about the location of the extracted installation packages, see IBM Knowledge Center.

2. Upgrade the GUI package. For upgrading the previously installed package, use the `rpm -Fvh` or `rpm -Uvh` options. The `rpm -Fvh` option upgrades the existing installed package, and the `rpm -Uvh` option installs the package and upgrades the package as well.

For upgrading on Ubuntu, use **dpkg -i.** By issuing the **dpkg -s gpfs.gui** command, you can check more properties about your management GUI installation. If the status in the command output is `install ok installed`, the upgrade was performed successfully. Consider the following points:

- On RHEL or SLES, issue the following command:

  ```
  rpm -Fvh gpfs.gui-5.0.3-0.noarch.rpm
  ```

- On Ubuntu, issue the following command:

  ```
  dpkg -i gpfs.gui_5.0.3-0_all.deb
  ```

3. If a new Java version is available, upgrade the package as shown in the following example:

   ```
   rpm -Fvh gpfs.java-5.0.3-0.x86_64.rpm
   ```

   Java packages are platform-dependent. For more information about the latest Java package that is required for each platform, see 2.2, “Manually installing the IBM Spectrum Scale GUI” on page 8.

4. If the minimum release level set for IBM Spectrum Scale is not same as the GUI version, the new GUI features might not be available. You can change the release level by issuing the **mmchconfig release=LATEST** command.

   Because changing the minimum release level affects the cluster behavior, see the **mmchconfig** command man page and other related topics before you make this configuration change. For more information, see IBM Knowledge Center.

5. Start the GUI by issuing the **systemctl start gpfsgui** command.

6. To ensure that the GUI and performance tool are started on the boot process, issue the following commands:

   ```
   systemctl enable gpfsgui.service
   systemctl enable pmsensor.service
   systemctl enable pmcollector.service
   ```

7. Issue the **systemctl status gpfsgui** command to verify the GUI service status.

8. Issue the **systemctl status pmcollector** and **systemctl status pmsensors** commands to verify the status of the performance tool.

2.5 Securing the IBM Spectrum Scale GUI

You can secure access to the GUI by using firewalls and HTTPS certificates, as described in this section.

2.5.1 Firewall recommendations and supported ports

Dedicating certain ports for firewalls helps to secure IBM Spectrum Scale management and installation GUIs. Different ports are used for securing the installation GUI and management GUI. The ports that must be used to secure the GUI are listed in Table 2-2.

Table 2-2 Firewall recommendations for the GUI

Port number	Function	Protocol
47080	Management GUI	HTTP, localhost only
47443	Management GUI IBM Spectrum Scale management API	HTTPS, localhost only
80	Management GUI	HTTP
443	Management GUI IBM Spectrum Scale management API	HTTPS
4444	Management GUI	Localhost only
4739, 9085, and 9084	Performance monitoring collector	N/A

If multiple GUI nodes are available in a cluster, the communication among those GUI nodes is carried out through the port 443.

The port 80 is open only to receive events if an older version than 4.2.3 of GPFS is used. It cannot be used to access the GUI or REST API. Ports 443 is internally forwarded to 47443, and port 80 is internally forwarded to 47080. This forwarding is done automatically by an `iptables` rule. The `iptables` rules are added when the `gpfsgui` service is started and are removed when it is stopped. Therefore, to access the GUI, ports such as 443, 47443, and 47080 must be opened.

The update mechanism for iptables can be disabled by setting the variable **`UPDATE_IPTABLES`** to `false`, which is stored at: `/etc/sysconfig/gpfsgui`. You must restart the GUI for the changes to take effect.

The iptables rules that are necessary for the port forwarding and to bind the non-root users to the privileged ports, are automatically checked every time when the GUI is started through the **`systemctl start gpfsgui`** command. The user does not have to configure anything manually for this process.

Note: The IBM Spectrum Scale GUI ports are not configurable. The GUI cannot coexist with a web server that uses the same ports.

The management GUI uses ZIMon to collect performance data. ZIMon collectors are normally deployed with the management GUI and sometimes on other systems in a federated configuration. Each ZIMon collector uses three ports, which can be configured in ZIMonCollector.cfg. The default ports are 4739, 9085, and 9084. The GUI is sending its queries on the ports 9084 and 9085 and these ports are accessible only from the localhost.

2.5.2 Creating and using an HTTPS certificate to secure communications between GUI web server and web browsers

The IBM Spectrum Scale system GUI supports self-signed and trusted certificates that are provided by a certificate authority (CA) to secure communications between the server and web browser.

During the GUI installation, an initial self-signed certificate is created to use for secure connections between the GUI web servers and web browsers. Based on the security requirements for your system, you can create a new self-signed certificate or install a signed certificate that is created by the CA. Self-signed certificates can generate web browser security warnings and might not comply with organizational security guidelines.

The trusted certificates are created by a third-party CA. These CAs ensure that certificates include the required security level for an organization based on purchase agreements. Trusted certificates often feature higher security controls for data encryption and do not cause browser security warnings. Trusted certificates are also stored in the WebSphere Liberty SSL keystore.

Major web browsers trust the CA-certified certificates by default; therefore, they can confirm that the certificate was received by the GUI server can be trusted. You can buy a signed certificate from a trusted third-party authority or create your own certificate and get it certified. You can use self-signed and trusted certificates. However, the use of a trusted is the preferred method because the browser trusts this certificate automatically without any manual interventions.

You can use the **Services** → **GUI** page in the GUI to install and use the certificates (see Figure 2-1).

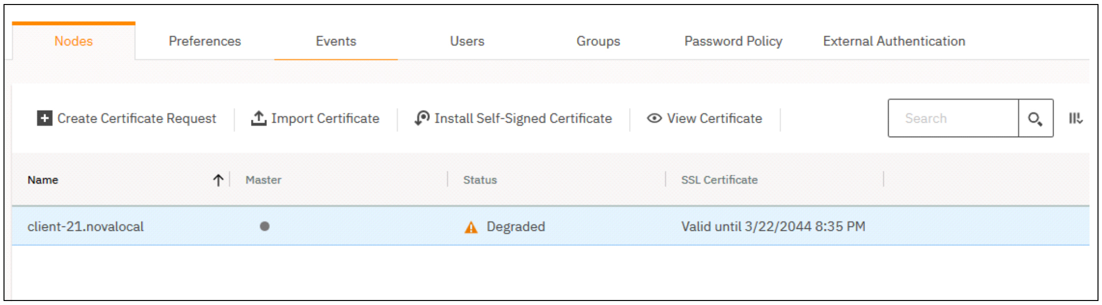

Name	Master	Status	SSL Certificate
client-21.novalocal	●	⚠ Degraded	Valid until 3/22/2044 8:35 PM

Figure 2-1 Services page

You can use the **Services** → **GUI** page in the GUI to perform the following tasks:

- Generate a self-signed certificate by using the Install Self-Signed Certificate option.
- Generate a certificate request and install it after getting it certified by the CA by using the Create Certificate Request option.

> **Note:** You can use new attributes for Subject Alternative Names, if the OpenSSL version on the GUI node is 1.1.1 or later.

- Install an issued certificate by using the Import Certificate option.
- View the details of the certificate that is applied on the local GUI node by using the View Certificate option.

When you export the certificate, the certificate is shown while accessing the GUI in the browser through HTTPS as shown in Figure 2-2.

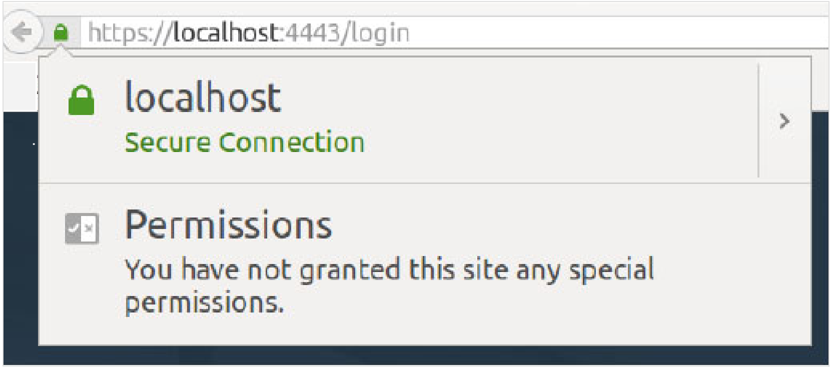
https://localhost:4443/login

localhost
Secure Connection

Permissions
You have not granted this site any special permissions.

Figure 2-2 SSL certificate

2.5.3 Root privilege considerations for the IBM Spectrum Scale GUI

IBM Spectrum Scale 4.2.3 or later no longer runs the GUI WebSphere Java process as “root” but as a user named *scalemgmt*. This method provides improved security because web applications running as root are vulnerable to security threats. The scalemgmt user is set up as a system account with no login privileges.

The GUI user still requires root privileges to perform the following tasks at the backend:

- Issue IBM Spectrum Scale CLI commands.
- Bind to the privilege ports 80 (HTTP) and 443 (HTTPS). The `iptables` rule that is used internally in the system forwards port 80 to 47080 and port 443 to 47443.

The system automatically creates the scalemgmt user. It does not require any configuration to be performed by the user.

Enabling the scalemgmt user to monitor and manage the system through the GUI

Because the root privileges are not available to the GUI user, the system enables the scalemgmt user to run the CLI commands. The GUI uses sudo wrappers to run the CLI commands. The GUI installation adds the `/etc/sudoers.d/scalemgmt_sudoers` file, which allows the scalemgmt user to run commands that match the `/usr/lpp/mmfs/bin/mm` wildcard.

2.6 Configuring the IBM Spectrum Scale GUI to use sudo wrappers

The GUI can be configured to run on a cluster where remote root access is disabled and sudo wrappers are used. On such a cluster, the GUI process still runs as root, but it issues SSH to other nodes by using a user name, for which sudo wrappers were configured.

Make the following configuration changes to use the IBM Spectrum Scale management GUI on a cluster where sudo wrappers are used:

1. Issue the **`mmchconfig sudoUser=gpfsadmin`** command to configure the user name.

2. Issue `/usr/lpp/mmfs/gui/cli/runtask DAEMON_CONFIGURATION` to refresh GUI configuration.

Passwordless SSH is set up between the root user on the node where the GUI is running on all the remote nodes in the cluster. The SSH calls are equivalent to SSH gpfsadmin@destination-node. Therefore, it is not necessary to set up passwordless SSH between gpfsadmin users on any two nodes.

The root user of the node where the GUI is running can do passwordless SSH to any other node using the gpfsadmin user login. Therefore, unidirectional access from the GUI node to the remote nodes as gpfsadmin user is sufficient.

> **Note:** If sudo wrappers are enabled on the cluster but GUI is not configured for it, the system raises an event.

2.7 Ensuring high availability of the GUI service

Multiple GUI nodes must be configured in the system to ensure high availability of the GUI service. You also must set up a CCR when you plan to configure multiple GUI nodes in the cluster. The CCR is used to store certain important configuration details that must be shared among all GUI nodes.

A GUI high availability configuration with two GUI nodes is shown in Figure 2-3.

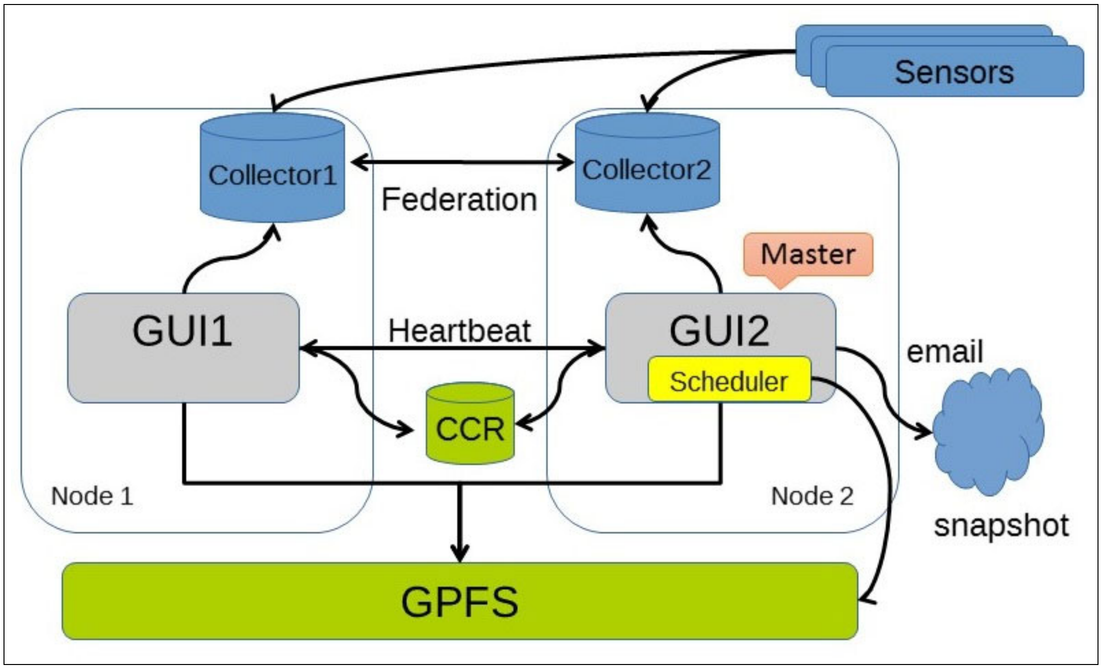

Figure 2-3 GUI high availability configuration

When configured for high availability, GUI nodes include the following main aspects:

- The GUI nodes are configured in the active/active configuration. Up to three GUI nodes can be configured in a cluster. All GUI nodes are fully functional and can be used in parallel.

- Each GUI has its own local configuration cache in PostgreSQL and collects configuration changes individually.
- One GUI node is elected as the *master node*. This GUI instance exclusively performs tasks that must be run only once in a cluster, such as running snapshot schedules and sending email and SNMP notifications.

  If services that are run on the master GUI node are configured, the environment for all the GUI nodes must support these services on all nodes. For example, ensure that access to SMTP and SNMP servers is possible from all GUI nodes and not only from the master GUI node.

  You can use the utility function that is shown in Example 2-3, which displays the current master GUI node.

Example 2-3   Master GUI node

```
[root@gpfsgui-11 ~]# /usr/lpp/mmfs/gui/cli/lsnode
Hostname             IP          Description     Role               Product version Connection status GPFS status Last updated
gpfsgui-11.novalocal 10.0.100.12 Master GUI Node management,storage 4.2.4.0         HEALTHY           HEALTHY     7/10/17 10:19 AM
gpfsgui-12.novalocal 10.0.100.13                 storage,ces        4.2.4.0         HEALTHY           HEALTHY     7/10/17 10:19 AM
gpfsgui-13.novalocal 10.0.100.14                 storage,ces        4.2.4.0         HEALTHY           HEALTHY     7/10/17 10:19 AM
```

- All GUI nodes are equal from the user’s perspective. If a GUI node fails, the user must manually connect to the other GUI. The master role fails over automatically, but no failover exists for the IP address of the other GUI server.
- Data that cannot be gathered from GPFS is stored in CCR as shared-cluster repository. This type of data includes GUI users, groups and roles, snapshot schedules, email notification settings, policy templates, and ACL templates.
- All GUI nodes must run on the same software level.
- If an external authentication method is used to store the GUI user details and authenticate them, such as AD or LDAP, the AD/LDAP configuration must be done on all GUI nodes to ensure high availability. If an internal authentication method is used, the GUI nodes get the user information from the CCR.
- To display the performance monitoring information, a performance monitoring collector must be installed on each GUI node, and these collectors must be in the federated mode. The data collection from the sensors can be configured in such a way that the details are sent to all collectors or only to a single collector.

  You can configure the collector redundancy by modifying the **`colRedundancy`** option in `/opt/IBM/zimon/defaults/ZIMonSensors.cfg`. For more information about how to ensure collector redundancy, see IBM Knowledge Center.
- The Mark as Read operation can be performed on events that are stored locally on the GUI node. The changes that are made to the events are not visible through the other GUI node.
- Each GUI has its own local configuration cache and collects configuration changes individually.
- A corrupted cache database affects only the local GUI; other GUIs continue working. Most of the configuration changes are simultaneously reported in the GUIs. Some configuration changes are gathered through the individually scheduled refresh tasks, which might result in displaying unsynchronized information.

2.8 Node classes used for the IBM Spectrum Scale GUI

The IBM Spectrum Scale GUI automatically creates the following node classes during the installation:

- `GUI_SERVERS`: Contains all nodes with a server license and all the GUI nodes.
- `GUI_MGMT_SERVERS`: Contains all the GUI nodes.

Each node on which the GUI services are started is added to these node classes.

> **Note:** These node classes must not be modified manually because the GUI regularly checks and possibly updates the node class members.

2.9 Modifying the IBM Spectrum Scale GUI property file

The GUI stores some settings that can be adjusted in the following properties file:

```
/usr/lpp/mmfs/gui/conf/gpfsgui.properties
```

The properties file is not maintained over upgrades; therefore, modifications to this file must be reapplied when the GUI is upgraded. Typically, this file does not need to be updated.

The important settings that can be modified in the properties file are listed in Table 2-3.

Table 2-3 Setting in the properties file

Setting	Description
KEEP_LOG_INTERVAL=168	Defines the number of hours the logs must be kept before they are discarded.
JDBC_DB_URL=jdbc:postgresql://localhost:5432/postgres	Sets the URL of the GUI owned postgres database.
MAX_ALLOWED_TIME_DIFF=60	Defines the maximum time difference (in seconds) that is allowed between the GUI node and any other cluster node before an event is triggered. It is important to synchronize the time across the cluster to ensure proper functioning of various features that are configured in the system. You can use NTP or some other similar function to synchronize time across the cluster.
ZIMonAddress=localhost, ZIMonPort=9084	Sets the host name and port where the performance collector service is running. The only supported configuration is hosting the collector on the node where the GUI is running.
GPFS_ADMIN=root	Specifies the current GPFS admin user. On a system with sudo wrappers enabled, it is automatically changed to a selected sudoUser.

2.10 Distributed GUI preferences

In contrast to the GUI properties file, the GUI preferences are shared between two GUIs through CCR and persist even after uninstalling or upgrading the GUI. These GUI preference files start with “_gui” and can be viewed by using the **`mmccr flist`** command. The local version of the GUI preferences is stored at `/var/lib/mmfs/gui`.

The distributed preferences contain various information that are needed in the distributed environment such as user repository, LDAP and certificate settings, user account templates, policies, and thresholds.

> **Note:** Distributed GUI preferences must not be edited manually.

You can also use the **Preferences** tab in the **Services** → **GUI** page to set the following options for the GUI node:

- Login message

  A message that can be displayed in the login page of the management GUI. Usually, this message is used to display some important information that must be shared with other users. For example, “Do not alter snapshot configuration”, “To get access to the system, please contact…… “, and so on. You can specify only up to 160 characters in the message.

- Session timeout

  The system automatically logs out the user after a specified period of inactivity.

- Display cluster name on the banner

  You can enter a name for the cluster and display it on the banner.

  > **Note:** The settings that are made under the Preferences tab are stored centrally in CCR. Therefore, the settings that are made for one GUI node are applicable to all GUI nodes of the cluster.

3

Understanding the GUI options

This chapter provides an overview of the basic options that are available in the IBM Spectrum Scale GUI and includes the following topics:

- 3.1, "Login" on page 24
- 3.2, "Navigation" on page 25
- 3.3, "Header area" on page 26
- 3.4, "Assistance for understanding the features associated with a GUI page" on page 28

3.1 Login

The users who are created by using the **Services** → **GUI** → **Users** page can log in to the GUI. When you log in for the first time after the installation, the GUI log in window guides you through the process of how to create the first GUI user.

Figure 3-1 shows the login window.

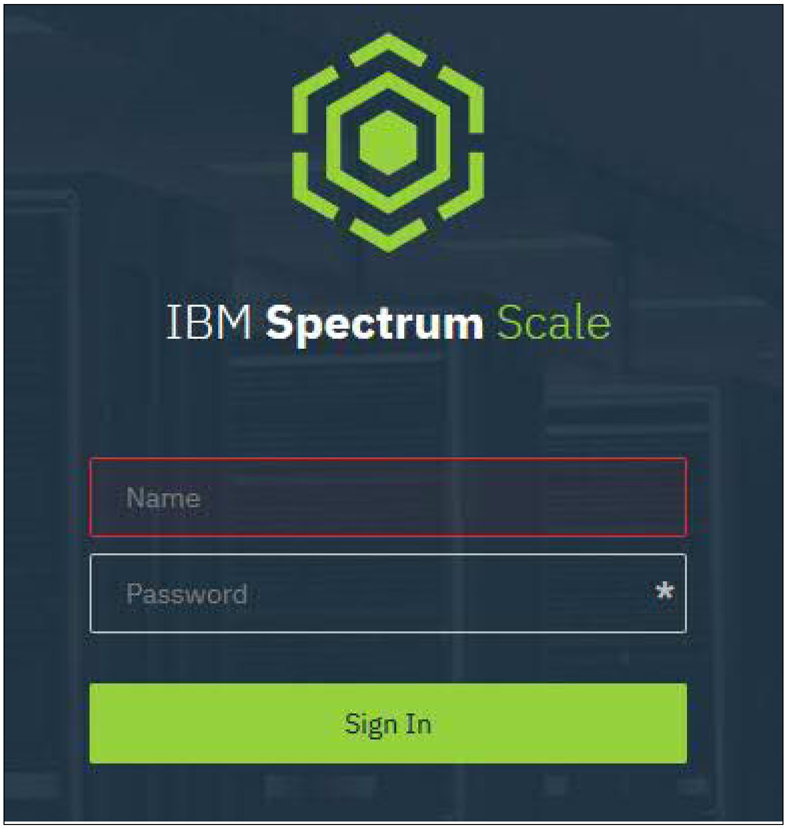

Figure 3-1 Login window

3.2 Navigation

You can navigate to the various GUI pages by using the navigation menu on the left side of the GUI window (see Figure 3-2). Each GUI page features a unique URL that you can use to directly access the page, bookmark pages, and start the GUI in-context.

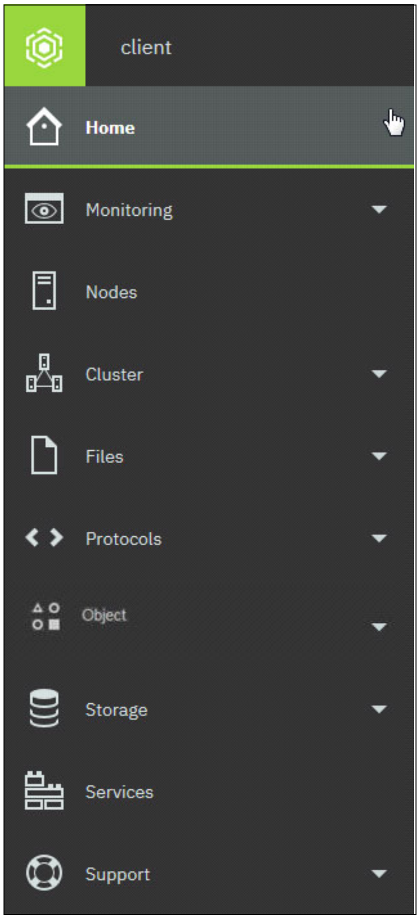

Figure 3-2 Navigation menu of the GUI

Note: The Object and Protocols menus are displayed only when these features are enabled in the cluster.

3.3 Header area

The header area provides the following details:

- A list of events of type *tips*, which provides recommendations to avoid certain issues in the future (see Figure 3-3).

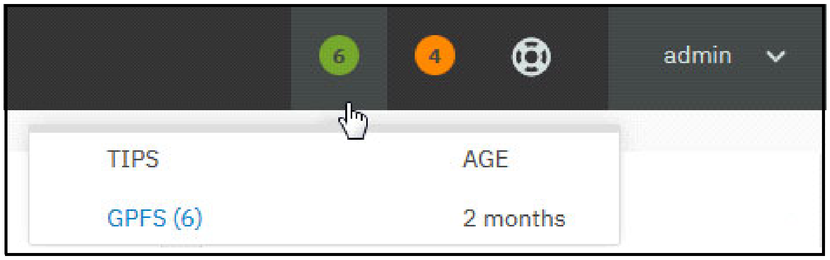

TIPS	AGE
GPFS (6)	2 months

Figure 3-3 Type “tips”

- The health status of various services, which displays only events that are in the Warning or Critical status (see Figure 3-4).

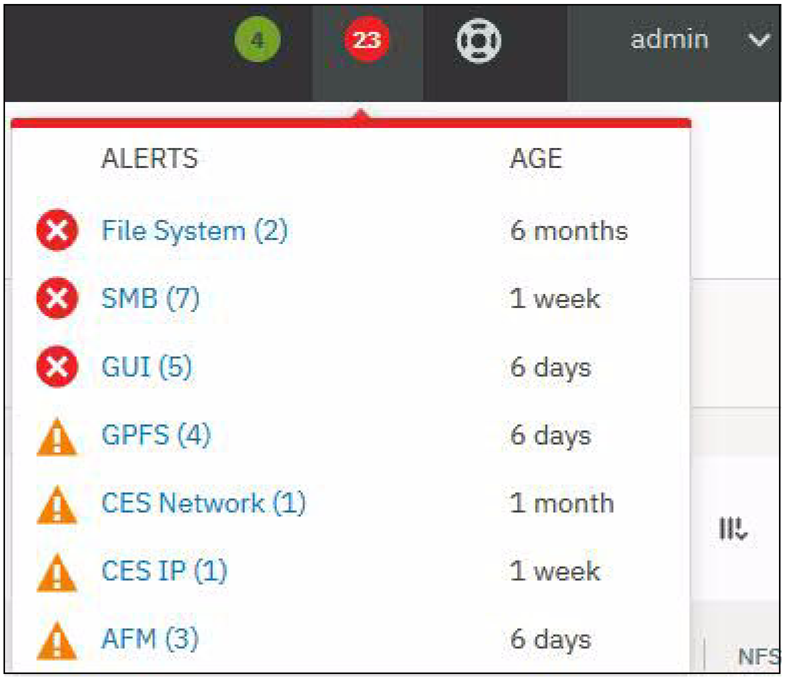

ALERTS	AGE
File System (2)	6 months
SMB (7)	1 week
GUI (5)	6 days
GPFS (4)	6 days
CES Network (1)	1 month
CES IP (1)	1 week
AFM (3)	6 days

Figure 3-4 Health status of even (Warning or Critical status only)

- A link to the context-sensitive help page (see Figure 3-5). This help file provides a detailed explanation of the features that are associated with the page. The context-sensitive help files are available in the Help menu, which is in the upper-right corner of the GUI page.

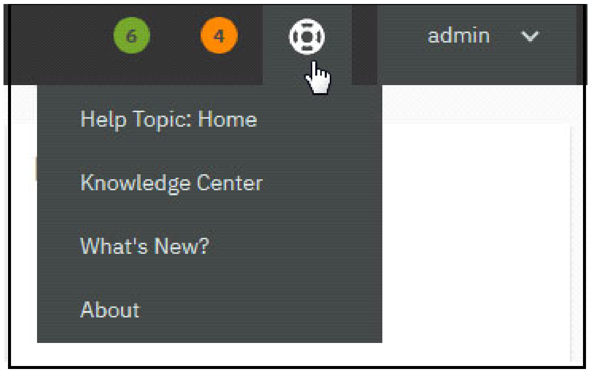

Figure 3-5 Link to the context-sensitive help page

- A link to the IBM Spectrum Scale Knowledge Center is also available in the Help menu, which is in the upper-right corner of the GUI page.
- The currently logged in user name, log out, and provide feedback options are available in the user menu (see Figure 3-6).

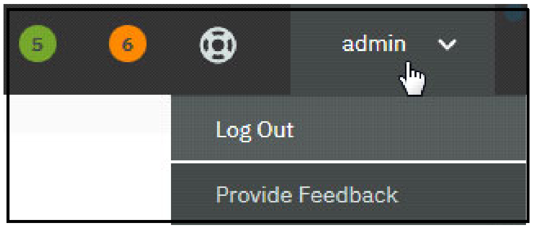

Figure 3-6 Login information

- Connection indicators that show active data transfers between browser and the GUI server (blue light indicates, “Loading”) and connection issues (yellow light indicates, “Disconnected”), as shown in Figure 3-7.

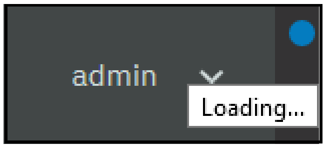

Figure 3-7 Connection indicators

3.4 Assistance for understanding the features associated with a GUI page

The following levels of assistance are available for the GUI users:

- Hover help

  When you hover the mouse over the tiny question mark next to the field label, the system displays a brief description of the feature that is associated with that field. Hover help is available only for the important and complex fields.

- Context-sensitive help

  Provides a detailed explanation of the features that are associated with the page. The context-sensitive help files are available in the Help menu, which is in the upper right corner of the GUI page.

- IBM Spectrum Scale Knowledge Center

  The IBM Knowledge Center provides the entire details of the product. A link to IBM Spectrum Scale Knowledge Center is also available in the help menu, which is in the upper-right corner of the GUI page.

4

Monitoring options available in the IBM Spectrum Scale GUI

This chapter describes the various monitoring options that are available with the IBM Spectrum Scale GUI and includes the following topics:

- 4.1, “Monitoring performance” on page 30
- 4.2, “Using the dashboard to view performance charts” on page 35
- 4.3, “Monitoring waiters” on page 36
- 4.4, “Monitoring capacity” on page 36
- 4.5, “Monitoring system health” on page 39
- 4.6, “Monitoring nodes” on page 41
- 4.7, “Monitoring Transparent Cloud Tiering” on page 44
- 4.8, “Monitoring Active File Management” on page 45
- 4.9, “Monitoring file systems” on page 48
- 4.10, “Monitoring filesets” on page 50
- 4.11, “Monitoring pools” on page 51
- 4.12, “Monitoring NSDs” on page 52
- 4.13, “Monitoring networks by using GUI” on page 53
- 4.14, “Monitoring remote cluster through GUI” on page 55
- 4.15, “Monitoring thresholds” on page 58
- 4.16, “Monitoring command audit log” on page 63

4.1 Monitoring performance

You can use the IBM Spectrum Scale GUI to monitor the status and historical trends of key indicators. You can then use this data to help make decisions more quickly and efficiently.

Table 4-1 lists the performance monitoring options that are available in the IBM Spectrum Scale GUI.

Table 4-1 Performance monitoring options

GUI page	Function
Monitoring → Statistics	Displays the performance of system resources and file and Object Storage in various performance charts. You can select the required charts and monitor the performance based on the filter criteria. You can also pan and zoom charts with detailed metrics and display past intervals.
Monitoring → Dashboards	Provides an easy-to-read and real-time user interface that shows a graphical representation of the status and historical trends of key performance indicators. This view helps users make more efficient decisions.
Nodes	Provides an easy way to monitor the performance, health status, and configuration aspects of all available nodes in the IBM Spectrum Scale cluster.
Cluster → Network	Provides the performance details, health status, and configuration aspects of network components.
Files → File Systems	Provides a detailed view of the performance, capacity, and health aspects of individual file systems.
Storage → NSDs	Provides a detailed view of the performance, capacity, and health aspects of individual Network Shared Disks (NSDs).
Storage → Pools	Provides a detailed view of the performance, capacity, and health aspects of storage pools.
Files → Active File Management	Provides a detailed view of the configuration, performance, and health status of AFM cache relationship, AFM disaster recovery (AFMDR) relationship, and gateway nodes.
Files → Filesets	Provides a detailed view of the fileset and its capacity details.
Files → Transparent Cloud Tiering	Provides insight into health, performance, and configuration of the Transparent Cloud Tiering feature.
Protocols → NFS Exports	Provides an overview of the performance aspects of the NFS export.
Protocols → SMB Shares	Provides an overview of the performance aspects of the SMB shares.

The performance and capacity data are collected with the help of the following components:

- Sensors: The *sensors* are placed on all the nodes, and they share the data with the collector. The sensors run on any node that is required to collect metrics.
- Collector: The *collector* collects data from the sensors. The metric collector must run at least on one node to gather metrics from all the nodes that are running the associated sensors. The metrics are stored in a database on the collector node.

The collector ensures aggregation of data when data gets older. The collector can run on any node in the system. You can configure multiple collectors in the system. To configure performance monitoring through the GUI, you must configure a collector on each GUI node.

The performance monitoring configuration for the GUI is shown in Figure 4-1.

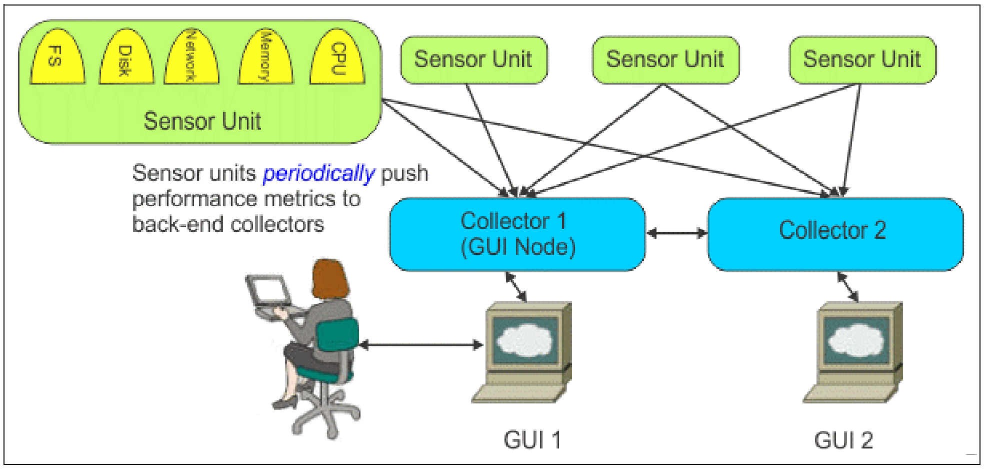

Figure 4-1 Performance monitoring configuration for the GUI

You can use the **Services** → **Performance Monitoring** page to configure sensors. You can also use the mmperfmon command to configure the performance data collection through the CLI. The GUI displays a subset of the available metrics that are available in the performance monitoring tool.

The performance monitoring tool installation can have a single collector, or can consist of multiple collectors to increase the scalability or the fault-tolerance of the performance monitoring system. This latter configuration is referred to as “federation”.

You can configure the system to monitor the performance of the following functional areas in the system:

- Network
- InfiniBand network
- System resources
- NSD server
- IBM Spectrum Scale client
- NFS
- SMB
- Object
- CTDB
- Transparent Cloud Tiering
- AFM
- Waiters

Note: The functional areas, such as NFS, SMB, Object, CTDB, and Transparent Cloud Tiering, are available only if the feature is enabled in the system.

4.1.1 Display options in the Statistics page of the GUI

The **Monitoring** → **Statistics** page is used for selecting the attributes based on which aspects of the performance of the system are monitored and comparing the performance based on the selected metrics. You can monitor the performance aspects of local cluster and remote clusters.

You can also use this page to monitor capacity. The customized charts that are marked as favorite charts can be selected n when you add widgets in the dashboard. You can display either or two charts at a time in the Statistics page.

The predefined performance charts and metrics help in investigating every node or any specific node that is collecting the metrics. Figure 4-2 shows various configuration options that are available in the Statistics page of the GUI.

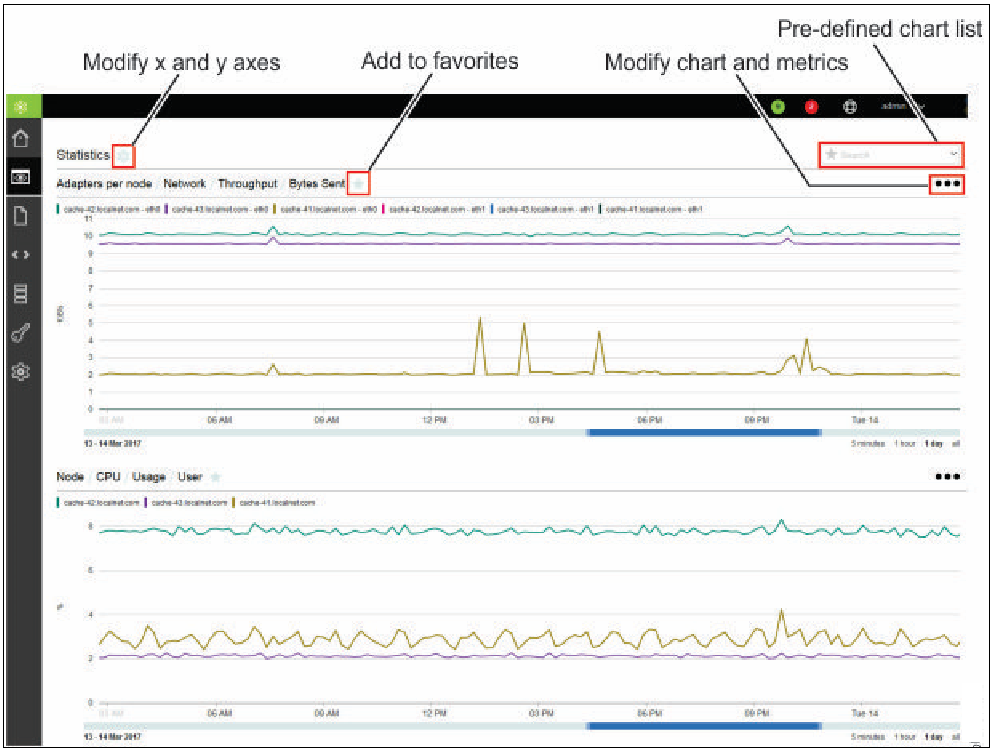

Figure 4-2 Statistics page in the IBM Spectrum Scale GUI

Predefined charts can be selected from a predefined chart list. You can add charts to the predefined list by clicking the **Favorites** button.

4.1.2 Display options in performance charts

The charting section displays the performance details based on various aspects. The GUI provides a rich set of controls to view performance charts. You can use these controls to perform the following actions on the charts that are displayed in the page:

- Zoom-in the chart by using the mouse wheel or resizing the timeline control. The y-axis is automatically adjusted during zooming.
- Drag the chart or the timeline control at the bottom. The y-axis is automatically adjusted during panning.

- Compare charts side by side. You can synchronize the y-axis and bind the x-axis. To modify the x and y axes of the chart, click the configuration symbol that is next to the title Statistics and select the wanted options.
- Link the timelines of the two charts together by using the display options that are available.

The Dashboard page helps to access all charts, which are predefined or custom-created favorites.

4.1.3 Select performance and capacity metrics

To monitor the performance and capacity of the system, use the predefined charts or select the appropriate metrics and create a chart, as shown in Figure 4-3.

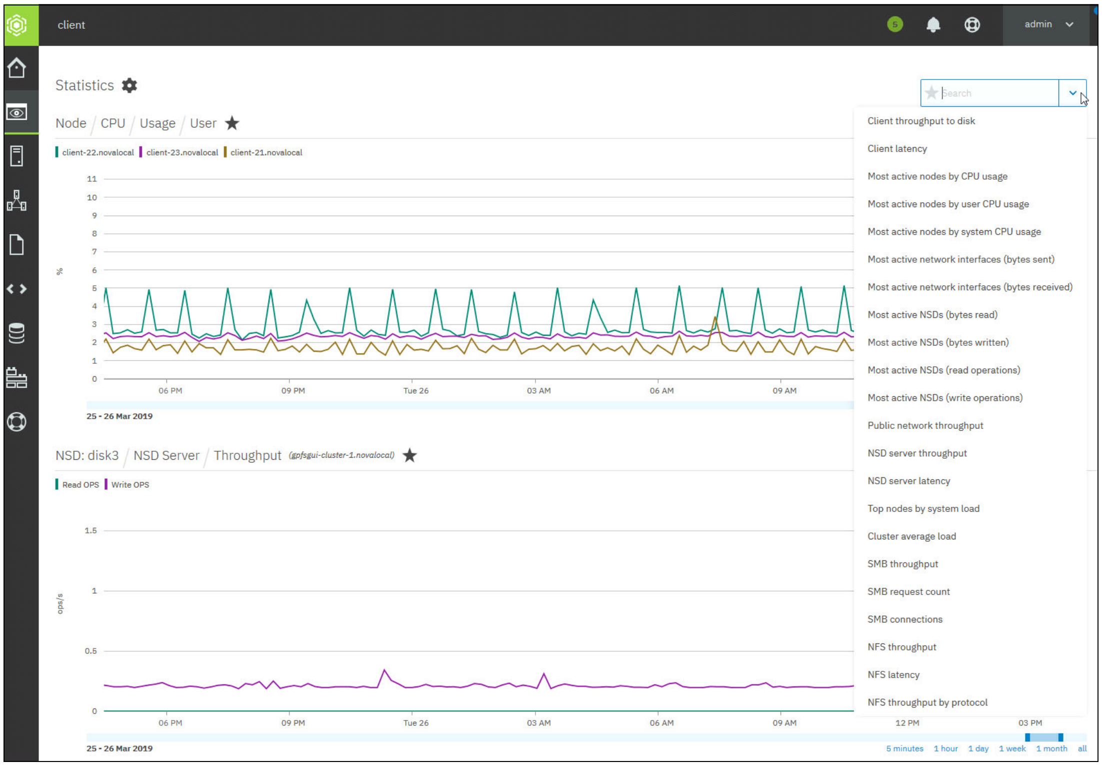
client
5
admin
Statistics
Search
Node / CPU / Usage / User
client-22.novalocal
client-23.novalocal
client-21.novalocal
%
11
10
9
8
7
6
5
4
3
2
1
0
06 PM
09 PM
Tue 26
03 AM
06 AM
09 AM
25 - 26 Mar 2019
NSD: disk3 / NSD Server / Throughput (gpfsgui-cluster-1.novalocal)
Read OPS
Write OPS
ops/s
1.5
1
0.5
0
06 PM
09 PM
Tue 26
03 AM
06 AM
09 AM
12 PM
03 PM
25 - 26 Mar 2019
5 minutes
1 hour
1 day
1 week
1 month
all
Client throughput to disk
Client latency
Most active nodes by CPU usage
Most active nodes by user CPU usage
Most active nodes by system CPU usage
Most active network interfaces (bytes sent)
Most active network interfaces (bytes received)
Most active NSDs (bytes read)
Most active NSDs (bytes written)
Most active NSDs (read operations)
Most active NSDs (write operations)
Public network throughput
NSD server throughput
NSD server latency
Top nodes by system load
Cluster average load
SMB throughput
SMB request count
SMB connections
NFS throughput
NFS latency
NFS throughput by protocol

Figure 4-3 List of predefined charts in the Statistics page

4.1.4 Create customized performance charts

Complete the following steps to create customized performance charts:

1. Click **Edit** in the menu to view the modification options. The performance and capacity options appear, as shown in Figure 4-4.

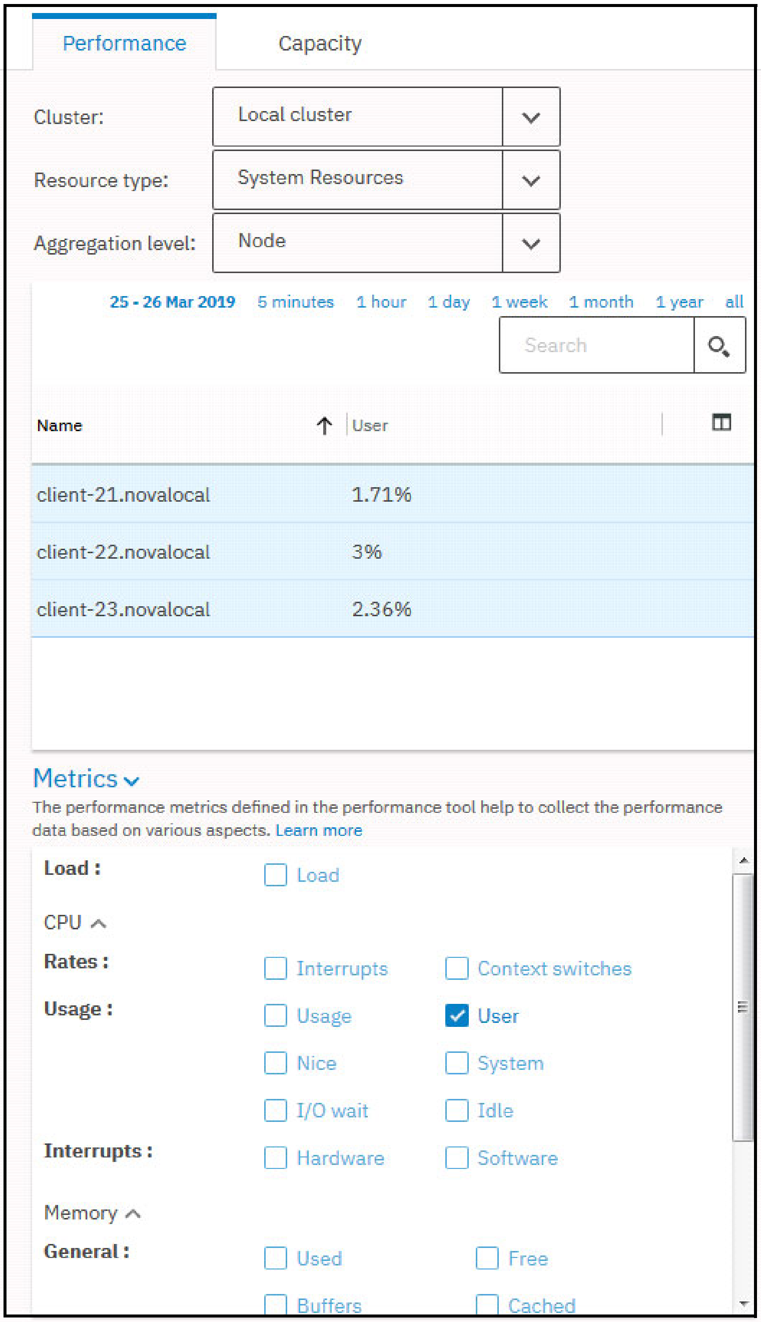

Cluster: Local cluster

Resource type: System Resources

Aggregation level: Node

Name	User
client-21.novalocal	1.71%
client-22.novalocal	3%
client-23.novalocal	2.36%

The performance metrics defined in the performance tool help to collect the performance data based on various aspects. Learn more

Load : □ Load

CPU

Rates : □ Interrupts □ Context switches

Usage : □ Usage ☑ User

□ Nice □ System

□ I/O wait □ Idle

Interrupts : □ Hardware □ Software

Memory

General : □ Used □ Free

□ Buffers □ Cached

Figure 4-4 Options to create a performance chart on the Statistics page

The performance metrics are grouped under the combination of resource types and aggregation levels. The resource types determine the area from which the data is taken to create the performance analysis and aggregation level determines the level at which the data is aggregated.

The aggregation levels that are available for selection varies based on the resource type.

2. Select whether you need to monitor performance of the local cluster or remote cluster from the **Cluster** field.
3. Select the required resource type from the **Resource type** field.
4. Select the aggregation level from the **Aggregation level** field.
5. Select the resource that you want to monitor.
6. Select the time frame for the performance data display.
7. Select the metrics that you want to be displayed on the performance chart.
8. Click **Apply** to apply the changes or **Close** to cancel the process.

4.2 Using the dashboard to view performance charts

The **Monitoring** → **Dashboard** page provides an easy-to-read, single-page, and real-time user interface that provides a quick overview of the system performance. The dashboard consists of several dashboard widgets and the associated favorite charts that can be displayed within a chosen layout.

The following important widget types are available in the dashboard:

- Statistics
- File system capacity by fileset
- System health events
- System overview
- Filesets with the largest growth rate in last week
- Timeline

Figure 4-5 highlights the configuration options that are available in edit mode of the dashboard.

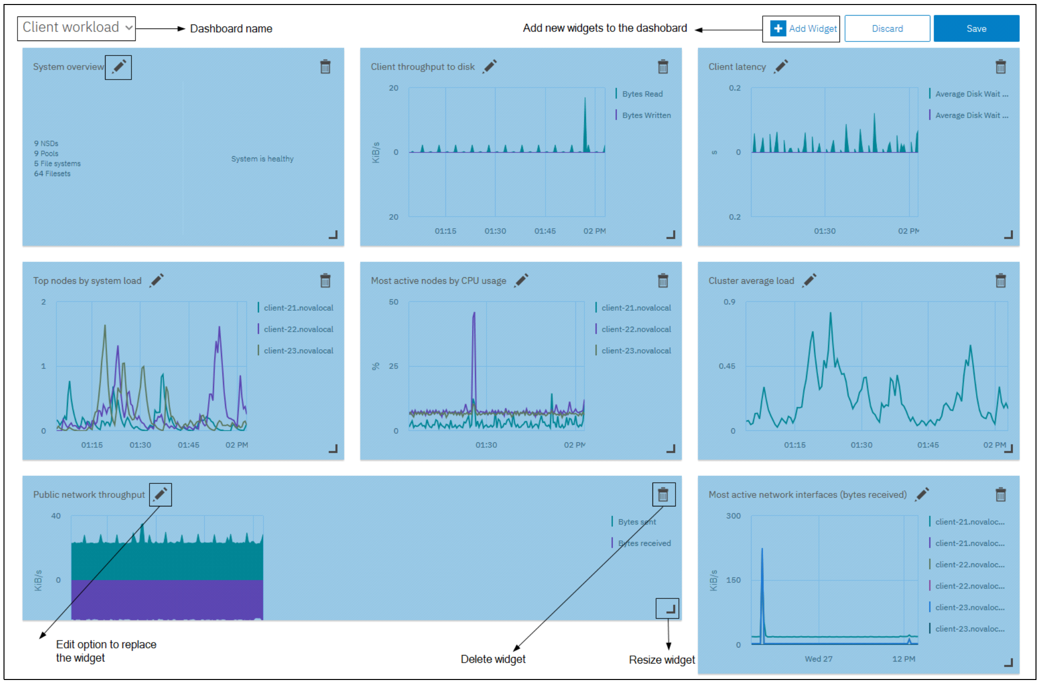

Dashboard name

Add new widgets to the dashboard

Edit option to replace the widget

Delete widget

Resize widget

Figure 4-5 Dashboard page in edit mode

4.2.1 Custom Dashboards

Select the **Create Dashboard** and **Delete Dashboard** options from the menu in the upper-right corner of the **Dashboard** page to create and delete dashboards. The dashboards are stored centrally in the cluster configuration repository (CCR). If several GUI nodes are configured in a cluster, all the saved dashboards are available to all GUI users on all nodes.

When you open the IBM Spectrum Scale GUI after the installation or upgrade, you can see the default dashboards that are included with the product. You also can modify or delete the default dashboards to suit your requirements.

4.2.2 Display options

Select **Display Options** from the menu that is available in the upper-right corner of the Dashboard GUI page to change the display options.

4.2.3 Widget options

Several dashboard widgets can be added in a selected dashboard layout. Select the **Edit Widgets** option from the menu that is available on the upper-right corner of the Dashboard GUI page to edit or remove widgets in the dashboard. You can also modify the size of the widget in the edit mode. Use the **Add Widget** option that is available in the edit mode to add widgets to the dashboard.

The widgets with type Performance list the charts that are marked as favorite charts in the Statistics page of the GUI. Favorite charts (along with the predefined charts) can be selected when you add widgets in the dashboard.

To create favorite charts, click the **Star** icon next to the chart title on the **Monitoring** → **Statistics** page.

4.3 Monitoring waiters

The metrics that are related to waiters are collected on the cluster level. Different wait time thresholds, such as 0.1s, 0.2s, 0.5s, 1s, 30s, and 60s can be selected. You can create charts to monitor the waiters by selecting **Waiters** as the resource type in the **Monitoring** → **Statistics** page.

4.4 Monitoring capacity

You can monitor the capacity of file system, pools, filesets, users, and user groups.

The capacity details that are displayed in the GUI are obtained from the following sources:

- GPFS quota database. The system collects the quota details for users, groups, and filesets daily and stores them in the postgres database.
- Performance monitoring tool. The GUI queries the performance monitoring capacity and displays capacity data in various pages in the GUI.

Based on the source of the capacity information, different procedures must be performed to enable capacity and quota data collection.

For GPFS quota database and performance monitoring tool-based capacity and quota collection, use the **Files** → **Quotas** page to enable quota data collection per file system and enforce quota limit checking. If quota is not enabled for a file system, the following results can occur:

- No capacity and inode data is collected for users, groups, and filesets.
- Quota limits for users, groups, and filesets cannot be defined.
- No alerts are sent, and the data writes are not restricted.

To enable capacity data collection from the performance monitoring tool, the `GPFSFilesetQuota` sensor must be enabled. For more information about how to enable the performance monitoring sensor for capacity data collection, see *Manual installation of IBM Spectrum Scale GUI in IBM Spectrum Scale: Concepts, Planning, and Installation Guide*.

4.4.1 Capacity data obtained from the GPFS quota database

The capacity and quota information that is collected from the GPFS quota database is displayed on the **Files** → **Quotas** and **Files** → **User Capacity** pages in the management GUI:

- **Files** → **Quotas** page

  Use quotas to control the allocation of files and data blocks in a file system. You can create default, user, group, and fileset quotas on the **Quotas** page.

  A *quota* is the amount of disk space and the amount of metadata that is assigned as upper limits for a specified user, group of users, or fileset. Use the **Actions** menu to create or modify quotas. The management GUI allows you to manage only capacity-related quota. The inode-related quota management is possible in the command-line interface only.

  You can specify a soft limit, a hard limit, or both. When you set a soft limit quota, a warning is sent to the administrator when the file system is close to reaching its storage limit. A grace period starts when the soft quota limit is reached. Data is written until the grace period expires, or until the hard quota limit is reached. Grace time resets when used capacity falls below the soft limit.

  If you set a hard limit quota, you cannot save data after the quota is reached. If the quota is exceeded, you must delete files or raise the quota limit to store more data.

  > **Note:** Consider the following points:
  >
  > - User or user group quotas for filesets are supported only if the **Per Fileset** option is enabled at the file system level. Use the command-line interface to set the option. For more information, see the **man** pages of **`mmcrfs`** and **`mmchfs`** commands.
  > - You must unmount a file system to change the quota enablement method from per file system to per fileset or vice versa.

  You can set default user quotas at the file system level rather than defining user quotas explicitly for each user. Default quota limits can be set for users. You can specify the general quota collection scope, such as per file system or per fileset to define whether the default quota must be defined at file system level or fileset level and set the default user quota. After this value is set, all child objects that are created under the file system or fileset are configured with the default soft and hard limits. You can assign a custom quota limit to individual child objects, but the default limits remain the same unless changed at the file system or fileset level.

  After reconfiguring quota settings, it is recommended to run the **`mmcheckquota`** command for the affected file system to verify the changes.

For more information about how to manage quotas, see the “Managing GPFS quotas” section of *IBM Spectrum Scale: Administration Guide*.

Capacity data from users, groups, and filesets with no quota limit set are not listed in the **Quotas** page. Use the **Files** → **User Capacity** page to see capacity information of such users and groups. Use the **Files** → **Filesets** page to view current and historic capacity information of filesets.

- **Files** → **User Capacity** page

  The **Files** → **User Capacity** page provides predefined capacity reports for users and groups. Although capacity information of file systems, pools, and filesets is available in the respective areas of the GUI, the **Files** → **User Capacity** page is the only place where information about used capacity per user or group is available.

  The User Capacity page depends on the quota accounting method of the file system. You must enable quota for a file system to display the user capacity data. If quota is not enabled, you can follow the fix procedure in the **Files** → **Quotas** page or use the `mmchfs <Device> -Q yes` CLI command to enable quota.

  Even if the capacity limits are not set, the User Capacity page shows data when the quota accounting is enabled, and users can start writing the data. This feature is different in the Quotas page, where only users and groups with quota limits that are defined are listed. The user and group capacity quota information is automatically collected once per day by the GUI.

  For users and user groups, you can see the total capacity and whether quotas are set for these objects. You can also see the percentage of soft limit and hard limit usage. When the hard limit is exceeded, no other files that belong to the respective user, user group, or fileset can be written. However, exceeding the hard limit allows a certain grace period before disallowing more file writes. Soft and hard limits for disk capacity are measured in units of kilobytes (KiB), megabytes (MiB), or gigabytes (GiB). Use the **Files** → **Quotas** page to change the quota limits.

4.4.2 Capacity data collected through the performance monitoring tool

The historical capacity data collection for file systems, pools, and filesets depend on the correctly configured data collection sensors for fileset quota and disk capacity. When the IBM Spectrum Scale system is installed through the installation toolkit, the capacity data collection is configured by default. In other cases, you need to enable capacity sensors manually.

If capacity data collection is not configured correctly, you can use `mmperfmon` CLI command or the **Services** → **Performance Monitoring** → **Sensors** page.

The **Services** → **Performance Monitoring** → **Sensors** page allows to view and edit the sensor settings. By default, the collection periods of capacity collection sensors are set to collect data with a period of up to one day. Therefore, it might take a while until the data is refreshed in the GUI.

The following sensors are collecting capacity-related information and are used by the GUI:

GPFSDiskCap NSD, Pool, and File System level capacity. Uses the `mmdf` command in the background and typically runs once per day because it is resource-intensive. Should be restricted to run on a single node only.

GPFSPool Pool and file system level capacity. Requires a mounted file system and typically runs every 5 minutes. Should be restricted to run on a single node only.

GPFSFilesetQuota Fileset capacity based on the quota collection mechanism. Typically, runs every hour. Should be restricted to run only on a single node.

GPFSFileset Inode space (independent fileset) capacity and limits. Typically runs every 5 minutes. Should be restricted to run only on a single node.

DiskFree Overall capacity and local node capacity. Can run on every node

4.4.3 Capacity information for file systems, pools, NSDs, and filesets

The **Monitoring** → **Statistics** page helps to create customized capacity reports for file systems, pools, and filesets. You can store these reports as favorites and add them to the dashboard.

The dedicated GUI pages combine information about configuration, health, performance, and capacity in one place. The following GUI pages provide the corresponding capacity views:

- **Files** → **File Systems**
- **Files** → **Filesets**
- **Storage** → **Pools**
- **Storage** → **NSDs**

The Filesets grid and details depend on quota that is obtained from the GPFS quota database and the performance monitoring sensor GPFSFilesetQuota. If quota is disabled, the system displays a warning message in the Filesets page.

4.4.4 Viewing the capacity for licensing

You can use capacity-based licensing to license IBM Spectrum Scale based on the storage capacity that is managed in an IBM Spectrum Scale cluster. The storage capacity to be licensed is the capacity in terabytes (TiB) under all NSDs in the IBM Spectrum Scale cluster. You can view the cluster capacity for licensing by selecting the **About** option from the menu that is available in the upper-right corner of the management GUI.

4.5 Monitoring system health

The system health monitoring options that are available in the IBM Spectrum Scale GUI are listed in Table 4-2.

Table 4-2 System health monitoring options available in the IBM Spectrum Scale GUI

GUI page	Function
Monitoring → **Events**	Lists the events that are reported in the system. You can monitor and troubleshoot errors on your system on the Events page.
Monitoring → **Tips**	Lists the events of type tips, which provide recommendations about certain events that might occur in the future.
Home	Provides the overall system health of the IBM Spectrum Scale system.
Monitoring	Displays the health status of nodes and lists the events that are reported at the node level.
Cluster → **Network**	Displays health status and configuration aspects of all available networks and interfaces that are part of the networks.

GUI page	Function
Files → **File Systems**	Displays the health status of file systems and lists the events that are reported at the file system level.
Files → **Filesets**	Displays the health status of filesets and lists the events that are reported at the fileset level.
Files → **Transparent Cloud Tiering**	Lists the events that are reported for the Transparent Cloud Tiering service. The GUI displays this page only if the Transparent Cloud Tiering feature is enabled in the system.
Files → **Active File Management**	Displays health status and lists events that are reported for AFM cache relationship, AFM disaster recovery (AFMDR) relationship, and gateway nodes.
Storage → **Pools**	Displays health status and lists events that are reported for storage pools.
Storage → **NSDs**	Displays health status of NSDs and lists the events that are reported at the NSD level.
Monitoring → **Command Audit Log**	Displays a record of various actions that are performed on the system. This information helps the system administrator to audit the commands and tasks that are performed by the administrators. These logs can also be used to troubleshoot issues that are reported in the system.
Health indicator (available in the upper-right corner of the GUI)	Displays the number of events with warning and critical status that are specific to each component.
System overview widget in the **Monitoring** → **Dashboard** page	Displays the number of events that were reported against each component.
System health events widget in the **Monitoring** → **Dashboard** page	Provides an overview of the warning and error events that are reported in the system.
Timeline widget in the **Monitoring** → **Dashboard** page	Displays the events that are reported in a selected time frame on the selected performance chart.

4.6 Monitoring nodes

The **Nodes** page provides options to monitor the performance, health status, and configuration aspects of the respective components that are shown in Figure 4-6.

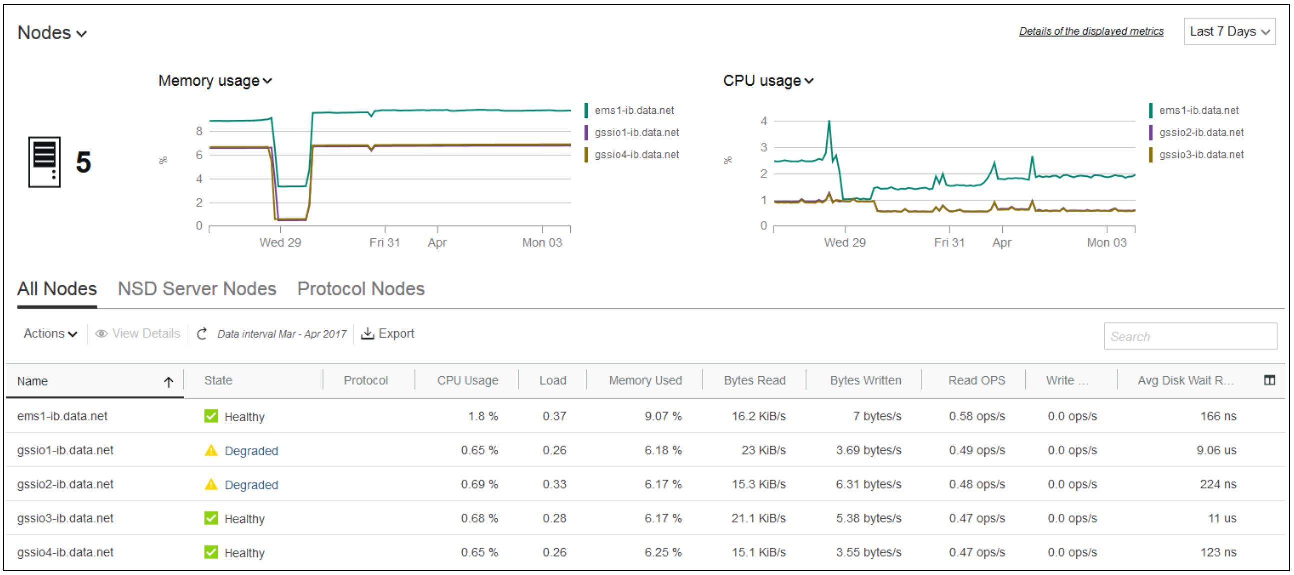

Name	State	Protocol	CPU Usage	Load	Memory Used	Bytes Read	Bytes Written	Read OPS	Write ...	Avg Disk Wait R...
ems1-ib.data.net	✓ Healthy		1.8 %	0.37	9.07 %	16.2 KiB/s	7 bytes/s	0.58 ops/s	0.0 ops/s	166 ns
gssio1-ib.data.net	⚠ Degraded		0.65 %	0.26	6.18 %	23 KiB/s	3.69 bytes/s	0.49 ops/s	0.0 ops/s	9.06 us
gssio2-ib.data.net	⚠ Degraded		0.69 %	0.33	6.17 %	15.3 KiB/s	6.31 bytes/s	0.48 ops/s	0.0 ops/s	224 ns
gssio3-ib.data.net	✓ Healthy		0.68 %	0.28	6.17 %	21.1 KiB/s	5.38 bytes/s	0.47 ops/s	0.0 ops/s	11 us
gssio4-ib.data.net	✓ Healthy		0.65 %	0.26	6.25 %	15.1 KiB/s	3.55 bytes/s	0.47 ops/s	0.0 ops/s	123 ns

Figure 4-6 Nodes page

The properties of a node display the status of various CES services, such as Object, NFS, and SMB and the authentication status of these services if they are enabled. It also displays other details, such as network status and information about attached NSDs and file systems.

4.6.1 Nodes tables

The following nodes tables provide a prefiltered view on nodes, each with specific information:

- All nodes

  Shows all nodes in the cluster and provides information about node roles, services, node health, and basic performance information about system and IBM Spectrum Scale client level.

- NSD server nodes

  Shows all nodes that are NSD servers with specific performance information that is related to NSDs. If no NSD servers are in the cluster, this table is not displayed.

- Protocol nodes

  Shows all protocol nodes. Specific performance and health information that is related to protocol services is displayed in this table. If no protocol nodes are in the cluster, this table is not displayed.

You can customize the nodes tables individually by adding or removing columns by using the **Customize Columns** feature.

You can use the **Set Attributes** option that is available in the Actions menu to set the node attributes, such as site, room, and rack on any of the views. You can set attributes of multiple nodes at a time. The attributes can be used to filter nodes in the nodes view.

The health status information of each service and component can have the following values:

- Healthy: The component is working as expected.
- Disabled: The component is not enabled.
- Suspended: When a CES is in the suspended state, most components also report as suspended.
- Starting: The component (or monitor) was recently started. This transient state is updated after the startup is complete.
- Unknown: Something is preventing the monitoring from determining the state of the component.
- Stopped: The component was intentionally stopped. This situation might occur briefly if a service is restarted because of a configuration change. It might also occur if a user issues the **`mmces service stop protocol`** command for a node.
- Degraded: A problem occurred with the component that is not a complete failure. This state does not cause the CES addresses to be reassigned.
- Failed: The monitoring detected a significant problem with the component that means it is unable to function correctly. This state causes the CES addresses of the node to be reassigned.
- Dependency failed: This state implies that a component has a dependency that is in a failed state; for example, an NFS or SMB service shows Dependency failed if authentication is failed.

4.6.2 Performance monitoring of nodes

The Nodes page provides the following options to analyze the performance of nodes:

- A *quick view* gives the number of nodes in the system and the overall performance of nodes based on CPU and memory usages. You can access this view by selecting the **Expand** button that is next to the title of the page. You can close this view if not required.

  Many graphs in the overview show the three nodes that feature the highest average performance metric over a past period. These graphs are refreshed regularly. The refresh intervals of the top three entities are depended on the displayed time frame:

  – Every minute for the 5-minute time frame
  – Every 15 minutes for the 1-hour time frame
  – Every six hours for the 24-hour time frame
  – Every two days for the 7-day time frame
  – Every seven days for the 30-day time frame
  – Every four months for the 365-day time frame

- A *nodes table* displays many different performance metrics. To find nodes with extreme values, you can sort the values that are displayed in the nodes table by different performance metrics. Click the performance metric in the table header to sort the data based on that metric.

  You can select the time range that determines the averaging of the values that are displayed in the table and the time range of the charts in the overview from the time range selector. This control is in the upper right corner. The metrics in the table do not update automatically. Use the **Refresh** button that is above the table to refresh the table content with more recent data.

- A *detailed view* of the performance and health aspects of individual nodes is available on the **Nodes** page. Select the node for which you need to view the performance details and select **View Details**. The system displays various performance charts in the right pane.

The detailed performance view helps to drill-down to various performance aspects. The following performance details can be obtained from each tab of the performance view:

- The Overview tab provides a performance chart for the following information:
  - Client IOPS
  - Client data rate
  - Server data rate
  - Server IOPS
  - Network
  - CPU
  - Load
  - Memory
- The Events tab can be used to monitor the events that are reported in the node. Similar to the Events page, you can also perform the operations like marking events as read and running fix procedure from the Events tab. Only the current issues are displayed in this view. The **Monitoring** → **Events** page displays the entire set of events that are reported in the system.
- The File Systems tab provides performance details of the file systems that are mounted on the node. The file system's read or write throughput, average read or write transactions size, and file system read or write latency are also available.

  You can also mount or unmount individual file systems or multiple file systems on the selected node. For more information, see 5.2, "Mounting a file system through the GUI" on page 66 and 5.3, "Unmounting a file system by using the GUI" on page 67.
- The NSD tab provides the status of the disks that are attached to the node. The NSD tab is displayed only when the node is configured as an NSD server.
- The SMB and NFS tabs provide the performance details of the SMB and NFS services that are hosted on the node. These tabs appear in the chart only when the node is configured as a protocol node.
- The Network tab displays the network performance details.
- The AFM tab displays the details of the AFM and AFM DR relationships for which the node is configured as a gateway node.
- The Properties tab provides an overview of the node-related attributes. You can also use the Prevent file system mounts option to allow or prevent file systems from mounting the node.

4.6.3 Creating and managing user-defined node classes

Node classes are used to group nodes. It helps you to select only the required set of nodes when you want to limit the scope of certain administrative tasks. The following types of node classes can be defined in the IBM Spectrum Scale system:

- System node classes
- User-defined node classes

The system node classes are hardcoded, but you can create user-defined node classes by using the **Nodes** → **Node Classes** → **Create Node Class** option in the IBM Spectrum Scale GUI. While creating a node class, consider the following points:

- The name of the new node class must be different from the name of the existing nodes or node classes.
- You can add individual nodes and other existing node classes in a new node class from the All Nodes and Node Classes tabs of the Create Node Class window.

- When you add a node class in the new node class, the nodes that are part of the existing node class become part of the new node class. When nodes are added or removed in the existing node class later, those changes also are applied to the new node class.

Use the **Modify** option to change the node class name and nodes and node classes that are part of a node class. You cannot modify system node classes.

Use the **Delete** option to delete the user-defined node class. You cannot delete the system node classes.

4.7 Monitoring Transparent Cloud Tiering

Transparent Cloud Tiering is a separately installable feature of IBM Spectrum Scale that provides a native cloud storage tier. It allows data center administrators to free on-premises storage capacity by moving out cooler data to the cloud storage. Freeing storage capacity helps to reduce capital and operational expenditures.

The Transparent Cloud Tiering feature uses the ILM policy query language semantics. The system administrators can define policies to tier data to a cloud storage.

On an IBM Spectrum Scale cluster with multiple storage tiers configured, this external cloud storage can be used as the cooler storage tier to store infrequently accessed data from a cool storage pool. Because of performance reasons, avoid moving any active or hot data to this external storage pool because it drives excessive data traffic that results in delays and application timeouts.

Transparent Cloud Tiering service features the following core functions:

- Migrate: Migrates the specified files or filesets to the cloud storage tier.
- Recall: Recalls the specified files or filesets from the cloud storage tier.
- Remove: Deletes the cloud storage tier.

The **Files** → **Transparent Cloud Tiering** page (see Figure 4-7) provides performance and health information charts about the Transparent Cloud Tiering service.

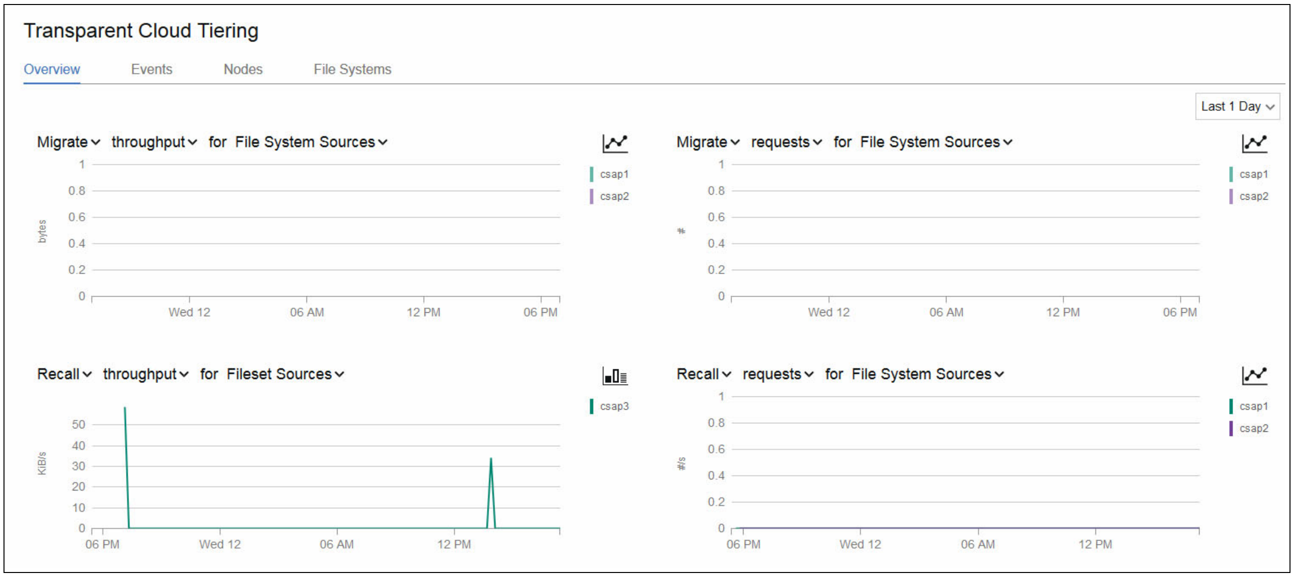
Transparent Cloud Tiering
Overview
Events
Nodes
File Systems
Last 1 Day
Migrate throughput for File System Sources
bytes
1
0.8
0.6
0.4
0.2
0
Wed 12
06 AM
12 PM
06 PM
csap1
csap2
Migrate requests for File System Sources
#
csap1
csap2
Recall throughput for Fileset Sources
KiB/s
50
40
30
20
10
06 PM
csap3
Recall requests for File System Sources
#/s
csap1
csap2

Figure 4-7 Transparent Cloud Tiering page

The following tabs are available on this page:

- Overview: Displays the aggregated data of all file systems and nodes that are associated with a particular cloud provider.
- Events: Displays events that are associated with the Transparent Cloud Tiering service.
- Nodes: Lists the node on which the cloud services are installed. The cloud services, such as Transparent Cloud Tiering, cloud data sharing, or both, can be activated on this node.
- File Systems: Displays the details of the file systems that are mapped with the Transparent Cloud Tiering service.

You can select a line chart or a bar chart to display the details. The line chart shows an average rate, whereas the bar chart shows aggregated data. For example, the aggregate view can be used by administrators to see how much data was transferred in one day.

4.8 Monitoring Active File Management

The **Files** → **Active File Management** page provides an easy way to monitor the performance, health status, and configuration aspects of the Active File Management and Active File Management disaster recovery relationships in the IBM Spectrum Scale cluster. It also provides details of the gateway nodes that are part of the Active File Management or Active File Management disaster recover relationships.

The GUI combines the following sources on the cache or primary side:

- Active File Management sensors from the performance monitoring tool
- Health status and events from the `mmhealth` component
- Active File Management configuration information

The Active File Management GUI architecture is shown in Figure 4-8.

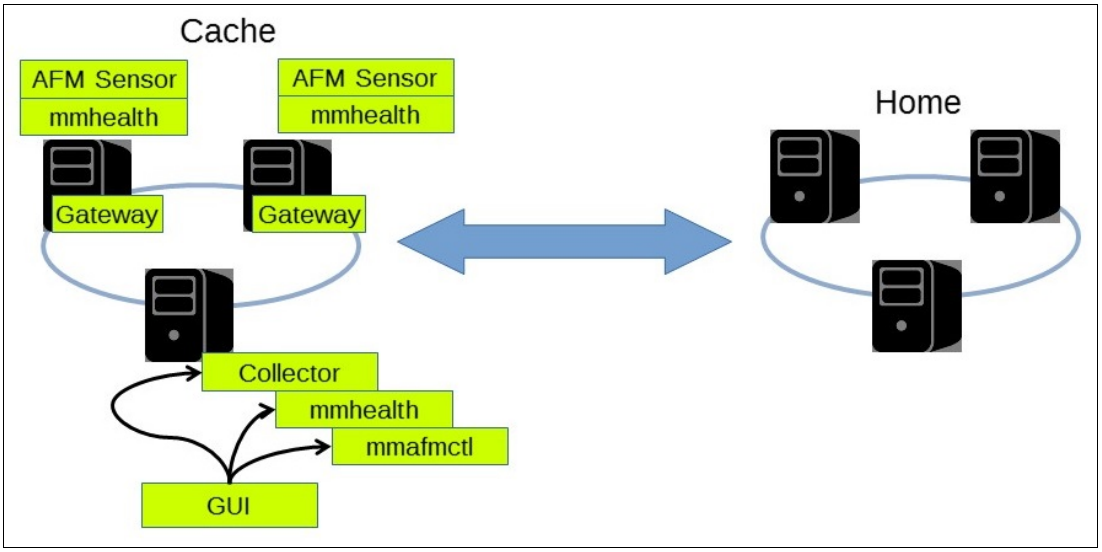

Figure 4-8 Active File Management GUI architecture

The **Active File Management** page of the GUI is shown in Figure 4-9.

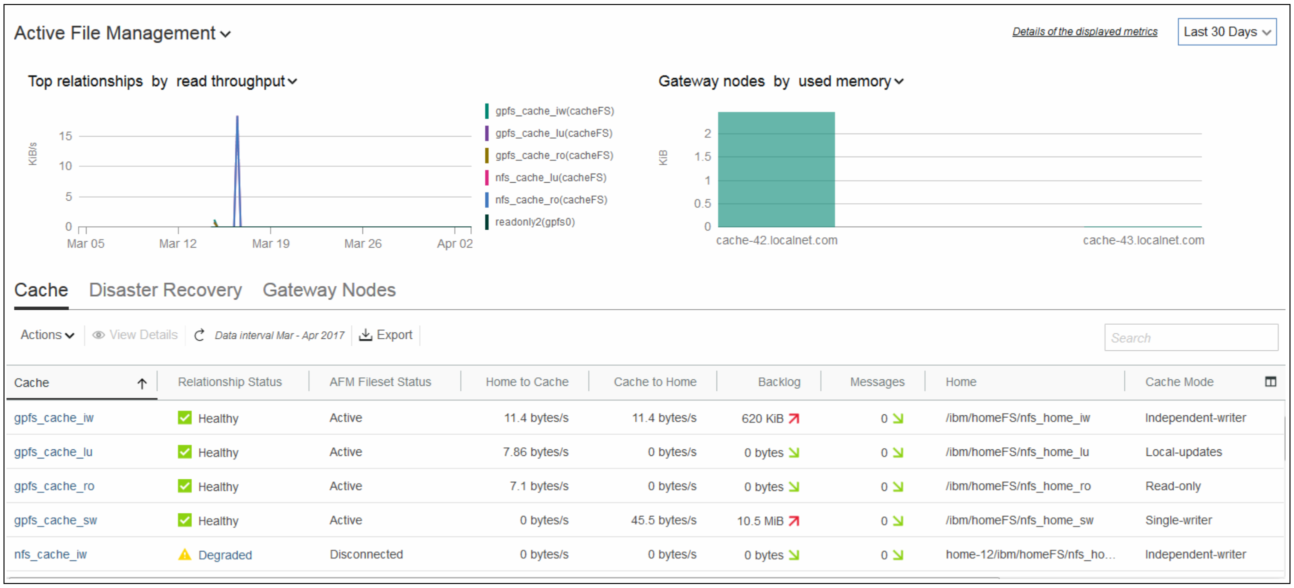

Cache ↑	Relationship Status	AFM Fileset Status	Home to Cache	Cache to Home	Backlog	Messages	Home	Cache Mode
gpfs_cache_iw	✓ Healthy	Active	11.4 bytes/s	11.4 bytes/s	620 KiB ↗	0 ↘	/ibm/homeFS/nfs_home_iw	Independent-writer
gpfs_cache_lu	✓ Healthy	Active	7.86 bytes/s	0 bytes/s	0 bytes ↘	0 ↘	/ibm/homeFS/nfs_home_lu	Local-updates
gpfs_cache_ro	✓ Healthy	Active	7.1 bytes/s	0 bytes/s	0 bytes ↘	0 ↘	/ibm/homeFS/nfs_home_ro	Read-only
gpfs_cache_sw	✓ Healthy	Active	0 bytes/s	45.5 bytes/s	10.5 MiB ↗	0 ↘	/ibm/homeFS/nfs_home_sw	Single-writer
nfs_cache_iw	⚠ Degraded	Disconnected	0 bytes/s	0 bytes/s	0 bytes ↘	0 ↘	home-12/ibm/homeFS/nfs_ho…	Independent-writer

Figure 4-9 The Active File Management GUI

The following options are available to monitor Active File Management and Active File Management disaster recovery relationships and gateway nodes:

- A *quick view* gives the details of top relationships between cache and home sites in an Active File Management or Active File Management disaster recovery relationship. It also provides the performance information for gateway nodes that are by used memory and the number of queued messages. The graphs that are displayed in the quick view are refreshed regularly. The refresh intervals depend on the following selected time frame:
  - Every minute for the 5-minute time frame
  - Every 15 minutes for the 1-hour time frame
  - Every 6 hours for the 24-hour time frame
  - Every two days for the 7-day time frame
  - Every seven days for the 30-day time frame
  - Every four months for the 365-day time frame
- Different performance metrics and configuration details display in the tabular format. The following tables are available:
  - Cache

    Provides the information about configuration, health, and performance of the Active File Management feature that is configured for data caching and replication.
  - Disaster Recovery

    Provides information about configuration, health, and performance of Active File Management disaster recovery configuration in the cluster.
  - Gateway Nodes

    Provide details of the nodes that are designated as the gateway node in the Active File Management or Active File Management disaster recovery configuration.

  To find an Active File Management or Active File Management disaster recovery relationship or a gateway node with extreme values, sort the values that are displayed on the table by different attributes. Click the performance metric in the table header to sort the data based on that metric.

You can select the time range that determines the averaging of the values that are displayed in the table and the time range of the charts in the overview from the time range selector, which is placed in the upper right corner. The metrics in the table do not update automatically. You can refresh the table with more recent data by clicking the **Refresh** button that is above the table.

- A *detailed view* of the performance and health aspects of the individual Active File Management or Active File Management disaster recovery relationship or gateway node is available. To see the detailed view, double-click the row that lists the relationship or gateway node of which you need to view the details, or select the item from the table and click **View Details**.

  The following details are available for each item:

  - Cache:
    - Overview: Provides the number of available cache inodes and displays charts that show the amount of data that is transferred, data backlog, and memory that is used for the queue.
    - Events: Provides details of the system health events that are reported for the Active File Management component.
    - Snapshots: Provides details of the snapshots that are available for the Active File Management fileset. The snapshots are taken for backup purposes. The snapshot that is taken in the Active File Management cache relationship is called a *peer snapshot*. It functions in the same way as the GPFS snapshots. When a snapshot is taken on the cache site, it also propagates the request to take a snapshot of the home.
    - Gateway Nodes: Provides details of the nodes that are configured as gateway node in the Active File Management configuration.
  - Disaster recovery:
    - Overview: Provides the number of available primary inodes and displays charts that show the amount of data that is transferred, data backlog, and memory that are used for the queue.
    - Events: Provides details of the system health events that are reported for the Active File Management component.
    - Snapshots: Provides details of the snapshots that are available for the AFM fileset. The snapshots that are taken in the Active File Management disaster recover are called *recovery point objective (RPO) snapshots*. These peer snapshots are taken at the same time on the primary and the secondary sites.
    - Gateway Nodes: Provides details of the nodes that are configured as gateway node in the Active File Management configuration.
  - Gateway Nodes:

    The details of gateway nodes are available under the following tabs. The same details are available in the **Nodes** page.

    The Overview tab provides performance chart for the following information:

    - Client IOPS
    - Client data rate
    - Server data rate
    - Server IOPS
    - Network
    - CPU
    - Load
    - Memory

The Events tab provides details of the events that are reported in the node. Similar to the Events page, you can perform the operations, such as marking events as read and running fix procedures from this events view. Only current issues are shown in this view. The **Monitoring** → **Events** page displays the entire set of events that are reported in the system.

The File Systems tab provides performance details of the file systems that are mounted on the node. The file systems' read or write throughput, average read or write transactions size, and file system read or write latency are also available. Use the **Mount File System** or **Unmount File System** option to mount or unmount individual file systems or multiple file systems on the selected node. The nodes on which the file system are mounted or unmounted can be selected individually from the list of nodes or based on node classes.

The NSD tab provides status of the disks that are attached to the node. The NSD tab appears only if the node is configured as an NSD server.

The SMB and NFS tabs provide the performance details of the SMB and NFS services that are hosted on the node. These tabs appear in the chart only if the node is configured as a protocol node.

The AFM tab provides details of the configuration and status of the Active File Management and Active File Management disaster recovery relationships for which the node is configured as the gateway node.

The Network tab displays the network performance details.

The Properties tab displays the basic attributes of the node. You can use the **Prevent file system mounts** option to specify whether you can prevent file systems from mounting on the node.

4.9 Monitoring file systems

The **Files** → **File Systems** page (see Figure 4-10) provides options to monitor the performance, health status, and configuration aspects of the all available file systems in the IBM Spectrum Scale cluster.

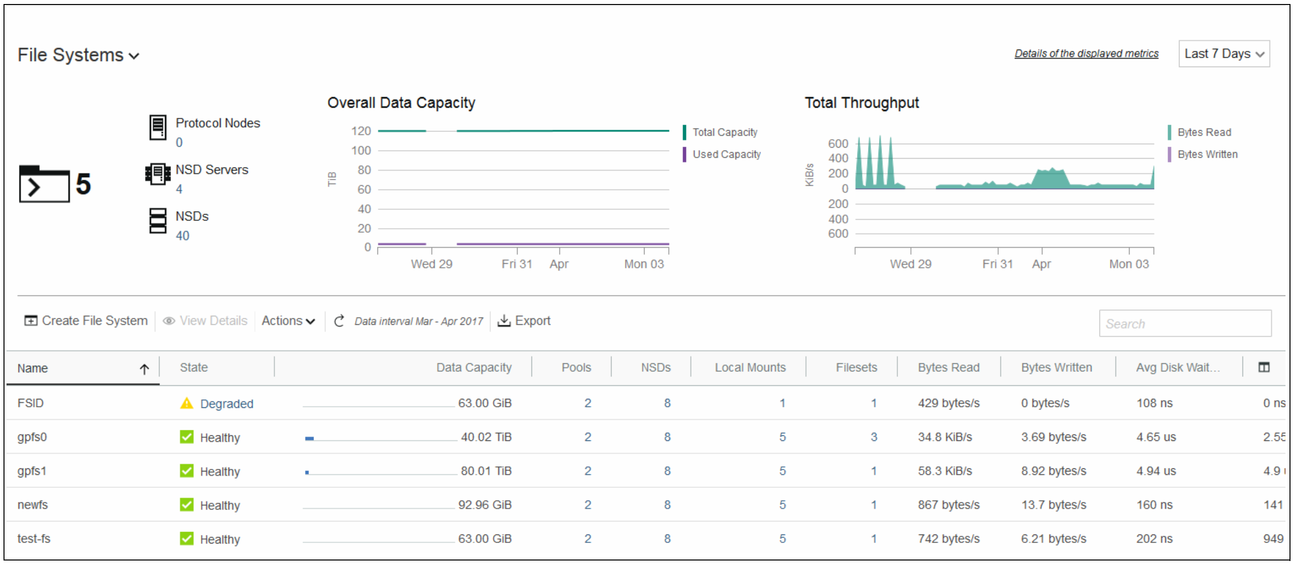

Name	State	Data Capacity	Pools	NSDs	Local Mounts	Filesets	Bytes Read	Bytes Written	Avg Disk Wait...	
FSID	Degraded	63.00 GiB	2	8	1	1	429 bytes/s	0 bytes/s	108 ns	0 ns
gpfs0	Healthy	40.02 TiB	2	8	5	3	34.8 KiB/s	3.69 bytes/s	4.65 us	2.55
gpfs1	Healthy	80.01 TiB	2	8	5	1	58.3 KiB/s	8.92 bytes/s	4.94 us	4.9
newfs	Healthy	92.96 GiB	2	8	5	1	867 bytes/s	13.7 bytes/s	160 ns	141
test-fs	Healthy	63.00 GiB	2	8	5	1	742 bytes/s	6.21 bytes/s	202 ns	949

Figure 4-10 File Systems page

The following options are available to analyze the file system performance:

- A *quick view* gives the number of protocol nodes, NSD servers, and NSDs that are part of the available file systems that are mounted on the GUI server. It also provides overall capacity and total throughput details of these file systems. You can access this view by selecting the **Expand** button next to the title of the page. You can close this view if it is not required.

  The graphs displayed in the quick view are refreshed regularly. The refresh intervals are depend on the following displayed time frames:

  - Every minute for the 5-minute time frame
  - Every 15 minutes for the 1-hour time frame
  - Every six hours for the 24-hour time frame
  - Every two days for the 7-day time frame
  - Every seven days for the 30-day time frame
  - Every four months for the 365-day time frame

- A *file systems table* displays many different performance metrics. To find file systems with extreme values, you can sort the values that are displayed in the file systems table by using the different performance metrics. Click the performance metric in the table header to sort the data based on that metric.

  You can select the time range that determines the averaging of the values that are displayed in the table and the time range of the charts in the overview from the time range selector. This control is in the upper-right corner. The metrics in the table do not update automatically. You can use the **Refresh** button that is above the table to refresh the table with more recent data.

- A *detailed view* of the performance and health aspects of individual file systems is available. To see the detailed view, double-click the file system for which you need to view the details, or select the file system and click **View Details**.

  The detailed performance view allows you to drill down to various performance aspects. The following performance details can be obtained from each tab of the performance view:

  - Overview: Provides an overview of the file system performance.
  - Events: System health events that are reported for the file system.
  - NSDs: Details of the NSDs that are part of the file system.
  - Pools: Details of the pools that are part of the file system.
  - Nodes: Details of the nodes on which the file system is mounted. You can also perform the tasks, such as mount file system, unmount file system, and specify whether to automatically mount file system if GPFS daemon starts or prevent any mounts of this file system on selected nodes.
  - Filesets: Details of the filesets that are part of the file system.
  - NFS: Details of the NFS exports that were created in the file system.
  - SMB: Details of the SMB shares that were created in the file system.
  - Object: Details of the IBM Spectrum Scale Object Storage on the file system.
  - Properties: Provides details of the file system attributes. You can also use the **Automatic mount** option to configure the automatic mount mode of the file system.

4.10 Monitoring filesets

The **Files** → **Filesets** page (see Figure 4-11) provides an easy way to monitor the performance, health status, and configuration aspects of all available filesets in the IBM Spectrum Scale cluster.

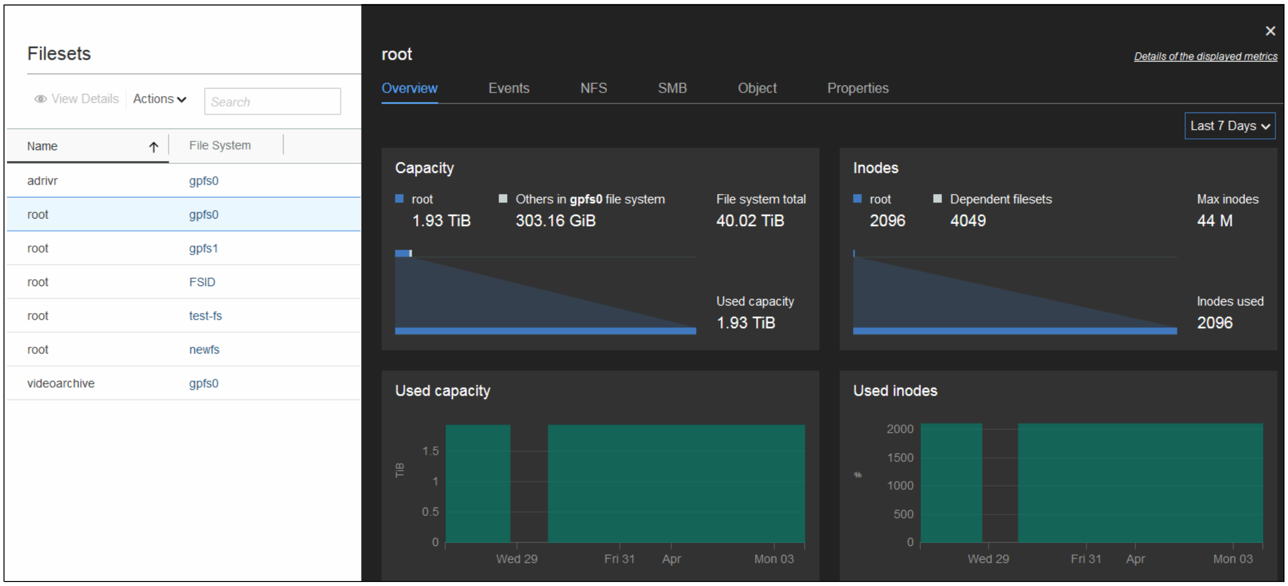

Filesets

Name	File System
adrivr	gpfs0
root	gpfs0
root	gpfs1
root	FSID
root	test-fs
root	newfs
videoarchive	gpfs0

root

Capacity: root 1.93 TiB; Others in **gpfs0** file system 303.16 GiB; File system total 40.02 TiB; Used capacity 1.93 TiB

Inodes: root 2096; Dependent filesets 4049; Max inodes 44 M; Inodes used 2096

Figure 4-11 Filesets page in the detailed view

The following options are available to view the fileset details:

- A *fileset table* displays the details of the filesets that are available in the system. You can sort the values that are displayed in the filesets table by different attributes. Click the column header in the table to sort the data based on that attribute.
- An *overview* section that displays the fileset size and growth rates in graphical format. If you enabled quota accounting and the GPFSFilesetQuota sensor is active, you can view reports on fileset size distribution, absolute and relative growth, and growth rates by size range. You can change the observed time frame from the upper right corner of the display. The fileset display interacts with the filesets table so that selection in the graphical display correlates to the selection in the table.
- A *detailed view* of the performance and health aspects of individual filesets is available. To see the detailed view, double-click the fileset for which you need to view the details, or select the fileset and click **View Details**.

  The detailed view allows you to drill down to various performance, health, and configuration aspects. The following details can be obtained from each tab of the performance view:

  – Overview: Provides an overview of the fileset capacity, inodes, and quota limits.

  – Events: System health events that are reported for the fileset.

  – NFS: Details about the NFS exports that were created in the fileset.

  – SMB: Details about the SMB shares that were created in the fileset.

  – Object: Details of the Object Storage policy that is mapped to the fileset. When objects are uploaded to a container, they are stored in the fileset that is associated with the container’s storage policy.

  – Properties: Provides details of the fileset attributes.

4.11 Monitoring pools

The **Storage** → **Pools** page (see Figure 4-12) provides options to monitor the performance, health status, and configuration aspects of the all available pools in the IBM Spectrum Scale cluster. The GUI shows a table of all internal pools in a cluster.

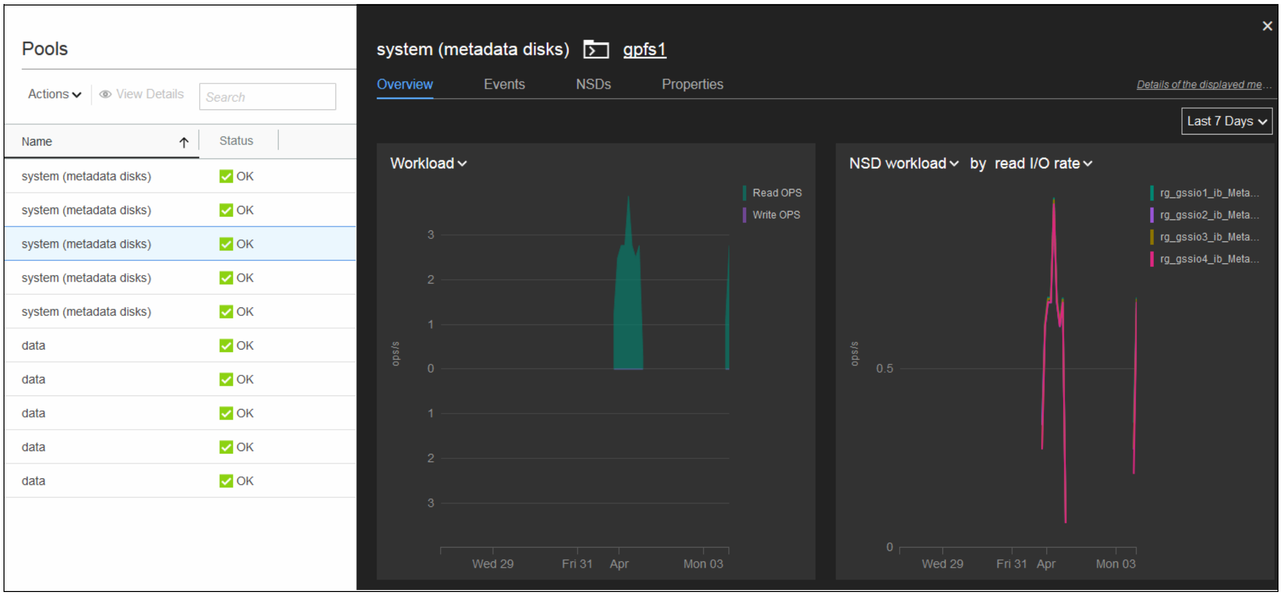

Pools

Name	Status
system (metadata disks)	✓ OK
system (metadata disks)	✓ OK
system (metadata disks)	✓ OK
system (metadata disks)	✓ OK
system (metadata disks)	✓ OK
data	✓ OK
data	✓ OK
data	✓ OK
data	✓ OK
data	✓ OK

Figure 4-12 Pools page in the detailed view

The systems pools contain metadata for the entire file system. Therefore, the GUI shows separate lines for the system pool, depending on the usage type of the NSDs in the system pool. If the system pool of a file system is used only for storing metadata, the GUI shows one row of details. If the system pool consists of NSDs that are of type *metadataOnly* and *dataOnly*, the GUI shows two rows with separate data. The detailed views for this pool also show separate performance and capacity information in the overview and NSD sections.

The following options are available to analyze the pools performance:

- A *pools table* displays many different performance metrics. To find pools with extreme values, you can sort the values that are displayed in the pools table by different performance metrics. Click the performance metric in the table header to sort the data based on that metric.

  You can select the time range that determines the averaging of the values that are displayed in the table and the time range of the charts in the overview from the time range selector. This control is in the upper right corner. The metrics in the table do not update automatically. You can refresh the table with more recent data by clicking the **Refresh** button that is above the table.

- A *detailed view* of the performance and health aspects of individual pools is available. To see the detailed view, double-click the pool for which you need to view the details, or select the pool and click **View Details**.

  The detailed performance view allows you to drill down to various performance aspects. The following performance details can be obtained from each tab of the performance view:

  – Overview: Provides an overview of the pools performance.
  – Events: Provide details of the system health events reported for the file system.
  – NSDs: Gives details of the NSDs that are part of the file system.
  – Properties: Provides an overview of the pool’s properties.

4.12 Monitoring NSDs

The **Storage** → **NSDs** page (see Figure 4-13) provides an easy way to monitor the performance, health status, and configuration aspects of the all NSDs that are available in the IBM Spectrum Scale cluster.

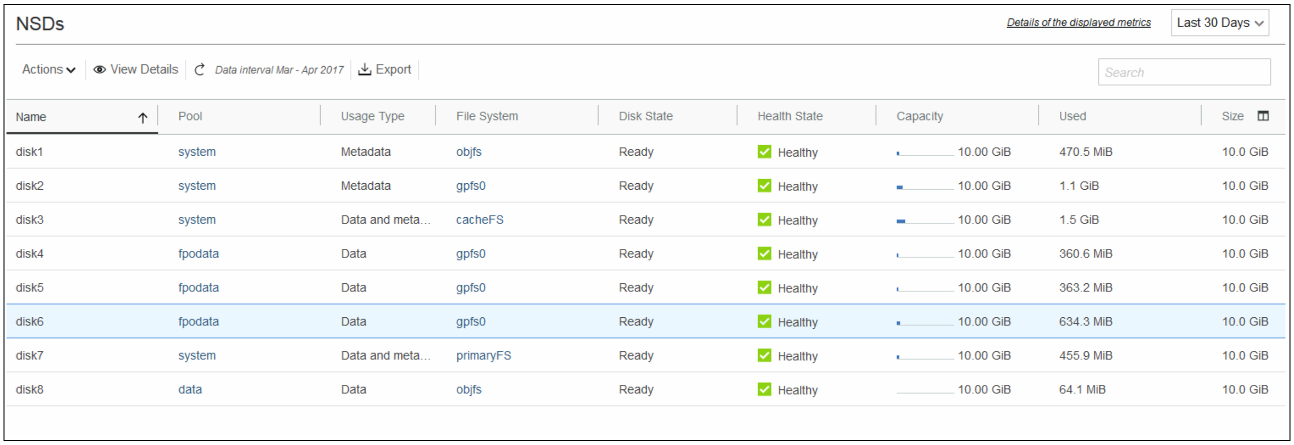
NSDs
Details of the displayed metrics
Last 30 Days
Actions
View Details
Data interval Mar - Apr 2017
Export
Search

Name	Pool	Usage Type	File System	Disk State	Health State	Capacity	Used	Size
disk1	system	Metadata	objfs	Ready	✓ Healthy	10.00 GiB	470.5 MiB	10.0 GiB
disk2	system	Metadata	gpfs0	Ready	✓ Healthy	10.00 GiB	1.1 GiB	10.0 GiB
disk3	system	Data and meta...	cacheFS	Ready	✓ Healthy	10.00 GiB	1.5 GiB	10.0 GiB
disk4	fpodata	Data	gpfs0	Ready	✓ Healthy	10.00 GiB	360.6 MiB	10.0 GiB
disk5	fpodata	Data	gpfs0	Ready	✓ Healthy	10.00 GiB	363.2 MiB	10.0 GiB
disk6	fpodata	Data	gpfs0	Ready	✓ Healthy	10.00 GiB	634.3 MiB	10.0 GiB
disk7	system	Data and meta...	primaryFS	Ready	✓ Healthy	10.00 GiB	455.9 MiB	10.0 GiB
disk8	data	Data	objfs	Ready	✓ Healthy	10.00 GiB	64.1 MiB	10.0 GiB

Figure 4-13 NSDs page

An NSD is a logical grouping of storage disks in a network on file storage systems. It provides a method for cluster-wide disk naming and high-speed access to data for applications that are running on nodes that do not have direct access to the disks.

The NSDs in the cluster might be physically attached to all nodes or serve their data through an NSD server that provides a virtual connection. You can specify up to eight NSD servers for each NSD. If one server fails, the next server in the list takes control from the failed node.

Each NSD server must have physical access to the same NSD. However, different servers can serve I/O to different non-intersecting sets of clients. The subnet functions in IBM Spectrum Scale determine which NSD server must serve a particular IBM Spectrum Scale client.

The following options are available in the NSDs page to analyze the NSD performance, health status, and configuration details:

- An *NSD table* displays the available NSDs and many different performance metrics. To find NSDs with extreme values, you can sort the values that are displayed in the table by different performance metrics. Click the performance metric in the table header to sort the data based on that metric.

  You can select the time range that determines the averaging of the values that are displayed in the table from the time range selector. This control is in the upper-right corner. The metrics in the table are refreshed based on the selected time frame. You can refresh it manually to view the latest data.

- A detailed view of the performance and health aspects of individual NSDs is also available in the NSDs page. Select the NSD for which you need to view the performance details and click **View Details**. The system displays details of the NSD on the right pane.

  The detailed view allows you to drill down to various performance and configuration aspects. The following details can be obtained from each tab of the detailed view:

  – Overview: Provides an overview of the NSD performance details.

  – Events: Reports system health events for the NSD.

- – Nodes: Gives details of the nodes that serve the NSDs.
- – Properties: Provides an overview of the NSD-related attributes. This tab does not provide any performance details.

Note: NSD performance metrics are not collected if the client is running on the NSD server. Therefore, the GUI does not display all SAN environments or workload from local clients.

4.13 Monitoring networks by using GUI

The **Cluster** → **Network** page (see Figure 4-14) provides an easy way to monitor the performance, health status, and configuration aspects of all available networks and interfaces that are part of the networks.

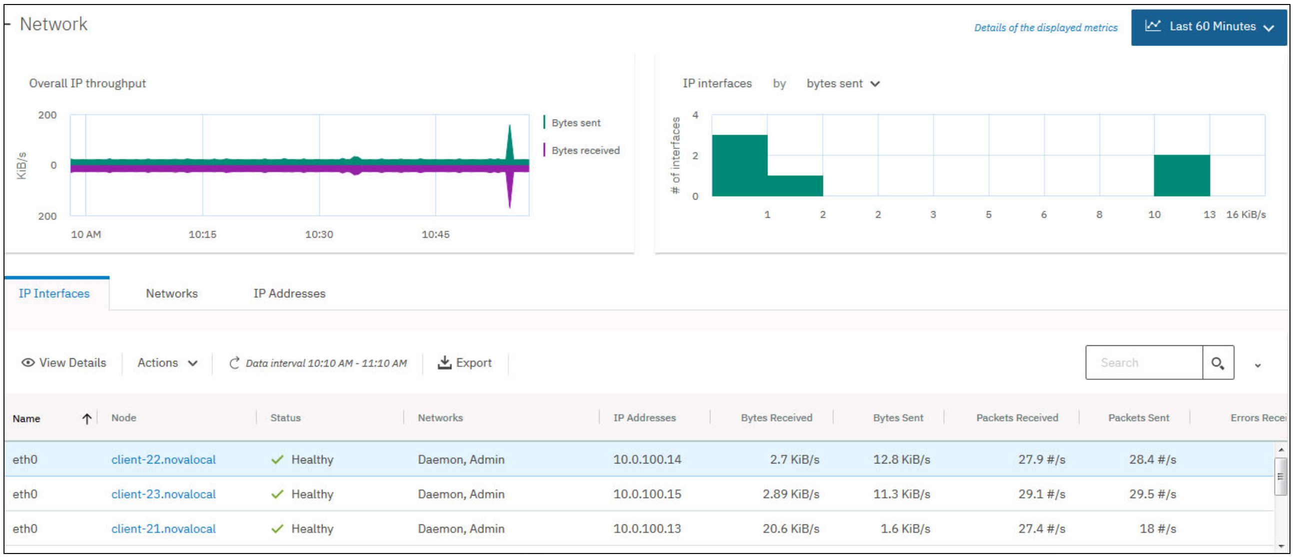

Name	Node	Status	Networks	IP Addresses	Bytes Received	Bytes Sent	Packets Received	Packets Sent	Errors Rece
eth0	client-22.novalocal	✓ Healthy	Daemon, Admin	10.0.100.14	2.7 KiB/s	12.8 KiB/s	27.9 #/s	28.4 #/s	
eth0	client-23.novalocal	✓ Healthy	Daemon, Admin	10.0.100.15	2.89 KiB/s	11.3 KiB/s	29.1 #/s	29.5 #/s	
eth0	client-21.novalocal	✓ Healthy	Daemon, Admin	10.0.100.13	20.6 KiB/s	1.6 KiB/s	27.4 #/s	18 #/s	

Figure 4-14 Network page

A dedicated network is used within the cluster for certain operations. For example, the system uses the administration network when an administration command is issued. It is also used for sharing administration-related information. This network is used for node-to-node communication within the cluster.

The daemon network is used for sharing file system or other resources data. Remote clusters also establish communication path through the daemon network. Similarly, the dedicated network types like CES network and external network can also be configured in the cluster.

The performance of network is monitored by monitoring the data transfer managed through the respective interfaces. The following types of network interfaces can be monitored through the GUI:

- IP interfaces on Ethernet and InfiniBand adapters.
- Remote Direct Memory Access (RDMA) interfaces on InfiniBand adapters with Open Fabrics Enterprise Distribution (OFED) drivers.

The GUI retrieves performance data from the performance monitoring tool. The IP-adapter-based metrics are taken from the Network sensor and the RDMA metrics are taken from the InfiniBand sensor. If no performance data appears in the GUI, verify that the monitoring tool is correctly set up and that these two sensors are enabled.

The Network page also exposes adapters and IPs that are not bound to a service to provide a full view of the network activity on a node.

The details of the networks and their components can be obtained both in graphical and tabular formats. The Network page provides the following options to analyze the performance and status of networks and adapters:

- A quick view that gives graphical representation of overall IP throughput, overall RDMA throughput, IP interfaces by bytes sent and received, and RDMA interfaces by bytes sent and received. You can access this view by selecting the expand button that is next to the title of the page. You can close this view if not required.

  Graphs in the overview are refreshed regularly. The refresh intervals of the top three entities are depended on the following displayed time frames:

  - Every minute for the 5-minutes time frame
  - Every 15 minutes for the 1-hour time frame
  - Every 6 hours for the 24-hour time frame
  - Every two days for the 7-day time frame
  - Every seven days for the 30-day time frame
  - Every four months for the 365-day time frame

  If you click a block in the IP interfaces charts, the corresponding details are displayed in the IP interfaces table. The table is filtered by the IP interfaces that are part of the selected block. You can remove the filter by clicking the link that appears above the table header row.

- A table that provides the following performance metrics that are available under the following tabs of the table:

  - IP Interfaces: Shows all network interfaces that are part of the Ethernet and InfiniBand networks in the cluster. To view performance details in graphical format or to see that the events reported against the individual adapter, select the adapter in the table and then select **Actions** → **View Details**.
  - RDMA Interfaces: Shows the details of the InfiniBand RDMA networks that are configured in the cluster. To view performance details in graphical format or to see that the events reported against the individual adapter, select the adapter in the table and then select **Actions** → **View Details**.

    The system displays the RDMA Interfaces tab only if there are RDMA interfaces available.
  - Networks: Shows all networks in the cluster and provides information on network types, health status, and number of nodes and adapters that are part of the network.
  - IP Addresses: Lists all IP addresses that are configured in the cluster.

  To find networks or adapters with extreme values, you can sort the values that are displayed in the tables by different performance metrics. Click the performance metric in the table header to sort the data based on that metric. You can select the time range that determines the averaging of the values that are displayed in the table and the time range of the charts in the overview from the time range selector, which is in the upper right corner. The metrics in the table do *not* update automatically. Click the **Refresh** button that is above the table to refresh the table content with more recent data.

- A detailed view of performance aspects and events reported against each adapter. To access this view, select the adapter in the table and then select **Actions** → **View Details**. The detailed view is available for both IP and RDMA interfaces.

4.14 Monitoring remote cluster through GUI

The IBM Spectrum Scale GUI can monitor and manage a single cluster. Cluster setups exist in which multiple clusters exchange data through AFM or cross cluster mounts. To provide consolidated monitoring of multiple clusters by using the IBM Spectrum Scale GUI, monitoring information can be exchanged among GUI nodes of different clusters.

By establishing a connection between the GUI nodes, both the clusters can monitor the other cluster. To enable remote monitoring capability among clusters, the release-level of the GUI nodes that are communicating with each other must be 5.0.0 or later.

To establish a connection with the remote cluster, complete the following steps:

1. Complete the following steps on the local cluster to raise the access request:
   a. Click **Cluster** → **Remote Connections**.
   b. Select the **Request Access** option that is available under the Outgoing Requests tab to raise the request for access.
   c. In the Request Remote Cluster Access dialog, enter an alias for the remote cluster name and specify the GUI nodes to which the local GUI node must establish the connection.
   d. If you know the credentials of the security administrator of the remote cluster, you can also add the user name and password of the remote cluster administrator and skip step 2.
   e. Click **Send** to submit the request.
2. Complete the following steps on the remote cluster to grant access:
   a. When the request for connection is received in, the GUI displays the details of the request in the **Cluster** → **Remote Connections** → **Incoming Requests** page.
   b. Select Grant Access to grant the permission and establish the connection.

Now, the requesting cluster GUI can monitor the remote cluster. To enable both clusters to monitor each other, repeat the procedure with reversed roles through the respective GUIs.

Note: Only the GUI user with Security Administrator role can grant access to the remote connection requests.

You can see the details of the connections established with the remote clusters under the Remote Cluster tab.

The remote cluster monitoring options that are available in the GUI are listed in Table 4-3.

Table 4-3 Remote cluster monitoring options available in GUI

GUI option	Description
Home	The Remote clusters grouping shows the following details: ► Number of remote clusters that are connected to the resource cluster. ► Number of file systems that are mounted on the local nodes. ► Number of local nodes on which the remote file systems are mounted.
Files → File Systems	The grid view provides the following remote cluster monitoring details: ► Whether the file system is mounted on a remote cluster. ► Capacity information. ► Number of local nodes on which the file system is mounted. ► Performance details. ► Pools, NSDs, filesets, and snapshots.
Files → File Systems → View Details → Remote Nodes	Provides the details of the remote cluster nodes where the local file system
Files → Filesets	The Remote Fileset column in the filesets grid shows whether the fileset belongs to a remote file system. The fileset table also displays the same level of details for both remote and local filesets; for example, capacity, parent file system, inodes, AFM role, and snapshots.
Files → Active File Management	When remote monitoring is enabled, you can view the following AFM details: ► On home and secondary, you can see the AFM relationships configuration, health status, and performance values of the Cache and Disaster Recovery grids. ► On the Overview tab of the detailed view, the available home and secondary inodes are available. ► On the Overview tab of the detailed view, the details such as NFS throughput, IOPs, and latency details are available, if the protocol is NFS.
Files → Quotas	When remote monitoring is enabled, you can view quota limits, capacity, and inode information for users, groups, and filesets of a file system that is mounted from a remote cluster. The user and group name resolution of the remote cluster is used in this view. It is not possible to change quota limits on a file system that is mounted from a remote cluster.

GUI option	Description
Cluster → **Remote Connections**	Provides the following options: ▶ Send a connection request to a remote cluster. ▶ Grant or reject the connection requests received from remote clusters. ▶ View the details of the remote clusters that are connected to the local cluster.
Monitoring → **Statistics** and **Monitoring** → **Dashboard**	You can create customized performance charts to monitor the remote cluster performance. For more information, see “Monitoring performance of the remote cluster”.

Monitoring performance of the remote cluster

You can monitor the performance of the remote cluster with the help of performance monitoring tools that are configured in the remote and local clusters. The performance details that are collected in the remote cluster is shared with the local cluster by using REST APIs.

After establishing the connection with the remote cluster by using the **Cluster** → **Remote Connections** page, you can access the performance details of the remote cluster from the following GUI pages:

- **Monitoring** → **Statistics**
- **Monitoring** → **Dashboard**
- **Files** → **File Systems**

To monitor performance details of the remote cluster in the Statistics page, you must create customized performance charts by completing the following steps:

1. Access the edit mode by clicking the icon that is available in the upper right corner of the performance chart and select **Edit**.
2. In the edit mode, select the remote cluster to be monitored from the Cluster field. You can select the l**ocal cluster** or **remote cluster** from this field.
3. Select **Resource** type. The data is taken from this area to create the performance analysis.
4. Select **Aggregation level.** This level determines the level at which the data is aggregated. The aggregation levels that are available for selection varies based on the resource type.
5. Select the entities that must be graphed. The table lists all entities that are available for the chosen resource type and aggregation level. When a metric is selected, you can also see the selected metrics in the same grid and use methods, such as sorting, filtering, or adjusting the time frame to select the entities that you want to select.
6. Select **Metrics**. These metrics are the type of data that must be included in the performance chart. The list of metrics that is available for selection varies based on the resource type and aggregation type.
7. Click **Apply** to create the customized chart.

After creating the customized performance chart, you can mark it as favorite charts so that they are displayed on the Dashboard page.

If a file system is mounted on the remote cluster nodes, the performance details of such remote cluster nodes are available in the Remote Nodes tab of the detailed view of file systems in the **Files** → **File Systems** page.

4.15 Monitoring thresholds

You can configure the IBM Spectrum Scale to raise events when certain thresholds are reached. Use the **Monitoring** → **Thresholds** page to define or modify thresholds for the data that is collected through the performance monitoring sensors.

You can set the following types of threshold levels for data collected through performance monitoring sensors:

- Warning level

  When the data that is monitored reaches the warning level, the system raises an event with severity Warning. When the observed value exceeds the current threshold level, the system removes the warning.

- Error level

  When the data that is monitored reaches the error level, the system raises an event with severity Error. When the observed value exceeds the current threshold level, the system removes the error state.

Certain types of thresholds are predefined in the system. The following predefined thresholds are available:

- Inode utilization at the fileset level
- Data pool capacity utilization
- Metadata pool capacity utilization
- Free memory utilization

Apart from the predefined thresholds, you can create user-defined thresholds for the data that is collected through the performance monitoring sensors.

You can use the **Monitoring** → **Thresholds** page in the GUI and the `mmhealth` command to manage predefined and user-defined thresholds.

Defining thresholds

Use the **Create Thresholds** option (see Figure 4-15) to define user-defined thresholds or to modify the predefined thresholds.

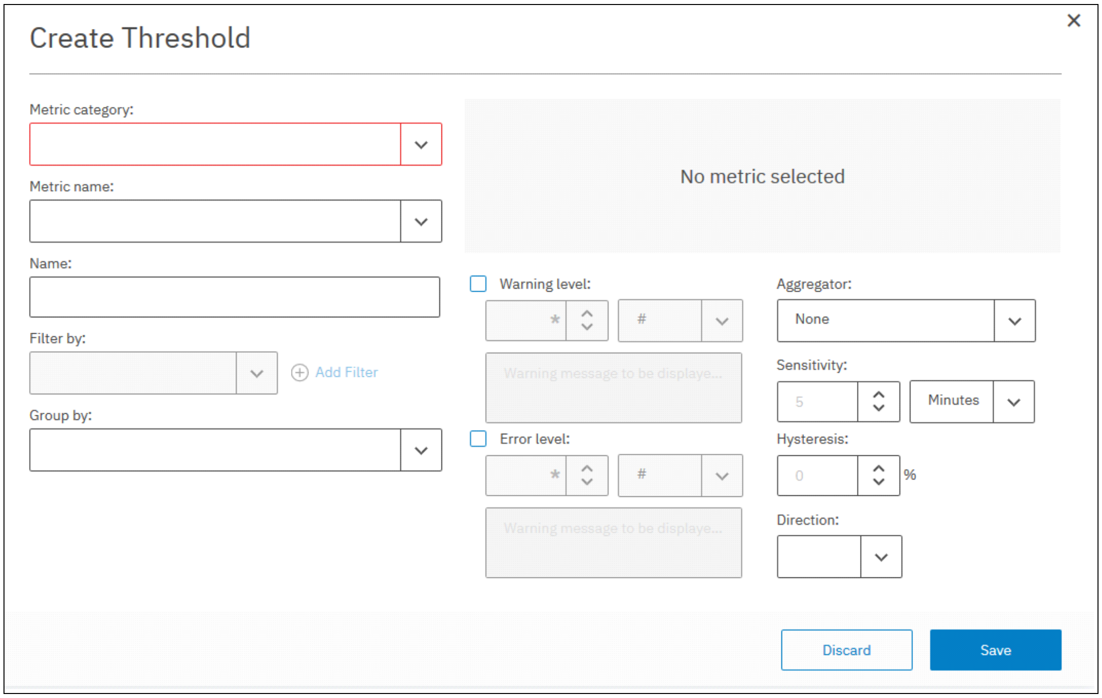

Create Threshold

Metric category:

Metric name:

Name:

Filter by:

Group by:

No metric selected

□ Warning level:

□ Error level:

Aggregator: None

Sensitivity: 5 Minutes

Hysteresis: 0 %

Direction:

×
Add Filter
Warning message to be displaye...
Warning message to be displaye...
Discard
Save

Figure 4-15 Create Threshold dialog

You can use the **Use as Template** option that is available in the Actions menu to use a defined threshold as the template to create a threshold. You can specify the following details in a threshold rule:

- Metric category

  Lists all performance monitoring sensors that are enabled in the system and thresholds that are derived by performing certain calculations on certain performance metrics. These derived thresholds are referred as *measurements*. The measurements category provides the flexibility to edit certain predefined threshold rules. The following measurements are available for selection:

  – `Fileset_inode`

    Inode capacity utilization at the fileset level. This level is calculated as shown in the following example:

    ```
    (sum(gpfs_fset_allocInodes)-sum(gpfs_fset_freeInodes))/sum(gpfs_fset_maxInod
    es)
    ```

  – `DataPool_capUtil`

    Data pool capacity utilization, which is calculated as shown in the following example:

    ```
    (sum(gpfs_pool_total_dataKB)-sum(gpfs_pool_free_dataKB))/sum(gpfs_pool_total
    _dataKB)
    ```

  - `MetaDataPool_capUtil`

    Metadata pool capacity utilization, which is calculated as shown in the following example:

    `(sum(gpfs_pool_total_metaKB)-sum(gpfs_pool_free_metaKB))/sum(gpfs_pool_total_metaKB)`

  - `FsLatency_diskWaitRd`

    File system latency for the read operations. Average disk wait time per read operation on the IBM Spectrum Scale client:

    `sum(gpfs_fs_tot_disk_wait_rd)/sum(gpfs_fs_read_ops)`

  - `FsLatency_diskWaitWr`

    File system latency for the write operations. Average disk wait time per write operation on the IBM Spectrum Scale client:

    `sum(gpfs_fs_tot_disk_wait_wr)/sum(gpfs_fs_write_ops)`

  - `SMBNodeLatency_read`

    SMB read latency at the node level:

    `avg(op_time)/avg(op_count)`

  - `SMBNodeLatency_write`

    SMB write latency at the node level:

    `avg(op_time)/avg(op_count)`

  - `NFSNodeLatency_read`

    NFS read latency at the node level:

    `sum(nfs_read_lat)/sum(nfs_read_ops)`

  - `NFSNodeLatency_write`

    NFS write latency at the node level:

    `sum(nfs_write_lat)/sum(nfs_write_ops)`

- Metric name

  The list of performance metrics that are available under the selected performance monitoring sensor or the measurement.

- Name

  User-defined name of the threshold rule.

- Filter by

  Defines the filter criteria for the threshold rule.

- Group by

  Groups the threshold values by the selected grouping criteria. If you select a value in this field, you must select an aggregator criterion in the **Aggregator** field. By default, there is no grouping, which means that the thresholds are evaluated based on the finest available key.

- Warning level

  Defines the threshold level for warning events to be raised for the selected metric. When the warning level is reached, the system raises an event with severity Warning. You can customize the warning message to specify the user action that is required to fix the issue.

- Error level

  Defines the threshold level for error events to be raised for the selected metric. When the error level is reached, the system raises an event with severity “Error”. You can customize the error message to specify the user action that is required to fix the issue.

- Aggregator

  When grouping is selected in the Group by field, an aggregator must be chosen to define the aggregation function. When the Rate aggregator is set, the grouping is automatically set to the finest available grouping.

- Sensitivity

  Defines the sample interval value. If a sensor is configured with interval period greater than 5 minutes, the sensitivity is set to the same value as sensors period. The minimum value allowed is 120 seconds. If a sensor is configured with interval period less than 120 seconds, the `--sensitivity` is set to 120 seconds.

- Hysteresis

  Defines the percentage of the observed value that must be under or over the current threshold level to switch back to the previous state. The default value is 0 percent. Hysteresis is used to avoid frequent state changes when the values are close to the threshold. The level must be set according to the volatility of the metric.

- Direction

  Defines whether the events and messages are sent when the value that is being monitored exceeds or goes below the threshold level.

You can also edit and delete a threshold rule.

Threshold configuration: A scenario

The user wants to configure a threshold rule to monitor the maximum disk capacity usage. The values against each field of the Create Threshold dialog and their respective functionality are listed in Table 4-4.

Table 4-4 Threshold rule configuration: A sample scenario

GUI fields	Value and functions
Metric Category	GPFSDiskCap Specifies that the threshold rule is going to be defined for the metrics that belong to the GPFSDiskCap sensor.
Metric name	Total Capacity The threshold rule is going to be defined to monitor the threshold levels of total capacity usage.
Name	Total capacity threshold By default, the performance monitoring metric name is used as the threshold rule name. Here, the default value is overwritten with “Total capacity threshold”.
Filter by	Cluster The values are filtered at the cluster level.

GUI fields	Value and functions
Group by:	File System Groups the selected metric by file system.
Aggregator	Maximum When maximum capacity exceeds the threshold level, the system raises the event. If the following values are selected, the nature of the threshold rule changes: ▶ Sum: When the sum of the metric values exceeds the threshold levels, the system raises the events. ▶ Average: When the average value exceeds the average, the system raises the events. ▶ Maximum: When the maximum value exceeds maximum level, the system raises the events. ▶ Minimum: When the minimum value exceeds the sum of or goes below the threshold levels, the system raises the events. ▶ Rate: When the rate exceeds the threshold value, the system raises the events. Rate is only added for the “finest” group by clause. If we wanted to get a rate for a “partial key”, this function is not supported; that is, when Rate is selected, the system automatically selects the best possible values in the grouping field.
Warning level	9 GiB The system raises an event with severity Warning when the total capacity usage exceeds 9 GiB.
Error level	10 GiB The system raises an event with severity level Error when the total capacity usage exceeds 10 GiB.
Sensitivity	24 hours The threshold value is being monitored once in a day.
Hysteresis	50% If the value is reduced below 4.5 GiB, the warning state is removed. Similarly, if the value is reduced below 5 GiB, the error state is removed.
Direction	High When the value that is being monitored exceeds the threshold limit, the system raises an event.

4.16 Monitoring command audit log

The audit log maintains a record of various actions that are performed on the system. This log helps the system administrator to audit the commands and tasks that are performed by the administrators. These logs can also be used to troubleshoot issues that are reported in the system.

You can monitor the command audit log from the **Monitoring** → **Command Audit Log** page.

You can use the **`Copy Command and Arguments`** option from the **Actions** menu to copy the command and arguments that are used in an operation.

5

Configuring and managing tasks by using the IBM Spectrum Scale GUI

This chapter lists various configuring and managing tasks that can be completed by using the IBM Spectrum Scale GUI and includes the following topics:

- 5.1, “Creating file systems” on page 66
- 5.2, “Mounting a file system through the GUI” on page 66
- 5.3, “Unmounting a file system by using the GUI” on page 67
- 5.4, “Creating filesets” on page 68
- 5.5, “Creating and managing snapshots” on page 68
- 5.6, “Deleting snapshots” on page 71
- 5.7, “Configuring quota” on page 71
- 5.8, “Information lifecycle management” on page 72
- 5.9, “Managing storage” on page 75
- 5.10, “Managing access control lists” on page 75
- 5.11, “Managing Object Storage, SMB shares, and NFS exports” on page 76
- 5.12, “Managing IBM services” on page 77
- 5.13, “Configuring role-based access for GUI users” on page 81

5.1 Creating file systems

After you create NSDs, a GPFS file system can be created. Use the **Create File System** option that is available in the **Files** → **File Systems** page to start the wizard that assists you to create a file system.

> **Note:** You cannot create file system from the GUI in an FPO-enabled environment. In an FPO environment, the Create File System option is disabled.

You can specify the following details while creating the file system by using the GUI:

- File system name.
- Storage pools that provide storage for the file system. You can create a system pool and multiple data pools for a file system. Storage pools provide storage for the file system. You can create a system pool and multiple data pools for a file system. A system pool can be used for storing metadata, data, or both. The data pool can be used only for storing data.
- Size of each inode.
- Maximum number of inodes for the root fileset.
- Data and metadata block size.
- Default replication policy for the file system by specifying the number of data and metadata copies that are required. You can later change the number of copies of data and metadata through the CLI by using the `mmchfs` command.
- NSDs that provide storage to the pools that are defined for the file system.
- Failure groups of NSDs and certain attributes that are important for the failure group definition.
- Round-robin order in which data must be written to the NSDs in a pool.
- Maximum number of IBM Spectrum Scale client nodes that can access the file system concurrently.
- Whether to enable quota for the file system. If yes, whether the user and group quota definitions must be set at the file system level or at the individual fileset level.
- Whether to enable DMAPI for the file system.
- The IBM Spectrum Scale release with which the file system features are compatible.
- Mount point and automatic mount mode of the file system.

5.2 Mounting a file system through the GUI

You can use the IBM Spectrum Scale GUI to mount individual file systems or multiple file systems on the selected nodes. Use one of the following pages in the GUI to mount a file system:

- **Files** → **File Systems**
- **Files** → **File Systems** → **View Details** → **Nodes**
- **Nodes** → **View Details** → **File Systems**

The GUI has the following options that are related to mounting the file system:

- Mount local file systems on nodes of the local IBM Spectrum Scale cluster.
- Mount remote file systems on local nodes.

- Select individual nodes, protocol nodes, or nodes by node class while selecting nodes on which the file system needs to be mounted.
- Prevent or allow file systems from mounting on individual nodes.

  To prevent file systems from mounting on a node, complete the following steps:

  a. Go to **Nodes**, and select the node on which you need to prevent or allow file system mounts.
  b. Click **Actions** → **Prevent Mounts**.
  c. Select the required option, and click **Prevent Mount** or **Allow Mount** based on the selection.
- Configure the **automatic mount** option.

  The **automatic mount** option determines whether to automatically mount file system on nodes when the GPFS daemon starts or when the file system is accessed for the first time. You can also specify whether to exclude individual nodes while enabling the **automatic mount** option.

  To enable automatic mount, complete the following steps:

  a. Go to **Files** → **File Systems**, and select the file system for which you need to enable automatic mount.
  b. Click **Actions** → **Configure Automatic Mount**.
  c. Select the required option from the list of automatic mount modes.
  d. Click **Configure**.

> **Note:** You can configure **automatic mount** option for a file system only after the file system is *unmounted* from all nodes; that is, you need to stop I/O on this file system to configure this option. However, you can include or exclude the individual nodes for automatic mount without unmounting the file system from all nodes.

5.3 Unmounting a file system by using the GUI

You can use the IBM Spectrum Scale GUI to unmount individual file systems or multiple file systems on the selected nodes. Use the following pages in the GUI to unmount a file system:

- **Files** → **File Systems**
- **Files** → **File Systems** → **View Details** → **Nodes**
- **Nodes** → **View Details** → **File Systems**

The following unmount features are supported in the GUI:

- Unmount the local file system from local nodes and remote nodes.
- Unmount the remote file system from the local nodes. When a local file system is unmounted from the remote nodes, the remote nodes can no longer be seen in the GUI. The **Files** → **File Systems** → **View Details** → **Remote Nodes** page lists the remote nodes that currently mount the selected file system. The selected file system can be a local or a remote file system, but the GUI allows you to unmount only local file systems from the remote nodes.
- Select individual nodes, protocol nodes, or nodes by node class while selecting nodes from which the file system needs to be unmounted.

- Specify whether to force unmount. Selecting the **Force unmount** option while unmounting the file system unmounts the file system, even if it is still busy performing the I/O operations. Forcing the unmount operation affects the outstanding operations and causes data integrity issues. The IBM Spectrum Scale system relies on the native **`unmount`** command to conduct the unmount operation.

  The semantics of forced unmount are platform-specific. On certain platforms, such as Linux, even when forced unmount is requested, the file system cannot be unmounted if it is still being referenced by the system kernel. To unmount a file system in such cases, identify and stop the processes that are referencing the file system. You can use system utilities like `lsof` and `fuser` for this process.

5.4 Creating filesets

Use the **Files** → **Filesets** page (see Figure 5-1) to create, manage, and monitor filesets.

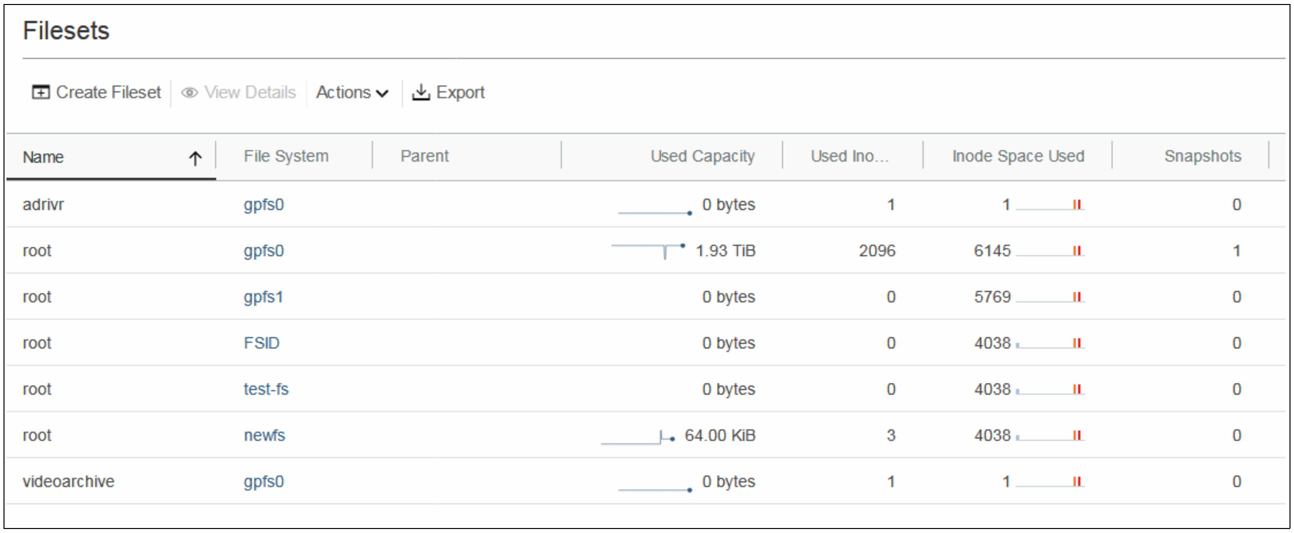

Name	File System	Parent	Used Capacity	Used Ino...	Inode Space Used	Snapshots
adrivr	gpfs0		0 bytes	1	1	0
root	gpfs0		1.93 TiB	2096	6145	1
root	gpfs1		0 bytes	0	5769	0
root	FSID		0 bytes	0	4038	0
root	test-fs		0 bytes	0	4038	0
root	newfs		64.00 KiB	3	4038	0
videoarchive	gpfs0		0 bytes	1	1	0

Figure 5-1 Filesets page

Use the **Files** → **Filesets** → **Create Fileset** option to create a fileset. You can create an independent or dependent fileset. You can specify the maximum number of inodes and the allocated number of inodes for an independent fileset. You can also specify access control lists for the fileset.

When Quota data collection is enabled, the GUI also provides information on fileset size and growth rates.

5.5 Creating and managing snapshots

Use the **Files** → **Snapshots** page to manage snapshots through the GUI. Snapshots can be used in environments where multiple recovery points are necessary. A snapshot can be taken of file system or fileset data, and then the data can be recovered from the snapshot if the production data becomes unavailable.

Note: Snapshots are read-only. You can change only the normal and active files and directories, not the snapshot.

When a snapshot of an independent fileset is taken, only nested dependent filesets are included in the snapshot.

5.5.1 Scheduling snapshot creation by using snapshot rules

You can manually create the snapshots or snapshot rules to automate the snapshot creation and retention.

To manually create a snapshot, complete the following steps:

1. Click **Create Snapshot** in the **Snapshots** page.
2. Enter the required details under the Manual tab of the Create Snapshot window.
3. Click **Create** after providing the details.

You can automate the snapshot creation and retention by creating a *snapshot rule*. That is, in a snapshot rule, you can specify a frequency at which the snapshots must be created and the number of snapshots that must be retained for a period. The system determines which snapshots are retained based on the retention policy. The retention policy helps to avoid unwanted storage of snapshots that result in the waste of storage resources.

The retention policy includes the following parameters:

- Frequency of snapshot creation.
- Number of most recent snapshots to be retained. The most recent snapshot is identified based on the frequency of snapshot creation.
- Number of days for which you must keep the latest snapshot of each day.
- Number of weeks for which you must keep the latest snapshot of each week.
- Number of months for which you must keep the latest snapshot of each month.

5.5.2 Example scenario for retention policy

Table 5-1 provides an example for the values that are specified against these parameters.

Table 5-1 Example for retention period

Snapshot deletion time	**Frequency**	**Minutes**	**# of most recent snapshots**	**Keep latest snapshot for...**			
				Hours	Days	Weeks	Months
2:30 AM	Hourly	1	2	2	6	2	3

Based on this retention rule, the snapshots that are shown in Table 5-2 are created and retained on 20 March 2016 at 06:10 AM.

Table 5-2 Time stamp of snapshots that are retained based on the retention policy

Time stamp	**Condition based on which snapshot is retained**
December 31 (Thursday, 23:01 AM)	Keep the latest snapshot for last 3 months

Time stamp	Condition based on which snapshot is retained
January 31 (Sunday, 23:01 AM)	Keep the latest snapshot for last 3 months
February 29 (Monday, 23:01 AM)	Keep the latest snapshot for last 3 months
March 5 (Saturday, 23:01 AM)	Keep the latest snapshot for last 2 weeks
March 12 (Saturday, 23:01 AM)	Keep the latest snapshot for last 2 weeks
March 14 (Monday, 23:01 AM)	Keep the latest snapshot for last 6 days
March 15 (Tuesday, 23:01 AM)	Keep the latest snapshot for last 6 days
March 16 (Wednesday, 23:01 AM)	Keep the latest snapshot for last 6 days
March 17 (Thursday, 23:01 AM)	Keep the latest snapshot for last 6 days
March 18 (Friday, 23:01 AM)	Keep the latest snapshot for last 6 days
March 19 (Saturday, 23:01 AM)	Keep the latest snapshot for last 6 days
March 20 (Sunday, 1:01 AM)	Keep two most recent snapshots.
March 20 (Sunday, 2:01 AM)	Keep two most recent snapshots

According to this rule, 13 snapshots are retained on 20 March 2016 at 06:10 AM.

To schedule snapshot creation and retention, complete the following steps:

1. Go to **Files** → **Snapshots**, and click **Create Snapshot**.
2. In the Create Snapshot window, enter the path of the file system or independent fileset for which you need to create snapshots.
3. In the **Snapshot name** field, enter the name of the snapshot.
4. Click **Snapshot Rules**.
5. Click **Create Rule** to schedule the snapshot creation and retention. The system displays the Create Snapshot Rule window.
6. In the **Name** field, enter the name of the snapshot scheduling rule.
7. In the Frequency field, select the frequency at which you need to create snapshot. You must enter some more details based on the value that is selected in the Frequency field. For example, if the selected value is **Multiple Times an Hour**, select the minutes of the hour in which you need to create snapshots.
8. In the **Retention** fields, specify the number of snapshots that must be retained in a time period.
9. In the **Prefix** field, specify a prefix to be added with the name of the snapshots that are created with this rule.
10. Click **OK** to save the changes.

Important naming information: If you do not specify a name for the snapshot, the default name is provided. The default snapshot ID is generated at creation time by using the format `@GMT-yyyy.MM.dd-HH.mm.ss`. If this option is provided and the `@GMT-date-time` format is omitted, this snapshot is not identifiable by Windows VSS, and the file restore is not possible by using that method.

Avoid white spaces, double and single quotation marks, parentheses (), the star (*), forward slash /, and backward slash \.

5.6 Deleting snapshots

To manually delete the snapshots, right-click the snapshot from the **Files** → **Snapshots** page and select **Delete**. The snapshots that are automatically created based on the snapshot creation rule are deleted automatically based on the retention period specified in the rule. When the condition for deletion is met, the GUI immediately starts to delete the snapshot candidates.

> **Note:** Snapshot capacity usage is not collected automatically because it can negatively affect the performance of the system. If you are trying to determine overall file system capacity, you must consider the capacity that is used by snapshot and manually aggregate the overall capacity usage.

5.7 Configuring quota

Use the **Files** → **Quotas** page to control the allocation of files and data blocks in a file system. You can create default, user, group, and fileset quotas by using the Quotas page.

You can also enable quota on file systems, set grace time defaults, and trigger a quota database repair action.

You can create new quotas and modify existing ones. A *quota* is the amount of disk space and the amount of metadata that is assigned as upper limits for a specified user, group of users, or fileset. Use the **Actions** menu to create or modify quotas. The management GUI allows you to manage the capacity-related quota only. The inode-related quota management is possible in the command-line interface only.

You can specify a soft limit, a hard limit, or both. When you set a soft limit quota, a warning is sent to the administrator when the file system is close to reaching its storage limit. A grace period starts when the soft quota limit is reached. Data is written until the grace period expires, or until the hard quota limit is reached. Grace time resets when the used capacity goes below the soft limit.

If you set a hard limit quota, you cannot save data after the quota is reached. If the quota is exceeded, you must delete the files or raise the quota limit to store more data. The grace period can be modified per device by using the `mmsetquota` command.

> **Note:** User or user group quotas for filesets are supported only if the Per Fileset option is enabled at the file system level.
>
> You must unmount a file system to change the quota enablement method from per file system to per fileset, or vice versa.

You can set the default user quotas at the file system level rather than defining user quotas explicitly for each user. Default quota limits can be set for users. You can specify the general quota collection scope, such as per file system or per fileset to define whether the default quota must be defined at file system level or fileset level and set the default user quota.

After this value is set, all child objects that are created under the file system or file set are configured with the default soft and hard limits. You can assign a custom quota limit to individual child objects, but the default limits remain the same unless changed at the file system or fileset level.

5.8 Information lifecycle management

The information lifecycle management (ILM) feature that is available in the IBM Spectrum Scale system facilitates automated tiered storage management. As part of the ILM, you must create a set of policies and rules that automatically determine where to physically store your data, regardless of its placement in the logical directory structure. The proper management of files ensures the efficient use and balance of premium and less expensive storage resources.

Use the **Files** → **Information Lifecycle** page (see Figure 5-2) in the IBM Spectrum Scale GUI to manage ILM rules and policies.

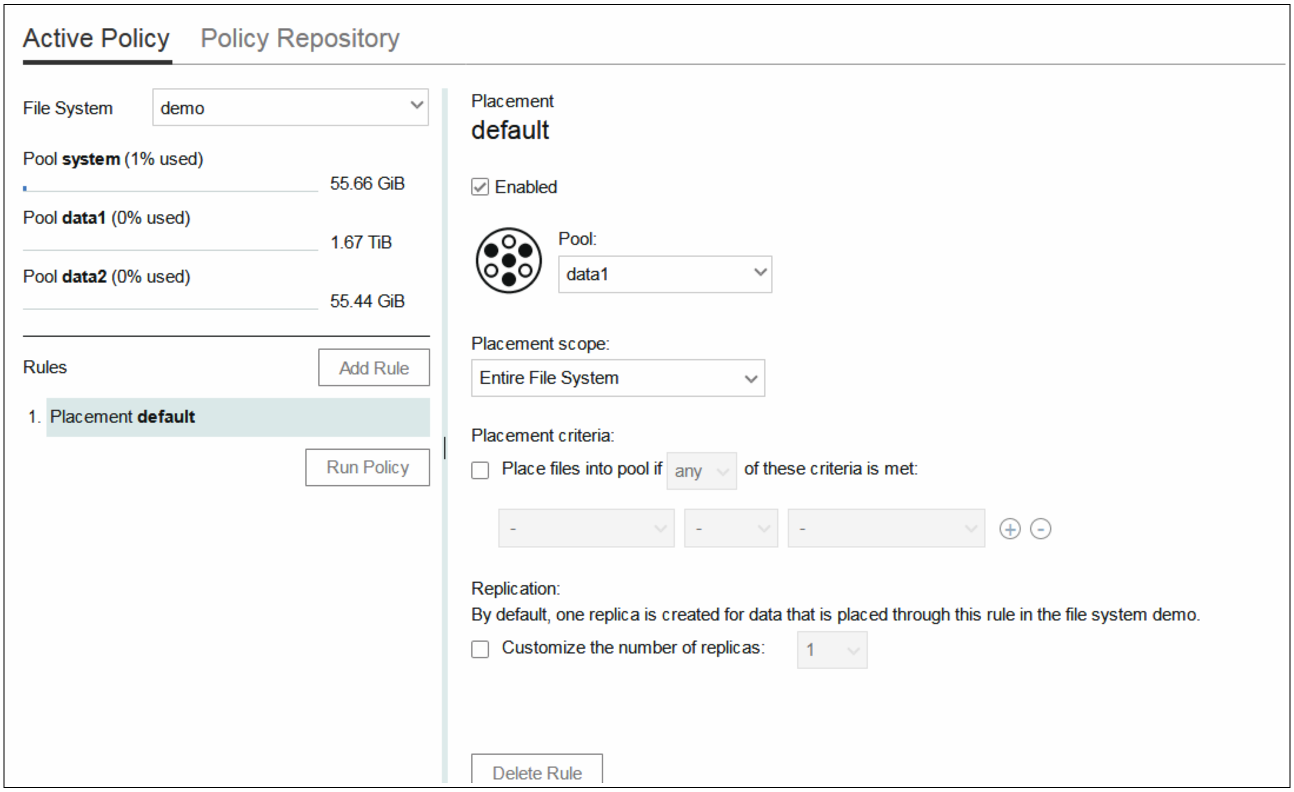

File System: demo

Pool **system** (1% used) 55.66 GiB

Pool **data1** (0% used) 1.67 TiB

Pool **data2** (0% used) 55.44 GiB

Rules

1. Placement **default**

Placement

default

☑ Enabled

Pool: data1

Placement scope: Entire File System

Placement criteria:

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

Replication:

By default, one replica is created for data that is placed through this rule in the file system demo.

□ Customize the number of replicas: 1

Figure 5-2 Information Lifecycle page

Policies and the rules are used to assign files to specific file system pools. A file system pool typically contains a set of volumes that provide a specific quality of service for a specific use, such as storing frequently accessed files on a premium or a pool with high performance and the non-frequently accessed files on a less expensive pool.

A policy is a set of rules that describes the lifecycle of user data that is based on the file's attributes. Each rule defines an operation or definition, such as placing new files into different pools or migrating files from one pool to another pool. A policy rule is an SQL-like statement that tells the file system what to do with a file in a specific file system pool if the file meets specific criteria. A rule can apply to any file within a file system or only to files within a specific fileset or group of filesets.

ILM rules include the following main functions:

- Initial file placement
- File management activities, such as migration of files from one storage pool to another pool, automatic file deletion, file compression, and file encryption

- File restoration

The active policy of a file system is available on the **Files** → **Information Lifecycle** → **Active Policy** page. Use the GUI to create more policies for a file system manually by editing the existing active policy, or by creating a policy and applying it as the active policy for the file system.

5.8.1 Creating and applying policy

Select **Policy Repository** to create a policy and define rules for it. You can also modify the created policies and apply a policy as the active policy for a file system. You must select **Active Policy** to see the active policy for a file system. You can also modify the active policy based on the requirement.

To create and apply a policy, complete the following steps:

1. Go to **Files** → **Information Lifecycle**, and select **Policy Repository**.
2. Click **Create Policy** and specify the required details.

   The policy is created. Next, you must add rules in the policy that manages the files in the system.
3. Click **Add Rule in the Policy Repository** and define rules with the required rule types. You can create multiple rules in a policy. You can drag the rules in the rules list to change the order in which the rules are applied in a policy. The **Add Rule** option supports only adding placement, migration, file compression, or deletion rules, or defining an external pool. To add encryption, exclusion, or list rules, you must modify the policy text by using the text editor.

   Optionally, you can use the text editor to edit policy text. Click **Policy Text** that is available in the upper-right corner of the GUI page to start the text editor. To work with list rules or less-frequently used policy rule syntax constructs, the policy text must be modified by using the text editor.
4. After editing the policy details, click **Apply Changes**.
5. If you want to apply a policy as the active policy for a file system, select the policy from the Policy Repository and then, select **Actions** → **Apply as Active Policy**. You can also change the active policy of the file system.

5.8.2 Editing a policy by using the text editor

To define or modify file placement, migration, file compression, deletion, or external pool rules, the GUI provides an easy to use graphical editing mode. For working with rules, such as encryption, exclusion, and list, you must manually edit the SQL policy text by using the text editor.

If only one rule is in the policy and it is not supported in the graphical editing mode, the entire policy can be displayed or modified by using the policy text editor.

5.8.3 Defining the policy run settings

You can define some of the policy run parameters that are used every time the ILM policy is run from the **Information Lifecycle** page in the GUI.

Note: The policy run settings that you set in the GUI are applicable when the policy execution is triggered by the default threshold callback or when using the Run Policy action in the GUI. These parameters are not applicable when a custom callback script is registered or when you run the policy by using the **`mmapplypolicy`** command in the CLI.

You can specify the following details that determine the policy run characteristics:

- Node that run the policies

  The ILM policy can run parallel on multiple nodes. The following types of nodes are available:

  - Nodes of a node class.
  - Default helper nodes. Nodes can be marked as helper nodes by using the **`defaultHelperNodes`** parameter of the **`mmchconfig`** command.
  - Manager nodes. These nodes are the nodes from which file system managers and token managers are selected.
  - Individual nodes.

- Local work directory

  The directory to be used for temporary storage during policy execution. This local directory, such as `/tmp`, is used on each helper node. A significant amount of temporary storage is required if the file system or directories contain many files.

- Global work directory

  A global directory to be used for temporary storage during policy execution. The specified directory must exist within a shared file system. It must also be mounted and available for writing and reading from each of the nodes. The use of a global work directory causes high performance and fault-tolerant protocols during policy execution.

- File selection algorithm. The following algorithm types are available:

  - Exact: Sorts all the candidate files by weight, then serially considers each file from the highest weight to the lowest weight by choosing feasible candidates for migration, deletion, or listing according to any applicable rule LIMITs and current storage-pool occupancy.
  - Fast: Uses a combination of statistical, heuristic, and parallel computing methods to favor higher weight candidate files, but the set of chosen candidates might be different than the exact method.
  - Best: Chooses the optimal method based on the rest of the input parameters.

- Average number of CPU cores per node.

  The number of threads and sort pipelines that each node runs during the parallel inode scan and policy evaluation.

- Number of threads per policy scan.

  The number of threads are created and dispatched within each **`mmapplypolicy`** process during the directory scan phase. The default is 24.

- Number of threads for policy execution.

  The number of threads that are created and dispatched within each **`mmapplypolicy`** process during the policy execution phase. The default value is 24.

- Maximum number of files per batch.

  Specifies how many files are passed for each invocation of the EXEC script. The default value is 100. If the number of files exceeds the value that is specified, the use of the **mmapplypolicy** command starts the external program multiple times.

5.8.4 Log files

The policy executions that are started by using the Run Policy action log the details in the `/var/log/cnlog/ilm` directory.

The policy executions can also be triggered based on a threshold that is managed by the callback handler, which is installed on the GUI node. Such policy execution details are logged in the `/var/adm/ras` directory and `/var/adm/ras/mmfs.log` file.

5.9 Managing storage

You can monitor and manage pools and NSDs of the IBM Spectrum Scale system by using the GUI options that are listed in Table 5-3.

Table 5-3 Storage management options

GUI option	Function
Storage → Pools	Provides an easy way to monitor the performance, health status, and configuration aspects of all of the pools that are available in the IBM Spectrum Scale cluster.
Storage → NSDs	Provides an easy way to monitor the performance, health status, and configuration aspects of all of the NSDs that are available in the IBM Spectrum Scale cluster.

5.10 Managing access control lists

Use the **Files → File System ACL** page to create access control lists (ACLs) for the files or directories in a file system. Access to the files and directories is managed through ACLs. It ensures that only authorized users can access directories and files. The IBM Spectrum Scale ACLs are stored in the NFSV4 ACL format.

An ACL is a list of permissions that are associated with a directory or file. It defines which users are allowed to access a particular directory or file. An access control entry in the ACL defines the permissions for a user or a group of users. An ACL usually consists of multiple entries. Each ACL has an owner that is associated with it who owns the file or directory for which the ACL is defined. Owners usually have full access to the files or directories that they own. If the directory contains files or subdirectories, the owner, owning group, and ACL cannot be modified.

You can define ACL templates to help the users to set default access control permissions for files and directories. The use of ACL template helps to save time and ensures that the correct standard and values for each ACL entry are used. You can use any of the predefined ACL templates to set the access rights to files and directories.

Only users with the DataAccess role can modify ACLs of files and directories on the **Files → File System ACL** page.

Users with the Administrator, SecurityAdmin, and DataAccess roles can edit ACL templates.

Users with the DataAccess role can modify the ACL of non-empty file system, filesets, and exports path by using the **Edit Access Control** option in the corresponding GUI pages.

Users with the Administrator and SecurityAdmin role are allowed to modify only the ACLs of file system root path, fileset link path, and export paths if they are empty.

5.11 Managing Object Storage, SMB shares, and NFS exports

The options that are available to configure, monitor, and manage the data exports through the NFS, SMB, and Object protocols are listed in Table 5-4.

Table 5-4 GUI options available for monitoring and managing protocol data exports

GUI page	Function
Protocols → NFS Exports	Create and manage NFS exports and add NFS clients. Protocols pages are displayed in the GUI only when the protocol feature is enabled on the cluster.
Services → NFS	Specify NFS server settings and start or stop NFS services.
Protocols → SMB Shares	Create and manage SMB shares. Protocols pages are displayed in the GUI only when the protocol feature is enabled on the cluster.
Services → SMB	Specify SMB server settings and start or stop SMB services.
Object → Accounts	Create and manage accounts and containers in the Object Storage. Object pages are displayed in the GUI only when the object feature is enabled on the cluster.
Object → Users	Create object users.
Object → Roles	Define roles for the object users.
Services → Object	View and change the object service status. You can define object administrator who can manage accounts in the object storage.

5.12 Managing IBM services

An IBM Spectrum Scale setup comprises various services. You can monitor, configure, and manage them through the newly introduced Services page (see Figure 5-3) in the IBM Spectrum Scale 5.0.3 management GUI. This page acts as the single place where you can view all the supported services, health status of each service. It also provides certain configuration options for some of the services.

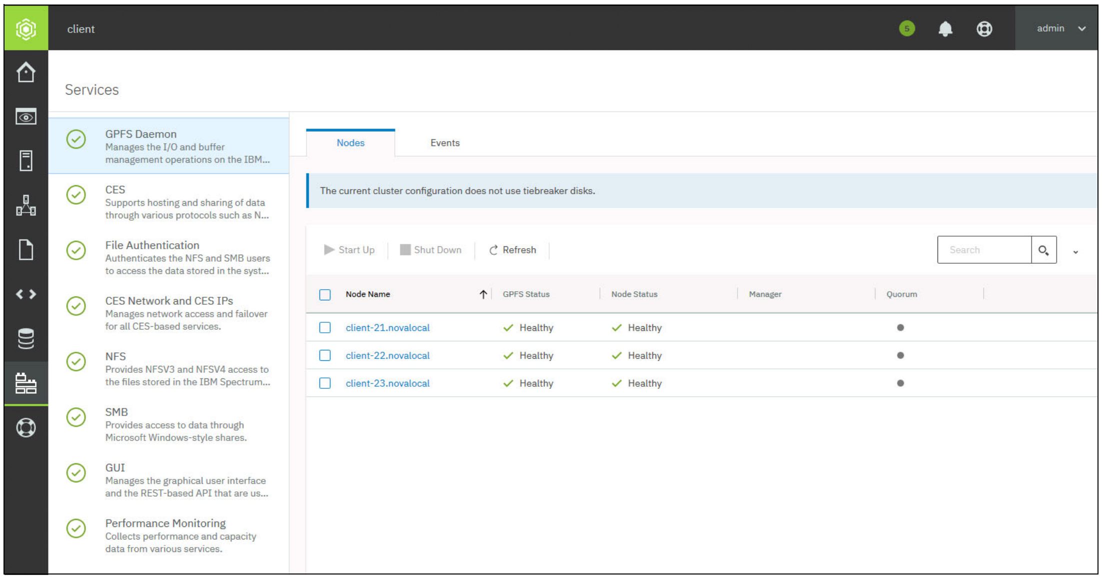

GPFS Daemon
Manages the I/O and buffer management operations on the IBM...

CES
Supports hosting and sharing of data through various protocols such as N...

File Authentication
Authenticates the NFS and SMB users to access the data stored in the syst...

CES Network and CES IPs
Manages network access and failover for all CES-based services.

NFS
Provides NFSV3 and NFSV4 access to the files stored in the IBM Spectrum...

SMB
Provides access to data through Microsoft Windows-style shares.

GUI
Manages the graphical user interface and the REST-based API that are us...

Performance Monitoring
Collects performance and capacity data from various services.

The current cluster configuration does not use tiebreaker disks.

Node Name	GPFS Status	Node Status	Manager	Quorum
client-21.novalocal	✓ Healthy	✓ Healthy		●
client-22.novalocal	✓ Healthy	✓ Healthy		●
client-23.novalocal	✓ Healthy	✓ Healthy		●

Figure 5-3 Services page to manage IBM Spectrum Scale services

The following sections describe the IBM Spectrum Scale services that can be managed through the Services page in the GUI.

5.12.1 GPFS daemon

The GPFS daemon performs all I/O operations and buffer management for GPFS. You can perform the following actions from the GPFS Daemon section:

- Start and shut down the GPFS services on the nodes.
- Monitor the status of GPFS service and the nodes on which the GPFS service is configured.
- Monitor the events that are raised against the GPFS service.

5.12.2 CES

The cluster export service (CES) provides highly available file and object services by using NFS, SMB, and Object protocols. The nodes that support these protocol services are referred as CES nodes. You can perform the following actions from the CES section:

- Stop CES service on a node when you are suspending the node.
- Start CES service on a node when you are resuming the node.
- Suspend and resume CES nodes.
- Monitor the status of the CES nodes and the protocol services hosted on the node.
- Monitor the events that are raised against the CES service.

5.12.3 CES network and CES IPs

Provides the details of the nodes that are part of the CES network. You can perform the following actions from the CES network section:

- View all the nodes that are part of the CES network.
- Add CES IP addresses by using the **Add CES IP** option that is available under the Addresses section.
- Monitor the events that are raised against the CES network service.
- Select the distribution policy to be used for assigning the CES IP address.

5.12.4 NFS

The NFS services must be configured on the system to use NFS protocol for data transfer between client and the IBM Spectrum Scale system. You can use NFSv3, NFSv4, or both to use for communication between server and client. You can perform the following actions from the NFS section:

- Start and Stop NFS service
- Monitor the health status of the NFS service that is configured on the CES nodes
- View the events that are raised against the NFS service
- Configure the lease lifetime, domain, and NFS protocol version at the NFS server level

5.12.5 SMB

You can monitor the SMB service and change global SMB configuration parameters from the GUI.

The following options are available under the respective tabs in the SMB section:

- SMB Service Status

  View the details and health status of the SMB service that is configured in the CES nodes. You can also start and stop the SMB service.
- Events

  Displays the events that are raised against the SMB service.
- Settings

  Provides options to configure the disk free quota, specify server description, set SMB server encryption node, and specify whether to restrict anonymous access.

5.12.6 Object

You must enable and start the object services to use the Object Storage facility.

You can monitor start, and stop object service on all or individual CES nodes, from the Object section in the Services page. In addition, object specific events and object administrator credential settings can be managed from this section.

5.12.7 File authentication

Use the File Authentication section to configure an authentication method or view the authentication method that is used for NFS and SMB users. You can also view the events that are reported against the authentication configuration.

The following file user authentication methods can be used to authenticate the user:

- Active Directory

  Uses Microsoft Active Directory as the authentication server. This method is used if you must authenticate SMB users to access the data through SMB shares. When you select AD as the authentication server, you must configure an ID mapping method to map the user IDs from the external domain with a set of internal user IDs.

- LDAP

  Uses an LDAP server to authenticate users. This method is the ideal method to use to authenticate the NFS protocol users to access the data through the NFS exports.

- NIS

  The NIS-based authentication is useful in NFS-only environment where NIS acts as an ID mapping server and is used for netgroups. When file access is configured with NIS, SMB access cannot be enabled.

- User-defined

  The user can select the authentication and ID mapping methods of their choice. It is the responsibility of the administrator of the client system to manage the authentication and ID mapping for file access to the IBM Spectrum Scale system.

5.12.8 Object Authentication section

The Object Authentication section shows the object user authentication configuration details and the events that are raised against the object authentication service in the cluster.

5.12.9 Hadoop connector

The IBM Spectrum Scale system provides access to the Hadoop Distributed File System Transparency (HDFS) clients. You can monitor the status of the HDFC transparency service from the Hadoop Connector section of the Services page.

5.12.10 GUI

The GUI service manages the GUI and REST APIs that are used to configure, monitor, and manage the IBM Spectrum Scale system. You can perform the following actions from the GUI section:

- Monitor the GUI node configuration in the cluster
- Configure a login message. This message appears in the login page of the GUI, which is typically used to display some important information that must be shared with the users who are attempting to log in.
- Configure session timeout.
- Create SSL certificate request.
- Install a self-signed certificate or a certificate that is issued by the certificate authority.

- View the certificate information of the GUI node.
- View the issues that are raised against the GUI service in the cluster.
- Manage GUI users, groups, and their password policy.
- Assign user roles for the GUI users.
- Configure an LDAP-based external configuration method for the GUI users.

For more information about the GUI user management, see "Configuring role-based access for GUI users" on page 81.

5.12.11 Performance monitoring

The performance monitoring tool collects metrics from various components of the IBM Spectrum Scale system and provides performance information.

The Performance Monitoring section of the Services page organizes the monitoring and configuration aspects of performance monitoring under the following tabs:

- Nodes

  Provides the nodes on which performance monitoring is enabled. You can also see the health status of these nodes and the performance monitoring sensors that are enabled on the node.
- Sensors

  Lists all the sensors that are available with the IBM Spectrum Scale system and provides the option to edit the sensor configuration. By clicking the **Edit** option, you can modify the data collection intervals and the scope of data collection. The data can be collected at all nodes, node group, or individual node level.
- **Collectors**

  Provides the health status of the performance monitoring collector that is configured in the system.
- **Events**

  Lists all events that are raised against the performance monitoring component.

5.12.12 File auditing

The file auditing logs data access to the files. Each file operation is generated as a local event on the node that serves the file operation.

The following details are available in the File Auditing section:

- Nodes: Provides file auditing service status per node.
- File Systems: Provides file auditing service status per file system and node.
- Events: Lists the events that are reported against the file auditing service.

5.12.13 Message queue

The message queue collects and stores events that are published by the producers that are running on nodes in an IBM Spectrum Scale cluster. This message queue provides a scalable infrastructure for file auditing services.

5.13 Configuring role-based access for GUI users

GUI administrators of the IBM Spectrum Scale system can monitor, configure, and manage the IBM Spectrum Scale system and are distinguished from the data users.

You can manage GUI users locally within the system and in an external authentication server, such as Microsoft Active Directory (AD) or Lightweight Directory Access Protocol Server (LDAP). By default, the IBM Spectrum Scale system uses an internal authentication repository for GUI users. Internal and external authentication methods can be configured in the system.

5.13.1 Managing GUI users locally in the IBM Spectrum Scale system

You can create users who can perform different administrative tasks on the system. Each user must be part of a user group or multiple groups that are defined on the system. When you create a user, you assign the user to one of the default user groups or to a custom user group. User groups are assigned with predefined roles that authorize the users within that group to a specific set of operations on the GUI.

Use the **Services** → **GUI** page to create users and add them to a user group.

Predefined roles are assigned to user groups to define the working scope within the GUI. If a user is assigned to more than one user group, the permissions are additive, not restrictive. The predefined role names cannot be changed.

The IBM Spectrum Scale GUI includes the following default user groups:

- Administrator

  Manages all functions on the system except those deals with managing users, user groups, and authentication.

- SecurityAdmin

  Manages all functions on the system, including managing users, user groups, and user authentication.

- SystemAdmin

  Manages clusters, nodes, alert logs, and authentication.

- StorageAdmin

  Manages disks, file systems, pools, filesets, and ILM policies.

- SnapAdmin

  Manages snapshots for file systems and filesets.

- DataAccess

  Controls access to data, such as managing access control lists.

- Monitor

  Monitors objects and system configuration, but cannot configure, modify, or manage the system or its resources.

- ProtocolAdmin

  Manages object storage and data export definitions of SMB and NFS protocols.

- UserAdmin

  Manages access for GUI users. Users who are part of this group have edit permissions only in the Access pages of the GUI.

When you log in to the system for the first time after the installation, the system lists the option to create the first GUI user.

Use the various controls that are available under the Password Policy tab of the GUI Users page to enforce strong passwords for the users. You can modify or expire password of the individual users or all the users that are created in the system. If the password is set as expired, the user is prompted to change the password at the next login.

Password policy modifications: Only users with the User Administrator role can modify the password policy of a user.

5.13.2 Assigning roles to user groups

Users who are part of Security Administrator and User Administrator user groups can create role-based user groups in which any users that are added to the group adopt the role that is assigned to that group.

Roles apply to users on the system and are based on the user group to which the user belongs. A user can be part of multiple user groups so that a single user can play multiple roles in the system.

You can assign the following roles to your user groups:

- Administrator

  Users can access all functions on the GUI except those deals with managing users and user groups.

- Security Administrator

  Users can access all functions on the GUI, including managing users and user groups.

- System Administrator

  Users can manage clusters, nodes, and alert logs.

- Storage Administrator

  Users can manage disks, file systems, pools, and filesets.

- Snapshot Administrator

  Users can manage snapshots for file systems and filesets.

- Monitor

  Users can view objects and system configuration, but cannot configure, modify, or manage the system or its resources.

- Data Access

  Users can perform the following tasks:

  - Edit owner, group, and ACL of any file or path on the **Files** → **File System ACL** → **Files and Directories** page.
  - Edit owner, group, and ACL for a non-empty directory of a file system, fileset, NFS export, or SMB share.
  - Create and delete object containers through the **Object** → **Accounts** page.

- Protocol Administrator

  Users manage object storage and data export definitions of SMB and NFS protocols.

- User Administrator

  Users manage GUI users and user groups.

> **Note:** Default groups are not created for the User Administrator user role in case the user is upgrading the IBM Spectrum Scale cluster from 4.2.0.x to a later release.

5.13.3 Managing GUI administrators in an external authentication server

By default, the IBM Spectrum Scale uses an internal authentication repository for the GUI administrators. You can configure an external authentication server either through GUI or CLI.

Configuring external authentication by using GUI

Use the **Configure External Authentication** option that is available under the External Authentication tab of **Services** → **GUI** page to configure an external LDAP-based authentication method for authenticating the GUI users. The Configure LDAP-Based External Authentication wizard assists you to configure the external authentication method (see Figure 5-4).

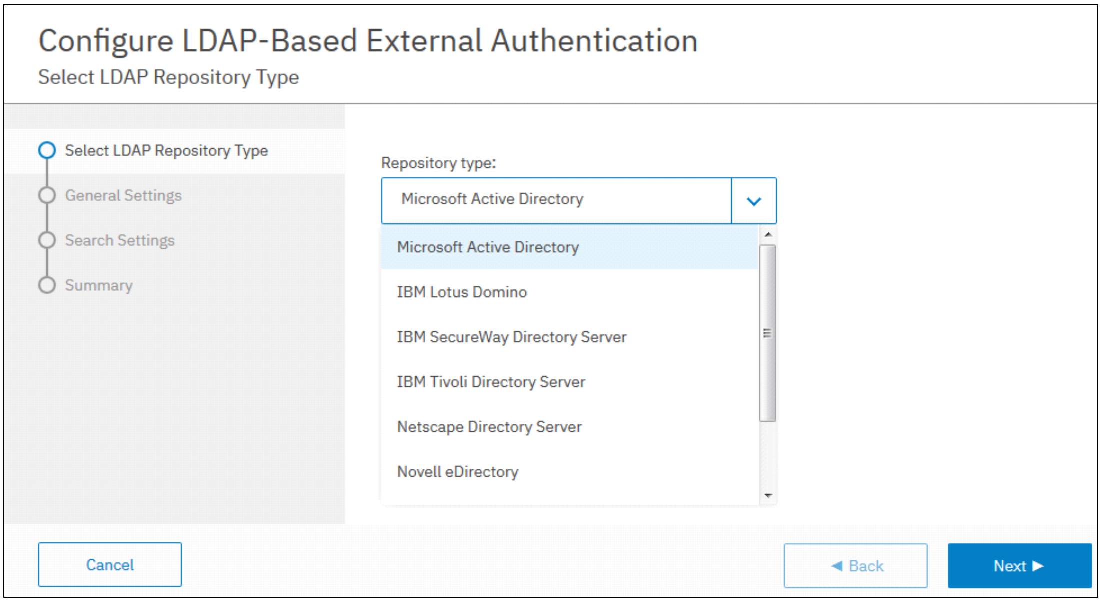

Configure LDAP-Based External Authentication

Select LDAP Repository Type

Repository type: Microsoft Active Directory

Select LDAP Repository Type
General Settings
Search Settings
Summary
Microsoft Active Directory
IBM Lotus Domino
IBM SecureWay Directory Server
IBM Tivoli Directory Server
Netscape Directory Server
Novell eDirectory
Cancel
◀ Back
Next ▶

Figure 5-4 Configure LDAP-Based External Authentication wizard

The user credentials are stored in an external repository. You can store the user credentials in the following repository types:

- Microsoft Active Directory
- IBM Lotus Domino
- IBM SecureWay Directory Server
- IBM Tivoli® Directory Server
- Netscape Directory Server
- Novell eDirectory
- Sun Java System Directory Server
- Custom (for example, OpenLdap)

Follow the wizard to complete the authentication service.

Use the **Test Connection** option that is available under the External Authentication tab to determine whether a user credential is available in the internal or external repository.

Configuring external authentication by using CLI

Complete the following steps to configure external authentication by using CLI:

1. Create your AD or LDAP configuration by issuing the **`mkldap`** command at the following location:

   ```
   /usr/lpp/mmfs/gui/cli/mkldap
   ```

   The use of this command writes the configuration automatically to `/opt/ibm/wlp/usr/servers/gpfsgui/ldap.xml`, which is then distributed across all GUI nodes. For secure AD or LDAP connection, ensure that the keystores are on the respective GUI nodes.

   The **`mkldap`** command accepts the parameters that are listed in Table 5-5.

Table 5-5 mkldap command parameters

Parameters	Description
id	Unique ID of the LDAP configuration.
--host	The IP address or host name of the LDAP server.
--baseDn	BaseDn string for the repository.
--bindDn	BindDn string for the authentication user.
--bindPassword	Password of the authentication user.
--port	Port number of the LDAP. Default is 389 or 636 over SSL.
--type	Repository type, such as “Microsoft Active Directory, ids, domino, secureway, iplanet, netscape, edirectory” or “custom”. Default value is “Microsoft Active Directory”.
--connecTimeout	Maximum time for establishing a connection with the LDAP server. Default value is 1m.
--searchTimeout	Maximum time for an LDAP server to respond before a request is canceled. Default value is 1m.
--keystore	Location with file name of the keystore file (.jks, .p12 or .pfx).
--keystorePassword	Password of the keystore.
--truststore	Location with file name of the truststore file (.jks, .p12 or .pfx).
--truststorePassword	Password of the truststore.
--userFilter	User filter for the LDAP repository.
--userIdMap	User ID map for the LDAP repository.
--groupFilter	Group filter for the LDAP repository.
--groupIdMap	Group ID map for the LDAP repository.
--groupMemberIdMap	Group member ID map for the LDAP repository.

The following examples show the use of the **mkldap** command:

- Example for standard AD:

  ```
  mkldap myad --host 9.155.106.19 --bindDn
  CN=Administrator,CN=Users,DC=mydomain,DC=local
  --baseDn CN=Users,DC=mydomain,DC=local
  ```

- Example for secure AD:

  ```
  mkldap mysecuread --host 9.155.106.19 --bindDn
  CN=Administrator,CN=Users,DC=mydomain,DC=local
  --baseDn CN=Users,DC=mydomain,DC=local --keystore /tmp/ad.jks
  ```

  If you specify multiple AD or LDAP servers, you might encounter a problem that a user with the same user name exists in multiple user repositories. This user cannot log in. To prevent this situation, you can specify LDAP filters for User Principal Names (UPN) for a selected server configuration.

- Example for a scenario where UPN filters are enabled:

  ```
  mkldap myfilteredad --host 9.155.106.19 --bindDn
  CN=Administrator,CN=Users,DC=mydomain,DC=local
  --baseDn CN=Users,DC=mydomain,DC=local --userFilter
  "(&(userPrincipalName=%v)(objectcategory=person))"
  --groupFilter "(&(cn=%v)(objectcategory=group))" --userIdMap
  "*:userPrincipalName"
  --groupIdMap "*:cn" --groupMemberIdMap "memberOf:member"
  ```

2. Map an existing AD or LDAP group to the SecurityAdmin GUI role as shown in the following example:

   ```
   /usr/lpp/mmfs/gui/cli/mkusergrp LDAPGroup --role securityadmin
   ```

   Now you can log in with your AD or LDAP user and create more group mappings through the GUI on the **Services** → **GUI** → **Users** page by using the **Create Group Mapping** option.

   If you want to remove the existing configurations, use the **rmldap** command. To see all specified LDAP configurations, issue the **lsldap** command.

Note: Configurations that are managed by using the **mkldap** and **rmldap** commands are not overwritten during the upgrade; that is, you do not need to back up the configuration data.

6

Troubleshooting options in the IBM Spectrum Scale GUI

This chapter describes the following troubleshooting options that are available in the IBM Spectrum Scale GUI:

- View an overall cluster status in the Home view.
- View health status of components, such as file systems, NSD, and nodes.
- View details of the events reported in the system.
- View options to download logs, trace files, and dumps to know more about the issues reported in the system.
- View the directed maintenance procedures to fix certain events reported in the system.

This chapter includes the following topics:

- 6.1, “Monitoring events” on page 88
- 6.2, “Monitoring tip events” on page 89
- 6.3, “Configuring event notifications” on page 89
- 6.4, “Configuring email notifications” on page 90
- 6.5, “Configuring SNMP manager” on page 91
- 6.6, “Collecting diagnostic data by using the GUI” on page 92
- 6.7, “Configuring IBM Call Home by using GUI” on page 94
- 6.8, “Directed maintenance procedures” on page 95
- 6.9, “Troubleshooting issues with capacity data displayed in the GUI” on page 96

6.1 Monitoring events

You can primarily use the **Monitoring** → **Events** page to review the entire set of events that are reported in the IBM Spectrum Scale system.

The following filter options are available in the Events page:

- Current Issues displays all unfixed errors and warnings.
- Unread Messages displays all unfixed errors and warnings and information messages that are not marked as read.
- All Events displays every event, no matter if it is fixed or marked as read.

The status icons help to quickly determine whether the event is informational, a warning, or an error. Click an event and select **Actions** → **Properties** to see the detailed information of that event. The event table displays the most recent events first.

6.1.1 Event deduplication

The events are raised against the respective component; for example, GPFS, NFS, and SMB. Some of these events might occur multiple times in the system. Such events are grouped under the Event Groups tab and the number of occurrences of the events are indicated in the Occurrences column. The Individual Events tab lists all the events irrespective of the multiple occurrences.

6.1.2 Event Filtering

The following filter options by event type are available as a drop-down list in the Events page:

- Current Issues: Displays all unfixed errors and warnings.
- Notices: Displays all transient messages of type “notice” that were not marked as read.
- Current State: Displays all events that define the current state of the entities, and excludes notices and historic events.
- All Events: Displays all messages, even historic messages and messages that are marked as read. This filter is not available in the Event Groups view because of performance implications.

A graphical view of events that are reported against each component is available. Clicking the graph displays only the relevant events in the grid view; that is, clicking a portion in the graphical view applies the corresponding filter on the search action and fetches only the relevant data in the Events table.

The status icons help to quickly determine whether the event is informational, a warning, or an error. Click an event and select **Actions** → **Properties** to see the detailed information of that event. The event table displays the most recent events first.

6.1.3 Marking events as read

You can mark events of type *Notice* as read to change the status of the event in the Events view. The status icons become gray when an error or warning is fixed or if it is marked as read.

6.1.4 Running fix procedure

Some issues can be resolved by running a fix procedure by clicking **Run Fix Procedure**. For more information, see 6.8, “Directed maintenance procedures” on page 95.

6.2 Monitoring tip events

You can monitor events of type “Tips” from the **Monitoring** → **Tips** page of the GUI. The tip events give recommendations to the user to avoid certain issues that might occur in the future. The system detects the entities with the tip event type as healthy. A tip disappears from the GUI when the problem behind the tip event is resolved.

Select **Actions** → **Properties** to view the details of the tip. After you review the tip, decide whether it requires attention or can be ignored. Select **Actions** → **Hide** to ignore the events that are not important. Select **Show** to mark the tips that require attention.

6.3 Configuring event notifications

The system can use Simple Network Management Protocol (SNMP) traps and emails to notify you when significant events are detected. Any combination of these notification methods can be used simultaneously. Use the **Monitoring** → **Event Notifications** page (see Figure 6-1) in the GUI to configure event notifications.

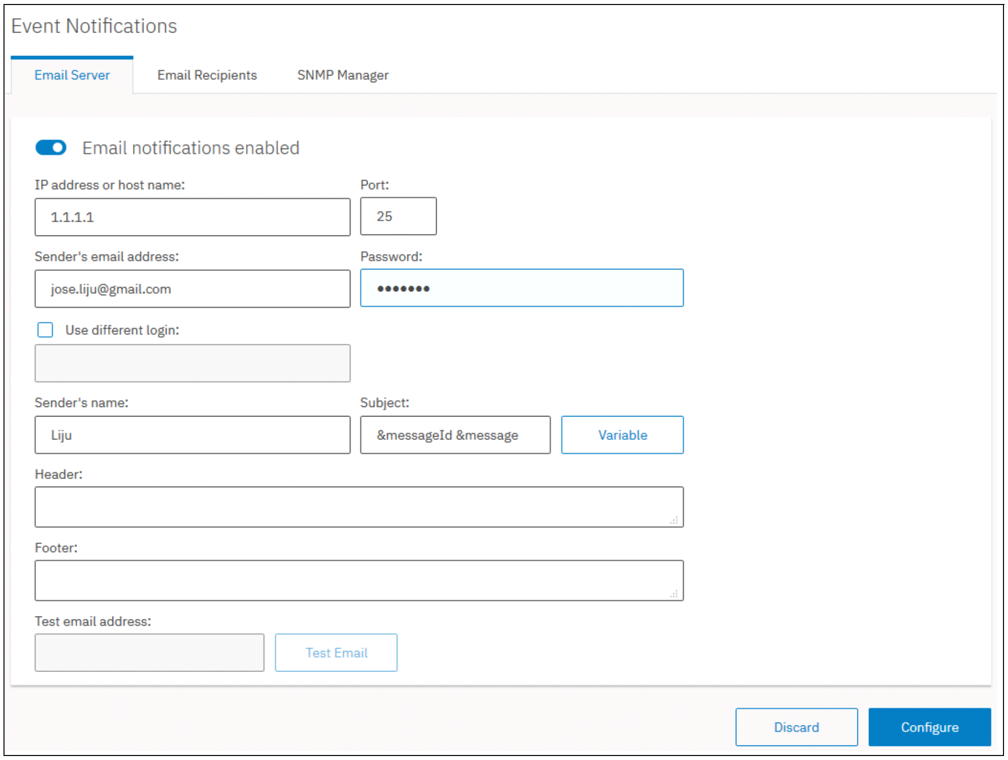

Event Notifications

Email Server | Email Recipients | SNMP Manager

● Email notifications enabled

IP address or host name: 1.1.1.1

Port: 25

Sender's email address: jose.liju@gmail.com

Password: •••••••

☐ Use different login:

Sender's name: Liju

Subject: &messageId &message

Header:

Footer:

Test email address:

Figure 6-1 Event Notifications page

Notifications are normally sent immediately after an event is raised.

6.4 Configuring email notifications

The email feature transmits operational and error-related data in the form of an event notification email. To configure an email server, complete the following steps:

1. Select **Monitoring** → **Event Notifications** → **Email Server**.
2. Select **Edit** and then, select **Enable email notifications**.
3. Enter the required details and, when you are ready, click **OK**.

Email notifications can be customized by setting a custom header and footer for the emails. You also can customize the subject by selecting and combining from the following variables:

- `&message`
- `&messageId`
- `&severity`
- `&dateAndTime`
- `&cluster`
- `&component`

Emails that contain the quota reports and other events that are reported in the following functional areas are sent to the recipients:

- AFM and AFM DR
- Authentication
- CES network
- Transparent Cloud Tiering
- NSD
- File system
- GPFS
- GUI
- Hadoop connector
- iSCSI
- Keystone
- Network
- NFS
- Object
- Performance monitoring
- SMB
- Object authentication
- Node
- CES

In the email notification method, you can also define whether a recipient must receive a report of events that are reported in the system. These reports are sent once daily. Based on the seriousness of the issue, a severity level is associated with each reported event in the system.

To create email recipients, select **Email Recipients** from the **Event Notifications** page, and then, click **Create Recipient**.

Note: You can change the email notification configuration or disable the email service at any time.

6.5 Configuring SNMP manager

SNMP is the standard protocol for managing networks and exchanging messages. The system can send SNMP messages that notify personnel about an event. You can use an SNMP manager to view the SNMP messages that the system sends.

With an SNMP manager, such as IBM Systems Director, you can view and act on the messages that the SNMP agent sends. The SNMP manager can send SNMP notifications, which are also known as *traps*, when an event occurs in the system.

Select **Monitoring** → **Event Notifications** → **SNMP Manager** to configure SNMP managers for event notifications. You can specify up to a maximum of six SNMP managers.

In the SNMP mode of event notification, one SNMP notification (trap) with object identifiers (OID) .1.3.6.1.4.1.2.6.212.10.0.1 is sent by the GUI for each event. The SNMP objects that are included in the event notifications are listed in Table 6-1.

Table 6-1 SNMP objects included in event notifications

OID	Description	Example
.1.3.6.1.4.1.2.6.212.10.1.1	Cluster ID	317908494245422510
.1.3.6.1.4.1.2.6.212.10.1.2	Entity type	NODE, FILESYSTEM
.1.3.6.1.4.1.2.6.212.10.1.3	Entity name	gss-11, fs01
.1.3.6.1.4.1.2.6.212.10.1.4	Component	NFS, FILESYSTEM, NSD
.1.3.6.1.4.1.2.6.212.10.1.5	Severity	INFO, TIP, WARNING, ERROR
.1.3.6.1.4.1.2.6.212.10.1.6	Date and time	17.02.2016 13:27:42.516
.1.3.6.1.4.1.2.6.212.10.1.7	Event name	nfs_active
.1.3.6.1.4.1.2.6.212.10.1.8	Message	NFS service is now active.
.1.3.6.1.4.1.2.6.212.10.1.9	Reporting node	The node where the problem is reported.

6.5.1 Understanding the SNMP OID ranges

The SNMP OID ranges are listed in Table 6-2.

Table 6-2 SNMP OID ranges

OID range	Description
.1.3.6.1.4.1.2.6.212	IBM Spectrum Scale
.1.3.6.1.4.1.2.6.212.10	IBM Spectrum Scale GUI
.1.3.6.1.4.1.2.6.212.10.0.1	IBM Spectrum Scale GUI event notification (trap)
.1.3.6.1.4.1.2.6.212.10.1.x	IBM Spectrum Scale GUI event notification parameters (objects)

The traps for the core IBM Spectrum Scale and the trap objects are not included in the SNMP notifications that are configured through the IBM Spectrum Scale management GUI.

6.5.2 Example for SNMP traps

Example 6-1 shows the SNMP event notification that is sent when performance monitoring sensor is shut down on a node.

Example 6-1 SNMP event notification: Performance monitoring sensor shutdown

```
SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::enterprises.2.6.212.10.0.1
SNMPv2-SMI::enterprises.2.6.212.10.1.1 = STRING: "317908494245422510"
SNMPv2-SMI::enterprises.2.6.212.10.1.2 = STRING: "NODE"
SNMPv2-SMI::enterprises.2.6.212.10.1.3 = STRING: "gss-11"
SNMPv2-SMI::enterprises.2.6.212.10.1.4 = STRING: "PERFMON"
SNMPv2-SMI::enterprises.2.6.212.10.1.5 = STRING: "ERROR"
SNMPv2-SMI::enterprises.2.6.212.10.1.6 = STRING: "18.02.2016 12:46:44.839"
SNMPv2-SMI::enterprises.2.6.212.10.1.7 = STRING: "pmsensors_down"
SNMPv2-SMI::enterprises.2.6.212.10.1.8 = STRING: "pmsensors service should be
started and is stopped"
SNMPv2-SMI::enterprises.2.6.212.10.1.9 = STRING: "gss-11"
```

Example 6-2 shows the SNMP event notification that is sent for an SNMP test message.

Example 6-2 SNMP event notification: SNMP test message

```
SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::enterprises.2.6.212.10.0.1
SNMPv2-SMI::enterprises.2.6.212.10.1.1 = STRING: "317908494245422510"
SNMPv2-SMI::enterprises.2.6.212.10.1.2 = STRING: "CLUSTER"
SNMPv2-SMI::enterprises.2.6.212.10.1.3 = STRING: "UNKNOWN"
SNMPv2-SMI::enterprises.2.6.212.10.1.4 = STRING: "GUI"
SNMPv2-SMI::enterprises.2.6.212.10.1.5 = STRING: "INFO"
SNMPv2-SMI::enterprises.2.6.212.10.1.6 = STRING: "18.02.2016 12:47:10.851"
SNMPv2-SMI::enterprises.2.6.212.10.1.7 = STRING: "snmp_test"
SNMPv2-SMI::enterprises.2.6.212.10.1.8 = STRING: "This is a SNMP test message."
SNMPv2-SMI::enterprises.2.6.212.10.1.9 = STRING: "gss-11"
```

6.5.3 SNMP MIBs

The SNMP Management Information Base (MIB) is a collection of definitions that define the properties of the managed objects.

The IBM Spectrum Scale GUI MIB OID range starts with 1.3.6.1.4.1.2.6.212.10. The OID range 1.3.6.1.4.1.2.6.212.10.0.1 denotes IBM Spectrum Scale GUI event notification (trap), and .1.3.6.1.4.1.2.6.212.10.1.x denotes IBM Spectrum Scale GUI event notification parameters (objects). While configuring SNMP, use the MIB file that is available at the following location of each GUI node:

```
/usr/lpp/mmfs/gui/IBM-SPECTRUM-SCALE-GUI-MIB.txt
```

6.6 Collecting diagnostic data by using the GUI

IBM Support might ask you to collect logs, trace files, and dump files from the system to help resolve a problem. You can perform this task by using the management GUI or the `gpfs.snap` command. Use the **Support** → **Diagnostic Data** page in the IBM Spectrum Scale GUI to collect details of the issues that are reported in the system.

The entire set of diagnostic data that is available in the system helps analyzing all types of IBM Spectrum Scale issues. Depending on the data selection criteria, these files can be large (gigabytes) and might take an hour to download.

The diagnostic data is collected from each node in a cluster. In a cluster with hundreds of nodes, downloading the diagnostic data might take a long time and the downloaded file can be large.

It is always better to reduce the size of the log file as you might need to send it to IBM Support to help fix the issues. You can reduce the size of the diagnostic data file by reducing the scope. The following options are available to reduce the scope of the diagnostic data:

- Include only affected functional areas
- Include only affected nodes
- Reduce the number of days for which the diagnostic data needs to be collected

The following three modes are available in the GUI to select the functional areas of the diagnostic data:

- Standard diagnostics

  The data that is collected in the standard diagnostics consists of the configuration, status, log files, dumps, and traces in the following functional areas:

  - Core IBM Spectrum Scale
  - Network
  - GUI
  - NFS
  - SMB
  - Object
  - Authentication
  - Cluster export services (CES)
  - Crash dumps

  You can download the diagnostic data for these functional areas at the following levels:

  - All nodes
  - Specific nodes
  - All nodes within one or more node classes

- Deadlock diagnostics

  The data that is collected in this category consists of the minimum amount of data that is needed to investigate a deadlock problem.

- Performance diagnostics

  The data that is collected in this category consists of the system performance details that are collected from performance monitoring tools. You can use this option only if it is requested by the IBM Support.

> **Note:** Instead of collecting the diagnostic data again, you can also use the diagnostic data that was previously collected. You can analyze the relevance of the historic data based on the date on which the issue is reported in the system. Ensure to delete the diagnostic data that is no longer needed to save disk space.

6.6.1 GUI logs

Run the **`gpfs.snap -N GUI_MGMT_SERVERS`** command or use the **Support** → **Diagnostic Data** page in the GUI to access the GUI logs to analyze GUI problems. These logs are available at the following location:

```
/var/log/cnlog/mgtsrv/
```

The following main log files are available:

- `mgtsrv-system-log-x` (most important)

  Logs everything that runs in background processes such as refresh tasks.
- `mgtsrv-trace-log-x`

  Logs everything that is directly triggered by the GUI user, such as starting an action, clicking a button, and running a GUI CLI command.
- `wlp-messages.log`

  Covers the underlying WebSphere Liberty. The log is mostly relevant during the start-up phase.
- `gpfsgui_trc.log`

  Logs problems related to incoming requests from the browser. Check this log if the GUI displays the following error message:

  ```
  Server was unable to process request.
  ```

6.7 Configuring IBM Call Home by using GUI

The Call Home feature provides a communication channel that automatically notifies the IBM service personnel about the issues reported in the system. You can also manually upload diagnostic data files and associate them with a PMR by using the GUI.

You can use the Call Home page in the GUI to perform the following tasks:

- Enable the Call Home feature on the cluster.
- Select one or more Call Home nodes that share the data with the IBM Support.
- Specify the contact information to be used by the IBM Support if any issues exist.
- Specify the proxy information that is needed to create a communication channel between the Call Home nodes and IBM support.
- Test connection with the IBM server.

Collecting data and sharing it with IBM Support

Call Home shares support information and your contact information with IBM on a scheduled basis. The IBM Support monitors the details that are shared by Call Home and any takes necessary action on any issues or potential issues.

Enabling Call Home reduces the response time for the IBM Support to address the issues. The GUI does not support to change the data gathering and sharing schedules.

You can also manually upload the diagnostic data that is collected through the **Support** → **Diagnostic Data** page in the GUI to share the diagnostic data to resolve a Problem Management Record (PMR).

To upload data manually, complete the following steps:

1. Go to **Support** → **Diagnostic Data**.
2. Collect diagnostic data that is based on the requirement. You can also use the previously collected data for the upload.
3. Select the relevant data set from the **Previously Collected Diagnostic Data** section and then, right-click and select **Upload to PMR**.
4. Select the **PMR** to which the data must be uploaded and then, click **Upload**.

6.8 Directed maintenance procedures

The directed maintenance procedures (DMPs) assist you to repair a problem when you select **Run Fix Procedure** on a selected event from the **Monitoring** → **Events** page. DMPs are present for only a few events that are reported in the system.

The available DMPs and the corresponding events are listed in Table 6-3.

Table 6-3 Directed maintenance procedures available for events

DMP	Event ID
Start NSD	`disk_down`
Start GPFS daemon	`gpfs_down`
Increase fileset space	`inode_error_high` and `inode_warn_high`
Start performance monitoring collector service	`pmcollector_down`
Start performance monitoring sensor service	`pmsensors_down`
Activate AFM performance monitoring sensors	`afm_sensors_inactive`
Activate NFS performance monitoring	`nfs_sensors_inactive`
Activate SMB performance monitoring	`smb_sensors_inactive`
Configure NFS sensor	`nfs_sensors_not_configured`
Configure SMB sensor	`smb_sensors_not_configured`
Mount file systems	`unmounted_fs_check`
Start GUI service on remote node	`gui_down`
Repair a failed GUI refresh task	`gui_refresh_task_failed`

6.9 Troubleshooting issues with capacity data displayed in the GUI

Because of the effect that capacity data collection can have on the system, different capacity values are collected on a different schedule and are provided by different system components. The following issues can arise from the multitude of schedules and subsystems that provide capacity data:

- Capacity in the file system view and the total amount of the capacity for pools and volumes view do not match.

  The capacity data in the file system view is collected every 10 minutes by performance monitoring collector, but the capacity data for pools and Network Shared Disks (NSD) are not updated. By default, NSD data is collected only once per day by the performance monitoring collector and this data is cached.

  Clicking the refresh icon gathers the last two records from performance monitoring tool and displays the last record values if they are not null. If the last record includes null values, the system displays the previous record. If the values of both records are null, the system displays “N/A” and the option for displaying a time chart is unavailable. The last update date is the record date that is fetched from performance monitoring tool if the values are not null.

- Capacity in the file system view and the total amount of used capacity for all filesets in that file system do not match.

  Differences exist in the collection schedule and the collection mechanism that contributes to the fact that the fileset capacities do not add up to the file system used capacity.

- Scheduling differences

  Capacity information that is shown for filesets in the GUI is collected once per hour by performance monitoring collector and displayed in the Filesets page. When you click the refresh icon, the information from the last record from performance monitoring is shown.

  If the last two records have null values, a “Not collected” warning for used capacity is shown. The file system capacity information on the file systems view is collected every 10 minutes by the performance monitoring collector. When you click the refresh icon, the information about the last record from performance monitoring is shown.

- Data collection differences

  Quota values show the sum of the size of all files and are reported asynchronously. The quota reporting does not consider metadata, snapshots, or capacity that cannot be allocated within a subblock. Therefore, the sum of the fileset quota values can be lower than the data shown in the file system view.

  You can use the **`mmlsfileset`** CLI command with the **`-d`** and **`-i`** options to view capacity information. The GUI does not provide a means to display these values because of the performance effect that results from data collection.

- The sum of all fileset inode values on the view quota window does not match the number of inodes that are displayed on the file system properties window.

  The quota value accounts for user-created inodes only, while the properties for the file system also display inodes that are used internally. Refresh the quota data to update these values.

- No capacity data is shown on a new system or for a newly created file system.

  Capacity data can appear with a delay of up to 1 day. The capacity data for file systems, NSDs, and pools is collected once per day because operation is resource-intensive. Line charts do not show a line if only a single data point exists. You can use the hover function to see the first data point in the chart.

- The management GUI displays negative fileset capacity or a highly used capacity, such as millions of Petabytes or 4,000,000,000 used inodes.

  This problem can be seen in the quota and filesets views. This problem is caused when the quota accounting is out of sync. To fix this error, issue the **`mmcheckquota`** command. This command recounts inode and capacity usage in a file system by user, user group, and fileset, and writes the collected data into the database. It also checks quota limits for users, user groups, and filesets in a file system. Running this command can affect the performance of I/O operations.

- No capacity data is displayed on the performance monitoring charts.

  Verify whether the sensors as described in 4.4.2, “Capacity data collected through the performance monitoring tool” on page 38 are enabled correctly. Ensure that all prerequisites are met; for example, quota is enabled for the GPFSFIlesetQuota sensor, or file systems are mounted on the nodes where the GPFSPools sensor runs.

For more information about how to enable the performance monitoring sensor for capacity data collection, see *Manual installation of IBM Spectrum Scale GUI in IBM Spectrum Scale: Concepts, Planning, and Installation Guide.*

A

Current limitations for the IBM Spectrum Scale GUI

IBM Spectrum Scale GUI features the following current limitations:

- The GUI runs on all Linux variants that are supported by IBM Spectrum Scale. However, the GUI does not run on AIX and Windows nodes.
- Up to 1000 nodes are supported.
- The GUI supports a subset of the CLI functionality. More capabilities are to be added in future releases of the product.
- The Object management pages do not support configurations with Keystone V2 API.
- One GUI manages a single cluster. A GUI instance can monitor a subset of performance and system health of remote clusters after creating a connection between the local and remote cluster GUI.
- In an IBM Spectrum Scale and ESS mixed support environment, the ESS GUI must manage the whole cluster to display the ESS-specific pages in the GUI.
- The GUI supports IBM Spectrum Scale release 4.2.0.0 or later.

  Issue the **`mmlsconfig`** command to see the value that is set for the `minReleaseLevel` attribute. Use the **`mmchconfig release=LATEST`** command and restart the GUI to make the management GUI fully operational at the new code level.

  Because changing the minimum release level affects the cluster behavior, see the **`mmchconfig`** command man page and other related topics before this configuration change is made.
- All IBM Spectrum Scale packages on one node must be of the same release. For example, do not mix the 5.0.3 GUI rpm with a 5.02 base rpm. However, GUI PTFs and efixes often can be applied without installing the corresponding PTF or efix of the base package. This design is helpful if you want to resolve a GUI issue without changing anything on the base layer.

Related publications

The publications that are listed in this section are considered particularly suitable for a more detailed discussion of the topics that are covered in this paper.

IBM Redbooks

The following IBM Redbooks publications provide more information about the topic in this document. Note that some publications that are referenced in this list might be available in softcopy only:

- *IBM Spectrum Scale Security*, REDP-5426
- *Implementing IBM Spectrum Scale*, REDP-5254

You can search for, view, download, or order these documents and other Redbooks, Redpapers, Web Docs, draft, and other materials at the following website:

ibm.com/redbooks

Other publications

The following publications are also relevant as further information sources:

- *IBM Spectrum Scale Version 4 Release 2.3 Administration Guide*, SA23-1455
- *IBM Spectrum Scale Version 4 Release 2.3 Concepts, Planning, and Installation Guide,* GA76-0441
- *IBM Spectrum Scale Version 4 Release 2.3 Problem Determination Guide,* GA76-0443

Online resources

The following websites are also relevant as further information sources:

- IBM Spectrum Scale GUI quick reference:

  https://www.ibm.com/support/knowledgecenter/STXKQY_5.0.3/com.ibm.spectrum.scale.v5r03.doc/bl1ins_quickrefforgui.htm

- Examples for GUI issues and their resolutions:

  https://www.ibm.com/support/knowledgecenter/STXKQY_5.0.3/com.ibm.spectrum.scale.v5r03.doc/bl1pdg_GUI_issues.htm

- IBM Spectrum Scale GUI videos:

  https://www.youtube.com/playlist?list=PLS7mekU2kxDrWbtK5AiVGPyF94xTv02xN

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

REDP-5458-01

ISBN 0738456454

Printed in U.S.A.